

# NEUROTECHNOLOGY FOR BIOMIMETIC ROBOTS



EDITED BY JOSEPH AYERS,  
JOEL L. DAVIS, AND ALAN RUDOLPH

# Neurotechnology for Biomimetic Robots



**This page intentionally left blank**

# Neurotechnology for Biomimetic Robots

edited by  
Joseph Ayers, Joel L. Davis, and Alan Rudolph

A Bradford Book  
The MIT Press  
Cambridge, Massachusetts  
London, England

© 2002 Massachusetts Institute of Technology

All rights reserved. No part of this book may be reproduced in any form by any electronic or mechanical means (including photocopying, recording, or information storage and retrieval) without permission in writing from the publisher.

This book was set in Palatino on 3B2 by Asco Typesetters, Hong Kong, and was printed and bound in the United States of America.

Library of Congress Cataloging-in-Publication Data

Neurotechnology for biomimetic robots / edited by Joseph Ayers, Joel L. Davis, and Alan Rudolph.

p. cm.

"A Bradford book."

Includes bibliographical references and index.

ISBN 0-262-01193-X (hc. : alk. paper)

1. Robotics. 2. Neural networks (Computer science) 3. Biosensors. I. Ayers, Joseph. II. Davis, Joel L., 1942- III. Rudolph, Alan.

TJ211 .N478 2002

629.8'92—dc21

2002016501

# Contents

	Preface	ix
	Introduction	xiii
<b>I</b>	<b>Neurotechnology</b>	<b>1</b>
	Overview: Advances in the Development of Biomimetic Sensors, Actuators, and Controllers	3
	<i>Alan I. Selverston</i>	
	<b><i>Biomimetic Sensors</i></b>	<b>11</b>
1	Biomimetic Flow and Contact/Bending MEMs Sensors	13
	<i>Nicol E. McGruer, George G. Adams, Thieu Q. Truong, Terrence G. Barnes, Xiaoqing Lu, and Juan Carlos Aceros</i>	
2	Fly-Inspired VLSI Vision Sensors	31
	<i>Reid R. Harrison</i>	
3	Visual Servo System Based on a Biologically Inspired Scanning Sensor	57
	<i>Stéphane Viollet and Nicolas Franceschini</i>	
4	Visual Sensors Using Eye Movements	73
	<i>Oliver Landolt</i>	
5	Using Biology to Guide Development of an Artificial Olfactory System	97
	<i>Joel White, Shalini Mall, and John S. Kauer</i>	
	<b><i>Biomimetic Actuators</i></b>	<b>115</b>
6	SMA Actuators Applied to Biomimetic Underwater Robots	117
	<i>Jan Witting, Koray Safak, and George G. Adams</i>	

7	Engineering a Muscle: An Approach to Artificial Muscle Based on Field-Activated Electroactive Polymers <i>Roy Kornbluh, Robert J. Full, Kenneth Meijer, Ron Pelrine, and Subramanian V. Shastri</i>	137
8	Polymer-Based MEMS Actuators for Biomimetics <i>Robert Horning and Burgess Johnson</i>	173
9	A Piezoelectrically Actuated Mesoscale Robot Quadruped <i>Michael Goldfarb, Michael Gogola, Gregory Fischer, Nikola Celanovic, Ivan Celanovic, and Ephraim Garcia</i>	187
	<b><i>Biomimetic Control Architectures</i></b>	205
10	Optimization of Swimming Locomotion by Genetic Algorithm <i>David Barrett</i>	207
11	Biomechanical Primitives and Heterarchical Control of Limb Motion in Tetrapods <i>Simon F. Giszter</i>	223
12	A Conservative Biomimetic Control Architecture for Autonomous Underwater Robots <i>Joseph Ayers</i>	241
<b>II</b>	<b>Vehicles</b>	261
	Overview: Vehicles in Neurobiology <i>Polly K. Pook</i>	263
	<b><i>Undulatory and Swimming Robots</i></b>	269
13	Snake Robots for Search and Rescue <i>Gavin S. P. Miller</i>	271
14	A Lamprey-Based Undulatory Vehicle <i>Cricket Wilbur, William Vorus, Yusong Cao, and Scott N. Currie</i>	285
15	Biomimetics in Action: Design and Performance of an Autonomous Robotic Fish <i>Jamie M. Anderson and Peter A. Kerrebrock</i>	297
16	Development of a Two-Joint Dolphin Robot <i>Motomu Nakashima and Kyosuke Ono</i>	309

17	Pectoral Fin Controllers <i>Naomi Kato</i>	325
	<b><i>Ambulatory Robots</i></b>	349
18	The Role of Reflexes versus Central Pattern Generators in Dynamical Legged Locomotion <i>Eric Klavins, Haldun Komsuoglu, Robert J. Full, and Daniel E. Koditschek</i>	351
19	A Decentralized, Biologically Based Network for Autonomous Control of (Hexapod) Walking <i>Holk Cruse, J. Dean, V. Dürr, Th. Kindermann, J. Schmitz, and M. Schumm</i>	383
20	Toward the Development of Agile and Mission-Capable Legged Robots <i>Roger D. Quinn, Gabriel M. Nelson, and Roy E. Ritzmann</i>	401
21	A Biologically Inspired Approach Toward Robust Real- World Locomotion in Legged Robots <i>Frank Kirchner, Dirk Spenneberg, and Ralf Linnemann</i>	419
	<b><i>Flying Robots</i></b>	449
22	The Organization of Insect Locomotion Systems and Computer-Based Flight Control in the Tobacco Hawkmoth <i>Manduca Sexta</i> <i>Gernot Wendler</i>	451
23	A Micromechanical Flying Insect Thorax <i>Ronald S. Fearing, Srinath Avadhanula, Domenico Campolo, Metin Sitti, Joseph Yan, and Robert Wood</i>	469
24	The Entomopter <i>Robert C. Michelson</i>	481
<b>III</b>	<b>Autonomous Behavior</b>	511
	Overview <i>Wendell G. Sykes</i>	513
25	Visual Homing without Image Matching <i>Ralf Möller</i>	517

26	Phonotaxis in Crickets and Robots <i>Barbara Webb and Reid R. Harrison</i>	533
27	Flow and Chemo-Sense for Robot and Lobster Guidance in Tracking Chemical Sources in Turbulence <i>Frank W. Grasso</i>	553
28	Sensory Fusion in Free-Flight Search Behavior of Fruit Flies <i>Michael Dickinson, Lance Tammero, and Michael Tarstino</i>	573
29	Visual Flight Control and Navigation in Honeybees: Applications to Robotics <i>Mandyam V. Srinivasan</i>	593
	Contributors	611
	Index	615

# Preface

What is *neurotechnology*? We submit that neurotechnology is the process of engineering devices that confer the performance advantages of animal systems on a new class of biomimetic machines. Thus neurotechnology is a concept that applies both to processes as well as to devices. As such, it requires a rather complete understanding of the biological systems both at the biomechanical and the physiological levels.

What is new about these robots? First, when compared to traditional robots, they are relatively small and agile. Second, in most cases, they are relatively cheap, relying on small processors and novel actuators. Third and most important, they have capabilities for dealing with novel real-world environments such as the woods, the sea bottom, and the air. In other words, they exhibit a behavioral set based on that of an animal model. Animals have evolved to occupy almost any environment in which we might want to operate a robot, save outer space. Thus the behavioral set of model animals provides a blueprint of how a biomimetic robot might operate in any environment.

Neurotechnology is a nascent field and most of the enabling technologies are also emerging. Several research areas can be identified:

## MYOMORPHIC ACTUATORS

Robots are traditionally built using geared motors (Kirchner, Spennberg, and Linnemann, chapter 21 of this volume), which obligate the designer to convert rotary motion to linear motion and make stiffness regulation problematical. All animals use muscle to cause motion, and the development of linear actuators that mimic muscle is key to the development of truly biomimetic robots.

In May 2000, the Neurotechnology for Biomimetic Robots conference reported the employment of several types of artificial muscle to actuate robots. For example, Witting, Safak, and Adams (chapter 6 of this volume) have successfully used the shape memory alloy (SMA) nitinol



to construct artificial muscle for underwater robots. However, SMA actuators use the bathing water to rapidly cool the muscle (increasing its dynamic range) and are problematical for operation in air. An alternative is flexible pneumatic chambers that undergo a length change when expanded and also allow proportional control (Quinn, Nelson, and Ritzmann, chapter 20 of this volume).

Two investigators reported myomorphic actuators that operated by electromechanical effects. Kornbluh and colleagues (chapter 7 of this volume) have developed a muscle based on field-activated electroactive polymers that feature high strains. These actuators generate force and movement by expansion rather than contraction. A contrasting approach is provided by the electrostatic polymer MEMS actuators developed by Horning and Johnson (chapter 8 of this volume), which, like biological muscle, act in tension. Thus antagonistic pairs of muscle can be utilized to control both joint movement and impedance. Several investigators describe the use of piezoelectric actuators that mediate high frequency movements and are ideal for miniature robots.

## NEUROMORPHIC SENSORS

Typical robots employ sensors that utilize different coding principles than animal sensors. Animal sensors universally rely on a labeled line code. Each sensory neuron encodes a particular sensory modality (light, pressure, chemical ligands, etc.) and senses that modality over a particular receptive field or orientation relative to the body. In addition, labeled lines may encode the intensity of the modality in a frequency-modulated code. Animal sensors often employ range fractionation where different labeled lines code for different ranges of a scalar variable such as inclination relative to vertical (Cohen, 1955).

McGruer and colleagues (chapter 1 of this volume) describe a novel application of MEMS cantilever microswitches to construct antennal and flow sensors. Harrison (chapter 2), as well as Viollet and Franceschini (chapter 3), describe optical flow sensors that are used in orientational control. Landolt (chapter 4) addresses the issue of moving motion sensors. White, Mall, and Kauer (chapter 5) describe the development of an electronic nose that discriminates odors using the same mechanisms as the vertebrate olfactory system.

A major challenge of neuromorphic sensor development is specification of labeled line codes and the establishment of behavioral releasers that can reactively trigger elements of the behavioral library of the vehicle. Releasers for orientational reflexes with regard to gravity, flow, collisions, and so on, lie at the very core of achievement of reactive autonomy.

## **BIOMIMETIC CONTROLLERS**

Mammalian systems rely on neocortical mechanisms to mediate skilled behavior acts. Many of the invertebrate models used for biomimetic robots, however, lack these higher-order computational systems. In these simpler systems, control is more highly distributed and relies more on genotypic specification rather than on learning. This volume contains descriptions of three types of control architectures. Barrett (chapter 10) describes the development of evolved controllers where a genetic algorithm is utilized to evolve control parameters relative to a fitness function. Giszter (chapter 11) addresses the organization of movement primitives around a set of modular organizers that may be combined to generate more complex controllers. Ayers (chapter 12) describes a controller based on the command neuron, coordinating neuron, central pattern generator model.

## **BEHAVIORAL LIBRARIES AND AUTONOMOUS BEHAVIOR**

Most animals can be characterized in terms of an ethogram or set of species-specific behavioral acts that define the response of the individual or group in its interactions with the environment. Defining behavior is problematical (Barlow, 1968) and is exacerbated when comparing animal and robotic behavioral acts.

At present, most robots are teleoperated. In other words, telemetry is utilized to present sensor status to a human operator, and the operator in turn alters the operational status of the vehicle. Animals, in contrast, exhibit reactive autonomous behavior resulting from continual sensor input. Animals select behavioral acts from a behavioral library based on specific releasers. Even teleoperated vehicles require some modicum of autonomous behavior. For example, if the vehicle strays from its effective communication path, it may need a set of behavioral acts to mediate reacquisition of lost communications links.

One can divide the behavior of robots into world-map-based robots that navigate relative to an internal representation of the surrounds, and purely reactive robots that navigate on the basis of sensor feedback from the environment. The papers presented in this conference addressed primarily the reactive component of autonomous behavior. Möller (chapter 25) modeled the role of visual landmarks in ant navigation using a robotic implementation. Grasso (chapter 27) modeled the olfactory navigation of a lobster mimic in odorant plumes. Dickinson, Tammero, and Tarstino (chapter 28) and Srinivasan (chapter 29) examined the role of optical flow in fly and bee navigation and saccades. Webb and Harrison (chapter 26) modeled phonotaxis in a robotic cricket and also examined the role of fusion of phonotactic inputs and optical flow.

Together, these contributions provide a new synthesis of neurotechnology that defines a pathway to truly autonomous robots. We submit that neurotechnology will evolve into the same relationship with neurobiology and biomechanics that biotechnology has to molecular biology.

Joseph Ayers and Joel L. Davis

## ACKNOWLEDGMENTS

The editors (J. A., J. D., and A. R.) would like to acknowledge the strong support of the Office of Naval Research (ONR) and the Defense Advanced Research Projects Agency (DARPA) for their generous support of the workshop that brought the authors together. In addition, the editors would like to thank these two agencies for their early recognition, leadership and support of the areas of biomorphic science and technology that are described in this volume.

## REFERENCES

- Barlow, G. W. (1968). Ethological units of behavior. In: D. Ingle (ed.), *The Central Nervous System and Fish Behavior*. Chicago: University of Chicago Press, pp. 217–232.
- Cohen, M. J. (1955). The function of receptors in the statocyst of the lobster (*Homarus americanus*). *J. Physiol.* 130: 9–33.

# Introduction

Living systems offer almost endless inspiration for the design and fabrication of robots. Evolution has overcome many of the challenges currently facing robotics, providing solutions that serve as local optima from which to study and extract useful principles. Robotics is inherently a multidisciplinary activity, requiring the integration of such fields as engineering, control theory, sensory fusion, and behavior. The effort to bring biomimetics to robotics further demands an understanding and extraction of useful componentry of living systems such as neural circuitry, biomechanics, unique soft and hard materials, and integration of multifunctional systems. In living systems, these come together to allow for highly functional behavior across an amazing scale of complexity and size. As our societal expectations grow for robots interacting with and substituting for human form and function, the hope increases that biomimetics will impact a greater functionality to robotic systems.

One of the key challenges of current robotics is mobility and maneuverability over complex and unplanned terrain. Most systems are currently limited to wheels and tracks that have very limited ability to climb or adapt to rough terrain. This provides a rich landscape for biomimetic roboticists: biological systems demonstrate incredible mobility and maneuverability over land, water, and in air. Many integrated structural and functional components enable this capability. As one example, appendages (legs, wings, arms) used by living systems are used to move and maneuver in environments, utilizing the combination of force dynamics (hydrodynamic, aerodynamic, gravitational) and neuronal control. This requires accurate sensing and perception of the body in the environment with subsequent interplay with neuronal and mechanical biosystems. These principles provide useful inspirations for design of mobile robotic systems, many of which are presented in this compilation. Significant advances in understanding the dynamics of the use of appendages from biomechanical and neuronal systems has led to a new age of highly mobile and maneuverable robots. This capability is critical for applications that, for example, reduce potential hazards

to humans, such as ordnance detection and removal, utility repair and maintenance, and surveillance over rough terrain.

Target recognition in highly complex backgrounds is another area of considerable activity in biomimetic robotics. Evolutionary pressures in nature have led to highly developed systems for finding food, identifying reproductive targets, and avoiding predation in both individuals and social groups. These principles have provided a rich ground for study and implementation in robotic systems. The development of highly specific and sensitive sensors that can be integrated into robotic platforms offers the potential for significant increase in functionality in this context. This can be achieved through organic sensing, such as in the better placement of a foot through tactile or visual sensors or detection of an obstacle to avoid or navigate. In addition, the integration of sensors in robotic payloads offers the modularity needed to address different users' needs. For example, highly sensitive detectors of explosive and other chemical compounds can now be integrated onto a mobile platform for detection of ordnance and other hazardous chemicals.

Such advances in robotics will certainly lead to wider deployment in a variety of applications. For the Department of Defense, the additional mission capabilities that robotics offer are continuing to be explored. These include applications that remove people from harm's way and those that complement military operations. Many of the technological advances envisioned for robotics (mobility, sensing, autonomy) are useful attributes for defense applications as well as the other significant uses for robotics. It is clear that the application of biological inspiration to the robotics field will enable many of these technological advances and result in the wider use of robotics.

Alan Rudolph

# *Biomimetic Sensors*

**This page intentionally left blank**

# I

# Neurotechnology



**This page intentionally left blank**

# Advances in the Development of Biomimetic Sensors, Actuators, and Controllers

Alan I. Selverston

Advances in bioengineering have led to a renewed interest in biologically inspired robots. The main goal of this work is to develop autonomous mobile robots that can operate in real-world environments, sometimes harsh, and carry out tasks that would be impossible or dangerous for humans. This précis covers some of the important work presented at the Neurotechnology for Biomimetic Robots conference held in Nahant, Massachusetts in May 2000. This research suggests that although reasonable progress is being made, the challenge of duplicating fundamental biological phenomena such as locomotion, sensation, and decision making with inanimate materials is considerable. Even more daunting is the conversion of the physiological mechanisms underlying these phenomena to electrical or mechanical engineering design methods. The problem is not only one of mismatched media, but also that the basic biological mechanisms underlying the phenomena we wish to duplicate in hardware are not yet fully explicable. A mechanical or electronic device that has the appearance of duplicating animal behavior, but in fact is working with a structural or computational plan that is completely nonbiological, is simply a toy. Devices that utilize biological principles such as olfactory coding, central pattern generation, or reflex arcs begin to capture the true meaning of the biomimetic approach. The Nahant meeting presented a mix of these two methods for the development of biomimetic robots.

As pointed out in the introduction to the meeting, "animals have evolved to occupy every environment where one might wish to operate a robot." In the course of this evolution, most animals have developed rather common mechanisms for generating behavior despite differences in particular environmental conditions. These mechanisms can be simplified into three fundamental components: sensing the physical environment and converting it to a neural code; central processing of this coded information; and motor code for generating movements. All of these interactions can be related in the following way: The components of biological behavior are all integrated and controlled by the neural

circuitry of the central nervous system (CNS). Buried within the CNS are ensembles of neurons called central pattern generators (CPGs) that give rise to rhythmic behaviors such as walking, flying, and chewing. The CPGs give meaning to the sensory inputs, which enter the CNS as massive parallel streams of firing patterns, by directing the muscles in a way that produces meaningful behavior. An important aspect of the neurons that constitute these circuits is that they are not static entities. The biophysical properties of each neuron, as well as most of the synapses, are in a continual state of flux as a result of their own activity and the actions of neuromodulatory substances that surround them. Nevertheless, identifiable neurons of invertebrates have demonstrated clearly that under precisely similar conditions, a particular neuron will behave the same way throughout the life of the animal—that is, it has the homeostatic mechanisms built in that allow its physiological properties to remain the same over a long period of time. Biological neurons process both digital and analog information, relying heavily on the plasticity of synapses to adapt the circuits to meet particular needs and also to store information (learn). Note that the behavior itself gives rise to feedback to the motor neurons and CNS as well as activating mechanoreceptors in the joints and muscles.

To apply this scheme to robots that are other than just toy walking machines, we should consider what we presently know about the biology of goal-directed behavior. If a robot is to be more than a toy, it should not only be autonomous (and we seem to be a long way from untethered robots) but also be able to accomplish something useful in the face of unforeseen obstacles (i.e., be able to make survival-based decisions).

Neurobiologists over the past thirty years or so have made enormous progress in describing the neurobiological mechanisms at both the systems and the cellular level. The field of neuroethology especially has advanced our understanding of how behavior per se is produced by the nervous system and is arguably the most influential discipline in terms of biomimetic robotics. But the extent to which biology can inform roboticists depends on paying attention to a biological schema such as I have illustrated and to the level of function one hopes to achieve in representing biological mechanisms in hardware. The argument is often made that trying to accomplish this goal may not be the best approach to take because there are not yet good analogues of biological media available. A hybrid mechanism—that is, one that is in part biologically based and in part based on engineering principles—may at present be the best route to achieve success in fabricating biologically inspired devices. But if this argument is pushed to its extreme, one finds that the inability to adequately mimic biological components constrains the performance of the robot to the performance criteria of a mechanical device. Put another way, not even DARPA (Defense Advanced Re-

search Projects Agency) is going to confuse these robotic machines with the living thing and much more development is necessary before such devices are mission ready.

Biomimetic robotics is a pointed example for the need to do much more basic research in neurobiology and engineering. In neurobiology, we actually know very little about the circuits that function in the brain in terms of neuron-to-neuron interactions. In engineering, there are still no electronic circuits capable of generating complex behaviors and only preliminary ideas for mimicking primary receptors capable of transduction with the range, sensitivity, and adaptability of biological sensors. One of the most serious deficits is developing actuators that have the complex properties of muscle. Although many attempts at manufacturing artificial muscle are underway, the control parameters available thus far are much different than those used by the nervous system in controlling biological muscle, and their mechanical properties are not even close in terms of elasticity, strength, and speed. Nevertheless, the idea of deriving "inspiration" from biology is popular, and little progress will be made in the field of biomimetic robotics unless we proceed with what is available.

## NEURAL CODING MECHANISMS AND ROBOTICS

In the area of neural engineering, the mechanisms for moving information around are based on hard physics. The coding process used by the nervous system is both analog and digital and relies heavily on biochemistry and ionic conductances. The timing of the digital events (spikes) and the channels through which these events are carried out play a fundamental role in how sensory information is sent to the nervous system; how the CNS deals with it in terms of recognition, memory formation, and action; and how rhythmic and episodic (unique) motor patterns are produced. There is no possibility of engineering these features into biomimetic robots at the moment. Robotic signaling is accomplished by electronic and digital circuitry. Sensors send electrical currents, proportional to the intensity of the stimulus, someplace into the robot's computer where it can be acted upon, usually in the form of simple feedback loops. Actuators receive patterned signals that make them shorten at the appropriate times. Controllers are generally small computers programmed to generate the motor patterns for different behaviors. The challenge of developing silicon brains that can interact meaningfully with the environment is still a long way off. Nevertheless, there are a host of robotic devices that can walk, swim, and even fly. And in addition, there are devices that can sense important environmental variables or obstacles, and small computers that can be used to exert some primitive control over these devices.

## Sensors

Biomimetic sensory arrays that can inform an autonomous robot about its external and internal environment would, in principle, respond to the same modalities found in animals. For specialized tasks such as locating toxins, weapons, or for general surveillance, sensors could be incorporated that are not found in biological systems. For long-term autonomous behavior, however, most of the sensors would have to function in a way similar to that of their biological counterparts.

The primary sensory receptor neurons of animals are highly specialized to convert the intensity and temporal properties of particular stimulus energies into neural codes that can be fed into the central nervous system as massive parallel streams of information. The receptor cells can be exquisitely sensitive—for example, photoreceptors can respond to a single quanta of light, chemoreceptors to single molecules, and mechanoreceptors in the human ear to oscillations smaller than the diameter of a hydrogen atom. Although the specificity of receptor cells is never absolute, they are generally classified as chemoreceptors, mechanoreceptors, photoreceptors, thermoreceptors, or electroreceptors. The biological receptor cells show an amazing variety of shapes and are often aided functionally by the presence of accessory structures such as pacinian corpuscles or cochlear hair cells. The biological transduction process converts stimulus energy to neural activity by setting up generator potentials that can be quite complex but in general have a threshold, a dynamic range that is proportional to the log of the stimulus intensity and a region where the cell is saturated. These localized potentials are due to changes in ionic conductances across the receptor cell membrane. This potential generally depolarizes the cell to the point where action potentials are elicited and transmitted to the CNS. So while the stimulus intensity is initially represented by the amplitude of the generator or receptor potential, it is recoded into the form of frequency-modulated action potentials. This method is the same for all sensory signaling in animals. How the brain interprets which modality a particular stream of action potentials represents depends upon which area of the brain receives the information. For every sensory modality there are reserved pathways (labeled lines) that have specific termination sites in the brain.

Two final points that distinguish biological sensory processing. Many sense organs do not simply supply the CNS with information in a centripetal direction, but receive feedback from downstream sensory receptors as well as centrifugal signals from the CNS. Efferent control of stretch receptors, for example, determines their sensitivity and demonstrates that the CNS actively participates in the flow of sensory information it receives. In addition, some abstraction of the sensory information occurs peripherally as a result of sensory filtering. The best

example of this is center-surround responses, which occur at the level of the retina as well as centrally.

At the Nahant meeting, sensors were proposed for biomimetic devices for almost all modalities. Microelectromechanical systems (MEMS) touch sensors were fabricated for sensing touch and water flow in underwater robotic vehicles (McGruer et al., chapter 1 of this volume). A bending sensor was fabricated, for example, that can be built into an artificial antenna for use in a walking robotic lobster. A low-power visual motion sensor based on very large-scale integration (VLSI) technology, and using fly vision for the basis of the design, was built by Harrison and his colleagues (chapter 2) and used to steer a wheeled robotic vehicle based on optical flow. For both of these devices the output was a time-varying voltage signal sent directly to the controllers of the two robotic platforms.

Viollet and Franceschini (chapter 3) have elaborated this idea beautifully by building a flying model aircraft whose position could be stabilized in the yaw plane by a small optical sensor that scans the visual field with variable angular speeds. The output voltage becomes a function of the position of some contrast feature in the field, so by making it part of a visuomotor control loop, it can stabilize the model plane with respect to some feature in the environment. Landolt (chapter 4) also presented the idea of biological eye movements to inspire a scanning visual sensor based on the oculomotor system, as yet untested. Here it was quite clear that much more in the way of biological principles must be found before such a visual sensing device can be used to provide visual input to a mobile robot.

Perhaps the most successful biologically inspired device was an artificial nose described by White, Mall, and Kauer (chapter 5). Based on the principles gleaned from a long-term study of olfaction in salamanders, they constructed a device that had an array of broadly tuned sensors and a pattern-recognition algorithm for processing the sensor output. Instead of a biological transduction mechanism, however, the artificial nose produced changes in the fluorescence properties of special polymers when exposed to certain compounds such as explosives—nonbiological but useful.

## **Actuators**

Although we have no idea what the voltages emanating from biomimetic sensory devices “represent” to the robot, they can nevertheless be part of actuator feedback loops that generate movement. From a biological point of view, we have the inverse situation from that described for sensory processing. Patterns of impulses generated by the CNS and modified by sensory inputs are carried in the motor nerves by similar labeled-line mechanisms to motor units in the periphery. A motor unit

is a bundle of muscle fibers and the single axon that innervates the bundle. The movements themselves produce sensory feedback (from the muscle stretch receptors) reflexively to the motor neurons and, via other ascending pathways, provide information about the movements to the brain. In addition, biological muscle, like sensory receptors, come in many physiological types in terms of contraction speed and resistance to fatigue. Muscles also demonstrate properties like elasticity and compliance that are hard to duplicate with hardware. But perhaps their most complex feature is the fine and coarse movements muscles can make under precise neural control. Unlike mechanical actuators, the control of contraction velocity and tension depends on the frequency of motor neuron firing and the number of motor units activated in each muscle respectively.

Unlike sensory systems, which can be mimicked in some cases quite well with hardware, nothing on the horizon comes close to duplicating the properties of biological muscle. Nevertheless, a large number of devices can produce movements and these devices can be turned on and off by software programs that mimic the motor centers of the brain and spinal cord. Research aimed toward developing more biological-like artificial muscle is a growing field and it is likely that more realistic actuators will be available in the not-too-distant future.

Several possible new approaches for use in mobile robots were put forward at the meeting. For example, Witting, Safak, and Adams (chapter 6 of this volume) described an actuator for the leg movements of an underwater lobster based on the thermal properties of nitinol, a Ni-Ti (nickel-titanium) shape memory alloy that shortens when a current that heats the wire is passed through it. The wire generates heat that must be dissipated, so it is well-suited for marine applications and does a creditable job in moving the legs of a mechanical lobster.

Another interesting approach is to use electroactive polymers as described by Kornbluh and colleagues (chapter 7). They discussed a dielectric elastomer that has high strain, high energy densities, and fast response times. Artificial muscles made of this material were successfully used to produce motion in a hexapod robot inspired by cockroach walking.

A similar polymer-based actuator was described by Horning and Johnson (chapter 8), which has many properties similar to biological muscle. Although the testing of artificial muscle made of this polymer looks promising, it has not been examined thoroughly enough to know if it would work in biomimetic robots under real-world conditions. The least biological, although possibly the most useful device presented, was a piezoelectrically actuated robot described by Goldfarb and his colleagues (chapter 9). Here, locomotion could be generated by vibrating the device at a frequency that causes the robot to move quite effec-

tively, but its lack of control parameters would suggest only limited usefulness.

## Controllers

Controllers of autonomous robots must not only supply patterned information to properly sequence the operation of actuators, but they also must be able to respond to sensory inputs in a way that can modify the motor patterns appropriately. Further, the controllers must function as a “brain” storing “memories” that can be referred to during normal operational activities. One should note that the biological brain devotes a large amount of its circuitry to generating smooth elegant movements with mechanically efficient trajectories. In addition to smoothness, the motor patterns must be able to adjust the motion of the robot constantly in response to environmental perturbations. The traditional view of robots—stiff, rigid machines that show jerky and rather stereotyped movements—is largely because of the lack of circuitry in the controller systems that might provide the information to mimic the graceful behavior of animals.

Control architectures based on known invertebrate circuits (none are known for vertebrates) are starting to replace the computer-generated patterns common to most locomotory robots. Such electronic brains are likely to revolutionize the control of walking machines, particularly if the muscle question can be solved. The Neurotechnology for Biomimetic Robots conference presented several new biologically inspired ideas for controllers. One approach to adjusting control parameters was suggested by Barrett (chapter 10). He used genetic algorithms to optimize body controllers for a robotic fish (RoboTuna). The method was described in detail, although its usefulness in fine tuning the performance of swimming movements has not been tested. Another route toward producing biological-like control mechanisms was described by Giszter (chapter 11). Here, force-field primitives were considered as the building blocks of movements and a framework for their use in muscle control was presented. The work at this stage is highly theoretical, however, and how it plays out practically, when tested experimentally, remains to be determined.

Certainly the most biological of the controller systems presented was the one proposed by the conference organizer, Joseph Ayers (chapter 12). Based on an experimentally derived central pattern generator (CPG) circuit for lobster locomotion, a computer program was designed to produce many forms of ambulation for a submersible robotic lobster. Here, the challenge of using a biologically inspired CPG to control nitinol actuators while at same time being able to respond to perturbations in the environment was met with a surprising degree of fidelity.



## CONCLUSIONS

It is clear that great strides are being made in biomimetic autonomous robots, but there is no question that the field is still in its infancy. From the point of view of biologists, building such robots represents one form of modeling, proving how well we understand a system. Much more informative than a computer animation, actually constructing a walking animal pushes our understanding of the underlying principles to the limit. The conference was successful in pointing out where development efforts would be best applied in order to realize a workable biomimetic autonomous robot capable of carrying out a successful mission.

# *Biomimetic Sensors*

**This page intentionally left blank**

# 1

## **Biomimetic Flow and Contact/Bending MEMS Sensors**

Nicol E. McGruer, George G. Adams, Thieu Q. Truong, Terrence G. Barnes, Xiaoqing Lu, and Juan Carlos Aceros

Contact/bending sensors and flow sensors have been designed, analyzed, fabricated, and tested using microelectromechanical systems (MEMS) technology. These sensors are to be used on ambulatory underwater robotic vehicles, such as lobsters and scorpions. This chapter will focus on the design and testing of contact/bending and flow sensors for the lobster robot. The contact/bending sensors are to be mounted on the lobster robot's flexible antennae in order to provide it with a sense of touch. The MEMS flow sensors will allow the lobster to sense water currents and thus to adapt its behavior accordingly. The family of MEMS sensors described here act as simple on/off switching devices and thus have the advantages of lost cost, low power consumption and simplicity in integration with the electronics.

One of the greatest threats that military troops face during coastal operations is the presence of sea mines. These mines can be located either in the surf zones along the coast or in the open water. The mines in the surf zones pose a serious threat to landing troops, while mines in the open water threaten ships. Unlike the situation with land mines, there are few devices in existence that can detect sea mines. Thus the ability to detect and destroy sea mines would be of great benefit to both military and civilian populations. To accomplish this objective and reduce the risk to human life, an inexpensive, ambulatory robot is being designed to detect mines on the ocean floor. Clearly, a traditional robot would have great difficulty in navigating along the sandy and rocky bottom of the ocean floor.

Because lobsters have evolved over hundreds of thousands of years to survive in this aquatic environment, the robot's design and behavior are patterned after those of a lobster. A real lobster is able to search for and recognize prey; the lobster robot will search for and classify objects as mines. In order for the robot to be biomimetic and survive in its underwater environment, it should have artificial senses so that it can explore with its antennae, feel the ground that it is walking on, sense if it is tilted at a large angle or flipped over, or sense extreme currents in the ocean and make appropriate adjustments.

Real lobsters navigate using their sight as well as their antennae. Experiments demonstrate that blindfolded lobsters [1] can navigate around obstacles using their antennae alone—when the antenna contacts an obstruction it bends; it is this bending that the lobster senses. The biomimetic lobster robot will mimic this behavior. Real lobsters detect flow through the bending of small hairs located in pits in its claws. The flow sensors on the lobster robot will mimic this behavior.

This chapter will focus on the contact/bending and flow sensors used in the lobster robot. A review of such existing sensors will be presented. The concept of microelectromechanical systems (MEMS) sensors will be introduced and used as the basis of the design of contact/bending and flow sensors. The analysis, fabrication, and testing of these MEMS devices will then be presented.

## **CONTACT/BENDING SENSORS**

### **Existing Bending Sensors**


Devices used to detect bending or deflection of a substrate have existed for many years. Whether an application calls for a bending, tactile, or force sensor, the principle of operation of these devices is to sense strain. Once strain is measured, other information such as the deflection or force can be calculated. All bending or force sensors that exist to date can be roughly categorized as thin-film metal foil strain sensors, piezo-resistive sensors, or capacitive array sensors.

With a conventional Wheatstone bridge, each of the three devices mentioned above can be used to measure strain. The advantages of these types of devices are their accuracy and their ability to measure a continuous range of strain. Their main disadvantages are that they need a continuous source of electrical power during operation and that the information collected requires processing due to the continuous range of sensitivities supplied.

### **NUMEM Contact/Bending Sensor**

It was decided to use bending sensors on the lobster antennae that closely resemble switches rather than traditional gauge sensors. During operation, the switch is not drawing current from the battery until the antenna is bent to a radius sufficient to activate the switch. Processing of the information from these switch sensors is much simpler than for traditional strain sensors. When a particular switch is activated, it can then activate a particular response.

The switch-type contact/bending sensors have been fabricated using the Northeastern University metal micromachining (NUMEM) process [2]. The contact/bending sensor is a gold cantilever beam (figure 1.1a)



**Figure 1.1** (a) A standard sensor with the source and drain. (b) Illustrates the closing of the switch by bending the substrate.


that is surface micromachined on a very flexible and thin silicon substrate. By attaching the thin silicon substrate to the antenna of the lobster robot, bending of the antenna will also bend the silicon substrate. Thus it is possible to design a switch that closes when the silicon substrate deflects to a certain curvature. Figure 1.1b illustrates how a switch will close by bending of the silicon substrate.

The sensors are to be placed inside the antenna of the lobster; therefore, the sensors have to be designed to fit in that small space. Because each cantilever switch can only detect one curvature, multiple switches are needed to detect a range of curvatures.


The design requirements for these contact/bending sensors are the ability to detect a range of radii of curvature and the capability of withstanding a stiction force as high as 100  $\mu\text{N}$ . The 100  $\mu\text{N}$  force was chosen because it was previously determined that electrostatic switches similar in design to these contact sensors required a release force of about 50 to 75  $\mu\text{N}$  to operate consistently [2]. The design constraints are that the beam thickness is 8  $\mu\text{m}$  and the switch gap is 0.25  $\mu\text{m}$ . By using these design requirements, design constraints, and a desired detectable radius of curvature, we can determine the width and length of each switch or beam.

**NUMEM Fabrication Process** The bending sensors will be fabricated using the NUMEM process [2]. The NUMEM process uses surface micromachining to deposit, pattern, etch, and plate various metals to fabricate a microswitch. The schematic shown in figure 1.2 illustrates the basic processing steps. Figure 1.2a shows the metal deposition, patterning of the photoresist, and dry etching steps used to produce the source and drain. Figure 1.2b shows the deposition of the copper sacrificial layer, patterning, and etching steps used to produce the base and tip of the sensors. Figure 1.2c shows the patterning of the beam, plating of the beam, and removal of the sacrificial layers to release the sensor.

Examples of standard contact/bending sensors, fabricated from the NUMEM process, can be seen in figure 1.3. In figure 1.3, the switches




**Figure 1.2** Schematic of NUMEM microfabrication process.



**Figure 1.3** SEM micrograph of two different length standard contact/bending sensors; each sensor measures a different curvature.

share the same source electrode located on the lower right and separate drains located on the upper left. There is a gate under both beams for electrical testing purposes.

**Antenna Design** The design of the lobster antenna is based on the desire to mimic a real lobster antenna, specifically the flexibility of the antenna and the function it serves as a contact sensing device. Other desired characteristics are that the antenna be relatively easy and inexpensive to produce. A drawing of the antenna design is shown in figure 1.4. Each antenna is composed of two halves of machined PVC sheets. Polycarbonate had been used initially but was replaced by PVC in order to eliminate some of the adhesive outgassing problems encountered with polycarbonate. Each of these sheets is 15 mils thick with a milled



**Figure 1.4** Polycarbonate antenna with the location of the contact/bending sensors and flexible circuit.

slot 5 mils deep. The bending sensors and the flexible circuit (with a polyimide substrate) used for the electrical connection are mounted to one half of the antenna. There are three sensors located at each of three locations as shown in figure 1.4. The second half of the antenna is bonded to the first half, sandwiching the sensors and the circuit. The milled slots on both halves of the antenna serve as housing to isolate the sensors and circuit from the sea water.

An antenna will be bent by contact with an obstacle. However, sufficient water flow will also cause antenna bending. How then does the lobster robot distinguish between bending due to contact and bending due to flow? The cross-section of an antenna of a real lobster is tapered. An analysis [3] has shown that it is the taper of the antennae that causes its deflected shape due to contact to differ significantly from its deflected shape due to flow.

One of the difficulties in processing is the use of 100  $\mu\text{m}$  silicon wafers in fabrication. These thin wafers are prone to breakage but are needed in order to allow the die to be easily bent by an antenna. A second problem is the amount of labor involved in the antenna assembly process. Processing on a plastic substrate is being explored to help with both problems. The concept is that the entire antenna tip, with all bending sensors, would be fabricated as a single unit. A number of these tips would be contained on each plastic wafer. Each tip would have the full contact and flow-sensing capabilities of the current antenna design, but at a much lower production cost, and with higher reliability. This fabrication method is being explored for use as the next-generation antennae and sensors.

**Radius-of-Curvature Calculations** The first step in designing both of these contact sensors is to develop the relationship between the length of the switch and the radius of curvature that it will detect. Because the length of these sensors is very small, it is assumed that the substrate below each switch deflects into a circular shape. The relationship between the switch length and its sensitivity can be obtained using






Figure 1.5 Radius of curvature for contact vs. switch length.

$$R^2 = L^2 + Y^2 \quad (1.1)$$

where  $R$  is the radius of curvature,  $L$  is length of the switch, and  $Y$  is the vertical location of the switch tip from the center of the circle. The value of  $Y$  in terms of  $R$  and the switch gap  $h$  ( $h$  is  $0.25 \mu\text{m}$ ) is

$$Y = R - h \quad (1.2)$$

Substituting equation (1.2) into equation (1.1) and neglecting  $h^2$  results in

$$L = \sqrt{2Rh} \quad (1.3)$$

which relates the detected radius of curvature to the switch length. Equation (1.3) is used to generate figure 1.5, which shows detected radius of curvature versus switch length. It shows that the sensitivity of the switch increases parabolically with switch length.

**Cantilever Beam Design** The base of the beam is attached to the source electrode, which is attached to the silicon oxide substrate (figure 1.1). The beam is approximately  $0.5 \mu\text{m}$  above the silicon oxide substrate and the tip height is approximately  $0.25 \mu\text{m}$ , resulting in a switch gap of about  $0.25 \mu\text{m}$ .

The theoretical calculations for the cantilever design were based entirely on small deflection Euler-Bernoulli beam theory. The deflection of the end of a cantilever beam with an end load is

$$h = \frac{PL^3}{3EI} \quad (1.4)$$

where  $h$  is the switch gap,  $P$  is the  $100 \mu\text{N}$  stiction load,  $L$  is the length of the switch,  $E$  is the elastic modulus of the beam, and  $I$  is the area moment of inertia given by

$$I = \frac{1}{12}wt^3 \quad (1.5)$$

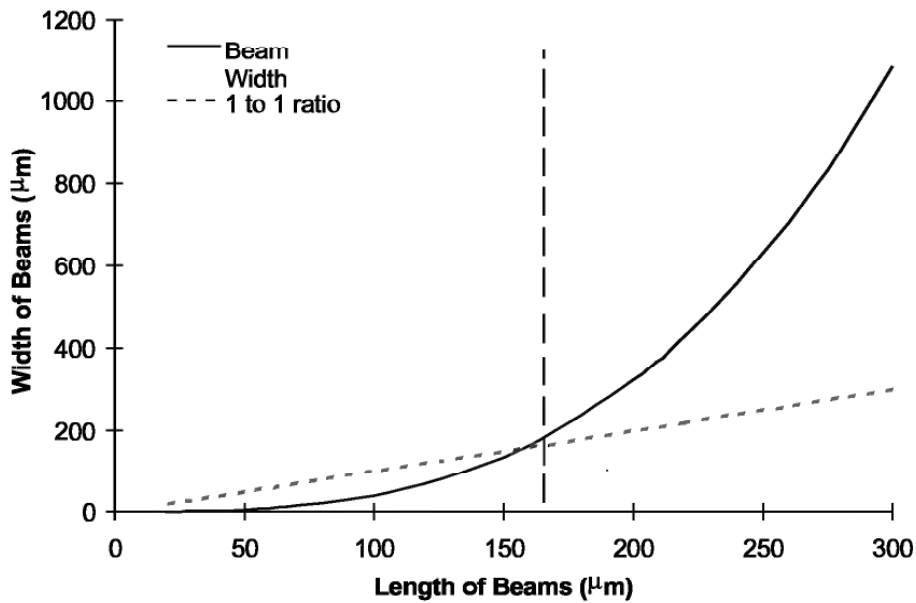


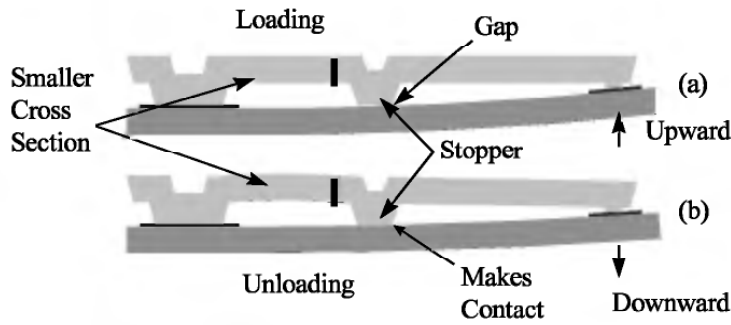
Figure 1.6 Plot of width vs. length for the standard cantilever sensors.

where  $w$  and  $t$  are the width and the thickness of the beam, respectively.

In equations 1.3 and 1.4, the two unknowns are  $L$  and  $I$ . Thus, if a specific sensitivity is desired, the length of the switch to achieve this sensitivity can be obtained from equation 1.3 (figure 1.5). Then the area moment of inertia ( $I$ ) and hence the width of this beam required to satisfy the stiction condition can be calculated using equations 1.4 and 1.5. These values of width versus length are plotted in figure 1.6.

In figure 1.6, the straight line represents a width-to-length ratio of one. The other line represents the width that is required to satisfy all of the design requirements and constraints for a given-length switch. The dashed vertical line indicates when the width of the switch becomes larger than the length of the switch. From figure 1.6, it was determined that a switch length greater than 200  $\mu\text{m}$  was not practical. For example, a 300  $\mu\text{m}$ -long switch would have to be more than 1 mm wide to satisfy the design requirements and constraints. Therefore, the standard cantilever sensor is limited to a length of less than 200  $\mu\text{m}$ , which corresponds to a radius of curvature of less than 80 mm.

**Cantilever Beam with Stopper Design** The stopper sensor design (figure 1.7) is similar to the standard sensor except that it uses a stopper to help extend its sensitivity. In operation, the stopper sensor is similar to the standard sensor except during the unloading process, when the stopper comes in contact with the substrate to pry the beam off the drain. The stopper comes into contact with the silicon oxide, making the beam stiffer in the unloading process by shortening the effective length of the



**Figure 1.7** (a) A stopper sensor electrically closing via substrate bending. (b) Stopper making contact with the substrate during unloading to pry the beam off the drain.

beam. During loading, the effective length of the switch is the length from the base to the tip. Another feature of the stopper sensor is that the cantilever beam switch has two different widths. The stopper is located at the junction of the thin beam to the thick beam.

One of the key concerns during the design process was what would happen if a sensor is bent far beyond the radius of curvature that it was designed to detect. If this happens, the switch may plastically yield or the contact force may be so high that the switch will stick shut, even after the substrate returns to its completely flat position. Calculations show that both the standard and stopper sensors will not yield when the substrate is bent to a radius as low as 3 cm. Also, the standard sensor will have a contact force directly proportional to the deflection of the tip. However, the thin beam on the stopper sensor has a very low stiffness and as a result, overbending will not cause a high contact force between the tip and drain.

The stopper sensor design allowed sensors to be designed with beam lengths up to 300  $\mu\text{m}$ . This beam length can measure a radius of curvature of 180 mm. Therefore, the stopper design increases the range of sensitivity that can be detected while still supplying sufficient retraction force to overcome stiction.

**Test Fixture** To test the accuracy of these contact/bending sensors, we created a fixture to activate the sensors and measure the radius of curvature. The fixture uses two stands or mounting fixtures that are attached to a linear slide. One stand holds a laser light source while the other holds the test specimen. The test specimen adheres to the matted polycarbonate sheet and is wire-bonded to a flexible circuit. The circuit is wired to light-emitting diodes (LEDs). With a small electrical source, the LEDs will light up when the sensors close.

To measure the radius of curvature of a specimen, the laser light is directed onto the specimen at a slight angle in a manner in which the reflection off the silicon oxide is visible on the surface of the stand holding the laser. The radius of curvature of the specimen is obtained

by moving the stands along the linear slide until the reflected light off the specimen is focused into a point. Once the reflected light is focused, the radius of curvature is twice the distance between the specimen and the surface it is focused on.

**Test Method** The testing of the sensors was limited to the testing of three sensitivities. For simplicity, the following notation is used:

- Switch 1 is the most sensitive sensor with an expected contact radius of 8 cm.
- Switch 2 is the second-most sensitive sensor with an expected contact radius of 5.1 cm.
- Switch 3 is the least sensitive sensor with an expected contact radius of 3.2 cm.

*Test 1:* The first test is to measure the radius of curvature at which switch 1 is activated and deactivated. This procedure was repeated ten times. The ten measurements were used to determine the accuracy and repeatability of switch 1.

*Test 2:* The second test measures the radius of curvature at which switch 1 is activated and then proceeds to activate switch 2 and record its radius of curvature for contact. Upon unloading, the radius of curvature at which switch 2 and switch 1 are deactivated was also recorded. This procedure was repeated ten times. The ten measurements were used to determine the accuracy and repeatability of switch 2. This test will also help determine if the performance of switch 1 is affected by overbending during the activation of switch 2. Test 2 is the first dynamic range test for switch 1.

*Test 3:* The third test is similar to test 2. The radius of curvature at which switch 1, switch 2, and switch 3 are activated was recorded in sequence. Upon unloading, the radius of curvature at which all three switches deactivated was also recorded.

This procedure was repeated ten times. The ten measurements were used to determine the accuracy and repeatability of switch 3. This test was also used to help determine if the performance of switch 1 and 2 were affected by the overbending that occurred during the activation of switch 3. Test 3 was also used to help determine the dynamic ranges of switch 1 and switch 2.

Tests 1, 2, and 3 were performed in sequence on the same test specimen.

**Test Results** The test data collected support the theoretical calculations. Figure 1.8 is a plot of the experimental data collected for one set of sensors along with the theoretically calculated values for the radius of curvature. Each data point represents an average of ten

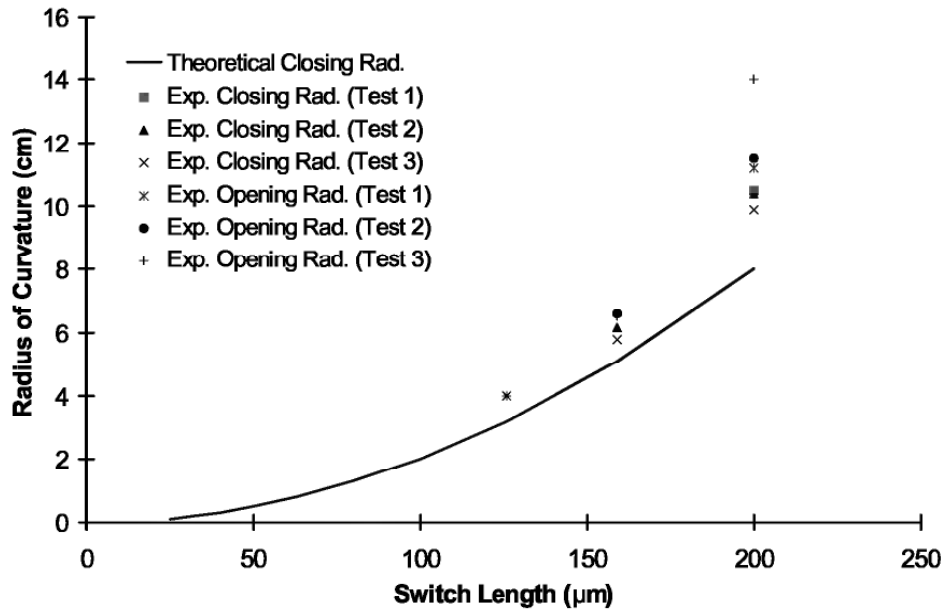


Figure 1.8 Plot of experimental values along with theoretically calculated values.

measurements. The range of the data points for the average closing radius is  $\pm 0.5$  cm. The range of the data points for the average opening radius is  $\pm 1.0$  cm. From figure 1.8, it is clear that stiction forces are present in these sensors after overbending. For example, the 200  $\mu\text{m}$  sensor activated at an average radius of about 10 cm but it is deactivated at an average radius of 14 cm. The 150  $\mu\text{m}$  sensor is activated and deactivated within a much smaller range.

The difference between the experimental closing radius and the theoretical closing radius can be attributed to the difference between the actual switch gap and the value used in the design and to systematic errors in the test setup. If the switch gap deviates from the theoretical value of 2500  $\text{\AA}$ , the closing radius for each sensor will change. Using equation 1.3 and the experimental closing radii values, the actual switch gap can be calculated to be approximately 2000  $\text{\AA}$ . The method of measuring the curvature of the silicon substrate may also contribute a 5% error to each measurement.

Results to date for the stopper sensors indicate that the switch gaps for these sensors are much larger than 2500  $\text{\AA}$ . As a result, the switches close at a much smaller radii than theoretically calculated. SEM (Scanning Electron Microscope) micrographs of the stopper sensors reveal that the cause of the larger switch gaps is an initial upward curvature of these sensors. It is believed that the curvature in these sensors is caused by residual stresses in the region where the thin beam is attached to the thick beam. This diagnosis is supported by SEM micrographs, which show that the majority of the bending in the beam occurs at this thin-to-thick beam junction. Away from this junction toward the

free end of the cantilever, the gap between the beam and the substrate increases linearly. A possible solution to alleviate the residual stresses at this junction is to change the electrical settings used to electroplate the sensors.

## FLOW SENSORS

### Existing Flow Sensors

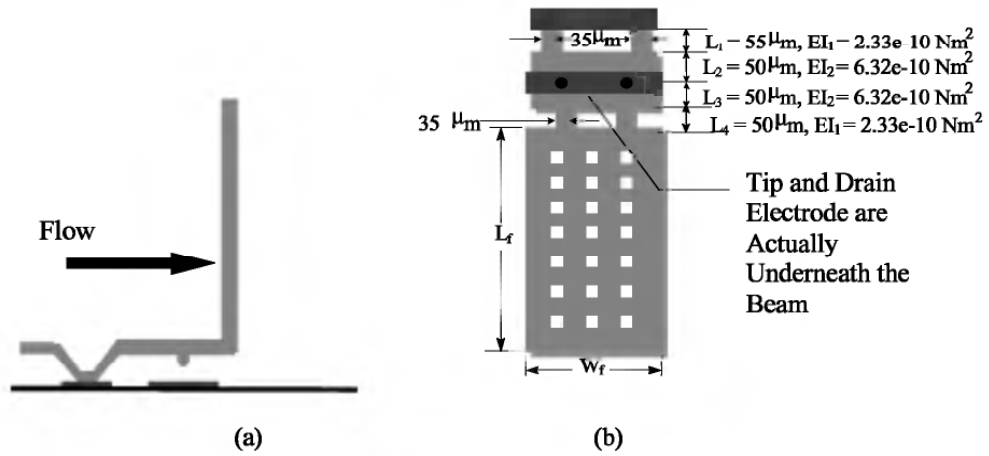
One type of existing flow sensor is based on a thermal anemometer. This device consists of a resistor that generates heat. Two sensors measure the temperatures upstream and downstream of the resistor. The asymmetry in the readings is used to determine the flow rate. Another design for a microflow sensor utilizes a flexible membrane that is deflected by the flow. A change in flow is detected by a change in capacitance between two electrodes. Another sensor uses fluid drag over an obstacle. This device uses a piezoresistor to convert the fluid flow to electrical information.

These sensors need more wiring and power than is desirable for use in the lobster robot and are generally more complicated than is needed for our application. Also, using any of these sensors as our flow sensor would be a step away from our goal to make the robot biomimetic. A lobster has small hairs that sense the water flowing around it [1]. Thus, for our application, a simple, on/off, 1-bit switch is most suitable.


MEMS sensors save space because they are very small (submillimeter). They are also economical to fabricate because they are made by surface micromachining. Furthermore, the NUMEM process has been selected to fabricate these flow sensors. Low power consumption and simplicity as an on/off 1-bit device are other advantages of these switch-based devices.

### Design of Paddle-Flow Sensor

A side-view and a top-view (before the sensor is bent at a right angle) of the lobster flow sensor are shown in figure 1.9 and will be referred to as the *paddle-flow sensor*. The water flows against the paddle as shown. Sufficient flow produces enough force to close the switch. The sensitivity of the switch is determined by the dimensions of the flow area, the beam thickness, the tip location, and the tip gap. The paddle switch is a complex cantilever beam having two different cross-sections with dimensions as shown in figure 1.9. By altering the dimensions of the paddle, the switch can be modified to detect different flow speeds. As shown in the SEM micrograph of figure 1.10, we have fabricated four switches of this type with water flow speed sensitivities of 0.5, 1.0, 3.4, and 5.4 m/sec.



**Figure 1.9** (a) Side view of a bent flow sensor. (b) Top view of a flow sensor prior to being bent.



**Figure 1.10** SEM micrograph of four paddle flow sensors.

In order to analyze the behavior of these switches, the first step is to calculate the moment required to make the switch close. The gap between the tip of the switch and the drain electrode, known as the tip gap ( $h$ ), was chosen to be  $2500 \text{ \AA}$ . For processing purposes, it is desirable to hold  $h$  constant. In order for the switch to close, there must be a sufficient moment applied to the cantilever such that the tip deflects  $2500 \text{ \AA}$ . Using superposition of the cantilever beam equations from strength of materials [4], and the known dimensions and deflection, the

moment necessary to close the switch can be found by

$$h = \frac{ML_1^2}{2EI_1} + \frac{ML_2^2}{2EI_2} + \frac{ML_1L_2}{EI_1} \quad (1.6)$$

where  $h = 2500 \text{ \AA}$ ,  $L_1 = 55 \text{ \mu m}$ ,  $L_2 = 50 \text{ \mu m}$ ,  $EI_1 = 2.33\text{e-}10 \text{ Nm}^2$ ,  $EI_2 = 6.32\text{e-}10 \text{ m}^2$ , and  $M$  is the moment needed to close the switch (Nm). Using this equation, we calculate the moment needed for contact to be  $1.23\text{e-}8 \text{ Nm}$ .

Next, the flow velocity of the seawater necessary to cause this moment must be determined. It is assumed that the switch will be placed near the tip of the lobster claw so that the boundary layer thickness is negligible and uniform flow contacts the paddle. Thus the moment is equal to the force from the flow on the paddle multiplied by half of the paddle length ( $L_f$ ). The force on the paddle is calculated from

$$F = \frac{1}{2}\rho AV^2 \quad (1.7)$$

where  $F$  is the force necessary to close the switch (N),  $\rho$  is the density of seawater ( $1030 \text{ kg/m}^3$ ),  $A$  is the total area of flow paddle ( $\text{m}^2$ ), and  $V$  is the velocity of current detected (m/sec).

In calculating the flow paddle area, the area of the release holes has been subtracted. The release holes are present for fabrication purposes and will be discussed later in this chapter. The above equation is derived from the Bernoulli equation, which is applicable for steady, inviscid, incompressible flow. These calculations were used to design the switches with four different sensitivities.

The four paddle switches, with sensitivities described above, are lined up side by side into an array of switches detecting flow from a single direction. Four arrays of switches are then arranged on a die such that they detect flow from four different directions.

One of the failure modes of this switch could be plastic deformation due to bending stress. Therefore, the maximum bending stress in the structure must be calculated. This occurs at the narrow points of the beam (figure 1.9b,  $L_1$  and  $L_4$  sections). Using strength of materials theory, the equation for the bending stress is

$$\sigma = \frac{Mc}{I}, \quad I = \frac{wt^3}{12} \quad (1.8)$$

where  $\sigma$  is the maximum bending stress (Pa),  $M$  is the moment caused by design flow ( $1.23\text{e-}8 \text{ Nm}$ ),  $c$  is half the thickness of the beam ( $4 \text{ \mu m}$ ),  $I$  is the cross-section area moment of inertia ( $2.98\text{e-}21 \text{ m}^4$ ),  $w$  is the width of beam at point of maximum stress ( $70 \text{ \mu m}$ ), and  $t$  is the thickness of beam ( $8 \text{ \mu m}$ ).

The maximum bending stress when a switch closes is  $16.5 \text{ MPa}$ . Because the yield stress of gold is approximately  $130 \text{ MPa}$ , the stress caused by the design flow speed should not cause yielding. Should



higher flow rates occur, the contact of the tip will offset further deflection and hinder plastic deformation from bending at the base of the cantilever. However, this will not help the bending stress at the 90° bend. The packaging of the switch might help this situation. If the switch is enclosed by some membrane with the paddle protruding through, an excessive flow would cause the paddle to deform the membrane and its elasticity would help to offset the moment. Similar effects would result from the paddle coming into contact with any material used as the packaging material.

When the flow velocity reaches the design velocity,  $V_c$ , the beam will bend down and the tip will touch the drain electrode on the substrate. As the flow increases beyond  $V_c$ , the elbow of the switch will continue to bend down and eventually the edge of the beam will contact the substrate. As the flow continues to increase, there will be an upward supporting force at the elbow point. This force will reduce the contact force at the tip. If it is too large, it could cause the tip to lift up off the drain electrode. Careful redesign of the switch dimensions will prevent this liftoff from occurring in the desired velocity range.

Another concern for the switch is stiction between the gold tip and the gold contact. The retraction force of the beam must be sufficient so that the beam lifts from the drain once the flow has slowed/stopped. The force needed to pull the beam off the contact is the force necessary to give the complex cantilever a 2500 Å deflection and is given by

$$h = F \left[ \frac{L_1^3}{3EI_1} + \frac{L_1^2 L_2}{EI_1} + \frac{L_2^2 L_1}{EI_1} + \frac{L_2^3}{3EI_2} \right] \quad (1.9)$$

where  $h$  is the tip gap (2500 Å),  $L_1$  and  $L_2$  are the beam lengths (m),  $EI_1$  and  $EI_2$  are the beam flexural rigidities (Nm<sup>2</sup>), and  $F$  is the retraction force (162 μN). From the experiment, it was found that a retraction force of about 100 μN was more than sufficient to lift the beam of a MEMS device upon unloading [2]. Our beams have a retraction force of about 162 μN. This should be more than stiff enough to lift the beam from the contact.

## Fabrication

The fabrication process is based on the Northeastern University metal micromachining (NUMEM) technology [1]. This process uses surface micromachining to deposit, pattern, etch, and electroplate various metals. The flow paddle needs to be lifted to a position that is perpendicular to the substrate without lifting the part of the beam where the tip is located. Thus a partial release is used so that while the paddle is being lifted by probes, the metal under the tip is holding the beam down. To accomplish this, the paddle was designed with etch holes to

allow the etchant underneath the paddle and to dissolve the sacrificial metal before the metal under the tip is etched.

This process was difficult to accomplish because of the inconsistency of the thickness of the sacrificial layer throughout the wafer. It was hard to determine an etch time that would consistently release the paddle without releasing the switch, so a new approach was examined. A mask was created so that, using careful alignment and photolithography, a layer of photoresist could be used to cover only the section of the switch that should not be released. Because the photoresist is leaving the paddle exposed, the etchant can flow into the release holes and dissolve the metal under the paddle while leaving the switch unreleased. Using microprobes, the flow paddle is lifted to its vertical position. After this step, the photoresist is etched away, then the remaining metal under the switch is etched away, finally exposing the tip and fully releasing the switch.

### **Packaging**

It is essential to not only have the circuitry but the switch tip and drain electrically isolated from the conductive seawater. Thus some sort of package is needed to keep the water away from the switch yet still allow the paddle to be exposed to the force of the external flow. One potential solution being investigated is an open-channel package.

A silicon shell with a slit is aligned and bonded on top of the flow sensor, with the sensor flow area protruding through the slit like a hair on a lobster. The small cavity that houses the switch is filled with a viscous fluid. Another possibility, specifically for the lobster, is to bond a hair onto the paddle and to place the switch in a pit on the lobster claw. The pit could be filled with a viscous fluid or covered with an elastic membrane to isolate the switch from the seawater while the hair protrudes out of the pit and is exposed to the flow. This solution would make the mechanical lobster more biomimetic, because lobsters actually use small hairs to sense flow. The packaging issue is under investigation; however, it is clear that the flow sensor must be isolated from its environment.

### **Testing**

A flow board has been constructed to test the behavior of the flow switches in deionized (DI) water. The DI water flows through into a flow meter where the flow can be measured and adjusted before flowing over the test specimen. The test specimen consists of a die with four paddle switches wire-bonded to a flexible circuit that is bonded to a strip of plastic to keep the specimen straight. The flexible circuit is

subsequently connected to an edge connector and then to an ammeter, where a voltage can be applied and a current can be measured to check if the switch is shut at a particular flow rate, which is read from the flow meter. The strip of plastic is wide enough so that it is held about halfway between the top and bottom of the Tygon hose. Consequently, the flow felt by the switch is that which is from the middle of the tube.

In order to properly measure the water flow, it should be fully developed both at the flow meter and at the test specimen. To calculate the length at which the flow of water becomes fully developed, it must be determined whether the flow is laminar or turbulent. This is determined by the Reynolds number ( $Re$ ). The Reynolds number for pipe flow is found by

$$Re = \frac{\rho VD}{\mu} \quad (1.10)$$

where  $Re$  is the Reynolds Number,  $\rho$  is the fluid (water) density (1000 kg/m<sup>3</sup>),  $V$  is the fluid velocity (m/sec),  $D$  is the pipe diameter (9.525e-3 m), and  $\mu$  is the fluid viscosity (1.12e-3 Nsec/m<sup>2</sup>).

A Reynolds number of 4000 or above indicates turbulent flow. The results are that all flows are turbulent. The maximum Reynolds number is about 45,900. Using this value, the minimum distance for fully developed flow for the fastest flow detectable by the switches can be determined by

$$l_e = 4.4D(Re)^{1/6} \quad (1.11)$$

where  $l_e$  is the length for fully developed flow (m). The length necessary for fully developed flow is about 0.25 meters—or about 10 inches. Therefore, the length of pipe from the faucet to the flow meter and the length from the valve to the test specimen each must be 10 or more inches in order to properly measure the flow rate.

Because the water out of the faucet only reaches about 1.2 gpm, this method of testing only works with the switches that measure smaller flows. For the two less-sensitive switches, a pump is needed to increase the flow to the point that these switches can close. In this configuration, the DI water is contained in a reservoir. The pump pulls the water out of the reservoir and into the meter. As before, the water flows out of the meter and into the testing device, which is held inside the Tygon tubing.

## Results

To date, only the low-flow-rate switches have been tested. The modified test device is in the process of being constructed. As for the low-flow-rate switches, the 1 m/sec (1.13 gpm) switch has been activated. The test device was inserted in the tube while the flow was being increased. Once the switch was activated, the flow was decreased to see

when the switch would open to check for hysteresis. This process was repeated at least 10 times for the same switch. The results agree well with the calculations. On average, the switch closed at about 1.1 gpm and lifted off at about 0.9 gpm. The voltage applied was 0.1 V and the electrical current was about 0.8 mA. The resistance of the switch and connecting leads is about 125  $\Omega$ . More low-flow switches will be tested and the higher-flow switches are in the process of being tested.

## CONCLUSIONS

Two types of MEMS contact/bending sensors have been designed, analyzed, fabricated, and tested. The standard sensor and the stopper sensor were designed to be placed in the antennae of a biomimetic robotic lobster. These sensors were designed based on theoretical calculations to account for stiction forces and to activate at certain curvatures. Testing to date suggests that the standard sensor operates as predicted. The stopper sensor was designed for improved sensitivity and increased dynamic range. Modification of the processing parameters should eliminate the residual stresses that are preventing effective operation of the stopper sensor.

The antennae for the lobster robot have also been designed, analyzed, fabricated, and tested. Each antenna is made from PVC and has three contact/bending sensors at each of three locations along its length. By obtaining curvature information from these three locations, the lobster robot will be able to distinguish bending due to contact from bending due to flow. These antennae are now working reliably. However, the next generation of antennae will be fabricated with the sensors and antenna as a single unit. This process is expected to increase reliability and reduce the time and cost involved in antenna assembly.

MEMS flow sensors for the biomimetic lobster have been designed, analyzed, fabricated, and tested. The structure is a nonuniform cantilever beam with a paddle region that is bent perpendicular to the flow. The paddle is made to mimic the hairs that lobsters use to sense flow. The area of the paddle is varied so that the switches close at different flow rates. The low-flow switches have been tested with some success and measurements agree with theoretical calculations. Packaging options, such as the open-channel package, are being explored and evaluated in order to seal the switch from the flow.

## ACKNOWLEDGMENTS

The authors are grateful to the Controlled Biological Systems Program of DARPA for their support of this work under Grant N00014-98-1-0381.

## REFERENCES

- [1] Ayers, J. (1995). A reactive ambulatory robot architecture for operation in current and surge. *Proceedings of the Autonomous Vehicles in Mine Countermeasures Symposium*. pp. 15–31.
- [2] Zavracky, P. M., McGruer, N. E., and Majumder, S. (1997). Micromachined switches fabricated using nickel surface micromachining. *IEEE/ASME J. Microelectromechanical Syst.* 6: 3–9.
- [3] Barnes, T. G., Truong, T. Q., Adams, G. G., and McGruer, N. E. (2001). Large deflection analysis of a biomimetic lobster robot antenna due to contact and flow. *ASME J. Appl. Mechanics* 68: 948–951.
- [4] Beer, F. P., and Johnston, E. R. Jr. (1992). *Mechanics of Materials*, 2d ed. New York: McGraw-Hill.

# 2

## Fly-Inspired VLSI Vision Sensors

Reid R. Harrison

Engineers have long looked to nature for inspiration. The diversity of life produced by five billion years of evolution provides countless existence proofs of organic machines with abilities that far surpass those of our own relatively crude automata. We have learned how to harness large amounts of energy and thus far exceed the capabilities of biological systems in some ways (e.g., supersonic flight, space travel, and global communications). However, biological information processing systems (i.e., brains) far outperform today's most advanced computers at tasks involving real-time pattern recognition and perception in complex, uncontrolled environments. If we take energy efficiency into account, the performance gap widens. The human brain dissipates 12 W of power, independent of mental activity. A modern microprocessor dissipates around 50 W, and is equivalent to a vanishingly small fraction of our brain's functionality.

Indeed, the human brain is a daunting goal for biologists and engineers alike. Our brain takes several years to fully develop, and contains between  $10^{10}$  and  $10^{11}$  neurons (nerve cells), each communicating with  $10^3$  other cells, on average. Brains of other animals (particularly invertebrates) are much smaller but still perform remarkably complex computations. Insect brains, for example, typically contain between  $10^5$  and  $10^6$  neurons, yet insects perform sophisticated information-processing tasks rapidly and efficiently.

We have attempted to extract computational principles from the visual system of the fly and apply these principles to an engineered system—an integrated, low-power visual motion sensor. As our engineering tool, we use very large-scale integration (VLSI) of silicon circuits—the most advanced information-processing substrate available today. In particular, we explore continuous-time (unclocked), continuous-value (analog) circuit architectures. This approach was pioneered by Mead and colleagues beginning in the 1980s (Mead, 1989).

The fly is an attractive target for biologically inspired approaches to engineering. Its brain and sensory systems have been studied for decades, so much is known about their operation. Of course, we are still decades (or centuries) away from understanding the entire system,

but a wealth of behavioral and electrophysiological data has led to the development of several models of information processing.

## THE VISUAL SYSTEM OF THE FLY

Vision is a vitally important sense for flying insects. In the housefly's brain, more than half of the 350,000 neurons are believed to have some role in visual processing. The fly's optic lobes contain motion-sensitive neurons that respond to moving stimuli over large portions of the visual field. Many of these neurons have been linked to specific visually guided behaviors that help the animal navigate through a complex environment in a robust manner (Egelhaaf and Borst, 1993).

Insects process visual motion information in a local, hierarchical manner. This information processing begins at the sensor—the retina. Despite the multilens construction of the compound eye, the pattern projected onto the underlying retina is a single image of the visual scene. Photoreceptors in the retina adapt to the ambient light level, and signal temporal deviations from this level. These signals are passed on to the next layer of cells, the lamina. Lamina cells generally show transient or high-pass responses, emphasizing temporal change (Weckström, Juusola, and Laughlin, 1992). The next stage of processing is the medulla, a layer of cells that are extremely difficult to study directly due to their small size. Indirect evidence suggests that local measures of motion (i.e., between adjacent photoreceptors) are computed here. These local, direction-selective motion estimates are integrated by large tangential cells in the lobular plate (Hausen and Egelhaaf, 1989). The housefly has 50–60 tangential cells in each hemisphere of its brain. These are the best-studied cells in the fly visual system, and much is known about their properties.

Lobular plate cells generally respond to motion over large parts of the visual field. Some of these cells seem to be matched filters for the optic flow patterns produced by rotation or translation along particular axes (Krapp and Hengstenberg, 1996). Some of these cells most likely control compensatory motor reflexes that prevent the fly from rotating during flight. Others are sensitive only to small objects moving across the visual field (Egelhaaf, 1985). It is believed that these “figure detection” cells allow the fly to locate nearby objects through motion parallax (Kimmerle, Warzecha, and Egelhaaf, 1997). All of these sensory abilities require that motion first be detected locally between every pair of photoreceptors.

### Photoreception

Each eye of the blowfly *Calliphora erythrocephala* consists of approximately 6000 individual lenses. Beneath each lens is a cluster of eight

light-sensitive cells. Each lens and its associated photoreceptors form a unit called an ommatidium. Six of the eight photoreceptors are used to implement neural superposition, a technique to increase the effective lens diameter by pooling the responses of neighboring ommatidia. The other two photoreceptors do not seem to be involved in the detection of motion. Mutants with these photoreceptors impaired cannot discriminate colors, but show no motion-related deficits (Heisenberg and Buchner, 1977). From an information-processing perspective, each ommatidium records one “pixel” of the external world’s image. Inter-ommatidial angular spacing is  $1.1\text{--}1.3^\circ$  (Land, 1997). This angular resolution is approximately 150 times worse than the  $0.008^\circ$  resolution in the foveated region of the human retina (Wandell, 1995). (This is roughly equivalent to having 20/3000 vision.) Although the ommatidia are arranged in a hexagonal array, it is useful to think of the equivalent array size in terms of the standard rectangular array used by computer monitors and digital cameras. Taking the square root of the number of ommatidia (6000), we see that *Calliphora*’s eye is roughly equivalent to a 77 by 77-pixel array covering one visual hemifield. The eye of the fruit fly *Drosophila* contains only 700 ommatidia, resulting in an equivalent array size of 26 by 26 (Land, 1997). Today’s cheap digital cameras provide 640 by 480-pixel images, and emerging photo-quality digital cameras provide 1800 by 1200 pixels or more—around two orders of magnitude more photoreceptors than a fly’s eye. Typical cameras concentrate these pixels into a  $40^\circ$  field of view, while each fly eye sees nearly a complete hemisphere.

It is remarkable that flies are capable of such impressive navigation when one considers their low-resolution eyes. This limited spatial acuity is a consequence of the compound eye design. In order to increase spatial acuity, more ommatidia are required. However, the resolving capability of each ommatidium is limited by diffraction, so each lens must also be made larger. If we wanted to build a compound eye with the acuity of the human fovea ( $0.008^\circ$ ), it would have a radius of 11.7 meters! The visual acuity of the largest insect eye in nature (that of the aeshnid dragonfly) reaches  $0.24^\circ$  in its most acute zone, still 30 times coarser than the human fovea (Land, 1997).

Although inferior to human eyes spatially, fly vision far exceeds ours temporally. Human vision is sensitive to temporal modulations up to 20 or 30 Hz, while fly photoreceptors respond to temporal frequencies as high as 300 Hz (Autrum, 1958).

## SIGNAL PROCESSING IN THE PERIPHERAL OPTIC LOBE

The laminar region, also called the first optic ganglion, contains cells that exhibit transient responses to step intensity changes. The large monopolar cells (LMCs) in this ganglion ignore the DC light level but



amplify temporal changes (Weckström, Juusola, and Laughlin, 1992). This high-pass response has been shown to optimize information transfer through this region (Laughlin, 1994). Laminar cells do not exhibit motion-specific responses. There is a strong retinotopic organization from the retina through the lamina to the next layer, the medulla. Every ommatidium has an associated neural “cartridge” beneath it in these underlying ganglia, suggesting many identical processing units operating in parallel (Strausfeld, 1976).

Cells in this second optic ganglion are extremely small and difficult to record from, and little is known about their structure or function. DeVoe recorded from medullar cells in *Calliphora* and reported a wide variety of response characteristics: transient temporal responses, sustained responses, directional motion responses, and nondirectional motion responses (DeVoe and Ockleford, 1976; DeVoe, 1980).

### **The Tangential Cells of the Lobular Plate**

The third optic ganglion is also known as the lobula-lobular plate complex. At this point in the optic lobe, the retinotopic organization ends with massive spatial convergence. Information from several thousand photoreceptors converges onto 50–60 tangential cells. These cells have broad dendritic trees that receive synaptic input from large regions of the medulla, resulting in large visual receptive fields (Hausen, 1982a, b; Hengstenberg, 1982; Hausen, 1984; Krapp and Hengstenberg, 1996).

A subset of these neurons were found to respond primarily to horizontal motion, and these cells were given names beginning with “H.” H1 is a spiking neuron that responds to back-to-front optic flow. HSS, HSE, and HSN are graded potential (nonspiking) neurons covering the southern, equatorial, and northern regions of the visual field, respectively. Collectively called the HS cells, these neurons are depolarized by full-field visual motion from the front to the back of the eye, and hyperpolarized by back-to-front motion. They have been shown to encode horizontal motion as effectively as the spiking H1 cell (Haag and Borst, 1997). Each HS cell integrates signals from an ipsilateral retinotopic array of elementary motion detectors (EMDs), units in the medulla that estimate local motion in small areas of the visual field. The HS cells synapse onto descending, spiking neurons, which relay information to the motor centers of the thoracic ganglion. Another class of neurons, the VS cells, respond to vertical motion. Recently, it has been shown that these HS and VS cells are not simply responsive to motion along one axis, but rather act as matched filters for complex patterns of optic flow that would be produced by body rotations (Krapp and Hengstenberg, 1996).

## Visually Guided Behaviors

Flies rely heavily on visual motion information to survive. In the fly, motion information is known to underlie many important behaviors including stabilization during flight, orienting toward small, rapidly moving objects (Egelhaaf and Borst, 1993), and estimating time-to-contact for safe landings (Borst and Bahde, 1988). Some motion-related tasks, like extending the legs for landing, can be executed less than 70 msec after stimulus presentation. Wagner reports a 30 msec reaction time for male flies chasing prospective mates (Wagner, 1986). The computational machinery performing this sensory processing is fast, small, low power, and robust.

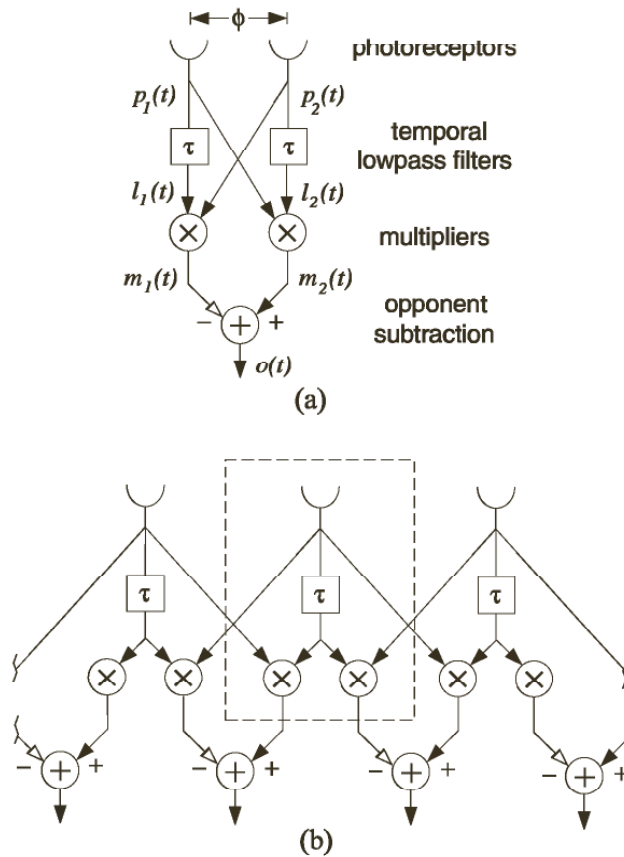
Flies use visual motion information to estimate self-rotation and to generate a compensatory torque response to maintain stability during flight. This well-studied behavior is known as the optomotor response. It is interesting from an engineering standpoint because it extracts relevant information from a dynamic, unstructured environment on the basis of passive sensors and uses this information to generate appropriate motor commands during flight. This system is implemented in biological hardware that is many orders of magnitude smaller and more power efficient than charge-coupled device (CCD) imagers coupled to a conventional digital microprocessor.

Much of the computation underlying the optomotor control system is performed by the HS cells (Geiger and Nüssel, 1981, 1982; Egelhaaf et al., 1988; Hausen and Wehrhahn, 1990; Egelhaaf and Borst, 1993). This well-studied system estimates rotation from optic flow and uses this information to produce a stabilizing torque with the wings (Götz, 1975; Warzecha and Egelhaaf, 1996).

## THE REICHARDT MOTION DETECTOR MODEL

Much is understood about the mechanisms used by flies to extract visual motion information from their retinal signals. The most-well-known model of fly motion detection is the Reichardt model, first proposed by Hassenstein and Reichardt in 1956 to explain visually mediated behaviors in a walking beetle (Hassenstein and Reichardt, 1956). This model is a type of correlation-based motion detector, which measures spatiotemporal correlations caused by moving objects.

This model has been successful at explaining both detailed electrophysiological responses of motion-sensitive neurons to visual stimuli (Egelhaaf and Borst, 1989; Zanker, 1990) and visually guided behavioral responses (Reichardt and Poggio, 1976; Reichardt and Egelhaaf, 1988; Borst, 1990; Warzecha and Egelhaaf, 1996). Modified versions of the Reichardt model have also been used to explain motion perception



**Figure 2.1** Reichardt motion detector architecture. (a) The signal from one photoreceptor is correlated with the delayed signal from an adjacent photoreceptor. Direction selectivity is increased by subtracting the responses of two half-detectors in opponency. (b) A 1-D array of Reichardt motion detectors, illustrating the repeated computational element. Subtraction currently performed off-chip for flexibility, but would be easy to implement given the current-mode outputs of the multipliers. A 2-D arrangement is possible with two additional multipliers in each cell.

properties in vertebrates, including humans (Borst and Egelhaaf, 1989; Clifford, Ibbotson, and Langley, 1997).

### Theoretical Analysis

We now describe the Reichardt motion detection algorithm, which will underlie our hardware implementation. As mentioned above, the Reichardt motion detector is a correlation-based algorithm, whose output is equivalent to the output of the spatiotemporal motion energy model proposed by Adelson and Bergen (1985) and others (van Santen and Sperling, 1985; Watson and Ahumada, 1985).

The basic idea of the Reichardt motion detector is to correlate the signal from one photoreceptor with the delayed signal from an adjacent photoreceptor (figure 2.1a). This delay-and-correlate algorithm produces a velocity-tuned response that is weakly directionally selective.

By subtracting the responses of two opponent half-detectors from each other, strong direction selectivity is achieved (Borst and Egelhaaf, 1990).

It is instructive to consider the case where the stimulus is a sinusoidal grating moving at velocity  $v$ . Image intensity  $i(x, t)$  can be expressed as

$$i(x, t) = I + \Delta I \sin[2\pi f_s(x + vt)] \quad (2.1)$$

where  $I$  is the mean intensity, and  $f_s$  is the spatial frequency. The contrast of the grating is  $\Delta I/I$ . At any single photoreceptor, this moving grating produces a temporal sinusoidal signal with a frequency  $f_t = vf_s$ . This allows us to rewrite equation 2.1 as

$$i(x, t) = I + \Delta I \sin(\omega_t t + \omega_s x) \quad (2.2)$$

where  $\omega_t = 2\pi f_t$  and  $\omega_s = 2\pi f_s$ . If two photoreceptors have an angular separation of  $\phi$ , then the signals measured by the photoreceptors can be expressed as

$$p_1(t) = |H(\omega_t)|\Delta I \sin\left(\omega_t t - \omega_s \frac{\phi}{2}\right) \quad (2.3)$$

$$p_2(t) = |H(\omega_t)|\Delta I \sin\left(\omega_t t + \omega_s \frac{\phi}{2}\right) \quad (2.4)$$

We introduce  $H(\omega_t)$  as the temporal frequency response of the photoreceptors. For simplicity, we ignore the phase contribution of  $H(\omega_t)$ , as it will be identical in  $p_1(t)$  and  $p_2(t)$ , and thus have no effect on perceived motion. We also assume that the photoreceptors have a high-pass behavior, which eliminates the DC component of illumination  $I$ . We model the photoreceptor response as

$$H(\omega_t) = \frac{Kj\omega_t\tau_H}{(j\omega_t\tau_H + 1)(j\omega_t\tau_{photo} + 1)} \quad (2.5)$$

where  $\tau_H$  is the time constant of the DC-blocking high-pass filter,  $\tau_{photo}$  is the time constant defining the photoreceptor bandwidth, and  $K$  is a constant of proportionality.

The delay required by the delay-and-correlate motion detector architecture is implemented using the phase lag inherent in a first-order low-pass filter. Low-pass filtering each photoreceptor signal yields

$$l_1(t) = \frac{|H(\omega_t)|\Delta I}{\sqrt{\tau^2\omega_t^2 + 1}} \sin\left(\omega_t t - \omega_s \frac{\phi}{2} - \tan^{-1} \tau\omega_t\right) \quad (2.6)$$

$$l_2(t) = \frac{|H(\omega_t)|\Delta I}{\sqrt{\tau^2\omega_t^2 + 1}} \sin\left(\omega_t t + \omega_s \frac{\phi}{2} - \tan^{-1} \tau\omega_t\right) \quad (2.7)$$

Correlation is accomplished by multiplying the phase-lagged signals with adjacent, nondelayed signals. The results are two ‘‘half-detector’’ responses:

$$m_1(t) = G[\cos(\omega_s \phi + P) - \cos(2\omega_t t - P)] \quad (2.8)$$

$$m_2(t) = G[\cos(\omega_s \phi - P) - \cos(2\omega_t t - P)] \quad (2.9)$$

where

$$G = \frac{(|H(\omega_t)|\Delta I)^2}{2\sqrt{\tau^2\omega_t^2 + 1}} \quad (2.10)$$

$$P = \tan^{-1} \tau\omega_t \quad (2.11)$$

Once these signals are subtracted in opponency, the final output becomes

$$o(t) = (\Delta I)^2 |H(\omega_t)|^2 \frac{\tau\omega_t}{\tau^2\omega_t^2 + 1} \sin \phi\omega_s \quad (2.12)$$

This describes the sensitivity of a Reichardt motion detector to a sinusoidal grating with a particular contrast, temporal frequency, and spatial frequency. Notice that the response is a separable function of these three parameters. We can rewrite this equation to make the dependency on the grating velocity  $v$  explicit:

$$o(t) = (\Delta I)^2 |H(\omega_s v)|^2 \frac{\tau\omega_s v}{\tau^2\omega_s^2 v^2 + 1} \sin \phi\omega_s \quad (2.13)$$

Although this response is direction selective (i.e., the sign of  $o[t]$  is equal to the sign of  $v$ ), it does not encode velocity independent of spatial frequency and contrast. Notice that the  $\sin \phi\omega_s$  term predicts spatial aliasing, as it becomes negative for  $1/2\phi < f_s < 1/\phi$ .

There are no time-dependent terms in this equation. This indicates a DC response to moving patterns. However, if the mean intensity of the image is not completely removed by prefiltering, or if the opponent subtraction is not perfectly balanced, oscillations at harmonics of the stimulus temporal frequency will be superimposed on the DC response. (For a complete analysis of the Reichardt motion detector in these non-ideal cases, see Egelhaaf, Borst, and Reichardt, 1989.) These oscillations may be reduced by using an array of Reichardt motion detectors (figure 2.1b) and summing their responses. This has the effect of integrating over different phases of the stimulus and canceling pattern-dependent oscillations, and has been proposed as a single model of tangential cell integration (Reichardt and Egelhaaf, 1988; Single and Borst, 1998).

## Hardware Implementations

Early attempts to implement the intensity-based Reichardt architecture in silicon used translinear, current-mode circuits (Andreou, Strohhahn, and Jenkins, 1991; Harrison and Koch, 1998). As we showed in the previous section, the response of these traditional Reichardt motion

sensors is affected strongly by contrast. Attempting to build contrast-independent Reichardt motion sensors, some have designed circuits that perform an initial binarization of the image based on temporal edges and then delay and correlate these digital signals (Moini et al., 1997; Jiang and Wu, 1999). These circuits would not be expected to perform well in noisy, low-contrast environments without additional image preprocessing. Another VLSI implementation involved continuous-level signal processing after the photoreceptors, but the final motion detector output was a binary value (Liu, 1997). Reichardt-inspired sensors have also been built in discrete hardware and used on mobile robots, although the particular implementation more closely resembled a feature-tracking, time-of-travel scheme (Pichon, Blanes, and Franceschini, 1989; Franceschini, Pichon, and Blanes, 1992).

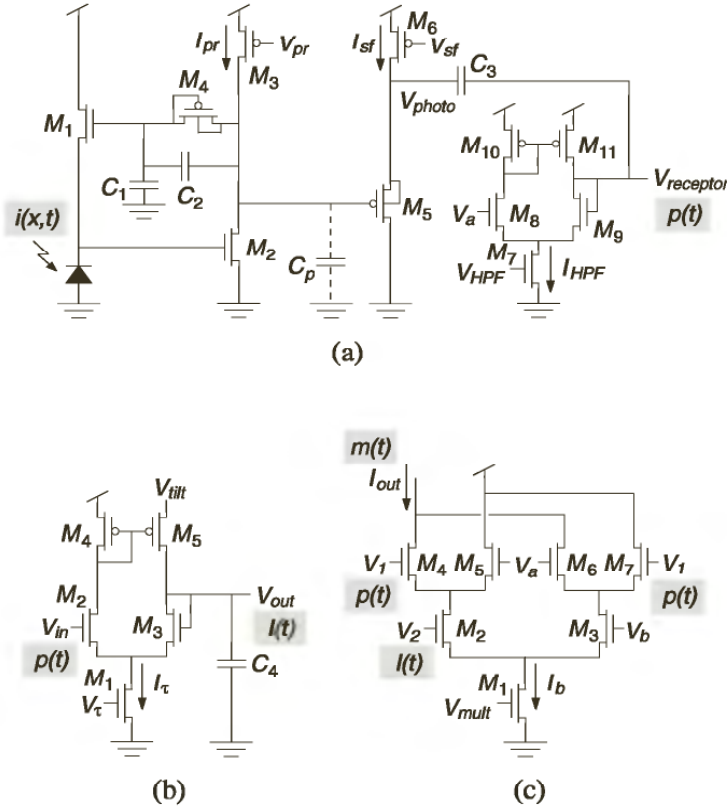
## CIRCUIT ARCHITECTURE

We implemented the Reichardt motion detector model in continuous-time VLSI circuitry. To the best of our knowledge, this is the closest approximation to this biological motion sensor that has been built.

We measure light intensity with an adaptive photoreceptor circuit developed by Delbrück and Mead (1996). This four-transistor circuit uses a substrate photodiode and source follower ( $M_1$ ) to convert incident light into a logarithmically encoded voltage (figure 2.2a). A high-gain amplifier ( $M_2$  and  $M_3$ ) and feedback network ( $C_1$  and  $C_2$ ) amplify the voltage signal by a factor of 18. The adaptive element ( $M_4$ ) acts as a nonlinear feedback element that conducts only if the voltage across it exceeds several hundred millivolts. This allows the photoreceptor to adapt to large changes in illumination. Thus we maintain a large dynamic range over a wide operating range. At low-bias current levels, the bandwidth of the photoreceptor is limited by the parasitic output capacitance  $C_p$ . (For a detailed discussion of this circuit, see Delbrück and Mead, 1996.)

The adaptive photoreceptor signal is sent to a  $g_m C$  high-pass filter (figure 2.2a). We use a source follower to provide a low-impedance driver, but in future designs we will leave this out and compensate for the increased output capacitance by increasing the photoreceptor bias current  $I_{pr}$ . We use a high-pass filter for two reasons. First, the AC coupling eliminates any systematic offsets caused by device variation in the adaptive photoreceptor. Second, by fixing the DC component of the signal to  $V_a$ , we can eliminate any common-mode effects later in the circuit.

The delay is accomplished with a first-order  $g_m C$  low-pass filter (figure 2.2b). The bias transistor in the circuit was made several times minimum size to improve time-constant matching across the chip. By



**Figure 2.2** Voltage-mode EMD subcircuits. Shaded labels indicate corresponding signals from figure 2.1a. (a) Adaptive photoreceptor ( $M_1$ – $M_4$ ,  $C_1$ – $C_2$ ) with source follower ( $M_5$ – $M_6$ ) and temporal highpass  $g_m C$  filter ( $M_7$ – $M_{11}$ ,  $C_3$ ) to remove the DC component of  $V_{photo}$ . (b) Temporal lowpass  $g_m C$  filter. This circuit’s phase lag acts as a delay. (c) Gilbert multiplier. This circuit multiplies delayed and nondelayed photoreceptor signals. The output is a current  $I_{out}$ , which allows for easy spatial summation.

operating this circuit at low current levels, we can achieve time constants useful for motion detection (10–100 msec) with reasonably sized capacitors (on the order of 1 pF).

Correlation is approximated by a Gilbert multiplier (figure 2.2c). The input  $V_2$  comes from the low-pass filter, and  $V_1$  comes from the high-pass-filtered photoreceptor from an adjacent pixel (figure 2.1b). The voltage  $V_a$  is the reference voltage used by the high-pass filter, and  $V_b$  is another DC bias voltage set a few tens of millivolts below  $V_a$ . We operate these field-effect transistors (FETs) in subthreshold, where their drain current  $I_d$ , ignoring channel-length modulation effects, is given by

$$I_d = I_0 e^{\kappa V_g / U_T} (e^{-V_s / U_T} - e^{-V_d / U_T}) \quad (2.14)$$

where  $I_0$  is a process-dependent constant,  $V_g$ ,  $V_s$ , and  $V_d$  are the gate, source, and drain voltages referenced to the bulk potential,  $\kappa$  is the gate efficiency factor (typically around 0.7), and  $U_T$  is the thermal voltage  $kT/q$  (Mead, 1989). Subthreshold FETs exhibit exponential behavior, much like the bipolar junction transistors (BJTs) with which the Gilbert

multiplier was originally built. We take a single-ended current-mode output from the circuit, which gives us

$$I_{out} = \frac{I_b}{2} + \frac{I_b}{2} \tanh\left[\frac{\kappa(V_1 - V_a)}{2U_T}\right] \tanh\left[\frac{\kappa(V_2 - V_b)}{2U_T}\right] \quad (2.15)$$

where  $I_b$  is the bias current. For small-signal inputs, this can be approximated as

$$I_{out} = \frac{I_b}{2} + \frac{\kappa^2 I_b}{8U_T^2} (V_1 - V_a)(V_2 - V_b) \quad (2.16)$$


For the multiplier to work properly, the common-mode voltage of the lower inputs ( $V_2$  and  $V_b$ ) must be lower than the common-mode voltage of the upper inputs ( $V_1$  and  $V_a$ ). Simulation results show that acceptable behavior is obtained with a difference of only 50 mV. In order to lower the DC level of the low-pass filter output, we lowered the source voltage of the output FET in the current mirror of the  $g_m C$  filter (see figure 2.2b). By placing the  $V_{tilt}$  bias a few tens of millivolts below  $V_{dd}$ , we lower the DC output level by  $(V_{dd} - V_{tilt})/\kappa$ . This source voltage “tilt” increases the time constant of the low-pass filter, but we can compensate by raising  $I_\tau$ . The difference in source voltages also creates an asymmetry in the up-going and down-going slew rates of the filter, but in practice this does not seem to have a significant effect on the overall circuit performance.

It can be shown from equation 2.15 that the circuit output saturates for differential inputs greater than about  $4U_T \approx 100$  mV. Rather than restrict our signals to this small linear region, we exploit the nonlinear behavior of the circuit to improve our motion detection algorithm. It has been shown that by adding saturating nonlinearities before the correlation stage, the contrast dependence of a Reichardt detector can be reduced (Egelhaaf and Borst, 1989).

Figure 2.3 shows the layout for one 1-D motion sensor, corresponding to the circuit element outlined in figure 2.1b. All experimental results shown below were measured from arrays of this circuit, which was fabricated in a 1.2  $\mu\text{m}$  double-poly, double-metal  $n$ -well CMOS process, yielding a pixel size of 61  $\mu\text{m}$  by 199  $\mu\text{m}$  with 32 transistors and 4 capacitors totaling 3.0 pF.

The outputs of all EMD pairs were summed to simulate the wide-field motion-sensitive neurons found in flies. We mounted a 2.6 mm lens over the chip, which gave the photoreceptors an angular spacing of 1.3° (similar to the 1–2° angular spacing observed in fly eyes), and a total field of view of 30° (much less than the fly’s eye, which sees almost an entire visual hemifield). The low-pass filter time constant was set to 50 msec, and the band-pass filters were set to pass frequencies between 0.5 Hz and 8 Hz. This circuit is characterized in detail in Harrison and Koch (2000).





**Figure 2.3** Voltage-mode motion detector layout. Cell measures  $61 \mu\text{m} \times 199 \mu\text{m}$  in a standard  $1.2 \mu\text{m}$  process with 32 transistors and 4 capacitors totaling 3.0 pF. In order to build a 2-D motion sensor, we need only add two more multiplier circuits, and some additional interpixel wiring.

## OPTOMOTOR CONTROL

Earlier, we introduced the optomotor response, where visual motion information is used as a feedback control signal to estimate and cancel self-rotation. This sensorimotor loop is perhaps the best-studied visually guided behavior of the fly. We will first describe optomotor experiments performed (by others) with flies and then describe real-time experiments where our VLSI sensor was compared directly against flies.

### Experiments Previously Performed on Flies

Warzecha and Egelhaaf (1996) recently characterized the optomotor behavior of the fly under closed-loop conditions. A female sheepfly (*Lucilia cuprina*, Calliphoridae) was rigidly attached to a meter that


measured yaw torque produced while the fly attempted to turn in response to visual stimuli (figure 2.4a), reducing the fly's behavior to a single degree of freedom. Vertical bars were presented to a large region of the fly's visual field, and could be drifted clockwise or counterclockwise. In closed-loop experiments, the fly's yaw torque was measured in real time and scaled by a constant gain term to yield angular velocity. This simulates the observed dominance of air friction in determining the instantaneous angular velocity in flies (Reichardt and Poggio, 1976). The fly's simulated angular velocity was subtracted from the angular velocity imposed by the experimenter. The resulting signal was used to control the drift rate of the visual stimulus. This simulated free-flight conditions, and allowed evaluation of the optomotor system performance.

The imposed motion schedule consisted of 3.75 sec of zero imposed motion, then 7.5 sec of clockwise rotation at  $44^\circ/\text{sec}$ . Figure 2.5a shows the torque data and resulting stimulus position for an individual trial. Figure 2.5b shows the averaged data over 139 trials in a total of five animals. (See Warzecha and Egelhaaf, 1996 for details on the experimental protocol.)

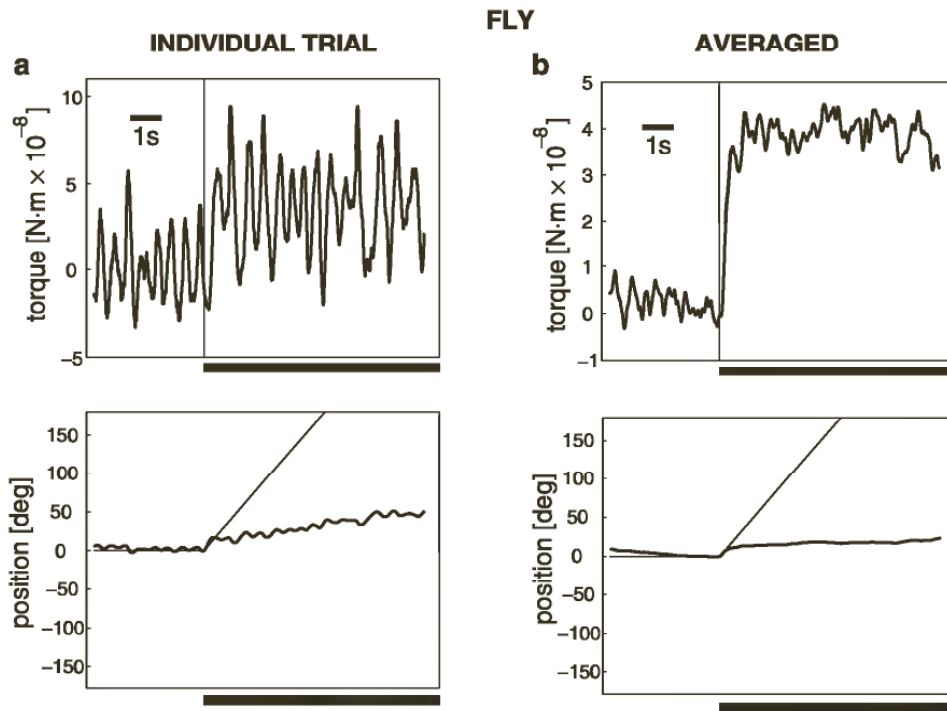
The fly is able to stabilize its flight and cancel out most of the imposed motion. Simulation results suggest that the nonmonotonic temporal frequency response of Reichardt motion detectors results in greater stability for the optomotor control system (Warzecha and Egelhaaf, 1996). The individual trials show an oscillatory component to the torque response around 2 Hz. This oscillation is not phase-locked to the stimulus because it is not present in the average torque trace. Oscillations are not observed under open-loop conditions, suggesting they arise from optomotor feedback (Geiger and Poggio, 1981; Warzecha and Egelhaaf, 1996). Notice that despite the large amplitude of the torque oscillations, the position trace is not dominated by this effect. This fluctuation amplitude, in terms of number of photoreceptors, is close to the amplitude observed in human microsaccades (Warzecha and Egelhaaf, 1996). Poggio and colleagues observed similar oscillations in closed-loop experiments and proposed that they arose from the 60–75 msec synaptic delay inherent in the fly visual system (Geiger and Poggio, 1981; Poggio and Reichardt, 1981).

### **Duplicating Experiments with the Silicon System**

We were able to directly replicate these experiments with our silicon analog of the optomotor system (figure 2.4b). To provide visual stimulation, we used an LED display with a 200 Hz refresh rate, which is currently being used to test flies in closed-loop experiments. The stimulation time schedule was identical to the fly experiments, but an



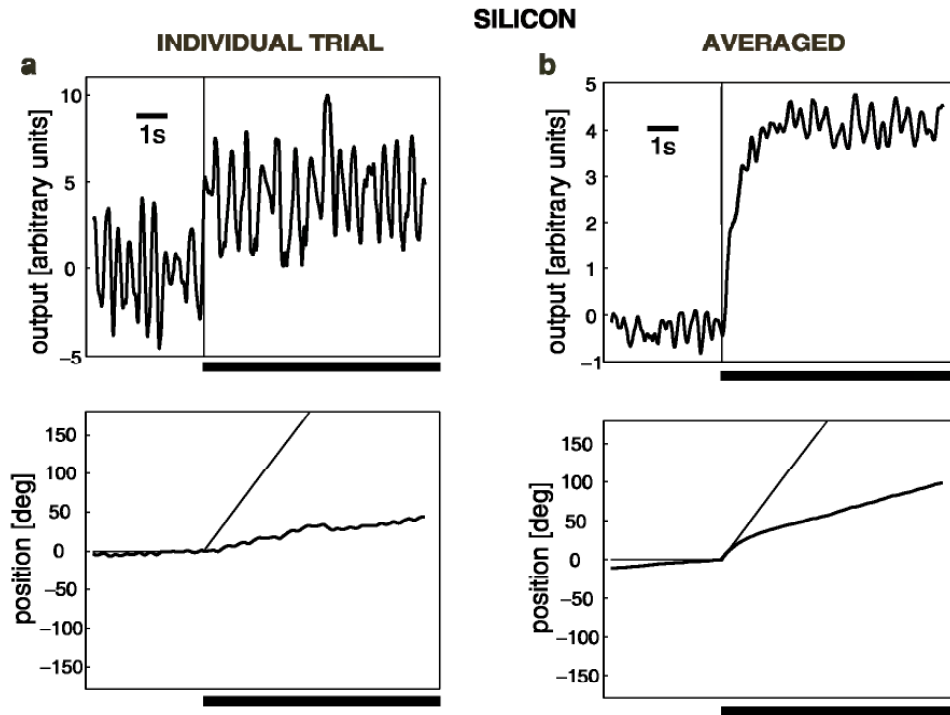
**Figure 2.4** Experimental methodology. (a) Setup used by Warzecha and Egelhaaf to measure the closed-loop torque response of the sheepfly *Lucilia* (Warzecha and Egelhaaf, 1996). The torque meter output is scaled to produce a measure of what the fly's self-motion would be if it were free to rotate. This self-motion is subtracted from the imposed motion to determine the pattern motion, creating the illusion of free flight in a room with distant walls. (Only rotation, not translation, is simulated.) (b) Setup used to replicate the closed-loop experiments with the silicon model. The output voltage from the circuit is used in place of the torque meter output voltage. The rest of the system is identical to (a).



**Figure 2.5** Fly's optomotor behavior. (a) Torque (top) and angular position (bottom) vs. time for an individual closed-loop trial with a fly. Dark horizontal bar indicates experimenter-imposed rotation. Thin lines on position trace indicate position in the open-loop case. Most of the imposed rotation is canceled out by the optomotor control system. Since the position is proportional to the integral of the torque (see text for details), large torque oscillations do not cause large position oscillations. (b) Averaged torque response and angular position trace for multiple trials ( $N = 139$ , 5 flies). The fly showed an average drift of 9.4% of the open-loop drift velocity, with position fluctuations of  $7.8^\circ$  (standard deviation) about this drift. Data in (a) and (b) redrawn from Warzecha and Egelhaaf, 1996.

angular velocity of  $50^\circ/\text{sec}$  was used. Our chip had a much smaller field of view ( $10^\circ$ ) than the fly, so we set the stimulus distance such that the EMD array saw approximately one wavelength of the pattern. The output signal from the silicon model of the HS cell was passed through an off-chip first-order low-pass filter with a time constant of 680 msec, modeling the behavior of the thoracic motor centers (Egelhaaf, 1987; Wolf and Heisenberg, 1990; Warzecha and Egelhaaf, 1996). The filtered output of the chip was treated exactly like the signal from the torque meter in the fly experiments, and closed-loop experiments were run in real time. Figure 2.6a shows torque and position data from the chip for an individual trial, and figure 2.6b shows the averaged response over 100 trials.

The silicon system shows the same ability to greatly cancel the imposed motion. The fly showed an average drift of 9.4% of the open-loop drift velocity, with position fluctuations of  $7.8^\circ$  standard deviation



**Figure 2.6** The optomotor behavior of our silicon system. (a) Chip output signal (analogous to torque) and position vs. time for the silicon system in an individual trial. (b) Averaged torque response and angular position trace for multiple trials ( $N = 100$ , 1 chip). The chip showed an average drift of 22% of the open-loop drift velocity, with position fluctuations of  $6.2^\circ$  (s.d.) about this drift.

(SD) about this drift. The chip showed an average drift of 22% of the open-loop drift velocity, with position fluctuations of  $6.2^\circ$  SD about this drift. Also, we observe the same 2 Hz oscillations in the individual trials. Because we did not build any explicit delay into our system, this demonstrates that the phase lags and nonlinearities in this simple model are sufficient to produce oscillations, even in the absence of additional synaptic delays.

We believe this hardware modeling approach will prove increasingly valuable in the future, as biological models of the neural circuitry underlying more complex and sophisticated behaviors arise. To simulate a sensorimotor system in software, one must construct *two* models: a model of the biological system, and a model of the world. The physical environment is an essential element in a sensorimotor feedback loop, so this world model must increase in detail as we study more advanced behaviors. Because animals interact with their three-dimensional environment in very dynamic ways, it may not be long before software simulations of sensorimotor systems require more computational resources to model the world than to model the neural circuitry of interest.

## APPLICATION TO AUTONOMOUS VEHICLE CONTROL


Optic-flow patterns produced by self-motion are one of the richest sources of navigation information available to a mobile creature (Gibson, 1950). As an animal moves through its environment, images of the outside world move across its retina in predictable ways. Objects being approached grow larger; objects left behind grow smaller. When moving forward, images of nearby objects move across the retina faster than images of distant objects. If a creature rotates in place, the entire visual scene moves across its retina at a rate that is independent of object distance. Much information can be gained from patterns of visual motion, even if no explicit object recognition is performed (Duchon, Kaelbling, and Warren, 1998). Indeed, motion parallax information is immune to camouflage that can defeat even the most sophisticated static pattern recognition scheme when object and background have similar textures. Humans have no difficulty detecting the structure of randomly patterned objects against identically patterned backgrounds from motion cues alone.

Using egomotion-induced optic flow for robot navigation is a computationally demanding sensory task. By its very nature, it must be done in real time. Most object recognition tasks are performed on static images, and often one can tolerate latencies of several seconds. But optic flow is available only while the robot is moving, and relevant information must be extracted in real time and fed back to the motor control system to steer the robot in the right direction. The rate of computation needed depends on the rate of robot motion, but typical real-world situations require times on the order of tens or hundreds of milliseconds.

Optic flow is also computationally demanding because, like other early vision tasks, it involves operations that must be performed identically on every pixel of an image. Local estimates of motion must be laboriously computed before the overall pattern is analyzed. This is a task that is ideally suited for parallel computation.

Measuring optic flow also involves large amounts of data. While audition involves one time-varying signal (two in the case of binaural audition), vision involves many time-varying signals. Rapid navigation requires many frames to be analyzed each second. This can tax even the most sophisticated microprocessor because it must deal with all the signals at once. If we divide the job to many processors, each dealing with one pixel and communicating with its immediate neighbors, the task becomes much easier.

In the past decade, researchers have been endowing mobile robots with biologically inspired (more specifically, insect-inspired) visual systems (Franceschini, Pichon, and Blanes, 1992; Srinivasan, Chahl, and Zhang, 1997; Huber, 1997; Lewis, 1998; Huber, Franz, and Bülthoff, 1999). These efforts have yielded promising results, but many problems



**Figure 2.7** Schematic of our optomotor system. A motion sensor chip is mounted facing forward on a two-wheeled robotic platform. Forward-facing motion sensors are largely blind to optic flow produced by forward translation, so we only measure rotation. The chip’s wide-field output is used as an estimate of self rotation, then lowpass filtered ( $\tau = 750$  ms) to stabilize the control loop. This signal is added to one motor and subtracted from the other, producing a compensatory rotation. A constant motor bias produces forward translatory motion.


still exist. Vision is a computationally intensive task, so powerful hardware is required to operate in real time. From an algorithmic viewpoint, the structure of visual scenes is often very complex, and it can be difficult to extract relevant information robustly.

As mentioned above, the fly uses visual motion information to stabilize its flight. Mismatch of body components or environmental disturbances may impart rotation on the animal, but sensory feedback is used to produce compensatory torque responses. This sensorimotor feedback is known as the optomotor system, and is one of the best-studied behaviors of the fly (Götz, 1975; Warzecha and Egelhaaf, 1996).

## Hardware Implementation

We constructed a hardware implementation of the optomotor system using a two-wheeled vehicle. We chose to build a physical motor system so we could evaluate our sensor’s performance in the real world. Figure 2.7 shows a schematic of the system. Our wide-field motion sensor estimates self-rotation, and this signal is used to produce a compensatory rotation by the drive motors.

We constructed a simple robotic platform on which we mounted the wide-field motion sensor (figure 2.8). The robot had two large wheels driven independently by two DC motors, and a free-turning wheel in the back to maintain balance. Each drive motor was controlled with a pulse-width modulation circuit that varied the duty cycle of a constant-amplitude square-wave voltage. By changing the duty cycle of the waveform, each motor could be driven at any speed up to a maximum.



**Figure 2.8** Photograph of the optomotor system. The lens is covering the motion sensor chip. Additional off-chip electronics have been constructed to drive the dc motors. The back wheels turn freely, and merely prevent the robot from falling over. The robot measures 13 cm  $\times$  19 cm  $\times$  22 cm, and is powered by on-board batteries.

If the motors were driven at different speeds, the robot would drive in a curved trajectory.

A large asymmetry was introduced into the robot's mechanics by connecting the left and right motors to their respective wheels with different gear ratios. The left motor was connected to the left wheel with a 1:5 gear ratio, while the right motor was connected to the right wheel with a 1:1 gear ratio. This caused the robot to drive in tight circles if both motors were driven at the same speed. This asymmetry was made extreme for the sake of experiment, but perfect symmetry is impossible to achieve in any physical robot. While two actuators may match perfectly in simulation, they will never match when built and tested in the real world. This difficulty is especially pronounced in outdoor terrain, where wheels or feet may slip in sand or mud. Legged robots are especially prone to walking in curved lines due to foot slip or terrain differences, even if they have been designed and constructed with high precision.

When open-loop control falls short, we must introduce sensory feedback to further improve performance. Optic flow information has the potential to guide a robot in a straight path, because any deviation



involves a yaw rotation, however slight. If yaw rotation can be estimated from optic flow reliably, we can use this as an error signal in a negative feedback loop in which the motors execute a compensatory rotation to null the sensory error signal.


We constructed a feedback loop of this type using our VLSI wide-field motion sensor. The sensor was mounted facing forward on the robot, oriented so it was sensitive to horizontal motion. We oriented the sensor facing straight ahead because translatory motion by the robot produces little optic flow in the direction of travel, while rotatory (yaw) motion produces uniform optic flow around the visual field parallel to the ground. Thus the optic flow in the forward region of the visual field will be dominated by the rotatory component. The hoverfly *Syrirtta pipiens* uses this strategy to stabilize its flight. When moving forward, the animal uses optic flow from the forward region of the visual field to estimate self rotation. This agile creature is also capable of flying sideways, and when doing so, it uses optic flow from the lateral visual fields to estimate self-rotation (Collett, 1980). Presumably, it is attempting to measure optic flow in the regions least contaminated with optic flow produced by its own translation.

The output of our motion sensor was a continuous, time-varying voltage. This signal was filtered by a first-order low-pass filter with a time constant of 750 msec. This is a simple model of the relationship between the output of a wide-field motion-sensitive neuron in the fly and the torque response produced by the wings (Egelhaaf, 1987; Warzecha and Egelhaaf, 1996). The filtered output of the motion sensor was added to the left motor command and subtracted from the right motor command (see figure 2.7). This has the effect of adding a rotatory component to the robot's trajectory. In the absence of visual feedback, both motors turned at the same rate (so one wheel turns five times faster than the other). Visual feedback slowed one wheel and sped up the other.

## Robot Experiments

Experiments were performed indoors in our laboratory, but the visual nature of the room was not altered in any way to accommodate the motion sensor. The room was a typical cluttered laboratory environment with many shady areas under tables. The robot's position was recorded 10–20 times per second with a magnetic field tracking system that returned location and orientation in three dimensions (Polhemus, Colchester, VT). The scale of experiments was limited by the range of this system, approximately a 70 cm by 140 cm area for highest accuracy.

The optic flow feedback proved capable of nearly eliminating the effect of physical asymmetry. Figure 2.9 shows one trial without visual feedback. The line shows the robot's path, and the circle indicates the




**Figure 2.9** Robot path with no sensory feedback. With the motion sensor disabled, the robot turns in circles due to the asymmetry in its mechanics. The circle denotes the ending location of the robot.

ending position. The robot is turning in tight circles. Figure 2.10 shows ten trials where visual feedback has been enabled. In general, the robot travels in straight lines. We purposely started the robot at different orientations to demonstrate that the sensor works well for general visual scenes around a room. When moving in straight lines, the robot traveled at a speed of approximately 20 cm/sec. Objects and walls were typically 0.2 to 1.5 meters away from the robot, depending on the direction.

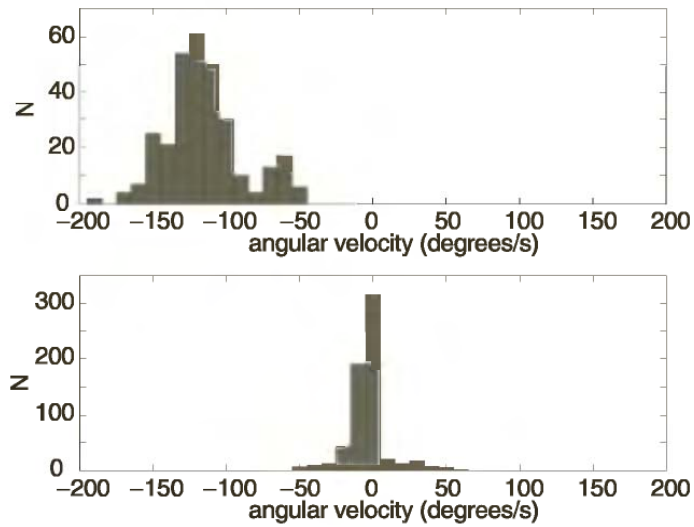
The angular velocity of the robot (yaw rate) was computed along each path by differentiating the robot's heading as recorded by the tracking system. Figure 2.11a shows a histogram of angular velocities for the trials without feedback and figure 2.11b shows all ten trials with visual feedback. The mean angular velocity in the open-loop case is  $-116^\circ/\text{sec}$ , while for the closed-loop case this decreased to  $-3.7^\circ/\text{sec}$ , an improvement by a factor of 31.

Occasionally, the feedback did fail to keep the course straight. A  $45^\circ$  turn is visible in figure 2.10, most likely caused by the sensor being oriented toward a relatively featureless part of the room, where no motion information is available. A larger field of view would reduce the likelihood of such occurrences. Also, the magnitude of the error depends on the degree of asymmetry in the gear ratios. In a more realistic situation with higher open-loop precision, it is likely that large closed-loop errors would be rare.

Our sensor is small and extremely low power, making it easily adaptable to mobile robot applications. While the power budget on a robot is usually dominated by motors, traditional CCD imagers consume



**Figure 2.10** Robot path with sensory feedback. Ten trials are shown where the motion sensor provided optomotor control to straighten the course of the robot despite a 5:1 mechanical asymmetry. Circles denote the ending location of the robot in each trial. The robot was exposed to different visual scenes during the trials.



**Figure 2.11** Histogram of angular velocities. (a) No visual feedback. The turning behavior of the robot is obvious. The mean angular velocity was  $-116$  deg/s. (b) Compilation of all ten trials with visual feedback. The mean angular velocity was greatly reduced to  $-3.7$  deg/s.

significant amounts of power, and digital microprocessors or digital signal processors (DSPs) capable of processing real-time video consume even more. On the Sojourner rover of the recent Pathfinder Mars mission, the CCD imagers alone consumed 0.75 W—5% of the total power budget at peak solar cell output. The CPU system consumed an additional 24%, and much of the CPU's time was devoted to processing static images while the rover was not moving (Matthies et al., 1995). By comparison, our EMD array consumed less than 5  $\mu$ W of power. Traditional imaging and image processing is expensive in terms of time, size, and power. Biologically inspired analog VLSI approaches to this problem can bring down the cost and make robot vision more practical.

## REFERENCES

- Adelson, E. H., and Bergen, J. R. (1985). Spatiotemporal energy models for the perception of motion. *J. Optical Soc. Am. [A]* 2: 284–299.
- Andreou, A. G., Strohbehm, K., and Jenkins, R. E. (1991). Silicon retina for motion computation. In *Proceedings of the 1991 IEEE International Symposium on Circuits and Systems*. pp. 1373–1376.
- Autrum, H. (1958). Electrophysiological analysis of the visual systems in insects. *Exp. Cell Res. Supplement* 5: 426–439.
- Borst, A. (1990). How do flies land? *BioScience* 40: 292–299.
- Borst, A., and Bahde, S. (1988). Spatio-temporal integration of motion. *Naturwissenschaften* 75: 265–267.
- Borst, A., and Egelhaaf, M. (1989). Principles of visual motion detection. *Trends Neurosci.* 12: 297–306.
- Borst, A., and Egelhaaf, M. (1990). Direction selectivity of blowfly motion-sensitive neurons is computed in a two-stage process. *Proc. Natl. Acad. Sci. U.S.A.* 87: 9363–9367.
- Clifford, C. W. G., Ibbotson, M. R., and Langley, K. (1997). An adaptive Reichardt detector model of motion adaptation in insects and mammals. *Vis. Neurosci.* 14: 741–749.
- Collett, T. S. (1980). Some operating rules for the optomotor system of a hoverfly during voluntary flight. *J. Comp. Physiol. [A]* 138: 271–282.
- Delbrück, T., and Mead, C. A. (1996). Analog VLSI phototransduction by continuous-time, adaptive, logarithmic photoreceptor circuits. CNS Memo 30, California Institute of Technology.
- DeVoe, R. D. (1980). Movement sensitivities of cells in the fly's medulla. *J. Comp. Physiol. [A]* 138: 93–119.
- DeVoe, R. D., and Ockleford, E. M. (1976). Intracellular responses from cells in the medulla of the fly, *Calliphora erythrocephala*. *Biol. Cybern.* 23: 13–24.
- Duchon, A. P., Kaelbling, L. P., and Warren, W. H. (1998). Ecological robotics. *Adaptive Behav.* 6: 473–507.
- Egelhaaf, M. (1985). On the neuronal basis of figure-ground discrimination by relative motion in the visual system of the fly: II. Figure-detection cells, a new class of visual interneurons. *Biol. Cybern.* 52: 195–209.

- Egelhaaf, M. (1987). Dynamic properties of two control-systems underlying visually guided turning in house-flies. *J. Comp. Physiol. [A]* 161: 777–783.
- Egelhaaf, M., and Borst, A. (1989). Transient and steady-state response properties of movement detectors. *J. Optical Soc. Am. [A]* 6: 116–127.
- Egelhaaf, M., and Borst, A. (1993). A look into the cockpit of the fly: Visual orientation, algorithms, and identified neurons. *J. Neurosci.* 13: 4563–4574.
- Egelhaaf, M., Borst, A., and Reichardt, W. (1989). Computational structure of a biological motion-detection system as revealed by local detector analysis in the fly's nervous system. *J. Optical Soc. Am. [A]* 6: 1070–1087.
- Egelhaaf, M., Hausen, K., Reichardt, W., and Wehrhahn, C. (1988). Visual course control in flies relies on neuronal computation of object and background motion. *Trends Neurosci.* 11: 351–358.
- Franceschini, N., Pichon, J. M., and Blanes, C. (1992). From insect vision to robot vision. *Philos. Trans. R. Soc. Lond. B Biol. Sci.* 337: 283–294.
- Geiger, G., and Nüssel, D. R. (1981). Visual orientation behaviour of flies after selective laser beam ablation of interneurons. *Nature* 293: 398–399.
- Geiger, G., and Nüssel, D. R. (1982). Visual processing of moving single objects and wide-field patterns in flies: Behavioural analysis after laser-surgical removal of interneurons. *Biol. Cybern.* 44: 141–149.
- Geiger, G., and Poggio, T. (1981). Asymptotic oscillations in the tracking behaviour of the fly *Musca domestica*. *Biol. Cybern.* 41: 197–201.
- Gibson, J. J. (1950). *The Perception of the Visual World*. Boston: Houghton-Mifflin.
- Götz, K. G. (1975). The optomotor equilibrium of the *Drosophila* navigation system. *J. Comp. Physiol. [A]* 99: 187–210.
- Haag, J., and Borst, A. (1997). Encoding of visual motion information and reliability in spiking and graded potential neurons. *J. Neurosci.* 17: 4809–4819.
- Harrison, R. R., and Koch, C. (1998). An analog (VLSI) model of the fly elementary motion detector. In M. I. Jordan, M. J. Kearns, and S. A. Solla (eds.), *Advances in Neural Information Processing Systems 10*. Cambridge: MIT Press, pp. 880–886.
- Harrison, R. R., and Koch, C. (2000). A robust analog VLSI Reichardt motion sensor. *Analog Integrated Circuits and Signal Processing* (in press).
- Hassenstein, B., and Reichardt, W. (1956). Systemtheoretische Analyse der Zeit-, Reihenfolgen-, und Vorzeichenauswertung bei der Bewegungsperzeption des Rüsselkäfers *Chlorophanus*. *Z. Naturforsch.* 11b: 513–524.
- Hausen, K. (1982a). Motion sensitive interneurons in the optomotor system of the fly: I. The horizontal cells: Structure and signals. *Biol. Cybern.* 45: 143–156.
- Hausen, K. (1982b). Motion sensitive interneurons in the optomotor system of the fly: II. The horizontal cells: Receptive field organization and response characteristics. *Biol. Cybern.* 46: 67–79.
- Hausen, K. (1984). The lobula-complex of the fly: Structure, function, and significance in behaviour. In M. A. Ali (ed.), *Photoreception and Vision in Invertebrates*. New York: Plenum, pp. 523–559.
- Hausen, K., and Egelhaaf, M. (1989). Neural mechanisms of visual course control in insects. In D. G. Stavenga and R. C. Hardie (eds.), *Facets of Vision*. Berlin: Springer-Verlag.

- Hausen, K., and Wehrhahn, C. (1990). Neural circuits mediating visual flight control in flies: II. Separation of two control-systems by microsurgical brain-lesions. *J. Neurosci.* 10: 351–360.
- Heisenberg, M., and Buchner, E. (1977). The rôle of retinula cell types in visual behavior of *Drosophila melanogaster*. *J. Comp. Physiol. [A]* 117: 127–162.
- Hengstenberg, R. (1982). Common visual response properties of giant vertical cells in the lobula plate of the blowfly *Calliphora*. *J. Comp. Physiol. [A]* 149: 179–193.
- Huber, S. A. (1997). Studies of the Visual Orientation Behavior in Flies Using the Artificial Life Approach. Ph.D. thesis, Universität Tübingen, Germany.
- Huber, S. A., Franz, M. O., and Bülthoff, H. H. (1999). On robots and flies: Modeling the visual orientation behavior of flies. *Robotics Autonomous Syst.* 29: 227–242.
- Jiang, H.-C., and Wu, C.-Y. (1999). A 2-D velocity- and direction-selective sensor with BJT-based silicon retina and temporal zero-crossing detector. *IEEE J. Solid State Circuits* 34: 241–247.
- Kimmerle, B., Warzecha, A.-K., and Egelhaaf, M. (1997). Object detection in the fly during simulated translatory flight. *J. Comp. Physiol. [A]* 181: 247–255.
- Krapp, H. G., and Hengstenberg, R. (1996). Estimation of self-motion by optic flow processing in single visual interneurons. *Nature* 384: 463–466.
- Land, M. F. (1997). Visual acuity in insects. *Annu. Rev. Entomol.* 42: 147–177.
- Laughlin, S. B. (1994). Matching coding, circuits, cells, and molecules to signals—general principles of retinal design in the fly’s eye. *Prog. Ret. Eye Res.* 13: 165–196.
- Lewis, M. A. (1998). Visual navigation in a robot using zig-zag behavior. In M. I. Jordan, M. J. Kearns, and S. A. Solla (eds.), *Advances in Neural Information Processing Systems 10*. Cambridge: MIT Press, pp. 822–828.
- Liu, S.-C. (1997). Neuromorphic Models of Visual and Motion Processing in the Fly Visual System. Ph.D. thesis, California Institute of Technology.
- Matthies, L., Gat, E., Harrison, R., Wilcox, B., Volpe, R., and Litwin, T. (1995). Mars microrover navigation: Performance evaluation and enhancement. *Autonomous Robots 2*: 291–311.
- Mead, C. (1989). *Analog VLSI and Neural Systems*. Menlo Park, Calif.: Addison-Wesley.
- Moini, A., Bouzerdoum, A., Eshraghian, K., Yakovleff, A., Nguyen, X. T., Blanksby, A., Beare, R., Abbott, D., and Bogner, R. E. (1997). An insect vision-based motion detection chip. *IEEE J. Solid-State Circuits* 32: 279–284.
- Pichon, J.-M., Blancs, C., and Franceschini, N. (1989). Visual guidance of a mobile robot equipped with a network of self-motion sensors. In *SPIE Mobile Robots IV*, vol. 1195, pp. 44–53.
- Poggio, T., and Reichardt, W. (1981). Visual fixation and tracking by flies: Mathematical properties of simple control systems. *Biol. Cybern.* 40: 101–112.
- Reichardt, W., and Egelhaaf, M. (1988). Properties of individual movement detectors as derived from behavioural experiments on the visual system of the fly. *Biol. Cybern.* 58: 287–294.
- Reichardt, W., and Poggio, T. (1976). Visual control of orientation behaviour in the fly. *Q. Rev. Biophys.* 9: 311–375.
- Single, S., and Borst, A. (1998). Dendritic integration and its role in computing image velocity. *Science* 281: 1848–1850.

- Srinivasan, M. V., Chahl, J. S., and Zhang, S. W. (1997). Robot navigation by visual dead-reckoning: Inspiration from insects. *Int. J. Pattern Recognition Artif. Intell.* 11: 35–47.
- Strausfeld, N. J. (1976). *Atlas of an Insect Brain*. Berlin: Springer.
- van Santen, J. P. H., and Sperling, G. (1985). Elaborated Reichardt detectors. *J. Optical Soc. Am. [A]* 2: 300–321.
- Wagner, H. (1986). Flight performance and visual control of flight of the free-flying housefly *Musca domestica* L.: II. Pursuit of targets. *Philos. Trans. R. Soc. Lond. B Biol. Sci.* 312: 553–579.
- Wandell, B. A. (1995). *Foundations of Vision*. Sunderland, Mass.: Sinauer Associates.
- Warzecha, A.-K., and Egelhaaf, M. (1996). Intrinsic properties of biological motion detectors prevent the optomotor control system from getting unstable. *Philos. Trans. R. Soc. Lond. B Biol. Sci.* 351: 1579–1591.
- Watson, A. B., and Ahumada, A. J. (1985). Model of human visual-motion sensing. *J. Optical Soc. Am. [A]* 2: 322–342.
- Weckström, M., Juusola, M., and Laughlin, S. B. (1992). Presynaptic enhancement of signal transients in photoreceptor terminals in the compound eye. *Proc. R. Soc. Lond. [B]* 250: 83–89.
- Wolf, R., and Heisenberg, M. (1990). Visual control of straight flight in *Drosophila melanogaster*. *J. Comp. Physiol. [A]* 167: 269–283.
- Zanker, J. M. (1990). On the directional sensitivity of motion detectors. *Biol. Cybern.* 62: 177–183.

# 3

## Visual Servo System Based on a Biologically Inspired Scanning Sensor

Stéphane Viollet and Nicolas Franceschini

The present study was carried out in the context of our previous studies using a biologically inspired minimalistic robotic approach, which were based on the idea of reconstructing natural visuomotor processes in order to better understand them [1–5]. Here, we show that an elementary visual sensor composed of only two pixels is able to accurately detect contrasting features within its visual field, provided it uses a particular scanning mode, such as that recently discovered in the compound eye of flies [6]. This low-cost, low-weight, low-complexity scanning sensor was found to deliver an output that is:

- Quasi linear with respect to the angular position of a contrasting object in its field of view
- Largely invariant with respect to the distance to the object and its level of contrast

The most novel feature of this visual sensor, which differs from other systems based on *scanning at constant angular speed* [4] and *pulse scanning* [7], is that it scans the visual field at an angular speed that varies gradually with time (variable-speed scanning). The sensor's output then becomes a graded function of the position of a contrast feature within its visual field, with the interesting consequence that if this sensor is incorporated into a visuomotor control loop, it is able to stabilize a robot such as a micro air vehicle with respect to the environmental features.

In the following two sections, we simulate the pure rotation at a variable speed of a pair of pixels in front of a stationary contrast feature (a light-to-dark edge) and show that the processing of the apparent motion yields a signal that gradually varies with the angular position of the edge in the visual field. In the section entitled Variable-Speed Scanning and Visual Position Sensing, we describe the structure of a miniature scanning sensor coupled with a local motion detector (LMD) based on the fly compound eye [8]. In the section entitled Visual Control of a Miniature, we describe the structure of a miniature Twin-Propeller Aircraft, aircraft equipped with the scanning sensor and the



implementation of a visual feedback loop designed to stabilize this model aircraft in yaw. The scanning sensor picks up an error signal, which is used to drive the two propellers differentially, causing the “seeing aircraft” to reorient and adjust its line of sight so as to remain constantly locked onto a nearby or distant contrasting target.

Contrary to the kind of yaw stabilization that can be achieved with a gyroscope having similar low-weight and low-cost characteristics, that described here is not only stiff but also devoid of drift. Some examples are given below, where the (tethered) aircraft either tracks a slowly moving dark edge or maintains fixation onto a stationary target for 17 minutes.

The scanning visual sensor and the visuomotor control loop we have developed are appropriate for:

- Detecting low-speed relative motion between a craft and its environment
- Rejecting environmental disturbances by maintaining visual fixation onto a nearby or distant target regardless of its contrast and distance within a given range
- And controlling the attitude of a craft optically with respect to environmental features

## SCANNING AT A VARIABLE ANGULAR SPEED

### Photoreceptor Model

The angular sensitivity of an insect lens–photoreceptor combination is usually approximated by a bell-shaped function [9]. In the present case, we use a Gaussian function, the bottom part of which is truncated below the threshold  $s_0$ .

The angular sensitivity  $S(\varphi)$  of a photosensor  $i$  is therefore given by:

$$S(\varphi) = \left( \frac{\text{abs}(G_i(\varphi) - s_0) + (G_i(\varphi) - s_0)}{2} \right), \quad (3.1)$$

with

$$G_i(\varphi) = \exp \left[ -\frac{\varphi^2}{2\sigma^2} \right] \quad (3.2)$$

The main parameters that characterize  $S(\varphi)$  are as follows:

- $l.v$ : total angular width of the visual field
- $\Delta\rho$ : angular width at half height

Both parameters can be expressed with respect to  $\sigma$  as follows:

$$\Delta\rho = 2\sigma \left[ 2 \ln \left( \frac{2}{1 + s_0} \right) \right]^{1/2} \quad (3.3)$$

$$Lv = 2\sigma \left[ 2 \ln \left( \frac{1}{s_0} \right) \right]^{1/2} \quad (3.4)$$

### Variable Angular Speed

As sketched in figure 3.1a, we have simulated the concerted rotation of two photosensors, 1 and 2, separated by a constant angle  $\Delta\varphi$ , called the interreceptor angle, in front of an elementary panorama consisting of a single segment AB. The pair of photosensors is assumed here to rotate clockwise at an angular speed  $\Omega$ , which decays exponentially with time. The angle  $\psi$  is given by (figure 3.1b):

$$\psi(t) = A \left( 1 - \exp \left( -\frac{t}{\tau} \right) \right) \quad (3.5)$$

Hence the angular speed  $\Omega(t)$  is:

$$\Omega(t) = \frac{A}{\tau} \exp \left( -\frac{t}{\tau} \right) \quad (3.6)$$

The position of the segment is defined by the points  $A(x_1, y_1)$  and  $B(x_2, y_2)$  with  $y_1 = y_2$ .

At any time, the output  $R_i(t)$  from photosensor  $i$  is obtained by integrating the angular sensitivity  $S(\varphi)$  weighted by the intensity  $I(\varphi)$ :

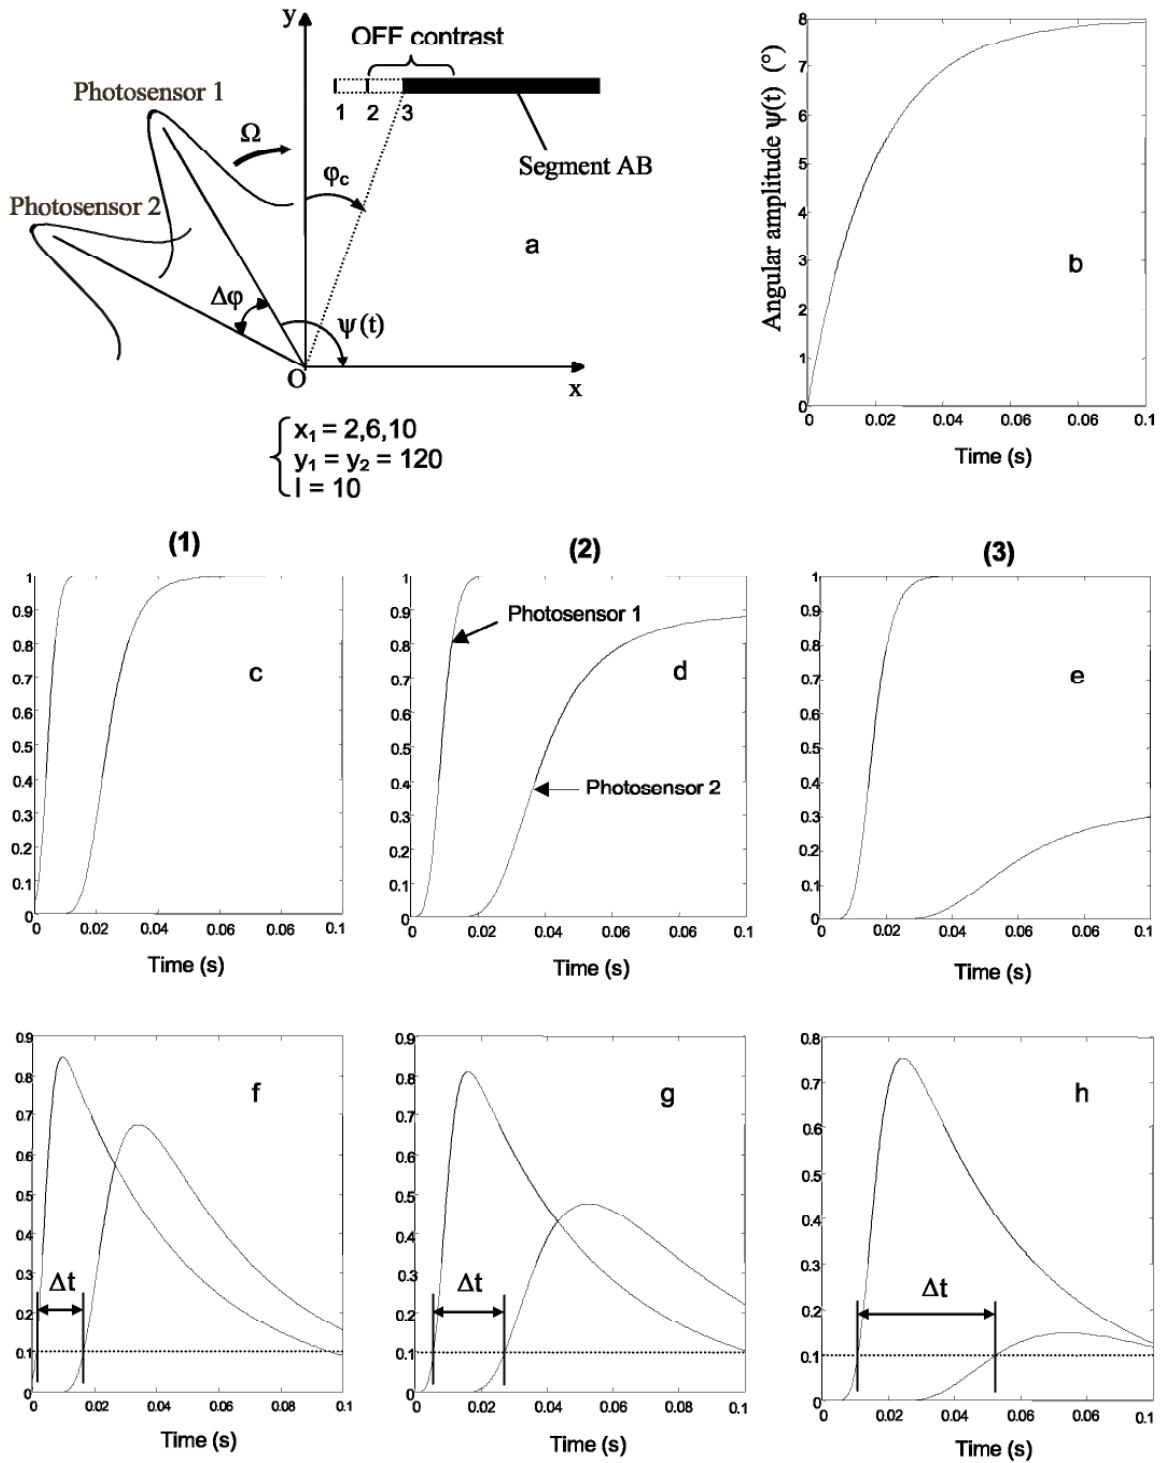
$$R_i(t) = \int_{-Lv/2}^{Lv/2} I(\varphi) S(\varphi) d\varphi \quad (3.7)$$

### Processing the Photoreceptor Signals

To calculate the output signal from each photoreceptor, we used a discrete version of equation 3.7. The signal sampled,  $R_i(kT)$ , with  $T$  in seconds, results from a convolution of the light intensity  $I$  with the Gaussian mask  $S(kTs)$  ( $Ts$  in degrees). In other words, at each value of  $\psi$ ,  $R_i(kT)$  results from the filtering of  $I$  by a filter with a Gaussian-shaped impulse response. In all these calculations, the Gaussian mask  $S(kTs)$  based on equation 3.1 was processed with a spatial sampling step  $Ts$  equal to  $0.004^\circ$ .  $R_i(kT)$  was normalized with respect to its extremal values:

- $R_i(kT) = 1$  when a black segment ( $I = 10$ ) covers the whole visual field  $Lv$ .
- $R_i(kT) = 0$  when the black segment is completely outside the visual field ( $I = 0$ ).

In order to show that the time lag between the two photosensors varies with the *position* of the contrast feature, we refer to the scheme of figure 3.1a, in which we put the dark edge at 3 positions (1, 2, 3) along



**Figure 3.1** Simulated response of two adjacent photosensors rotating clockwise at a variable angular speed  $W$ , and encountering a dark edge placed at various positions 1, 2, 3 (a). (b) Variation of angle  $\psi(t)$ . (c, d, e) Output from photosensors 1 and 2, depending of the position of the dark edge (1, 2, 3). (f, g, h) High-pass filtered version of (c, d, e) (cutoff frequency: 2 Hz) and thresholding of the output signal from each photosensor.  $Dr = 2^\circ$ ;  $Dj = 4^\circ$ ;  $Lv = 5.2^\circ$ ;  $s_0 = 0.01$ ;  $A = 2Dj$ ;  $t = 20$  ms;  $T = 0.01$  s.




Figure 3.2 Structure of the LMD.

the x-axis. For each of these stationary positions, figure 3.1c, d, e gives the output from photosensors 1 and 2 as they turn clockwise at a variable angular speed  $\Omega$ . In figure 3.1f, g, h, the output from each photosensor was passed through an analog high-pass filter with a cutoff frequency of 2 Hz.

Figure 3.1f, g, h shows that the time lag  $\Delta t$  between the differentiated and thresholded outputs of the two photosensors indeed varies gradually with the angular position  $\varphi_c$  of the contrast edge within the visual field.


### MEASURING ANGULAR SPEED WITH AN LMD

The local motion detector (LMD)—also called the elementary motion detector (EMD)—is an analog circuit yielding an output signal that decreases gradually when the time lag  $\Delta t$  between its two inputs increases. Like many insect motion-detecting neurons [8], the analog LMD is directionally selective. We designed an LMD that detects only light-to-dark transitions (“off contrasts”) when the scanning direction is clockwise. The use of the derivative in the first LMD processing step is essential, as it makes it possible to:

- Eliminate the DC level of the photosensors
- Discriminate between positive (on) contrasts and negative (off) contrasts

A block diagram of the LMD is shown in figure 3.2.

The signal output  $V_o$  results from multiplying an exponentially decaying pulse by a delayed fast pulse.  $V_o$  varies inversely with  $\Delta t$  and hence grows larger with  $\Omega$ . This pulse-based scheme for an analog motion detector was originally designed [10] on the basis of electrophysiological findings obtained at our laboratory on flies [8], and an array of motion detectors of this kind has been previously used onboard a visually guided autonomous robot capable of avoiding obstacles at 50 cm/sec [1–2].



**Figure 3.3** Sketch of the complete visual scanning sensor in front of a dark object posted up on the wall.

It was predicted from the simulations shown in figure 3.1 that if an LMD is driven by a pair of photosensors rotating at a variable angular speed, its output  $V_o$  will vary gradually with the angular position  $j_c$  of a dark edge in the visual field of the scanning sensor.

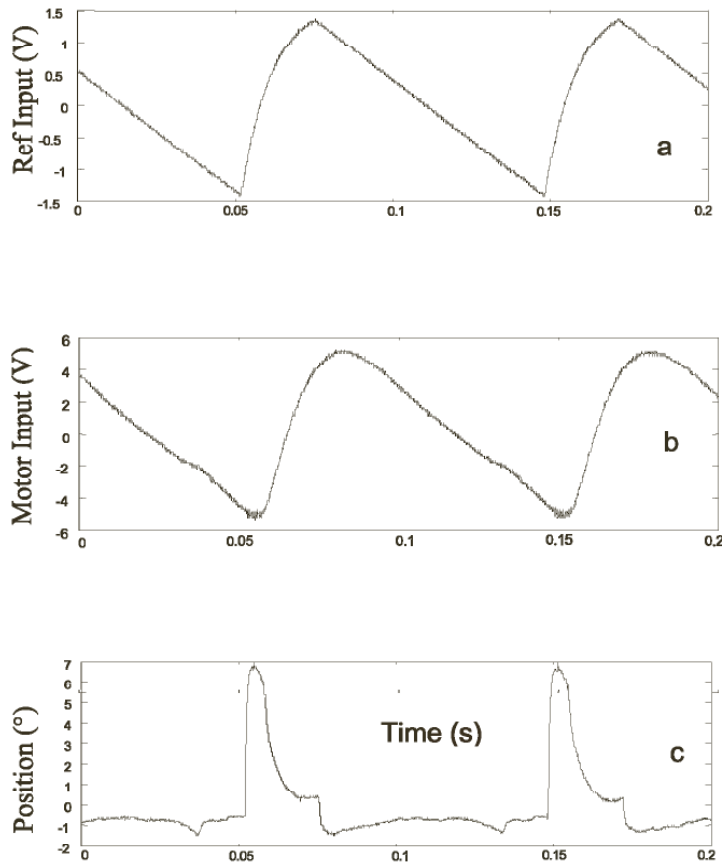
## VARIABLE-SPEED SCANNING AND VISUAL POSITION SENSING

### The Scanning Sensor

We built a miniature scanner, the components of which are sketched in figure 3.3. A dual photosensor and a lens (focal length: 8 mm), mounted opposite each other on a blackened Perspex drum (diameter 15 mm), form a miniature “camera eye,” which is driven by a DC micromotor (diameter 10 mm). The angle of divergence  $Dj$  between the visual axes of the two photoreceptors (interreceptor angle) is:  $Dj = 4^\circ$ . The angular position of the drum (and hence, that of the sensor’s line of sight) is monitored by a *magneto-resistive sensor* responding to the angular position of a micromagnet glued to the hub of the drum. This sensor is part of the loop of an accessory position-servo, which controls the eye so that it will follow any imposed signal.

This servo eye is positioned by a composite periodic signal at 10 Hz (cf. figure 3.4a), which eventually leads to two scanning phases (cf. figure 3.4c):

1. During the first phase (duration: 25 msec), the angular speed  $\Omega$  of the eye is made to decrease quasi exponentially with time.



**Figure 3.4** Measured periodic signals from the position servo of the scanning eye. (a) Reference input voltage imposed upon the position servo. (b) Actual motor input voltage. (c) Resulting orientation of the “eye” as monitored by the magnetoresistive sensor.

2. During the longer “return phase” (duration 75 msec), the eye is made to return to its original position at a quasi-constant speed.

The amplitude  $\Delta\xi$  of the scan is expressed in terms of the interreceptor angle  $\Delta\varphi$  by means of the scanning factor [4]  $\alpha$ :

$$\alpha = \frac{\Delta\xi}{\Delta\varphi}$$

The LMD responds only during the first part of the scan (short phase), delivering an output that varies in a graded manner with the position of the contrasting edge within its visual field. Imposing a slow and constant angular speed during the return phase improves the robustness of the sensor, making it responsive to environmental features far beyond the simple edge considered here.

### Sensor Output Versus Contrast and Distance

The complete scanning sensor was mounted vertically onto the shaft of a resolver and a 16-bit resolver-to-digital converter was used to

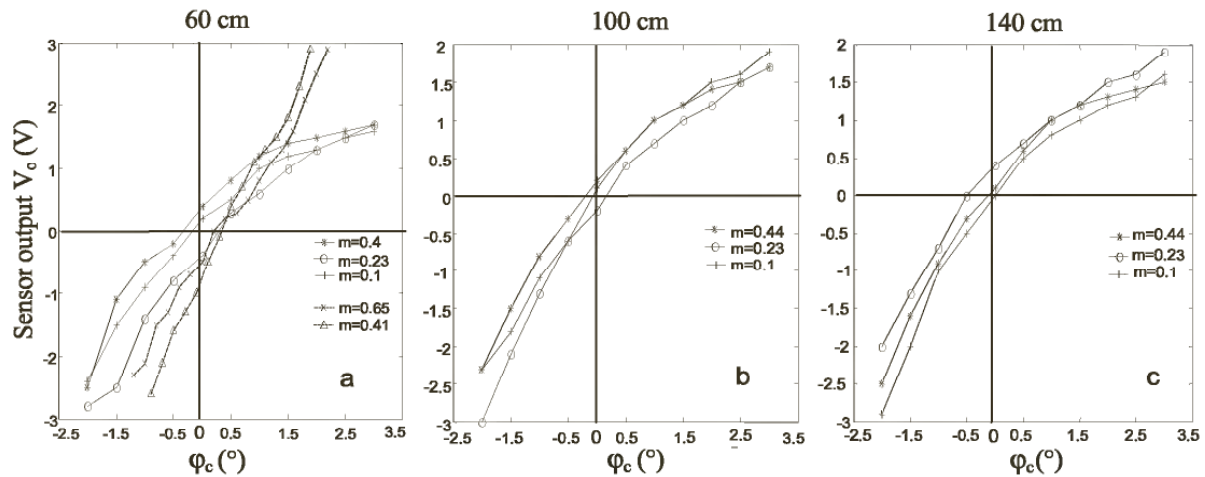
accurately determine the angular orientation of the sensor. The two photoreceptor outputs were connected to an LMD via soft microwires. We varied the orientation of the scanning sensor manually and at each orientation recorded the voltage output  $V_o$  from the LMD and the angular position  $\varphi_c$  of the sensor.

The object was a contrast step made of gray paper that stood out from the background. The contrast  $m$  was determined by measuring the relative illuminance of the paper ( $I_1$ ) and its background ( $I_2$ ) and calculating:

$$m = \frac{I_1 - I_2}{I_1 + I_2}$$

Contrast was measured in situ with a linear photo device having the same spectral sensitivity as the dual photosensor used.

Figure 3.5a, b, c shows that the sensor output varies with the position of a contrasting edge within its visual field—as predicted, based on the results in the introductory section and “Measuring Angular Speed with an LMD.” The responses are quasi linear with respect to the yaw orientation of the scanning sensor and largely invariant with respect to the contrast  $m$  and the distance  $D$  of the dark edge [5]. Subtracting a reference signal  $r = 5.5$  V from all the data gives an odd function, which is a useful feature for incorporating the scanning sensor into an optomotor loop. Moreover, invariance with respect to the distance—within a limited range—is an important feature for the dynamics of an optomotor loop of this kind, because it does not introduce any gain variations liable to cause instability [11]. The sensor was designed to be



**Figure 3.5** Voltage output  $V_o$  from the scanning sensor (cf figure 3.3) as a function of its orientation  $\varphi_c$  and the contrast  $m$  of the pattern. Accuracy:  $\pm 50$  mV. (a) Response at  $D = 60$  cm with respect to a dark edge (solid) and a dark stripe 10 mm in width (dotted). (b) Response at  $D = 100$  cm to a dark edge of various contrasts. (c) Response at  $D = 140$  cm to the same dark edges. ( $\Delta\varphi = 4^\circ$ ,  $\alpha = 2$ ).

robust against any spurious modulations caused by neon room illumination (100 Hz intensity modulations).

### VISUAL CONTROL OF A MINIATURE TWIN-PROPELLER AIRCRAFT

A sketch of the tethered model aircraft equipped with its scanning sensor is shown in figure 3.6.

This aircraft is made of carbon, wood, Fiberglas, and expanded polystyrene (EPS). Its total weight is 100 g, including the engines, scanning sensor, gyroscope, and the complete electronics. The onboard battery (metal-hydride 7.2 V-0.6 A.h) weighs another 90 g. In order to minimize the inertia around the yaw axis, the DC engines are positioned close to the axis of rotation and transmit their power to their respective propellers (diameter 13 cm) via a 12 cm-long carbon fiber shaft and bevel gear (reduction ratio 1:5). The SMD electronic components




Figure 3.6 Model aircraft equipped with its scanning sensor.



occupy five decks of printed circuit boards interconnected via miniature bus connectors, each board (diameter 55 mm) being dedicated to a particular task, as follows:

- Board 1: power conditioning
- Board 2: RC receiver and PWM signal generator
- Board 3: LMD
- Board 4: optomotor controller
- Board 5: scanner servo system

### Dynamic Modeling of the Twin-Propeller Aircraft

We analyzed the dynamics of the twin-propeller system around its yaw axis from its response to a step input. In order to convey a step input voltage to the autonomous plant, we designed a remote control system operating at 433 MHz (1000 baud). At the transmitter, the throttle and yaw commands were adjusted by means of “digital” joysticks. Onboard the aircraft, these signals were received by a miniature receiver and dispatched to two digital potentiometers via a PIC16F84 microcontroller. These potentiometers (8-bit digital-to-analog converters) generated a differential signal  $u_d$  to adjust the yaw and a signal  $u_c$  to adjust the common throttle. In order to make sure that the two thrusters reached the same nominal speed (i.e., the same “operating point”) during each experimental run, the DC level  $u_c$  was memorized in the microcontroller flash memory and called back via a third remote control channel. The differential propeller driving system is shown in figure 3.7.

The whole aircraft was fixed onto the shaft of the resolver (cf. section entitled The Scanning Sensor), which was selected because of its negligible friction and inertia. The 16-bit resolver-to-digital converter was

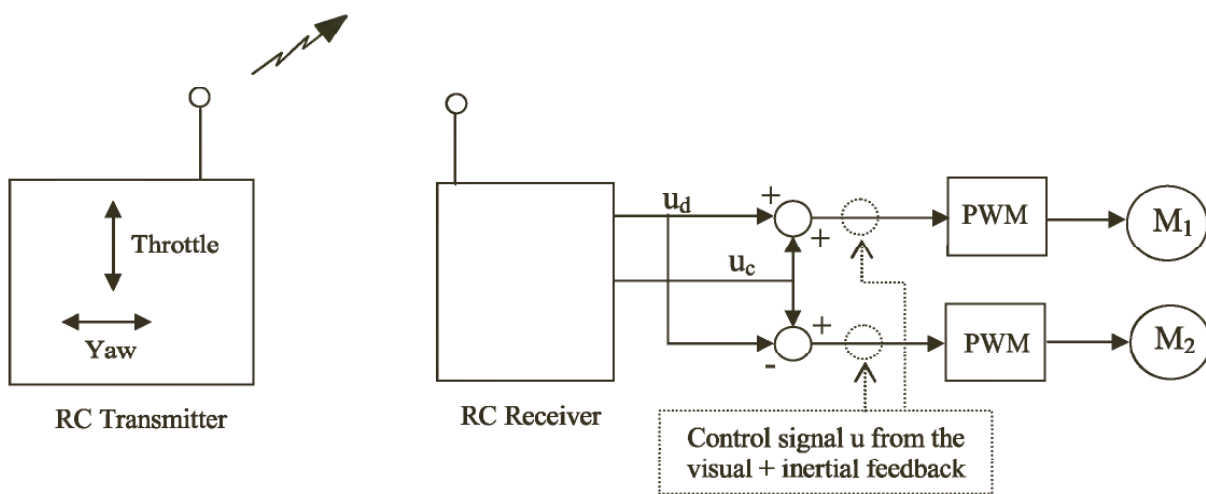


Figure 3.7 Block diagram of the differential system driving the propellers.

then used to measure the angular speed of the aircraft in yaw. This tachometer has a scale factor of  $2.27 \cdot 10^{-3}$  V/deg/s, a scale factor error of 10% and a bandwidth of 125 Hz.

The model  $G(s)$  of the aircraft obtained from its step response is as follows:

$$G(s) = \frac{\theta_a}{U_d} = \frac{2700}{(0.026s^2 + 0.232s + 1)}$$

where  $\theta_a$  is the angular speed (deg/s) of the aircraft and  $U_d$  is in Volts.

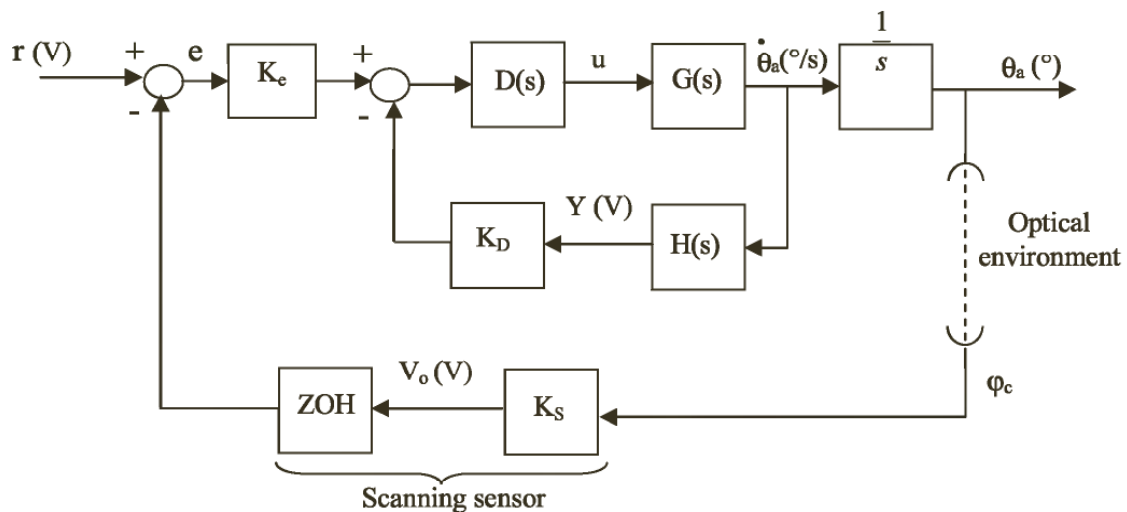
### The Visual Inertial Feedback System

The complete position feedback system we have designed is shown in figure 3.8.


The scanning sensor is modeled by a simple gain  $K_s \cong 1$  V/deg (cf. figure 3.5), and its output signal was maintained by a zero-order hold until the completion of each scan ( $T = 100$  msec). Because the transfer function from the controlling yaw input to the yaw angle has three main poles, we incorporated a velocity feedback [12] using a miniature piezo gyroscope. The latter was modeled by the following transfer function  $H(s)$ :

$$H(s) = \frac{Y}{\dot{\theta}_a} = \frac{6.7 \cdot 10^{-3}}{(0.0032s + 1)}$$

The gain  $K_D = 5$  was chosen so as to give a closed-loop step response with a settling time of 100 msec with a minor overshoot. The controller  $D(s)$  was designed in particular to compensate for the limited bandwidth (50 Hz) of the gyroscope.  $D(s)$  raised the phase margin further, compensating for the phase lag introduced by the low-pass filter  $H(s)$ ,



**Figure 3.8** Block diagram of the closed-loop visual control system incorporating the scanning visual sensor (outer loop) and a gyroscope (inner loop).



**Figure 3.9** Frequency response of  $B_o(s)$ .

and thus increasing the damping. The transfer function of  $D(s)$  is given by:

$$D(s) = 0.3 \left( \frac{0.08s + 1}{0.014s + 1} \right)$$

and

$$K_e = 6 \cdot 10^{-2}$$

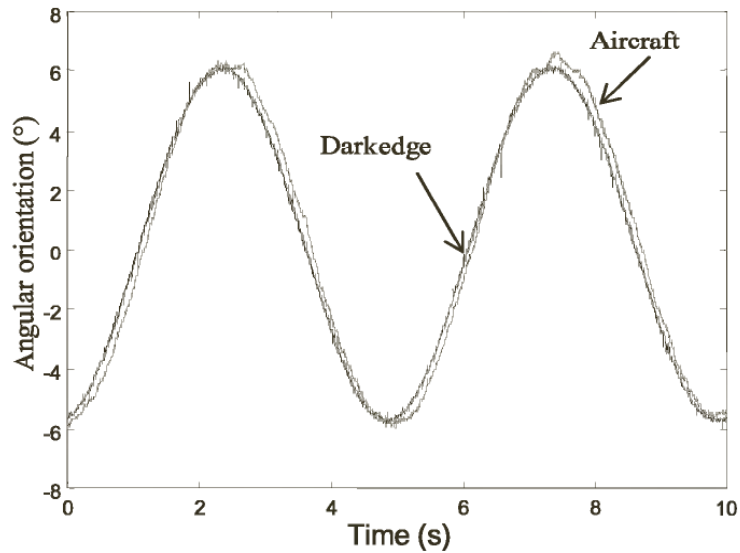
The frequency response of the open-loop transfer  $B_o(s) = \frac{V_o(s)}{r(s)}$  is shown in figure 3.9.

### Target Tracking and Visual Fixation


We assessed the closed-loop performance of the visual-inertial feedback system on the test bed described in figure 3.6. Both the angular position of the dark edge (input) and the angular orientation in yaw of the model aircraft (output) were measured with potentiometers, as depicted in figure 3.6.

Figure 3.10 illustrates the tracking performance of the visual feedback loop with respect to a sinusoidal variation of the position of a dark edge in the visual field of the scanning sensor.

The response was found to lag behind the command by about 100 msec—that is, by hardly more than the scan period—which turned out to be the main limiting factor. The maximum linear speed of the dark edge that the model aircraft was able to track was about 25 cm/sec at a distance of 1 m ( $14^\circ/\text{sec}$ ). The histogram in figure 3.11 shows that the



**Figure 3.10** Dynamic performance of the overall twin-propeller system (cf figure 4.6) tracking a dark edge (contrast  $m = 0.2$ ) located at a distance of 1 m in front of the eye and oscillating at 0.2 Hz perpendicularly to its line of sight.  $r = 5.5$  V.



**Figure 3.11** Distribution of the angular orientation of the model aircraft, which is constantly fixated on a stationary dark edge ( $m = 0.2$ ) for 17 minutes. Number of samples = 50000; acquisition time = 1000s; standard deviation =  $0.22^\circ$ .

aircraft steadily controls its two propellers so as to gaze steadily on a stationary target during a long 17 min experiment.

## CONCLUSION

In this chapter, we described the use of a low-level, low-complexity, power-lean miniature sensor designed to provide a micro air vehicle

with visual stabilization and tracking possibilities. We established that a 100 gram twin-propeller aircraft equipped with this miniature scanning eye can track an object with various levels of contrast (down to 10%) at various distances (up to 1.5 m), using its onboard electronics and energy in a completely autonomous manner.

The aim here was not to achieve a high-performance closed-loop system such as those based on a CCD camera, which recognize and track a target at high speeds. Such systems rely on more-complex control laws and require larger computational resources [11, 13]. Our low-level visual system is not based on any “understanding of the environment,” and consequently cannot detect and track a specific target among many others. Yet this moderately complex, inexpensive ( $\cong$  400 USD) system can be used:

- To stabilize a micro air vehicle around a given axis of rotation, where the time constant is small (the small time constants of micro air vehicles make it difficult for them to be stabilized by a human operator)
- To replace a gyroscope by maintaining an angular position with respect to environmental features—for a long time and with no drift
- To track a target visually with a large shape, contrast, and distance invariance, using a single pair of photosensors scanning concomitantly

It is worth noting that the optical position–servo method we have developed is based on a *position sensor* that relies on a *motion detector*. An essential message of this chapter is that a simple, low-amplitude scanning of two photoreceptors can provide positional information far beyond that expected from their angular sampling basis. Extending these results to a 1-D or 2-D photoreceptor array will retain this interesting property while enlarging the field of view.

## ACKNOWLEDGMENTS

We thank M. Boyron for his expert technical assistance and I. Daveniere and J. Blanc for improving the English manuscript. This research was supported by CNRS (Life Sciences, Engineering Sciences, and Programme on Microsystems). Stéphane Viollet received a predoctoral fellowship from Ministry of National Education and Scientific Research (MENRS).

## REFERENCES

- [1] Pichon, J. M., Blanes, C., and Franceschini, N. (1989). Visual guidance of a mobile robot equipped with a network of self-motion sensors. In *Mobile Robots 4*. Bellingham, Wash.: SPIE, pp. 44–53.
- [2] Franceschini, N., Pichon, J. M., and Blanes, C. (1992). From insect vision to robot vision. *Philos. Trans. R. Soc. Lond. B Biol. Sci.* 337: 283–294.

- [3] Mura, F., and Franceschini, N. (1994). Visual control of altitude and speed in a flying agent. In D. Cliff, P. Husbands, J.-A. Meyer, and S. W. Wilson (eds.), *From Animals to Animats 3*. Cambridge: MIT Press, pp. 91–99.
- [4] Mura, F., and Franceschini, N. (1996). Biologically inspired retinal scanning enhances motion perception of a mobile robot. In A. Bourjault (ed.), *Proceedings first Europe-Asia Congress on Mechatronics 3*. Besançon, France: pp. 934–940.
- [5] Viollet, S., and Franceschini, N. (1999). Biologically inspired visual scanning sensor for stabilization and tracking. In *Proceedings IEEE IROS '99 1*. Kyongju, Korea: pp. 204–209.
- [6] Franceschini, N., and Chagneux, R. (1997). Repetitive scanning in the fly compound eye. In Wässle and Elsner (eds.), *Proceedings Göttingen Neurobiology Conference*.
- [7] Mura, F., and Shimoyama, I. (1998). Visual guidance of a small mobile robot using active, biologically inspired eye movements. In *Proceedings of IEEE Robotics and Automation*. Leuven, Belgium: pp. 1859–1864.
- [8] Franceschini, N., Riehle, A., and Le Nestour, A. (1989). Directionally selective motion detection by insect neurons. In Stavenga and Hardie, (eds.), *Facets of Vision*. Berlin: Springer, pp. 360–390.
- [9] Hardie, R. C. (1985). Functional organisation of fly retina. In D. Ottoson (ed.), *Progress in Sensory Physiology 5*. Berlin: Springer, pp. 1–79.
- [10] Franceschini, N., Blanes, C., and Oufar, L. (1986). Appareil de mesure passif et sans contact de la vitesse d'un objet quelconque. *Dossier Technique ANVAR/DVAR N° 51 549*, Paris.
- [11] Corke, P. I. (1996). Dynamic issues in robot visual-servo systems. In Giraldo and Hirschinger, (eds.), *Robotics Research*. Berlin: Springer, pp. 488–498.
- [12] Franklin, G. F., Powell, J. D., and Emami-Naeini, A. (1994). *Feedback control of dynamic systems*. Addison Wesley.
- [13] Ishii, I., Yoshihiro, Y., and Ishikawa, M. (1996). Target tracking algorithm for 1 ms visual feedback system using massively parallel processing. In *Proceedings IEEE Robotics and Automation*. Minneapolis: pp. 2309–2314.

**This page intentionally left blank**

# 4

## Visual Sensors Using Eye Movements

Oliver Landolt

Visual processing in robotics applications lags far behind the visual systems of even simple animals. Close examination reveals that the operating principles underlying biological visual systems differ fundamentally from conventional electronic image sensing and processing hardware. One important difference, frequently overlooked even within the field of biologically inspired engineering, is that most animals can move their eyes with respect to their head and/or body. Such animals display a repertoire of eye movements designed to improve the visual data acquisition process and save neural hardware. By contrast, the visual axis of electronic visual systems is typically pointed in a fixed direction with respect to the supporting platform, or can move only slowly compared to the time scale of visual events. The present chapter argues that eye movements are an important aspect of visual sensing and should be built into electronic implementations of visual sensing systems. The next section presents an overview of the major types of eye movements found in humans and discusses their importance to vision. The two subsequent sections describe engineered systems incorporating similar principles. Last, the Conclusion discusses to which degree the biological domain can inspire the architecture and algorithms of synthetic vision systems.

### EYE MOVEMENTS IN HUMANS

Eye movements observed in humans fall into two broad categories, namely, saccades and target-holding movements. Saccades consist of fast (on the order of  $1000^\circ/\text{sec}$ ), essentially ballistic jumps from one fixation point to another. Humans usually look at visual scenes in an endless succession of saccades and fixations (up to 3 per second) concentrated on areas of interest [1]. The main purpose of saccades is to bring an image location of interest into the high-acuity foveal region of the retina for detailed analysis. Another purpose of saccades is to bring the eyeball rapidly back to the center after drifting away—for example,



in pursuit of a moving object. The most obvious benefit of a moving fovea is economical: To obtain foveal resolution over a wide field of view without eye movements, a tremendous number of photoreceptors would be needed, as well as a matching amount of cortical visual processing hardware. Similar economical considerations are highly relevant for electronic visual sensors as well.

Unlike saccades, target-holding eye movements are smooth and are designed to stabilize the retinal image. Several categories are distinguished depending on the nature of the stimulus controlling the movement. The vestibulo-ocular reflex (VOR) is driven by inertial measurements of head movements, performed by the vestibular system. The VOR counteracts the effect of head movements by generating opposite eye movements, thereby greatly reducing the velocity of retina image slip. The peak velocity of VOR movements is on the order of  $500^\circ/\text{sec}$ . The opto-kinetic reflex (OKR) is similar in purpose but controlled by vision itself. The OKR compensates for large-field image slip up to about  $80^\circ/\text{sec}$ , such as produced by the subject's own body movements with respect to the environment. Smooth pursuit consists of tracking a small moving object in order to keep it in the fovea. It operates up to  $40^\circ/\text{sec}$ . Target-holding movements are important for the reduction of motion blur [2]. Because of their limited temporal bandwidth, photoreceptors can follow light intensity fluctuations caused by a moving image only up to some velocity, dependent on the spatial features of the image. At higher velocities, the image appears blurred. Total or partial compensation of image slip increases the range of conditions under which a target can be seen clearly.

In addition to the large-scale eye movements described so far, it is known that even during attempts to maintain perfect fixation, humans actually generate small involuntary eye movements [3]. These movements typically alternate between smooth drift and microsaccades, with a median delay on the order of 0.6 sec between the onset of successive saccades. The amplitude of these movements is on the order of a few minutes of arc. It has been established that the function of these involuntary movements is to maintain permanent vision. If the image of a target is kept perfectly stable on the retina—by mounting the target directly on the eyeball for instance—then visual perception of the target vanishes in a matter of seconds [3]. Instead, observers report seeing a gray fuzzy field, or sometimes even extreme blackness. In other cases, observers describe intermittent perception of a dim, unsharp, or partial image of the target. Reportedly, the target may transiently reappear clearly on events such as illumination flashes, or by applying a movement of sufficient amplitude (on the order of  $2'$  of arc) to the target. Failure of vision in the case of a perfectly stabilized retina image is explained by the fact that the visual system detects and encodes light intensity fluctuations over time—and space—rather than steady light

levels. Therefore, permanent visual perception is possible only if the image keeps moving. Small-amplitude involuntary eye movements provide the conditions for sustained vision.

Although it may seem odd at first that humans rely on a visual system requiring eye movements to operate properly, this architecture turns out to have interesting properties. In natural viewing conditions, spatial patterns that remain perfectly stable on the retina are most likely artifacts from the visual system itself, such as the blind spot, shadows of blood vessels, afterimages due to adaptation, or fixed-pattern noise due to photoreceptor mismatch. By detecting only temporal fluctuations caused by eye movements, such artifacts can in principle be discounted. In contrast with biological retinas, most electronic image sensors—notably cameras—measure the DC light intensity received by every pixel and are thus affected by fixed-pattern noise caused by photoreceptor mismatch. To deal with this effect, the fixed pattern must be measured in a calibration step, stored in some type of memory and discounted at image readout. By building electronic visual sensors capable of shifting their optical axis, the benefits of the biological approach can be exploited without sacrificing sustained vision. Another prospective advantage of small-amplitude eye movements is resolution enhancement. Shifting a photoreceptor in space provides access to a continuum of image points instead of a discrete grid. In principle, it is possible to use such scanning movements to detect features much narrower than photoreceptor spacing. It is not likely that humans exploit this potential because photoreceptor spacing in the fovea is so narrow that the effective resolution is actually limited by the imaging optics [3]. However, there exist animals known to rely on scanning for visual data acquisition, presumably for resolution enhancement purposes [4–5].

Given the benefits provided by moving eyes to their animal users, perhaps gaze control could improve electronic visual sensors as well and contribute to bridging the performance gap between biological and electronic visual systems. For the development of practical visual systems exploiting eye movements to the same extent as animals, several technological issues must be addressed.

1. A mechanical/optical solution must be found to the problem of rapidly shifting the visual axis. Device size and power consumption must remain compatible with small robotics applications. For performance comparable to humans, the device should be capable of executing saccades within about 50–100 msec with an accuracy on the order of  $1^\circ$ , over a total angle of  $60\text{--}90^\circ$ .
2. Image-sensing principles must be developed, which take advantage of eye movements to achieve capabilities such as fixed-pattern noise rejection or resolution enhancement.

3. High-level control strategies must be developed to achieve the built-in functionality of the human oculomotor system, such as target-holding reflexes and saccadic exploration. Conflicts between subsystems implementing different functions must be resolved.
4. Efficient computation of coordinate mappings between visual space and motor space must be implemented. Motor space is defined by the degrees of freedom of the device shifting the visual axis. Fast and accurate coordinate mapping is an essential part of the control system of a moving eye.

The author and his colleagues have been addressing these issues over several years and have proposed and implemented original solutions in the frame of two independent developments of visual sensing systems exploiting eye movements. Custom hardware solutions were favored, which rely on operating principles inspired by neurobiology [6]. This choice is motivated partly by interest in exploring the potential of this approach, and partly by anticipation of strong constraints in terms of cost, size, speed, and power consumption in many applications of eye movements. Arguably, bioinspired approaches should provide superior solutions—in the long run—in classes of problems where biology excels but computers do poorly. Indeed, neural structures have been optimized over millions of years of evolution by selection of the best designs for survival.

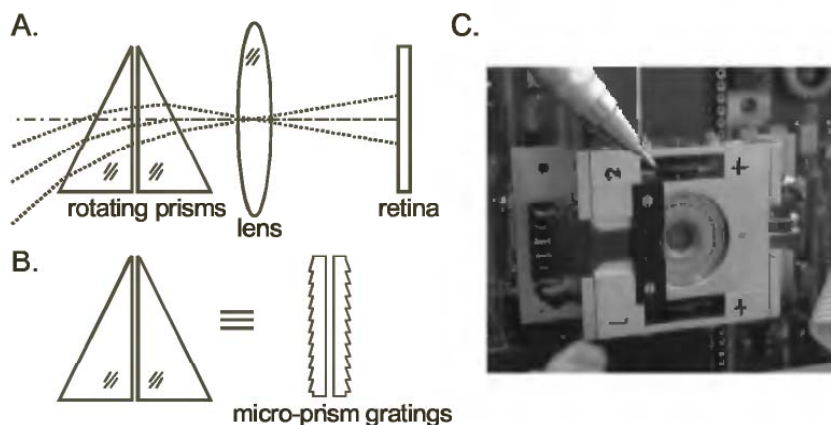
## **OCULOMOTOR SYSTEM BASED ON MICROPRISM GRATINGS**

### **Functionality**

An oculomotor system has been built that implements two types of eye movements found in humans—namely, saccades and smooth pursuit [7]. Rotations of the visual axis are produced by a light deflection device mounted in front of a fixed-image sensor. The system otherwise contains four analog very large-scale integrated circuits (VLSI) based on bioinspired principles. The primary purpose of this system is to demonstrate the validity of novel principles underlying the optical front end, the related control strategy, and its custom hardware implementation. However, some additional elements are likely to be required before practical applications can be addressed. In particular, difficult visual processing problems—such as identification of saccade targets in natural scenes or discrimination of moving objects from their background—have been circumvented by assuming visual scenes made of light patches over a dark background, whereby luminance directly identifies objects of interest. The functionality of the oculomotor system can be extended without alteration of its core by inserting more-sophisticated visual-processing hardware between the retina chips and the rest of the control system.

## Light Deflection Device

To rotate the visual axis, we have selected an approach consisting of keeping the retina chip and the focusing lens fixed, while deflecting light using refractive elements mounted in front of the lens (figure 4.1A). The light deflection device consists of two transparent and flat plastic disks with a microprism grating on one side, mounted perpendicularly to the optical axis of the lens. The angle of the microprisms is  $30^\circ$  and the refraction index is 1.49. The disks have an overall diameter of 7 mm and a thickness on the order of 1 mm. Each disk can rotate without restriction around this axis, independently from the other. As a whole, each microprism grating acts on light essentially like a single bulky prism, except that it takes much less space and weight (figure 4.1B). Although a single fixed prism cannot have an adjustable deflection angle, with two mobile prisms in the light path, any magnitude and direction of deflection can be obtained within some boundaries. Their contributions may combine either constructively or destructively depending on the relative prism orientations. With the selected material and geometry, the available range extends up to  $\pm 45^\circ$  horizontally and vertically. The relationship between prism orientations and the direction and magnitude of deflection has been derived in detail [8]. This relationship turns out to be strongly nonlinear and there are generally two combinations of prism orientations achieving the same overall deflection. Therefore, control of this device involves the computation of nontrivial coordinate mappings. However, the decisive advantage of this device over many other solutions is that only two small passive optical elements have to move. This feature enables fast movements with moderately powered motors and avoids electrical connections to moving parts. The measured delay between onset of a saccade and



**Figure 4.1** (A) Light deflection device principle. (B) Replacement of conventional prisms by micro-prism gratings. (C) Photograph of a mechanical device including the gratings, orientation sensors and motors. The overall device size is 56 mm  $\times$  34 mm  $\times$  16 mm.

stabilization on the target ranges from 45 msec to 100 msec depending on the saccade magnitude. Optical aberrations introduced by the prisms degrade the sharpness of the image. The degradation is negligible at low deflection angles, but at the extreme deflection of 45°, the image of an ideal spot extends over a distance of 300–400  $\mu\text{m}$ . However, when the device is used in conjunction with a typical electronic retina, this degradation is not significant up to deflections of 30–35°. The reason is that these image sensors are characterized by large pixel spacing (typically 50–100  $\mu\text{m}$ ) due to the presence of focal-plane electronic processing.

### Control System Architecture

The oculomotor system consists of three modules as shown on the block schematic in figure 4.2. The moving-eye module contains the light deflection device described in the previous section with its two motors. The orientations of the two microprism gratings are monitored by means of coding wheels and optical detectors. The motors are controlled by two independent position servo loops keeping each grating in the orientation specified by the content of a register. These registers are the points of convergence between the saccadic exploration module and the smooth-pursuit module. A saccade is generated by loading a new orientation setting in the registers, whereas smooth movements are generated by applying incremental changes to the stored settings.




Figure 4.2 Oculomotor system architecture.


The gratings track these changes within the delay imposed by the dynamics of the low-level control loops.

The light deflection device is mounted on top of a retina chip with a focusing lens dimensioned to achieve a relatively narrow field of view of  $12^\circ$ . Because this narrow field can be moved over a large angle, this retina can be considered as a high-resolution “spotlight” gathering image details sequentially in time. In this respect, it is similar to the fovea found in humans and other animals. The retina chip contains an array of 35 by 35 pixels turning the image into electrical signals encoded in the timing of pulses. The firing rate of a given pixel is proportional to the ratio between the incident light intensity and the local average intensity in the neighborhood of the pixel. This integrated circuit (described in detail in [9]) was not designed specifically for the oculomotor system. The pixels respond to steady illumination instead of only temporal transients as the human retina. However, this discrepancy does not significantly affect system operation. The data delivered by this retina chip is used internally for smooth pursuit control. It also constitutes the main output delivered by the oculomotor system to higher levels of processing.

### **Smooth-Pursuit Module**

The smooth-pursuit module contains an element detecting the location of a target on the retina image. The target location is transmitted to a so-called incremental control chip implementing nonlinear coordinate mapping between visual and motor space. This chip determines the direction and relative magnitudes of prism orientation increments required to slightly shift the image of the target toward the center of the retina. For this purpose, it also takes the current orientations of the two prisms into account, because the appropriate action depends on the present state of the system. The requested actions are executed by the moving-eye module, thereby causing the image of the target to shift on the retina. Corrective actions are integrated over time as a result of continuous operation of this visual control loop, whereby a target initially away from the center will progressively shift toward the center of the retina and then remain there. If the target moves, the eye tracks it until the next saccade causes locking onto a new target, or until the object leaves the field accessible to the moving eye.

Information coding and processing within the incremental control chip are based on principles inspired by biological neural structures. Throughout the chip, variables are represented by the location of a patch of activity within a topologically organized array of cells. This coding convention is called *place coding* [8] to emphasize the fact that location—rather than only amplitude—determines the represented




**Figure 4.3** Distribution of stimulation intensity encoding an angle of  $81^\circ$  on a map of 8 nodes whose preferred directions are uniformly distributed at intervals of  $45^\circ$ .


quantity. This approach is closely related to population coding [10], which is widespread in biological neural structures. At the input of the incremental control chip, the angular position of the visual target with respect to the center of the retina is encoded by a map consisting of eight nodes. Each of these nodes is related to a particular preferred direction. These directions are uniformly spaced at intervals of  $45^\circ$ . In order to indicate that the target is located at a given arbitrary angle (such as  $81^\circ$ ), the closest matching input cells of the chip ( $45^\circ$  and  $90^\circ$ ) must be stimulated by pulse streams of some frequency. The relative frequencies at which these nodes are driven indicate how closely the represented direction matches their preferred direction (figure 4.3). The convention chosen for the definition of place coding is that the value represented by a given pattern of stimulation is the centroid of the pattern [8]. A direction of  $81^\circ$  can therefore be represented by applying 20% of the total stimulation to the node related to  $45^\circ$ , whereas 80% of the stimulation is applied to the node whose preferred direction is  $90^\circ$ .

With a place-coding scheme, computation can be performed by networks of passive interconnections—called *networks of links*—which embody the relationship between an input map and an output map [8]. To implement functions of several variables, multiple input maps can be merged by means of an array of fuzzy logic AND gates. The resulting intermediate map can then be projected onto the output map by means of a network of links. Fuzzy logic AND gates can be implemented by an extremely simple circuit consisting of only one metal-oxide-semiconductor (MOS) transistor per input. The incremental control chip is a direct application of this approach. Its input stage consists of three place-coding maps representing the target location and the current orientations of the two prisms. These maps receive stimulation encoded as the frequency of brief pulses. Two networks of links project the three input maps onto two intermediate maps (figure 4.4). A third network of links performs an additional computational step resulting in the final maps. These maps represent control actions that are transmitted to the moving-eye module in the form of pulses. The architecture of the chip has been derived manually by studying the properties of the control function to be implemented. The topology of the net-





**Figure 4.4** Architecture of the incremental control chip.

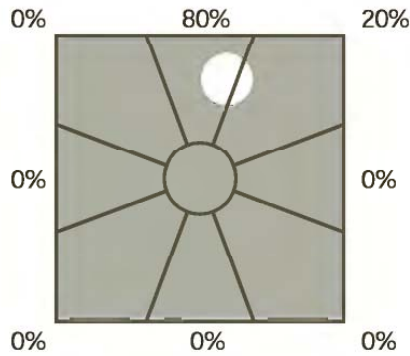


**Figure 4.5** Photograph of the incremental control chip. The shown area is  $2\text{ mm} \times 1\text{ mm}$ .

works of links has been generated automatically by software on the basis of a behavioral model of the oculomotor system.

Place coding may at first seem complicated compared to ordinary analog or digital coding, but it actually lends itself well to compact and power-efficient hardware implementation of nonlinear mappings such as the oculomotor system requires. The core of the incremental control chip (figure 4.5) uses  $2\text{ mm}^2$  of silicon area in a  $0.7\text{ }\mu\text{m}$  CMOS process, and consumes  $30\text{ }\mu\text{W}$  under a supply voltage of  $5\text{ V}$ . Conversion from digital coding to place coding on a chip is simple and efficient [8]. Besides, it is possible in some cases to get place coding directly from a sensor. In the smooth-pursuit module, the input map representing the angular location of the target is derived from the visual data delivered by the retina chip. The surface of the retina is divided into eight radial





**Figure 4.6** Division of the retina surface into radial slices resulting in place coding of the angular position of the target.

slices and one central area (figure 4.6). The pulses generated by all pixels within a radial slice are merged into a single signal, whereby the average frequency of the resulting pulse stream is roughly the sum of spiking frequencies of all pixels within that slice. This pulse stream stimulates one node of the target location map, namely, the node whose preferred direction matches the orientation of the radial slice. If a light spot shining on the retina is located away from the center, firing of pixels illuminated by the spot causes stimulation of a small fragment of the map. The location of the most active node provides coarse information about the target location. If the spot is large enough to overlap two slices, two neighboring nodes will be stimulated at a rate depending on the area of the intersection between the spot and the retina slice. The stimulation pattern received by the chip thereby provides continuous information about the spot location. This representation of the angular position may be slightly distorted, but is accurate enough to guarantee that the control actions applied under permanent feedback cause the spot to ultimately reach the center of the retina.

### **Saccadic-Exploration Module**

The saccadic-exploration module contains a second instance of the same retina chip as the moving-eye module. However, the visual axis of this retina is fixed and the associated focusing optics is such that the field of view extends over most of the area accessible by eye movements. This retina provides a coarse image of the surroundings for the selection of salient (i.e., interesting) locations. Therefore, it is comparable to the peripheral area of the human retina. Under the simplifying assumption that the visual scene is made of light patches over a dark background, luminance can be used as a measure of saliency. Therefore, the raw image delivered by the retina can be considered as a distribution of saliency without further processing. For practical applications, more-sophisticated and possibly task-dependent measures of saliency could

**Table 4.1** Saccade statistics in the presence of four equally salient points after 3 minutes of recording

Visual location	Number of hits	Average fixation time
A	436	133 ms
B	443	88.2 ms
C	418	106.4 ms
D	434	86.9 ms

be implemented [11–12]. The related processing could be carried out by an additional chip taking the retina image as input and providing a saliency map at its output (as suggested in figure 4.2).

The distribution of saliency derived from the retina image, encoded as mean firing rates on an array of cells, is applied to a so-called saccadic control chip. This chip has the function of selecting saccade targets based on the distribution of saliency and determining the timing of saccades between locations of interest. The general idea is to select the most salient point at any time as the target of the next saccade. To prevent neglect in the case of multiple salient points, incoming saliency signals are integrated over time, and the next saccade target is the location with the highest integral. Saliency signals originating from the currently attended location and its immediate vicinity are ignored and the corresponding integrators are let to slowly leak instead of letting them keep increasing. Thereby, whenever a given point becomes the target of a saccade, it will remain the most salient location only for a limited amount of time before another location wins the competition. As a result of the integration process, any location with nonzero saliency will become the saccade target sooner or later. However, the number of visits to more-salient locations is more frequent because the corresponding integrals increase at a higher rate. Moreover, fixation tends to be maintained for a longer time at more-salient locations. The network generating saccade dynamics has been implemented in analog hardware as part of the saccadic control chip (described in detail in [8]). Saccade dynamics have been measured on the chip alone under controlled electrical stimulations. For illustration, statistics obtained by recording saccades over three minutes in the case of four equally salient points are given in table 4.1. The four points, arbitrarily labeled from A to D, are spaced well apart in the visual field. Similar measurements were performed on the entire oculomotor system under optical stimulation as well (see System-Level Measurements section). Besides generating saliency-driven saccades, the chip also implements coordinate mapping between visual and motor spaces. This mapping differs from the smooth-pursuit module, because visual space is fixed instead of


being bound to the visual axis of the moving eye. Moreover, the saccadic control chip has to compute absolute prism orientations instead of increments with respect to the present location. The hardware implementation of this coordinate mapping relies on place coding as in the previous module.

### **Module Coordination**

Due to the coexistence of two independent control modules, the control system must incorporate some amount of coordination between them to avoid interference. For instance, during a saccade, the smooth-pursuit module would attempt to compensate for the observed retina image slip if no mechanism prevented it from doing so. In order to handle this situation, the saccadic-exploration module inhibits the smooth-pursuit module and resets the incremental control chip during every saccade. After the saccade, the smooth-pursuit module starts tracking the new target without memory of its previous state. Another coordination mechanism has been incorporated to handle boundary problems during smooth pursuit. If a tracked target crosses the boundary of the visual field accessible to the moving eye, this situation is detected by the smooth-pursuit module, which triggers a signal, causing the saccadic control chip to generate a saccade back to the center. This mechanism is represented in figure 4.2 by an arrow labeled “nystagmus,” by analogy with the reflex found in humans and other animals.

### **System-Level Measurements**

The complete oculomotor system (figure 4.7) consists of three independent printed circuit boards corresponding to the three modules described in the Control System Architecture section. Most experiments were carried out with the system mounted on an optical bench in front of a dark screen punctuated by light-emitting diodes (LEDs). The direction of the visual axis was permanently monitored by observing internal signals within the moving-eye module. Image data from either retina chip was acquired by recording all transitions on the output bus of the chip by means of a logic analyzer as long as memory would allow, then processing the data off-line in order to count the number of spikes emitted by each pixel. Stimulation patterns in portions of the system using place coding were acquired by a similar method. A first round of experiments consisted of qualitative observations of the behavior of the system while turning on either the saccadic-exploration module, or the smooth-pursuit module, or both. With the saccadic module only, when shown a scene consisting of a single LED turned on, the moving eye points toward this LED most of the time. Brief saccades toward any point in the background occur occasionally. The



**Figure 4.7** Photograph of the oculomotor system. The length of the largest printed circuit board is about 25 cm.

fraction of time spent watching the background depends on the contrast between target saliency and background saliency [7]. The experimental setup had a dark background intended to keep background saliency low, therefore the system spent only a small percentage of the time on the background. It is perfectly appropriate that the oculomotor system pays some attention to the background if its saliency is not strictly zero. This behavior is due to the temporal integration of saliency built into the saccadic control chip (Saccadic-Exploration Module section). With two LEDs turned on, saccades alternate from one target to the other in an essentially periodic fashion—except for infrequent saccades toward background locations as in the previous case. With additional LEDs, all intended targets are visited frequently, but no obvious periodicity can be observed. The system exhibits only minor preferences between targets of the same nominal intensity, consistent with observations made on the saccadic control chip alone (table 4.1). The rate at which saccades occur depends on the rate at which retina pixels emit spikes. This rate can be tuned by altering bias conditions of the retina chip. In most experiments, the retina was tuned to achieve about 3–5 saccades per second, well within the range supported by the underlying mechanics. Saccade accuracy was measured by activating several targets at various locations on the visual field. Immediately after every saccade, the location of the target was measured by computing the centroid of its image on the retina chip mounted underneath the light deflection device. The distance between the centroid of the target and the center of the retina was considered to be the error. In all cases,

the error is less than  $2^\circ$  of visual angle. The error vector tends to be identical upon successive returns to the same target, suggesting that the accuracy of saccades is limited by built-in errors rather than noise. The major causes are rounding distortions inherent to the operation of networks of links, and map alignment errors due to the fact that the visual axes of the two retinas are not exactly identical.

When the smooth-pursuit module is activated alone, the oculomotor system tracks a light source waved manually within its field of view. Because this module lacks a wide-angle acquisition mechanism, tracking is initiated only after the target enters the visual field of the moving eye. Smooth pursuit operates within a range of velocity limited by the bandwidth of the retina chip and the dynamics of the low-level control loops. For bright targets, tracking velocities up to  $50^\circ/\text{sec}$  have been observed. The dynamics of the feedback loop depend on target contrast because the system uses luminance directly as a feature to identify the target. Lighter targets cause the retina to fire at a higher rate, thereby causing more frequent updates of the control action fed back to the mechanical parts. When both modules are active simultaneously, the observed behavior is essentially a combination of the first two cases. When a fixed scene of LEDs is shown, saccades occur as before, except that targets end up accurately centered on the retina. The reason is that the smooth-pursuit module cancels the residual error after saccades under closed-loop visual control. Conversely, in the presence of a slowly moving target within the wide-angle visual field of the system, an initial saccade toward the target is triggered 1–2 sec after its initial appearance, then the target is tracked by the smooth-pursuit module.

## VISUAL SENSOR RELYING ON MECHANICAL VIBRATIONS

### Concept

In the Eye Movements in Humans section, it was argued that small-amplitude movements of the visual axis can be used as a means to overcome fixed-pattern noise inherent to photoreceptors, and as a means to gain access to the image in continuity rather than on a discrete grid. A visual sensor exploiting this technique is described in the present section. The general idea consists of applying continuous oscillatory scanning movements to the image focused on the surface of a photoreceptor array. The amplitude of the oscillation should be on the order of pixel spacing, and the frequency should be high enough that the image cannot change significantly over a single scanning cycle. A frequency of a few hundred cycles per second is expected to be sufficient in most cases. In principle, the oscillation does not need to be periodical, but periodicity simplifies signal processing. As a result of these oscillations, spatial variations of light intensity in the image turn into temporal

fluctuations of light intensity at every photoreceptor. Knowing the pattern of motion applied to the image, elementary spatial image features such as edges or textures can be detected by processing these temporal signals. For instance, by sweeping the image of a thin line—such as a dark cable over the background of a clear sky—over a photoreceptor, an impulse of photocurrent will be observed, even if the line is much thinner than pixel spacing. The temporal signature of a dark line is a sharp transition from light to dark followed soon by another transition from dark to light. In addition to the mere existence of this line, its orientation can also be determined by relating the occurrence of temporal changes in photocurrent with the instantaneous direction of movement applied to the image at the same time. No fluctuation will occur while the photoreceptor is shifted parallel to the line, whereas the steepest transitions will be detected when the scanning path crosses the line perpendicularly. If the photoreceptor scans the same area at different angles, the general orientation of the underlying pattern can be determined by identifying the directions of scanning producing the most or the least intensity fluctuations. Finally, the thickness of the line can be determined by measuring the time elapsed between the falling and the rising transition, and scaling it by the velocity of the scanning movement over this time interval.

With this principle, every single pixel on the image sensor has the functionality of a local spatial-feature detector, provided that the image is kept in permanent movement and that this movement is known at all times. Obviously, the whole signal-processing chain leading from raw photocurrents to a detailed interpretation of the visual scene cannot be incorporated into every pixel of a single chip. Instead, in the system described herein, pixel functionality is limited to the detection of significant temporal transitions. At the occurrence of such a transition, a pixel immediately fires a brief pulse. At the periphery of the pixel array, pulses are tagged by the coordinates of the firing pixel and transmitted off-chip in real time. External hardware is used to identify temporal patterns of pulses within each pixel and relate them to the instantaneous scanning movement applied to the image. Image-feature maps resulting from parallel operation of several temporal-pattern detectors can be used as the input of higher-level processing such as landmark recognition or depth estimation. The outcome of this process can be used for autonomous vehicle navigation or other machine vision applications.

The format in which visual data is encoded by such a sensor is very different from existing cameras. In fact, rebuilding the original image from pulse timings for display to human observers would be rather resource consuming, although possible in principle. Instead, the point of this device is to act as a front end to a visual system extracting relevant information from an image in order to provide a machine with visual-

sensing capabilities. In this respect, the proposed scheme has several interesting properties. First of all, spatial features are detected without relying on comparisons between different photoreceptors. Therefore, the issue of fixed-pattern noise is virtually eliminated. Spatial aliasing is also eliminated because scanning movements are continuous. However, temporal aliasing is still present as in most other image sensors because scanning occurs in discrete cycles. Temporal coding of image data in the timing of spikes lends itself to efficient implementation of feature detectors, either with custom VLSI chips or with off-the-shelf digital hardware. Another property of this coding scheme is that pixels communicate information only to the extent that the image contains significant features in their scanning area. This strategy makes more efficient use of the available communication bandwidth than systematic serial readout [13–14]. On the down side, it is clear that every pixel requires substantially more silicon area than a pixel of a conventional camera because some amount of signal processing is carried out locally. This drawback applies to all other image sensors with focal-plane processing as well. However, in the specific design described herein, this limitation is largely offset by the fact that the effective resolution and the amount of information extracted from the image by a scanning pixel is much larger than a fixed pixel can achieve.

An implementation of the approach outlined in this introduction is under development as of this writing. Some design details and intermediate results are presented in the remainder of this section.

### **Light Deflection**

Two different methods for applying mechanical oscillations to an image have been considered. Periodic image movements at a constant velocity on a circular path can be achieved by spinning a tilted mirror in front of the focusing optics (figure 4.8). The mirror must be mounted on the shaft of a motor, which should be tilted at an angle of about  $45^\circ$  with respect to the optical axis of the lens. If the mirror is not exactly perpendicular to the shaft but tilted by a small angle  $\varepsilon$ , rotation of the motor will cause the reflective surface to wobble, thereby causing the image to travel a circular path with a radius of  $2\varepsilon$  in viewing angle. A prototype device using this principle has been designed and built. The measured mirror angle  $\varepsilon$  was  $0.56^\circ$ . A DC motor spinning up to 19,000 rpm was selected in order to provide scanning frequencies in excess of 300 Hz. A magnetic encoder was coupled to the motor in order to indicate the orientation of the mirror at all times.


The spinning-mirror device provides accurate control over the scanning path. However, the size and power consumption of the motor are significant compared to the image-sensing chip. For applications



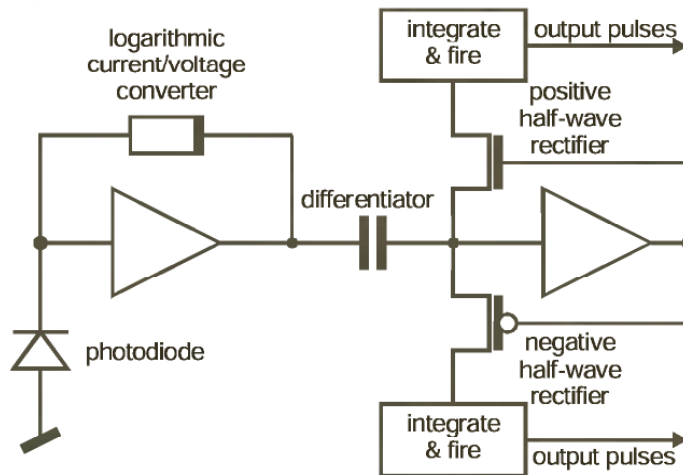
**Figure 4.8** Photograph of a device producing circular scanning by spinning a tilted mirror. The device size is 38 mm × 38 mm × 38 mm.

where space and power are an issue, an alternative device has been developed, where the scanning pattern is produced by displacements of the lens focusing the image onto the chip. In this device, the lens is mounted on springs allowing lateral X-Y displacements but maintaining constant spacing between the lens and the chip. If the system is mounted onto a vibrating platform such as a vehicle driving on a rough surface, the mechanical energy available in the vicinity of the resonance frequency of the lens/spring system will cause scanning movements. To be effective, the amplitude of these movements must be on the order of pixel spacing on the chip—that is, a few tens of microns. In applications where the permanent availability of environmental vibrations is not guaranteed, small piezoelectric actuators could be added to the device. These actuators would need to be turned on only when external vibrations are insufficient. The shape of the scanning path will depend on the relative magnitudes and phases of vibrations applied to the X and Y axes, and on the resonance frequency matching between these axes. As the scanning path will vary over time depending on environmental vibratory conditions, it is necessary to monitor the position of the lens and use this information in the interpretation of the signals generated by the visual-sensing chip. For this purpose, the lens position is monitored by capacitive measurements between the lens socket and surrounding fixed electrodes. A prototype scanning device operating on the principle described herein has been manufactured (figure 4.9). Its dimensions are 34.5 mm by 34.5 mm by 8.2 mm, and its measured resonance frequency is 645 Hz. This frequency is reached with the mass of just the lens and its socket. It can be reduced as needed by attaching additional mass to the lens socket.





**Figure 4.9** Photograph of a scanning device powered by environmental vibrations. The device size is 34.5 mm × 34.5 mm × 8.2 mm.



**Figure 4.10** Block schematic of a single pixel.

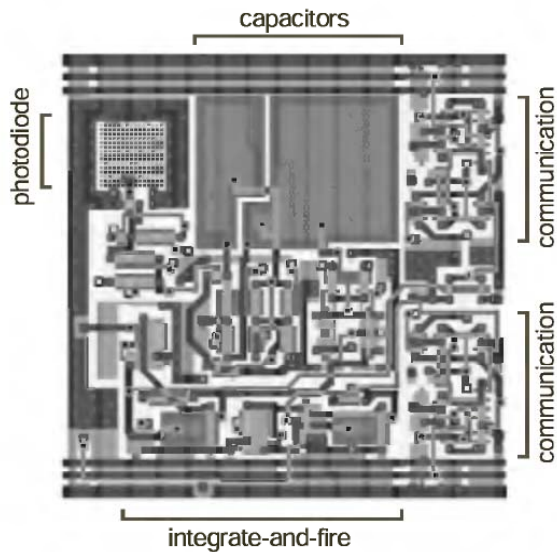
### Image Sensor

An integrated circuit implementing an array of 32 by 32 pixels has been designed and manufactured. A block schematic of a single pixel is shown in figure 4.10. In the first stage of signal processing, the current delivered by a photodiode is applied to a logarithmic current-to-voltage converter. The same visual scene under different illumination levels produces images differing only by a scaling factor in intensity. After logarithmic transform and differentiation, the temporal waveforms produced by scanning are essentially independent of illumination level. Besides logarithmic compression, this circuit also enhances the temporal bandwidth of the photoreceptor with respect to a passive solution. A large bandwidth is crucial to the operation of this visual sensor be-


cause scanning must occur at a much larger frequency than typical fluctuations within the image itself. In addition, photoreceptor bandwidth determines the effective resolution of the visual system. If the photocurrent impulse caused by crossing a thin line is too brief, it will be filtered out and therefore the line will not be detected. It can be shown that at 10 kHz or above, the limiting factor of resolution is no longer bandwidth [15].

The signal resulting from logarithmic compression is differentiated with respect to time and half-wave rectified, whereby both the positive and the negative fraction are retained separately. Current signals delivered at both outputs of the rectifier are sent to independent non-leaky integrate-and-fire circuits, where the charge is accumulated over time until the resulting voltage reaches a threshold. At this point, the integrate-and-fire block emits a short pulse, resets its integrator, and resumes operation. Whenever the scanning path of a photoreceptor crosses a sharp edge, causing an amplitude change exceeding the built-in threshold, at least one spike is reliably generated at this point at every scanning cycle. In another prototypical case where an area of the image contains only a weak intensity gradient instead of a sharp edge, the temporal waveform contains only low-amplitude fluctuations proportional to the magnitude of the gradient. In this case, it takes several scanning cycles—in inverse proportion to the gradient magnitude—before a spike can be generated, and this spike may occur any time the intensity is changing. In the general case, it can be shown that the probability of spiking at a particular phase of a periodical scanning cycle is proportional to the gradient of the image at this point.

The integrated circuit has been manufactured using a CMOS process with a feature size of 0.6  $\mu\text{m}$ , two levels of polysilicon and three levels of metal. A single pixel has a size of 68.5  $\mu\text{m}$  by 68.5  $\mu\text{m}$ , including a photodiode of 10  $\mu\text{m}$  by 10  $\mu\text{m}$ . The entire chip area is about 10  $\text{mm}^2$ . The layout of a single pixel is shown in figure 4.11. With a supply voltage of 3 V, the measured steady current consumed by the analog stages of the chip ranges between 22  $\mu\text{A}$  and 26  $\mu\text{A}$  depending on incident illumination. In addition, the digital communication bus transmitting pulses off-chip consumes about 1 mA when operating at a rate of 1.2 Mpulses/sec. Consumption of this block is roughly proportional to the data rate. Besides the main pixel array, an additional test pixel was incorporated on the chip, which can be stimulated electrically instead of optically for accurate control over experimental conditions. The bandwidth of the photoreceptor circuit was measured by sweeping the frequency of a sine wave at the electrical input emulating the photodiode current, and observing the waveform at a test point located immediately before the rectifier. The bandwidth depends on the DC photocurrent level. Measurement results within the current range of interest are plotted in figure 4.12. The bandwidth exceeds 10 kHz for

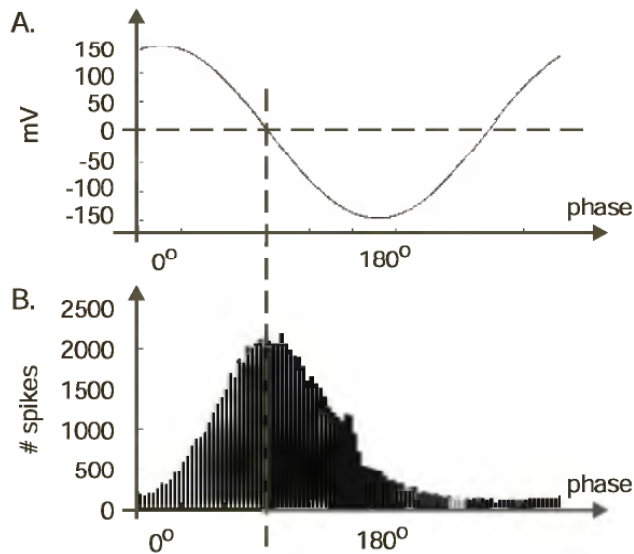


**Figure 4.11** Layout of a single pixel ( $68.5 \mu\text{m} \times 68.5 \mu\text{m}$ ).



**Figure 4.12** Measured bandwidth of the photoreceptor circuit versus photocurrent intensity.

photocurrent levels of 4 pA and above. With the lens built into the devices described in the previous section, this photocurrent level is reached in indoor illumination conditions. Another experiment aimed at verifying that spiking probability is proportional to the gradient of the input signal. For this purpose, a sine wave of constant frequency was applied to the input of the test pixel while recording the times at which pulses are emitted at one of its outputs. Timing of zero-crossings of the input sine wave were recorded simultaneously on the same instrument. Time within each period of the input signal was split into 100 bins, and a histogram was built by counting the number of spikes occurring in each bin over many cycles. As expected, the histogram has the shape of a half-wave rectified sine wave (figure 4.13) with a  $90^\circ$  phase shift with respect to the input signal. Although the chip de-



**Figure 4.13** Measured pixel firing histogram with a sinusoidal input. (A) Input signal. (B) Number of spikes versus phase of the input signal.

scribed in this paragraph has much of the expected functionality, it turned out that intrinsic noise within the pixel circuits caused excessive background firing, thereby making it unfit for acquisition of satisfactory image data. This problem is being addressed by a redesign in progress as of this writing.

## CONCLUSION

Providing robots or vehicles with sufficient visual capabilities to confer them some degree of autonomy in a wide range of environments is a difficult problem. The stumbling block is not light sensing, which is very well mastered in electronics, but rather processing of visual information. Existing cameras provide high quality images fit for display, but when it comes to automatically determining what the image represents, solutions available today typically require either excessive computing time, or an amount of hardware whose volume, weight, and/or power consumption is incompatible with most practical applications. In sharp contrast with this situation, even simple biological creatures possess visual systems powerful enough to let them navigate through their environment and take care of their business. The principles underlying the operation of their visual system are fundamentally different from current mainstream electronic image sensors and associated computational means. The performance gap is such that despite steady progress in integrated circuits manufacturing technology, it is questionable whether the traditional architecture based on a huge pixel array, serial readout, and a low number of powerful processors will ever reach the performance level found in biological organisms.

The field of bioinspired—or neuromorphic—engineering is founded on the premise that principles underlying biological computational structures can be used as the basis of electronic designs and can confer them a comparable level of performance. More than a decade after the emergence of this field [6], it is somewhat embarrassing to admit that few visual sensor designs following this approach have ever reached the stage of a commercial product [16]. It seems appropriate to inquire about the causes of this slow development. The relative modesty of means invested into this approach—compared to mainstream electronics—comes to mind. Other circumstances might also slow down neuromorphic engineering, such as a drop in popularity of analog circuit design, compared to digital VLSI, software or communication systems design. However, the most serious limitation might be the current degree of understanding of biological structures. The most common abstractions used today to grasp the functionality of neural structures implementing visual processing might simply miss essential points. For instance, although many neural models consider only the mean firing rate of neurons as relevant, it is now commonly recognized that the detailed timing of spikes also matters in many cases. Without stepping from a mean-firing-rate abstraction to a more detailed description, the operation of some neural circuits simply cannot be understood. In the field of vision, it is possible that exceedingly reduced abstractions hinder the discovery of the key features that really confer biological architectures their power. One of the missing dimensions might be the fact that an eye is not just an array of photoreceptors passively measuring incoming light, but a subsystem in permanent motion whose movements are the basis of a reliable visual data acquisition process. It is the author's belief that the premise of bioinspired engineering is correct, but that radically new perspectives on the operation of visual systems are still to be discovered, and that such perspectives will eventually lead to powerful hardware implementations of such systems. The work described in the present chapter is meant to be a modest step in this direction by attempting to consider the electronic and mechanical aspects of visual sensing within the same designs.

## ACKNOWLEDGMENTS

The oculomotor system has been developed at CSEM SA in Neuchâtel, Switzerland with basic research funding from the Swiss government. Besides the author, the project team includes Patrick Debergh, Friedrich Heitger, Stève Gyger, and Eduardo Franzi. Additional contributions by Johann Bergqvist, Lorenzo Zago, William Beaudot, Philippe Venier, Alessandro Mortara, and Eric Vittoz are gratefully acknowledged.

The vibrating visual sensor is under development at Caltech in Pasadena, California under funding by DARPA/ONR and the Center

for Neuromorphic Systems Engineering, as part of the National Science Foundation Engineering Research Center program. Besides the author, the project team includes Ania Mitros, Theron Stanford, and Christof Koch.

The author is very grateful to Ania Mitros for revising the manuscript and suggesting numerous improvements.

## REFERENCES

- [1] Yarbus, A. (1967). *Movements of the Eyes*. New York: Plenum Press.
- [2] Land, M. F. (1999). Motion and vision: why animals move their eyes. *J. Comp. Physiol. [A]* 185: 341–352.
- [3] Ditchburn, R. W. (1973). *Eye-Movements and Visual Perception*. Oxford: Oxford University Press.
- [4] Land, M. F. (1972). Mechanisms of orientation and pattern recognition by jumping spiders (Salticidae). In R. Wehner (ed.), *Information Processing in the Visual Systems of Arthropods*. Berlin: Springer Verlag, pp. 231–247.
- [5] Hengstenberg, R. (1972). Eye movements in the housefly *Musca domestica*. In R. Wehner (ed.), *Information Processing in the Visual Systems of Arthropods*. Berlin: Springer Verlag, pp. 93–96.
- [6] Mead, C. (1989). *Analog VLSI and Neural Systems*. Reading, Mass.: Addison-Wesley.
- [7] Landolt, O., and Gyger, S. (2000). An oculo-motor system with multi-chip neuromorphic analog VLSI control. In S. A. Solla, T. K. Leen, and K.-R. Miller (eds.), *Advances in Neural Information Processing Systems 12*. Cambridge: MIT Press, pp. 710–716.
- [8] Landolt, O. (1998). *Place Coding in Analog VLSI—A Neuromorphic Approach to Computation*. Dordrecht: Kluwer Academic Publishers.
- [9] Venier, P. (1997). A contrast sensitive silicon retina based on conductance modulation in a diffusion network. In *Proceedings Sixth International Conference on Microelectronics for Neural Networks and Fuzzy Systems*, Dresden.
- [10] Pouget, A. (1998). Statistically efficient estimation using population coding. *Neural Comput.* 10: 373–401.
- [11] Itti, L., and Koch, C. (2000). A saliency-based search mechanism for overt and covert shifts of visual attention. *Vision Res.* vol. 40, nos. 10–12: 1489–1506.
- [12] Itti, L. (2000). Models of bottom-up and top-down visual attention. Ph.D. thesis, California Institute of Technology.
- [13] Mortara, A., Vittoz, E., and Venier, P. (1995). A communication scheme for analog VLSI perceptive systems. *IEEE J. Solid-State Circuits* vol. 30, no. 6 (June): 660–669.
- [14] Boahen, K. (1996). Retinomorph vision systems: II. Communication channel design. In *Proceedings IEEE International Symposium on Circuits and Systems (ISCAS'96)*, Atlanta.
- [15] Landolt, O., Mitros, A., and Koch, C. (2001). Visual sensor with resolution enhancement by mechanical vibrations. In *Proceedings Nineteenth Conference on Advanced Research in VLSI*, Salt Lake City.
- [16] Arreguit, X., van Schaik, F. A., Bauduin, F., Bidiville, M., and Raeber, E. (1996). A CMOS motion detector system for pointing devices. In *Proceedings of the International Solid-State Circuits Conference '96*, San Francisco.

**This page intentionally left blank**

# 5

## Using Biology to Guide Development of an Artificial Olfactory System

Joel White, Shalini Mall, and John S. Kauer

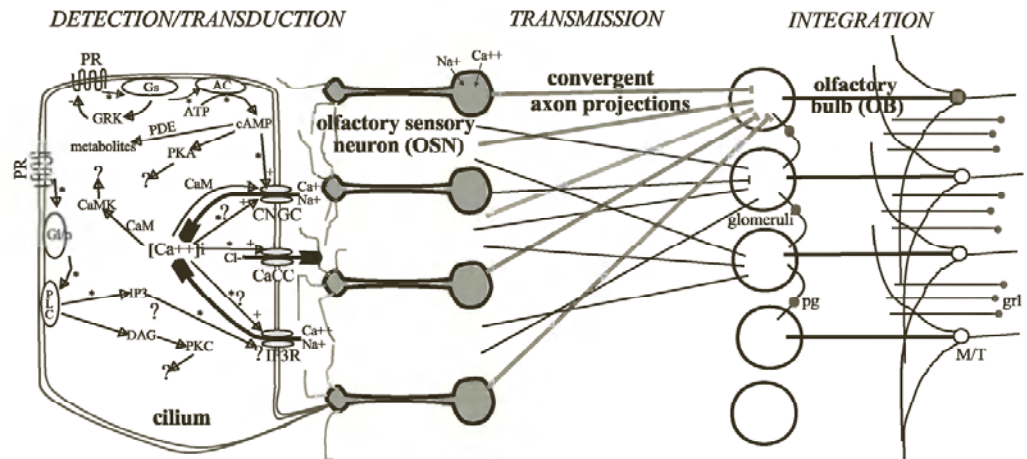
### DISTRIBUTED PROCESSING AND OLFACTORY FUNCTION

It has long been thought that odor properties of vapor phase compounds are encoded and represented by olfactory systems of terrestrial animals using processes that are distributed in time and space (see Adrian, 1953; Moulton and Beidler, 1967; Kauer, 1987; Hildebrand and Shepherd, 1997; Laurent, 1999; Christensen and White, 2000; Kauer and White, 2001). Strong evidence supporting this view has accumulated from many anatomical, electrophysiological, biochemical, behavioral, and molecular biological observations. A summary of some of the details of how the spatial and temporal hypothesis may work is briefly described below. This description of the olfactory process is presented in order to give a short review of its functional components as they pertain to the development of our artificial olfactory system device (figure 5.1). For a more detailed discussion of the many biological components of this process, there are recent reviews that cover various aspects of the physiology, biochemistry, and molecular biology of these events (see Ache and Zhainazarov, 1995; Buck, 1996; Hildebrand and Shepherd, 1997; Schild and Restrepo, 1998; Mori, 1999; Laurent, 1999; Nagao, and Yoshihara, Christensen and White, 2000).

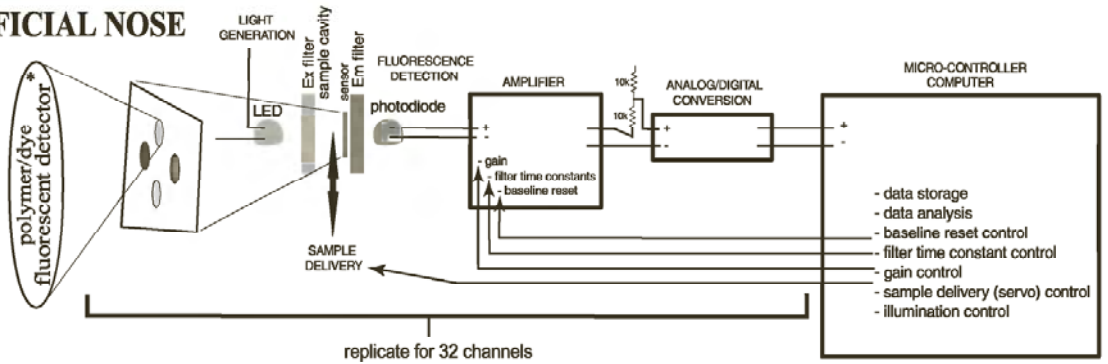
1. For air-breathing vertebrates, odorant stimulus vapors most often consist of mixtures of chemicals rather than pure compounds, with the amount of each chemical species in the air depending on its quantity in the mixture and on its vapor pressure. Odorant vapors are taken into the nasal cavity by an active, rhythmic, sniffing process, which is repetitive and usually in synchrony with the inspiration phase of the respiratory cycle. How the vapor phase molecules are distributed to the sensory cells in the nose depends on a number of complex and still rather poorly defined events that include the sniffing process itself (such as rise time, fall time, and duration of the sniff, and on the flow rate and tidal volume of the inspiration) as well as on the aerodynamic properties of the complex shape of the nasal cavity itself. In animals



## BIOLOGICAL OLFACTORY SYSTEM



## ARTIFICIAL NOSE



**Figure 5.1** (Top) Diagram of some of the major attributes of the peripheral olfactory system. Enlarged cilium at left shows schematic details of some of the biochemical cascades in the odorant transduction process ranging from odorants interacting with seven transmembrane receptors (PR) through G-proteins ( $G_s$ ,  $G_i$ ;  $G_o$ ), adenylate cyclase (AC) to generate camp which opens cyclic nucleotide gated channels (CNGC) and phospholipase (PLC) to generate inositol triphosphate ( $IP_3$ ) and diacyl glycerol (DAG) to open  $IP_3$ -gated channels (see Schild and Restrepo 1998 for more details). Olfactory sensory neurons project to OB glomeruli such that OSNs expressing a particular receptor converge on 1–2 glomeruli. Periglomerular (pg), mitral/tufted (M/T), and granule (grl) cells are the major components of the OB circuits (see Hildebrand and Shepherd 1997; Christensen and White 2000). (Bottom) Components of the artificial olfactory system beneath their approximate correlates in the biological pathway. Polymer fluorescent dye (sensors) detectors are illuminated by light emitting diodes (LEDs) and fluorescent changes are detected by photodiodes which lead into amplifiers and A/D converters to generate spatial/temporal arrays of signals to be analyzed by the microcontroller computer.

with good senses of smell (so-called macrosmats), such as dogs, the nasal cavity is highly convoluted, affording a large surface area of sensory epithelium. The sensory surface area is further elaborated by hair-like ciliary extensions on the luminal ends of the sensory neurons.

2. Once the odorants have entered and have been distributed around the nasal cavity lined with olfactory sensory neurons (OSNs), the molecules dissolve into the mucus overlying the OSNs before they interact with molecular receptors in the membranes of the cilia.

3. Odorant molecules bind to proteins that are part of a large family of 7 transmembrane, G-protein-linked receptors. There are estimated to be as many as 1000 different types of these receptors in mammals.

4. Binding of the odorant to the receptors generates a series of biochemical events in at least one (adenylate cyclase, generating cyclic AMP) and possibly another (phospholipase C generating inositol triphosphate) second messenger cascade that ultimately lead to the opening of channels in the OSN cell membrane. Channel opening usually leads to depolarization of the olfactory sensory neuron, although a few examples of hyperpolarizing potentials have also been reported.

5. Depolarization of the OSN generates temporally patterned trains of action potentials in each of the OSNs that respond to the presented odorant.


6. OSNs expressing a particular receptor are distributed within broadly defined zones in the olfactory epithelium. Within a zone, these OSNs are distributed randomly among a number of other receptor types.

7. Upon stimulation with any one odorant compound, many, but not all, OSNs respond. Thus, even a monomolecular odorant stimulus interacts with a number of different odor receptor types. In addition, there are many individual OSNs expressing any one receptor type.

8. As a result of events 6 and 7, an odorant stimulus activates large numbers of many different types of OSNs, generating temporally patterned firing that depends on the timing of arrival of the odorant pulse, on the intrinsic response properties of each OSN, and on the distribution of the stimulus by the physical properties of the sniff and aerodynamics of the nasal cavity. The widely distributed spatial/temporal response patterns arising from these events are the first step in odorant encoding across the OSN population.

9. OSNs that express one receptor type project their axons in a convergent manner onto, usually, two glomerular synaptic termination sites in the olfactory bulb (OB).

10. Stimulation with an odorant generates spatial and temporal patterns of glomerular activation and spatial and temporal patterns of OB output neuron (mitral and tufted cells) activation that are shaped by



**Figure 5.2** Simplified summary of the combinatorial hypothesis of odorant representation by the olfactory pathway of air-breathing vertebrates. A stylized odorant molecule (upper left) has “attributes” (undefined) that are represented by subpopulations of OSNs and the glomeruli and bulbar circuits to which they connect. This is the spatial aspect of odorant representation in which different subsets of the circuit respond to different elemental odorant properties. The temporal aspect of the coding process, manifest in the firing patterns of the OSNs and bulbar cells, is not shown.

the incoming patterns from the OSNs as well as by the local circuits within the OB. (figure 5.2)

11. The spatial and temporal patterns arising from the OB circuitry project onto a number of higher-order olfactory pathway nuclei about which relatively little is known with regard to odorant encoding mechanisms.


## RATIONALE FOR DEVELOPING AN ARTIFICIAL DEVICE

### Modeling


We have based the development of our artificial nose device on principles of olfactory function (as described above) that we and others have studied in biological preparations. In our lab, we have focused on the olfactory system of the salamander and have used data collected over many years from this preparation to assemble a mathematical model of the major events in the peripheral olfactory pathway (White et al., 1990, 1992, 1996, 1998; White and Kauer, 1999) (figure 5.3). The ability of this model to capture and compute a number of the first-order spatial and temporal events in the olfactory system formed the basis for the hardware implementation of the many biological principles that we have incorporated into the artificial olfactory system described here.

I.

A




B



**Figure 5.3** Computer simulation of olfactory system model. (I.A) Schematic diagram of the simplified, compartmental representation of the components of the peripheral olfactory system. The numbers indicate the number of each component in the model. REC, receptor cells; GLM, glomeruli; PGL, periglomerular cells; MIT, mitral/tufted cells; GRL, granule cells. (I.B) Diagram shows how the 6000 receptor cells are spatially represented and how they connect in a convergent fashion onto the OB. (II) Comparisons of computed intracellular recordings from mitral cells simulated in the model after low intensity orthodromic electrical stimulation (II.A model); high intensity orthodromic (II.B model); and antidromic stimulation (II.C model) compared with intracellular records from salamander mitral/tufted cells (A. salamander, B. salamander, C. salamander). (III) Same as (II) But using simulated odorant stimulation. Note correspondence between the complex changes in temporal firing patterns in the model and those seen in salamander mitral/tufted neurons.

II.


**A Low Intensity Orthodromic**



b. Salamander




**C Antidromic**



b. Salamander



**B High Intensity Orthodromic**




b. Salamander




III.


**A Low Intensity Odor**



b. Salamander



**B High Intensity Odor**



b. Salamander




Figure 5.3 (continued)

Figure 5.3.I.A shows the components of the olfactory pathway represented in the model. These components are connected together in the model with appropriate ratios of cells, connection polarities, and synaptic weights as derived from the known circuit connections in the biological system. Figure 5.3.I.B shows the distributions of the receptor elements in the olfactory epithelium and connections with elements of the OB. Figure 5.3.II and III show the computed outputs of the model after simulated electrical (II) and simulated odorant stimulation (III). These computed intracellular recordings from the output neurons are compared with actual intracellular recordings from salamander olfactory bulb mitral/tufted cells.


Our ability to capture a number of the complexities of odorant representation in this computational model was one of the forces driving development of the artificial olfactory system described in the next section. This hardware implementation of the biological system extends the representation of the circuits in the model by including odorant delivery (sniffing), chemically sensitive detectors, spatial and temporal pattern-recognition algorithms, and an odorant descriptor (spoken word) output.

## DESCRIPTION OF THE ARTIFICIAL DEVICE

### Overview

The concept of an “artificial nose” was first introduced by Persaud and Dodd (1982) and a number of artificial nose devices have been developed since that original description. All of these devices incorporate two main features characteristic of artificial noses: 1. an array of broadly specific sensors and 2. a pattern recognition method for processing sensor output. The artificial olfactory system that we describe here (Tufts Medical School Nose, TMSN; figure 5.4) builds on this early artificial nose concept and expands on it by incorporating numerous additional attributes of the biological system to enhance its sensitivity, discriminative ability, and functionality.

The TMSN uses an array of chemical sensors composed of dyes, polymers, and dye/polymer mixtures that change their fluorescence upon exposure to brief pulses of airborne compounds (i.e., *odorants*). Odorants are delivered to the sensor array via a suction fan and valving system that generate brief inhalation pulses (i.e., “sniffs” of 0.5–2.5 sec duration) for sampling the ambient environment. The change in fluorescence produced by each sensor over time is amplified and digitized, and these temporal data are used for vapor detection and identification. Data acquisition and processing are controlled by an embedded microcomputer (right side of figure 5.4a). User control of the device is via a touch screen panel on the microcomputer. The modes of analyte de-



**Figure 5.4** (a) Overall schematic of major components of the TMSN artificial olfactory system. (b) Details of the sensing chamber with inlet airflow at bottom left; the photo-sensing chamber where changes in sensor fluorescence are detected by photodiodes which generates spatial/temporal patterns that are sent to the computer for analysis. Each step of the process from input sniffing to output analysis is under feedback computer control.



livery and the data analysis algorithms for odor detection and discrimination are based in large part on processing steps thought to occur in the biological olfactory system.

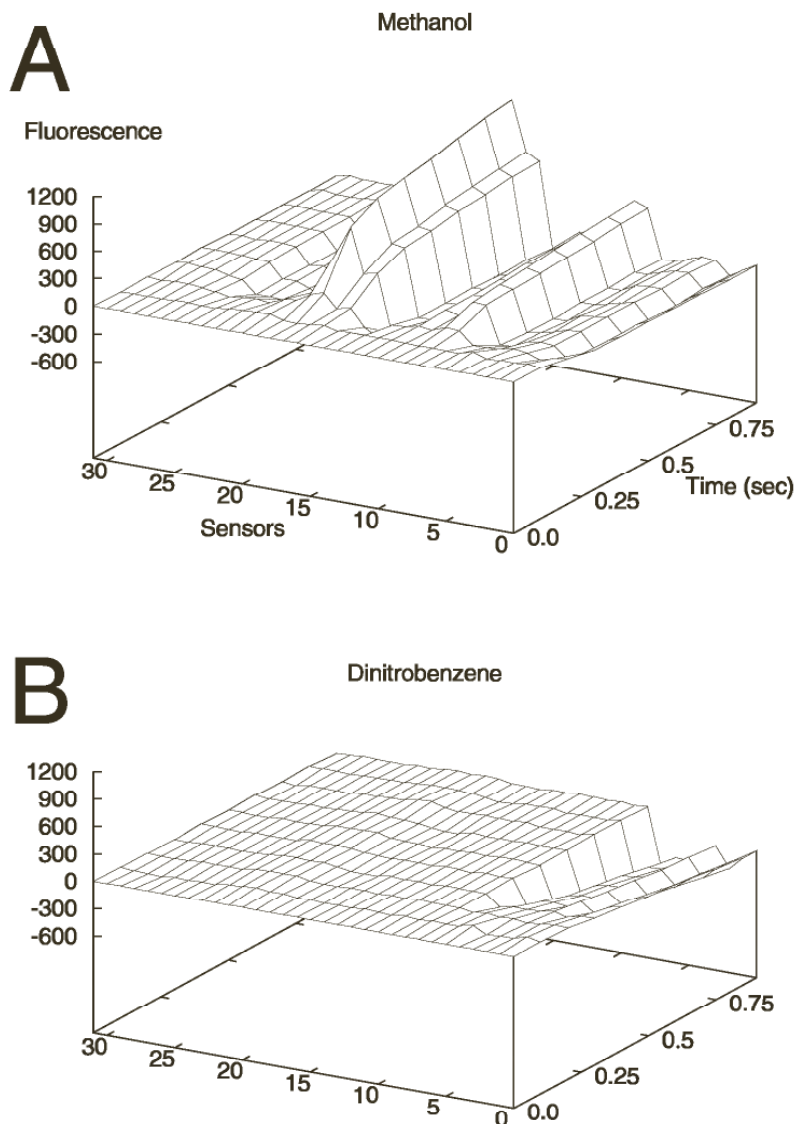
## Sensors

The device contains 32 sensors that can be illuminated and observed at 16 different excitation and emission wavelengths. The sensors are located in a *head stage* box (left side of figure 5.4a) and are placed along a narrow chamber (figure 5.4b) through which ambient air is drawn. The optical elements for illuminating and monitoring the sensors are positioned along the sides of the chamber. Excitation light is produced by LEDs providing wavelengths appropriate for the sensing materials being used. The excitation light passes through optical filters (which reduce “bleed” through to the emission filters) and illuminates the sensors. The sensors fluoresce at wavelengths longer than those used for excitation—this light passes through emission filters (at wavelengths appropriate for the particular materials) and is monitored by photodiodes. Thus, there is an LED/photodiode pair for each sensor channel. Electric current produced by each photodiode is converted to voltage, which is then amplified. The amplified signal from each sensor is digitized over time at 12-bits resolution by the embedded microprocessor. The sensors produce rapid changes in fluorescence (< 2 sec) with pulses of odor (see below).

Sensors are constructed from fluorescent dyes, intrinsically fluorescent polymers, and combinations of fluorescent dyes with nonfluorescent polymers. All sensing materials are deposited onto fibrous substrates with large surface areas. In our system, this method of sensor construction yields larger signal amplitudes and improved signal-to-noise compared to thin film sensors, such as those produced by spin-coating onto glass coverslips.

For our land mine detection efforts, eight sensors incorporate intrinsically fluorescing polymers developed by Swager and colleagues (Yang and Swager, 1998) for detecting DNT, TNT, and other nitro-aromatic compounds. Although sensors constructed with these polymers are sensitive to nitro-aromatic compounds, they can also show changes in fluorescence to high concentrations of other nonexplosive compounds (e.g., methanol; figure 5.5; sensors 1–8 at right on X-axis).

The remaining 24 sensors in the TMSN array were developed in our laboratory and include solvatochromic dyes (Nile red [Barnard and Walt, 1991] and other dyes), used either alone or in combination with various polymers examined at several different wavelengths. These sensors are broadly responsive, each producing signals to a variety of organic vapors. Sensors are selected and wavelength optimized by testing with pure compounds containing different functional groups



**Figure 5.5** Spatial/temporal patterns generated by methanol (A) and dinitrobenzene (B). Sensor number is on the X-axis; duration of sniff on the Z-axis (into the page); and change in fluorescence above and below baseline is on the Y-axis.

(e.g., methanol, acetone, amyl acetate, octane, xylene, tetramethylethylenediamine). The variety of sensor types in the device allows the potential detection and discrimination of a range of chemical compounds using the data processing and pattern-recognition algorithms described below.

### Vapor Sampling

The sampling method of the TMSN is inspired by the sniffing behavior of animals. Odorants are presented in a pulsatile fashion (sniffs) to the sensor array through a fan and valve arrangement. In the device shown

in figure 5.4a, there are two fans. One draws odorous air over the sensors through the odor inlet port (“inhale”); the other draws clean air from behind the sensor head, blowing it over the sensors and out the odor port (“exhale”).

Valves are positioned by servomotors, which are controlled by the microprocessor. Sniff timing can be dynamically controlled through feedback (see next section). Sniff volume is controlled by the valves and by the power of the fans, where more powerful fans produce higher air velocity at the odor inlet port and therefore higher flows and larger sample volumes.

All aspects of data acquisition and timing in the TMSN are controlled by the embedded microprocessor (right side of figure 5.4a). To reduce sensor photobleaching, the LEDs are illuminated only long enough for amplifier settling and signal digitization. An important aspect of the amplifier electronics is the ability to reset the amplifier to zero at the beginning of a sniff. Using this method, a large input due to a bright sensor can be offset to zero, allowing only the changes in fluorescence to be amplified and digitized. In this way, small changes in a large fluorescence background can still be measured with high gain. The digitized data are stored in memory on the embedded microprocessor for data processing and pattern recognition.

### **Description of Signals**

Examples of the temporal signals produced by the entire array of 32 sensors (X-axis) upon exposure to various odorants at high concentration are shown in figure 5.5. Upon exposure to odorant vapors, sensor signals evolve rapidly over time. In general, sensors are broadly responsive, showing fluorescent change to a variety of odorants. This fluorescent change can vary in time, course, and amplitude. When viewed across sensors, these signals produce a spatiotemporal signature of the odor (figure 5.5). This signature varies with the odor stimulus and enables accurate odor detection and identification (next section).

### **Pattern Recognition**

Most artificial nose systems evaluate the odorant/sensor interaction by measuring sensor output once equilibrium between the analyte vapor and sensor materials has been reached. It is not advantageous to wait for equilibrium in a sensing system requiring rapid, repeated sampling, as is necessary for land mine detection. In order to sample rapidly, our device takes brief sniffs so that equilibrium between the analyte vapor and the sensor material is not always attained. We have developed a number of data processing strategies that accommodate these dynamic, nonequilibrium responses.

A *one-of-k* algorithm is used when the number of target odors is small and their identities defined. This algorithm is useful for investigating the sensitivity and discrimination ability of the device for pure compounds in the laboratory and we have used it in the field (see below). For this method, samples of the target odors are acquired as described above. Sensors not producing appreciable signal are removed from analysis by the algorithm. This eliminates noise that those sensors might have contributed to the overall pattern. Temporal signals from the remaining sensors for each target odorant are used as templates for testing. During an odor test, a sample is acquired and the sensor signals are matched against the templates of the target odors. A match score is determined by calculating the sum-of-squares of the difference between each data point for the target and the test data. The target that is the closest match is identified as the unknown odorant.


### **Feedback Control of Data Acquisition**

Another aspect of TMSN operation that is under development is an “intelligent” mode, whereby sampling and data acquisition parameters for a sniff are tuned based on the results of one or more previous sniffs. In a fashion similar to the way animals behave, the device thus adapts its sniffing for the particular odor environment being investigated. For example, a slowly developing sensor signal may not generate sufficient sensor signal for discrimination during a brief sniff. Taking a longer second sniff would allow the signal to develop more fully, improving discrimination.

Our ultimate goal is a system that learns the nuances of a particular land mine area, much the way a trained dog appears to adapt to particular field situations. The device will adapt both its method of sampling and its knowledge base of odor signatures.

### **Test Results**

**DNT Detection under Controlled Conditions** We evaluated the detection performance of the TMSN in carefully controlled tests performed at the Institute for Biological Detection Systems (IBDS) at Auburn University in collaboration with Dr. Paul Waggoner. These tests were carried out in the same apparatus used to test odor generalization and thresholds in dogs for a number of compounds, including explosives and mine-related materials. The tests of our artificial olfactory system serve to evaluate the performance of the sensing elements as well as the data-processing algorithms. Explosives fate and transport studies (George et al., 1999) suggest that 2,4-DNT is likely to be the predominant mine-related compound in the soil and in the air above land mines—therefore, we have focused on testing the response of the




**Figure 5.6** TMSN performance over a range of DNT concentrations. Performance was tested under controlled conditions in IBDS dog testing apparatus. Delivered analytes were methanol (1:20 dilution), dinitrotoluene (at various concentrations), and background air. Responses to methanol were discriminated with 100% accuracy and are not shown. The TMSN discriminated among these stimuli in the automatic mode. Threshold for TMSN performance for DNT dropped to chance (1 of 3) between 1–2 and 4–5 ppb.

artificial system to DNT. As a basis for comparison, recent data indicate that behavioral detection thresholds to DNT for dogs tested in the IBDS apparatus range from 0.15 to 1 ppb (for different dogs), with false alarm rates at about 20–30% (Waggoner, unpublished).

For the IBDS tests, vapors from DNT and other compounds at various concentrations were generated by an air-dilution olfactometer (Hartell et al., 1998). DNT concentrations were measured explicitly after the odor trials by gas chromatography/mass spectroscopy. Compounds were presented in random order to the device every 25 sec. Test runs contained at least 10 presentations of each odor at each concentration. The IBDS computer system then triggered data acquisition by the TMSN. During these tests, odor detection and discrimination occurred in real time after each odorant presentation using the one-of-k algorithm described above. In this test situation, the TMSN was running in a “hands-off” mode, making discriminations without operator intervention. Data were also logged to a laptop computer for off-line analysis.

Results from a DNT concentration study using the paradigm described above are summarized in figure 5.6. Data for each DNT concentration were acquired in separate sets of runs. For each set of runs, target patterns for DNT, a 1:20 dilution of methanol (a possible interferent for some sensors—see figure 5.5), and blank air were established. Test runs of all three conditions (DNT, methanol, and blank) were then applied in random order. Only results for DNT and blank air are



**Figure 5.7** TMSN performance across a scan line over a TMA5 anti-tank mine at Fort Leonard Wood, MO. Samples were taken along a trajectory that ran north/south over an anti-tank mine marked by two flags about 2 feet apart. Response at each location (6" intervals) on the X-axis is shown by degree of correct identification (Y-axis) after training at a site over the mine and at a control site off to the side.

shown in figure 5.6—detection and discrimination of methanol vapor was 100% correct.

As seen in figure 5.6, TMSN showed 100% DNT detection and discrimination from methanol down to 4 ppb, with no false alarms. Performance fell to chance (33%) for 2 ppb DNT. The TMSN detection limits for DNT therefore appear to be within an order of magnitude of dog thresholds tested in the same apparatus.

**Field Results from Fort Leonard Wood** Land mine detection tests of the TMSN have been conducted at the DARPA land mine facility located at Fort Leonard Wood, Missouri. The facility consists of a large grassy field in which a variety of antitank and antipersonnel land mines have been buried. In a “calibration” area of the field, the mines are marked with flags indicating their type and position. We have tested the TMSN over a TMA5 antitank mine in this calibration area.

As shown in figure 5.7, the TMSN identified a signal between the flags marking the land mine position. It gave no or little signal south of the flags. No trials were taken north of the flags. The bars represent the average of five tests at each position. These results do not show that our device can unequivocally detect land mine chemical signatures but suggest that we are approaching the sensitivity and discriminability necessary to approach such a complex, low-concentration task. Additional tests are needed to ensure that the signals are specific to land mines and show acceptable levels of false negative rates.

## SUMMARY

We have designed an artificial system that mimics a variety of different attributes of the biological olfactory system including controlled, pulsatile odorant delivery, the use of a variety of cross-reactive detectors, generation of spatial and temporal response patterns, and pattern-recognition algorithms motivated by analysis of biological circuits. We have defined a set of target odorants (those associated with land mine chemical signatures) that the device is designed to detect and we have tested the device in both controlled laboratory and open-field settings with moderate success. For this particular, difficult, low-concentration chemical detection task, we are encouraged that the use of a biological understanding of olfactory function to guide artificial system development has met with sufficient success to warrant continued investigation.

## ACKNOWLEDGMENTS

This work was supported by grants from NIH (NINDCD), the Office of Naval Research (ONR), and the Defense Advanced Research Projects Agency (DARPA). We thank Dr. Paul Waggoner and his colleagues of Auburn University IBDS for graciously allowing us to test our device in the highly controlled and calibrated conditions of their dog-training apparatus.

## REFERENCES

- Ache, B. W., and Zhainazarov, A. (1995). Dual second-messenger pathways in olfactory transduction. (Review, 48 refs.) *Curr. Opin. Neurobiol.* 5: 461–466.
- Adrian, E. D. (1953). Sensory messages and sensation. The response of the olfactory organ to different smells. *Acta Physiol. Scand.* 29: 5–14.
- Barnard, S. M., and Walt, D. R. (1991). A fibre-optic chemical sensor with discrete sensing sites. *Nature* 353: 338–340.
- Buck, L. B. (1996). Information coding in the vertebrate olfactory system. (Review, 187 refs.) *Annu. Rev. Neurosci.* 19: 517–544.
- Christensen, T. A., and White, J. (2000). Representation of olfactory information in the brain. In T. E. Finger, W. L. Silver, and D. Restrepo (eds.), *The Neurobiology of Taste and Smell*. New York: Wiley-Liss, pp. 197–228.
- George, V., Jenkins, T., Leggett, D., Cragin, J., Phelan, J., Oxley, J., and Pennington, J. (1999). Progress on determining the vapor signature of a buried landmine. In *Proc. 13th Annu. Int. Symp. Aerospace/Defense Sensing, Sim. Controls.* pp. 258–269.
- Hartell, M., Myers, L., Waggoner, L., Hallowell, S., and Petrousky, J. (1998). Design and testing of a quantitative vapor delivery system. In *Proc. 5th Int. Symp. on Anal. Det. Explosives.*
- Hildebrand, J. G., and Shepherd, G. M. (1997). Mechanisms of Olfactory discrimination: Converging evidence for common principles across phyla. *Annu. Rev. Neurosci.* 20: 595–631.

- Kauer, J. S. (1987). Coding in the olfactory system. In T. E. Finger and W. L. Silver (eds.), *The Neurobiology of Taste and Smell*. New York: John Wiley and Sons, pp. 205–231.
- Kauer, J. S., and White, J. (2001). Imaging and Coding in the olfactory system. *Annu. Rev. Neurosci.*
- Laurent, G. (1999). A systems perspective on early olfactory coding. *Science* 286: 723–728.
- Mori, K., Nagao, H., and Yoshihara, Y. (1999). The olfactory bulb: Coding and processing of odor molecule information. (Review, 51 refs.) *Science* 286: 711–715.
- Moulton, D. G., and Beidler, L. M. (1967). Structure and function in the peripheral olfactory system. *Physiol. Rev.* 47: 1–52.
- Persaud, K., and Dodd, G. (1982). Analysis of discrimination mechanisms in the mammalian olfactory system using a model nose. *Nature* 299: 352–355.
- Schild, D., and Restrepo, D. (1998). Transduction mechanisms in vertebrate olfactory receptor cells. [Review] [377]. refs. *Physiol. Rev.* 78: 429–466.
- White, J., Dickinson, T., Walt, D., and Kauer, J. S. (1998). An olfactory neuronal network for vapor recognition in an artificial nose. *Biol. Cybern.* 78: 245–251.
- White, J., Hamilton, K. A., Neff, S. R., and Kauer, J. S. (1992). Emergent properties of odor information coding in a representational model of the salamander olfactory bulb. *J. Neurosci.* 12: 1772–1780.
- White, J., and Kauer, J. S. (1999). Odor recognition in an artificial nose by spatio-temporal processing using an olfactory neuronal network. *Neurocomputing* 26–27: 919–924.
- White, J., Kauer, J. S., Dickinson, T. A., and Walt, D. R. (1996). Rapid analyte recognition in a device based on optical sensors and the olfactory system. *Anal. Chem.* 68: 2191–2202.
- White, J., Neff, S. N., Cinelli, A. R., and Kauer, J. S. (1990). Modelling the salamander olfactory bulb: Single cell and network interactions. *Neurosci. Abstracts* 16: 403.
- Yang, J.-S., and Swager, T. M. (1998). Fluorescent porous polymer films as TNT chemosensors: Electronic and structural effects. *J. Am. Chem. Soc.* 120: 11864–11873.



**This page intentionally left blank**

# *Biomimetic Actuators*

**This page intentionally left blank**

# 6

## SMA Actuators Applied to Biomimetic Underwater Robots

Jan Witting, Koray Safak, and George G. Adams

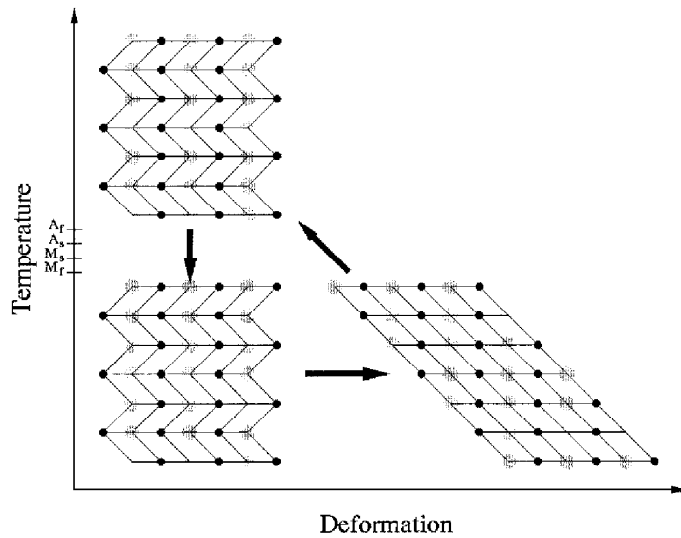
The work on Ni-Ti (nickel-titanium) shape memory alloy actuators presented here arises from the need for a biomimetic linear actuator, in the context of building a biomimetic ambulatory underwater robot. The model animal used is the American lobster *Homarus americanus*, chosen due to its behavioral robustness in dealing with the complex benthic habitat identified as the target environment for the vehicle. As reported elsewhere in this book Chapter 12, the biomimetic aspect of our effort consists of not only reverse kinematic analysis of the gait patterns of *Homarus*, but also extends to looking at the underlying control signals produced by the nervous system. Finally, the response of the muscles to these signals, in the form of electromyographical recordings, have also been used to generate the robot's control architecture. Realizing that the gait of an ambulator is the sum of control signals from the nervous system and the response of the actuators (muscles) to those signals, it becomes very important to match the physical actuator characteristics to the biological model. Timing parameters of control signals from invertebrate motor neurons have coevolved with muscles to produce motion with appropriate velocity and displacement. Because the control signals are recorded directly from the animal, the time domain of the actuator responses to these signals must approximate that of the animal muscles. Ni-Ti shape memory alloy (SMA) was identified as a promising material from which to build an artificial muscle fulfilling the requirements outlined above (Ohkata and Suzuki, 1998). The suitability of Ni-Ti-based artificial muscles for the construction of a biomimetic underwater ambulatory robot was tested in a series of experiments. Stress-strain relationship of the muscles was experimentally measured, together with contraction and relaxation velocities under different load and temperature conditions. Based on this data, a full-body postural model has been developed; this model will be used to aid in the construction of behavioral libraries, as well as fine tuning the control signals to minimize robot power consumption.

## BEHAVIOR OF NI-TI SHAPE MEMORY ALLOY

Understanding the behavior of nitinol wire is the necessary first step toward the construction, modeling, and characterization of the nitinol-actuated biomimetic underwater robot. Nitinol is one of the alloy materials exhibiting the shape memory effect (SME), which enables the material to recover residual strains of up to 8% upon deformation. The SME is due to a unique property of temperature- and stress-induced phase transformation. SMA materials show two important phenomena: The SME, in which the material exhibits large residual strains and fully recovers during a phase transformation initiated by a temperature rise; and a pseudoelastic effect, in which the specimen exhibits large residual strains, which are fully recovered upon unloading in a hysteresis loop. The ability of SMA materials to recover large amounts of plastic strain is due to the martensitic phase transformation initiated in certain conditions of stress and temperature (Wayman and Duerig, 1990).

The SMA material at zero stress and high temperature exists in the austenite crystalline phase, a body-centered cubic (BCC) structure. Upon decreasing the temperature, the material undergoes a self-accommodating crystal transformation into a martensite phase, which is usually a face-centered cubic (FCC) structure. Although there are some exceptions, in most SMA materials (Ni-Ti included), the transformation temperatures occur in the order  $M_f < M_s < A_s < A_f$ , or martensite-finish, martensite-start, austenite-start, austenite-finish temperatures, respectively. To illustrate the process, let us consider an SMA material transformed from the austenite phase at zero stress by cooling through  $M_s$  and  $M_f$  temperatures. This material has multiple variants and crystalline twins present, all with crystallographically equivalent but differing orientations. When the stress level is increased, multiple martensitic twins start converting into stress-preferred twins, which is called *detwinning* (figure 6.1). Detwinning converts the multiple crystal variants into a single variant, the direction of which is determined by alignment of habit planes with the axis of loading. During this process, the stress rises very slightly compared to the applied strain. Upon unloading, the backward transformation (to multiple variants) does not take place, but only a small elastic portion of the strain is recovered, leaving the specimen with a large amount of plastic strain. The plastic strain is recovered by heating the material above the  $A_f$  temperature, which results in a conversion to austenite with no variants. This returns the material to its original shape configuration (before the martensite transformation).

The mechanism behind the pseudoelastic effect (figure 6.2) is somewhat different. Starting with an SMA specimen entirely in the austenite phase ( $T > M_s$ ) at a constant temperature, there is a critical stress level





**Figure 6.1** The shape memory effect (SME). At high temperatures and zero-stress, SMA material exists in the austenite phase. Upon decreasing the temperature, it converts into the “twinned” martensite phase. Deformation at small stress levels results in “detwinning” and yields high amounts of strain, which can be fully recovered by heating the material above the  $A_f$  temperature.

that initiates the austenite to martensite transformation. At this stress level, the transformed phase will be of a single variant with completely detwinned martensite. If the material is unloaded at a temperature  $T \geq A_f$ , the detwinned martensite is transformed back into the austenite phase because of the instability of the martensite phase at these temperatures. Large amounts of plastic strain recovery, due to the pseudo-elastic effect, yield a characteristic hysteresis in the stress-strain diagram (figure 6.2b), similar to the stress-strain relationship of biological muscle.


## CONSTRUCTION OF NITINOL MUSCLE MODULES

The SME exhibited by Ni-Ti SMA can be utilized to construct an artificial muscle (Ohkata and Suzuki, 1998). The principle of operation is simple: Ni-Ti (or nitinol, Buehler et al., 1963) wire is heated by passing an electrical current through it. The generated heat brings about the martensitic transformation, producing a strain in the wire. The strain is conducted to a mechanical system of levers, constituting the robot’s leg. To achieve this in practice, the nitinol wire (250  $\mu\text{m}$  Flexinol brand SMA wire, Dynalloy Inc.) undergoes some preparatory steps. After annealing the wire (500°C, 30 min) in an argon atmosphere, the wire is acid etched in a mixture of  $\text{HNO}_3$  and HF to remove an oxide coating present on the wire’s surface. Removing the oxide coating on the wire provides more-uniform electrical resistance throughout the surface and



**Figure 6.2** (a) Pseudo-elastic effect. At high temperatures ( $T > M_s$ ), application of stress causes the material to convert into austenite phase. At this stress level, the transformed phase will be a single variant with completely detwinned martensite. If the specimen is unloaded at a temperature  $T > A_f$ , the detwinned martensite is transformed back into the austenite phase because of the instability of the martensite phase at these temperatures. Large amounts of plastic strain recovery due to the pseudo-elastic effect yields a characteristic hysteresis loop in the stress-strain diagram (b).

provides for better electrical connections. Acid-etched wires are cut to appropriate lengths for the assembly. The wire is then run through a PTFE insulating sleeve (figure 6.3) to provide both thermal and electrical insulation from the surrounding water. The two ends of the nitinol wires are then passed through a stainless steel crimp (302 SS) and crimped with a pneumatic crimping tool to form a terminal loop of nitinol in either end of the muscle. Kevlar strings are then attached to these loops to provide the muscles a means of a mechanical attachment to the robot legs. The electrical connectors are soldered to the stainless



**Figure 6.3** Schematic of nitinol muscle module.

steel crimps. Finally, the crimped ends are covered with shrink-tubing and potted in a urethane sealant to insulate the terminals from water, and to keep water from entering the polytetrafluoroethylene (PTFE) insulation. The muscle modules act in antagonistic pairs attached to a leg constructed from carbon composite tubing and Delrin (for a description of the legs, see the section entitled Biomimetic Underwater Robot Posture Control with Nitinol Actuation).

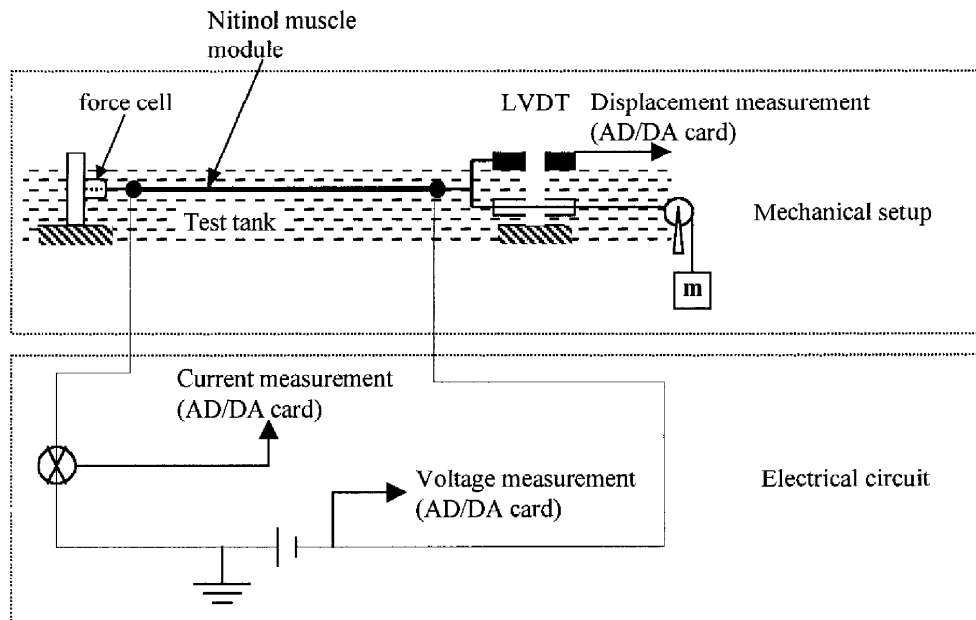
## EXPERIMENTAL CHARACTERIZATION OF NITINOL ARTIFICIAL MUSCLES

Although a number of constitutive models for SMAs exist, their success in predicting nitinol behavior in this application is very limited (Liang and Rogers, 1990; Brinson, 1993). The presence of the PTFE insulation, in combination with the surrounding water mass, creates a very complicated pathway for heat dissipation. In addition, we were unable to accurately measure nitinol temperature inside the PTFE sleeve. Because nitinol temperature is a direct input to all constitutive models, these models were effectively unavailable to us. Thus, an empirical approach was taken, and the behavior of nitinol was studied under various stress/strain and temperature conditions. In effect, a comprehensive look-up table was generated, both for the purposes of generating appropriate input signals as well as for constructing a full-vehicle dynamic postural model of the vehicle.

### Methods

To conduct experiments on assembled artificial-muscle modules, a constant water flow-through test bench was constructed. The system consists of a linear displacement transducer in line with the module and a hanging weight producing stress in the muscle module (figure 6.4). The control signal produced by a program running on an Apple Macintosh IIci drives a custom pulse-width-modulated current supply (Massa Products, Hingham, Mass.). With this system, the muscle input






**Figure 6.4** Schematic of the experimental setup. Nitinol muscle module, contracting in a temperature controlled water bath, lifts the mass  $m$ . During the action the displacement of the movement is recorded by the LVDT, while the actuation signal parameters are recorded via an AD/DA card.

voltage, current, and duty cycle could be continuously varied—the input signal consisting of trains of pulses separated by short quiescent intervals (figure 6.5). Displacement produced by the contraction of the muscle module was recorded by a National Instruments LabView-based virtual instrument (VI) running on an Apple Macintosh G3. The same VI recorded input voltage and current, all analog signals running through a National Instruments PCI/16x A/D acquisition card. Stress of the nitinol was varied by changing the mass of the hanging weight, and the water bath temperature was modulated by an in-line water chiller. With this apparatus, constant load-cyclic displacement tests at different electrical power levels and at different water temperatures were conducted.

### Experimental Procedure

The wire was activated in a cyclic manner—that is, for a predetermined activation (on) time followed by a predetermined relaxation (off) time (figure 6.5). We selected on/off times of sufficient duration so that the wire reached its steady-state minimum displacement level when activated by electric current and again reached its steady-state maximum displacement when the power was turned off. Tests were repeated with loads ranging from 50 to 1000 grams. Nitinol strain was calculated by dividing observed displacement with the austenite length of the wire.



**Figure 6.5** PWM current board output profile. The PWM current driver board produces a current output as shown here.  $T_d$  indicates the duration of the activation and  $T_p$  indicates the period of the individual pulses. Activation time is followed by a relaxation time, during which no electrical power is applied. Total time of activation and relaxation is called the period,  $T$ . The amplitude of the applied electrical current is  $I$ . The board used for the experiments provides a pulse period of 0.01 sec, which corresponds to a pulse frequency of 100 Hz.


Nitinol load-deformation behavior is a function of material temperature. Because the wire was heated by passing through an electrical current, the temperature is a function of applied electrical power and time. The amplitude and duration of the applied electrical power plays an important role in determining the amount of transformation to austenite and hence on the amount of strain recovered upon deformation. The pulse width modulation (PWM) frequency of the current driver board is 100 Hz. To avoid aliasing the data by the PWM pulses, the data acquisition frequency was set to 500 Hz.

For measuring the stress/strain relationship, we loaded the wire with the weights of 50, 100, 200, ..., 1000 grams and applied the electrical power. Signal was delivered at 7 volts, duration of the activation 1 sec, and total period 3 sec. Tests were repeated with pulse widths of 20%, 40%, 60%, and 80%. Displacement, voltage, and current were measured for 15 seconds (5 pulses). The relationship between power, displacement, and ambient water temperature was measured in a similar manner. In this set of experiments, the mass of the weight used was a constant 500 g. Water temperature was varied with an interval of 4°C (10, 14, 18, 22, 26°C). For each temperature, the muscle was activated with a pulse train (1 sec pulse, 3-sec period) in which both voltage and pulse-width duty cycle were varied at 1 V and 10% increments, respectively. For each pulse parameter combination, displacement, voltage, and current were measured for 15 seconds (5 pulses). For each displacement measurement, a corresponding power input level was calculated as  $P = V * I * \% \text{ PWM}$ .

## Results

The displacement, voltage, current, and force measurements were recorded into separate data files for each case of applied weight and PWM. Additional processing of the experimental data was done in the MATLAB environment. The original frequency of the sampling was 500 Hz; this sampling rate was selected to avoid aliasing with the frequency component of the PWM pulses. The frequency component due to the PWM does not convey important information regarding the behavior of the nitinol. It simply contaminates the displacement, current, and force signals with the high-frequency and small-amplitude oscillations. Therefore, we eliminate these pulses by first filtering at a lower cut-off frequency and then down-sampling. We filtered the signals at the cutoff frequency of 10 Hz with a 10-point FIR filter and then down-sampled to 25 Hz.


In a typical experiment, a 500 gram weight was attached to the nitinol wire and actuated repeatedly (figure 6.6). The force-time results show a slight variation in the force exerted on the wire (figure 6.6, filtered signal, 80% PWM), mainly due to the pulley friction and the




**Figure 6.6** A typical relationship of displacement, load and current in time. A 5 N force (imposed by a 500 g load) resists muscle module contraction and stretches the muscle back to full length. The observed rise times are typical, in the order of 250 ms to 90% contraction. The variation in the load is due to frictional losses and inertial forces when the mass is moved by the muscle module.

inertial effects present during cycling. We take into account this variation in the force by taking the mean of the measured force and accepting it as the actual force exerted on the nitinol. In the displacement-time record (figure 6.6), the upper peaks correspond to the maximum length during the relaxation period, whereas the lower peaks correspond to the minimum length during the activation period. By averaging the maximums over the 5 activation periods and dividing by the austenite length of the wire, the strain for the cold state of the wire at this applied force can be calculated. Similarly, averaging the minimums over the 5 activation periods and dividing by the austenite length gives the strain for the hot state of the wire (figure 6.7). For any constant force, the nitinol wire length cycles between the points determined by the hot and cold strains on the cyclic stress-strain curves. The maximum strain observed in the muscle modules was close to 6% at maximum stress of 200 MPa. These curves were used for the analysis of actuator motion, robot postures, and robot locomotion.

The transient behavior of the nitinol activation and relaxation was also measured. The recorded displacement versus time graphs were analyzed to obtain the time-dependent change in length of the nitinol wire during both activation and relaxation phases. The effect of applied force on the activation and relaxation times were obtained from the displacement versus time graphs. A 63% maximum displacement time was taken as the activation and relaxation time criterion. The results (figure 6.8) show that nitinol activation time varied between about 75




**Figure 6.7** Nitinol stress vs. strain graph for martensitic (cold) and austenitic (hot) states. The two lines necessarily meet at unity, however the first 0.01% strain required so little force we were unable to record it.



**Figure 6.8** Nitinol activation/relaxation time constants vs. stress, PWM = 80%.

and 175 msec with increased applied force. On the other hand, the nitinol relaxation time constant decreased with the applied force, within the range of 0.5–1.2 seconds.

There is a clear relationship between contraction velocity and amplitude and ambient water temperature (figure 6.9). For a given temperature, there is a range of power levels that yields some intermediate contraction amplitude, something that is useful in control of postural states of the robot. The power level required for full contraction varied from a low of 3.5 W (at 26°C) to a high of 7 W (at 10°C, figure 6.9). The resulting data set can be used as a look-up table to predict the resulting displacement from a given power level signal used. Such a predictive relationship is shown between water temperature and contraction amplitude (figure 6.10). A 90% displacement level was chosen. A best-fit exponential regression line was fitted through points less than 90% of full displacement level for each temperature ( $r^2$  from 0.87 to 0.95). The 90% displacement x-intercepts (the corresponding power levels) were plotted against the respective temperatures (figure 6.10). A third-degree polynomial fit produces an  $R^2$  of 0.999, providing a relationship that can be readily used in software to yield appropriate power levels under different environmental conditions. Similarly, an empirical relationship between the intermediate displacement states can be established. Figure 6.11 demonstrates three such intermediate states at three different power levels under otherwise identical conditions.




**Figure 6.9** Muscle module displacement at different power levels at three different temperatures. As power is increased, the displacement increases, initially exponentially. Upon reaching approximately 90% of full displacement the rate of increase declines dramatically. Each data point represents an independent observation. Tests were conducted with a 5 N force. Measurements were taken at five temperatures, only three are shown for clarity.


## BIOMIMETIC UNDERWATER ROBOT POSTURE CONTROL WITH NITINOL ACTUATION

High water currents and rough surface characteristics of the sea floor make locomotion and static stability a complex task for lobsters. To adapt to these conditions, the animals can realize many posture states in the main body pitch, roll, and height planes. Mimicking the posture adaptation competencies of *Homarus americanus* in the robot is important in order to give the robot the necessary behaviors to negotiate the same environment. In the case of the robot, achieving these postural states means the coordinated activation of the robot's eight identical legs (figure 6.12a, b).


The leg assemblies consist of 3-DOF linkages, actuated by nitinol muscle modules. Antagonistic pairs of nitinol are used to provide rotary motion to the joints in both directions.  $L_m$  and  $L_d$  denote the equal length of the middle and distal leg segments, respectively. Motion of the protraction-retraction joint provides the robot with necessary propulsive force during locomotion. The nominal motion range of this joint is  $40^\circ$ . The elevation-depression joint moves the leg segment up and down to provide or break with ground contact. Its motion range is within a  $30^\circ$  envelope. The extension-flexion joint remains fixed during



**Figure 6.10** The power level required to achieve a 90% displacement at different ambient water temperatures. An exponential regression was fitted through points up to the 90% displacement point in figure 9. The regression equations ( $r^2$  from 0.85 to 0.95) were used to calculate the power levels shown in this figure. A very good fit is found with a third-degree polynomial equation. The good fit ensures that this relationship can be incorporated in the control of the ambulator legs.



**Figure 6.11** Partial contraction under different power levels, under constant (500 g) load. Power was modulated by changing the duty cycle of the actuation signal. The ability to grade contractions is essential to achieve postural changes in the ambulator.




**Figure 6.12** (a) schematic of the functioning of the leg assembly.  $L_m$  and  $L_d$  denote the equal length of the middle and distal leg segments. (b) a photograph of the assembled 8-legged ambulator.

normal walking, but it drives the robot mainly during the lateral (side-ways) motion. This joint is capable of moving within a  $35^\circ$  angular range.

For a given posture behavior, all the robot leg joint angles need to be specified. To achieve this, the muscle modules should be activated in a graded fashion, which keeps the nitinol wires at an intermediate level of contraction. By keeping the wires at intermediate contraction levels, we can fix the joint angles and support the robot at a desired postural configuration. The posture control problem requires the solution of a





**Figure 6.13** Schematic of robot postures. The main body postures of the lobster robot are height, pitch and roll adaptations. In the figure,  $h_B$ ,  $\theta_B$ ,  $\beta_B$  correspond respectively to the height, pitch angle and roll angle of the coordinate axes attached to the robot body at its center of gravity. These postures are achieved with activation of the individual leg joints in an organized manner.

highly redundant system. The input to the problem is the desired final body posture and the output is the necessary power levels that need to be applied to all the nitinol muscle modules (figure 6.13). The inverse kinematics problem—that is, for a given desired posture, finding the angular positions of the joints—does not have a unique solution. Power consumption was used as a criterion for selecting the best possible solution of muscle activations in a set of all the possible combinations. For any statically stable body posture, we assume all eight legs touch the ground and the only force they carry is the ground reaction force due to robot's underwater weight (neglecting friction). This reaction force causes moments at each of the leg joints and these moments are balanced by the nitinol muscle forces. Using the torque balance equations and the experimentally obtained nitinol stress-strain-power relations, we solve for the static equilibrium joint angular positions of each leg, which then leads us to determine the final robot postural configuration.

The controllers installed onboard the vehicle are capable of providing three distinct levels of actuation power, which is analogous to the muscular recruitment levels of the real animal. We included this limitation in our posture analyses by using power levels: zero, one-third, two-thirds, and one. The full power level corresponds to the power level to produce the maximum contraction of the wire, and the others (one-third, two-thirds) are taken to be exact proportions of the full power. In the experiments, we determined that 7 W is sufficient to

**Table 6.1** Muscle activations and power requirements for roll postures

Rolls			Legs 1–4	Legs 5–8	Power
$\beta_{BDes}$	$\beta_B$	$h_B/L_d$	Elv/Dpr	Elv/Dpr	(W)
$-10^\circ$	$-11^\circ$	0.44	0.00/1.00	1.00/0.00	96
$-5^\circ$	$-6^\circ$	0.22	0.00/0.67	0.67/0.00	85
$5^\circ$	$6^\circ$	0.26	0.67/0.00	0.00/0.67	85
$10^\circ$	$11^\circ$	0.48	1.00/0.00	0.00/1.00	96

provide full contraction of the 90 mm muscle modules at the coldest temperature ( $10^\circ\text{C}$ ). For calculating the total power required for a given posture (table 6.1), we assume a 4 W per wire power consumption, corresponding to operation in approximately  $17^\circ\text{C}$  water temperature.

### Height Posture Results

For the height postures, we considered all the combinations of elevation and depression activations. No power was applied to the protraction and retraction muscles. We considered two sets of possible combinations for the extension and flexion joints. In one set, we selected extension and flexion activations to be equal to each other and tried all the activation levels. In the second set, we selected extension activation to zero and tried all the activation levels for the flexion muscle. All of these different combinations make a total of 32 cases. The desired levels of height are determined to be:


$$h_B = [0, 0.2, 0.4, 0.6, 0.8, 0.9]L_d$$

Example results for the height postures are shown in figure 6.14. The best possible activations and resulting actual heights are shown in table 6.1. In table 6.2,  $h_{BDes}$  and  $h_B$  correspond to the desired and actual values of the height. The maximum power consuming height is 0.55, which requires 64 Watts.

### Pitch and Roll Postures

For better static stability of the robot under arbitrary water currents, the robot can realize various pitch and roll postures. Pitch and roll postures depend mainly on the elevation-depression joint angles. For the pitch postures, the robot elevation-depression angles are given a gradient from back to front. For the roll postures, an angular gradient is given to the elevation-depression joints between one side and the other. The protraction-retraction joint does not carry much force, but it should be given some stiffness by coactivating both the protraction and the

## Ext:0.67 Flx:0.67



**Figure 6.14** Height postures, extension and flexion activations. At constant amounts of extension and flexion activations, varying the elevation and depression activations result in varying height postures. The maximum height is obtained by activating the depressor at full level and keeping elevator unactivated.

**Table 6.2** Muscle activations and required power levels for height postures

Heights						Power (W)
$h_{Bdes}/L_d$	$h_B/L_d$	Extension	Flexion	Elevation	Depression	
0.00	0.00	0.00	0.00	0.00	0.00	0
0.20	0.20	0.00	0.67	0.00	0.00	21
0.40	0.37	0.00	0.67	0.00	0.33	32
0.60	0.55	0.67	0.67	0.00	0.67	64
0.80	0.78	0.00	0.33	0.00	1.00	43
0.90	0.90	0.00	0.67	0.00	1.00	53

retraction muscles at a low level of activation. Similarly, extension and flexion joints are coactivated in a similar manner. The desired range of pitch and roll angles were specified to be:

$$\theta_B = [-20^\circ, -10^\circ, 10^\circ, 20^\circ]$$

$$\beta_B = [-10^\circ, -5^\circ, 5^\circ, 10^\circ]$$

For the pitch and roll postures, we found that protraction and retraction activations should be set to 0.33 to provide the joints with some amount of stiffness. Similarly, extension and flexion activations were set to 0.67. The pitch posture is symmetrical about the XY axis attached to the robot body, which simplifies the analysis. We calculated the

**Table 6.3** Muscle activations and required power levels for pitch postures

Pitches			Leg1	Leg2	Leg3	Leg4	Power
$\theta_{BDes}$	$\theta_B$	$h_B/L_d$	Elv/Dpr	Elv/Dpr	Elv/Dpr	Elv/Dpr	(W)
-20°	-19°	0.44	0.00/1.00	0.33/0.67	0.67/0.67	1.00/0.00	99
-10°	-10°	0.38	0.33/0.67	0.33/0.67	0.67/0.67	0.67/0.33	99
10°	13°	0.41	0.67/0.00	0.67/0.67	0.33/0.67	0.00/0.67	93
20°	18°	0.52	1.00/0.00	0.67/0.67	0.33/0.67	0.00/1.00	99

pitch postures obtained by various combinations of elevation and depression muscle activations and selected the best possible ones in terms of the power consumption (table 6.2).

Similarly for roll, we calculated the posture configurations for combinations of muscle module activations. Protraction-retraction and extension-flexion muscles were kept at two-thirds activation level, and the elevation and depression activations were similarly varied for different legs. Among these combinations, the ones that resulted in roll angles closest to the desired roll angles and requiring the minimum power (table 6.3) were selected. The output of the model will be used to obtain the desired vehicle roll and pitch angles at minimum power consumption levels.

## CONCLUSIONS

Although far from an ideal material for an artificial muscle, this study shows nitinol SMA to be a viable candidate for the construction of an underwater ambulatory robot. The advantages offered by SMAs in general are very high energy density, compact size, and low weight. Major drawbacks are relatively high power consumption and difficulties in controllability. Given the nature of nitinol as a thermal engine, it is necessary that a predictive relationship between actuation power and muscle displacement be established. This is needed to enable control in variable thermal environments as well as obtaining partial activation (graded contractions), and was established empirically (figures 6.10 and 6.11). The predictable behavior shown in this study in the relationship of displacement and power, and ambient temperature goes a long way in enabling real-world controllability of nitinol-based muscles, and by extension, nitinol-based vehicles. The addition of a thermal insulation layer in the form of a PTFE tube has allowed us to minimize power consumption to less than 4 W per module in warm ambient water. In colder ambient water, more insulation can be used, the trade-off being the lengthened relaxation time in warmer ambient water temperature.

Real lobsters use their postural behavior to adapt to such environmental conditions as high water currents and surge. Using the nitinol parameter space, we measured for the artificial muscle. A model mimicking the height, pitch, and roll posture adaptations of the lobster was constructed for use in the control of our biomimetic underwater vehicle. We modeled the static posture behavior of the underwater vehicle. The identical eight legs of the vehicle are supported by nitinol muscle modules and carry the net underwater weight of the vehicle. The friction force between the legs and the ground is neglected. Using the model, we obtained posture results for possible combinations of muscle activations. For any desired posture, we selected the best possible set of muscle activations based on the minimum power consumption criterion. Results show that by using the discrete actuation powers, it is possible to closely obtain the desired height, pitch, and roll postures. The main disadvantage of nitinol actuation is its high power consumption. Calculations show that the vehicle consumes as high as 99 watts of electrical power during some of the pitch and roll postures.

In the current study, we have taken the pre-strain in the muscle wires to be constant with some slack present at the unactivated states of the wires. But, preliminary studies with the no-slack case for the height postures have shown that it is possible to use about half of the power required for the slack case. Further efforts can be directed toward reducing the power consumption of the vehicle by taking into account the amount of pre-strain in the wires, the friction force between the legs and the ground, or reducing the underwater weight of the vehicle.

Using an artificial muscle in a biomimetic robot requires the careful matching of the response characteristics of the muscle to the control signals. In this study, we have sought to evaluate the suitability of nitinol SMA for constructing an artificial muscle for a very specific application—in our case, underwater use. The results show that nitinol does fulfil the requirements that the biological model we have adopted imposes upon the artificial muscles. The high power consumption and the heat dissipation limitations essentially exclude the use of nitinol in terrestrial applications. This is likely to be the case with new materials coming forth with applications as linear actuators as well. The match between the specific application and the actuator performance parameters has to be established, probably on a case-by-case basis.

## REFERENCES

- Brinson, L. C. (1993). One-Dimensional constitutive behavior of shape memory alloys: Thermomechanical derivation with non-constant material functions and redefined martensite internal variable. *J. Intell. Material Syst. Struct.* 4: 229–244.
- Buehler, W. J., Gilfrich, J. W., and Wiley, R. C. (1963). *J. Appl. Physiol.* 34: 1475.

- Liang, C., and Rogers, C. A. (1990). One-Dimensional thermomechanical constitutive relations for shape memory materials. *J. Intell. Material Syst. Struct.* 1, no. 2: 207–234.
- Ohkata, I., and Suzuki, Y. (1998). In K. Otsuka and C. M. Wayman (eds.), *Shape Memory Materials*. Cambridge: Cambridge University Press, pp. 240–266.
- Saburi, T. (1998). In K. Otsuka and C. M. Wayman (eds.), *Shape Memory Materials*. Cambridge: Cambridge University Press, pp. 49–96.
- Wayman, C. M., and Duerig, T. W. (1990). An introduction to martensite and shape memory. In T. W. Duerig, K. N. Melton, D. Stöckel, and C. M. Wayman (eds.), *Engineering Aspects of Shape Memory Alloys*. Boston: Butterworth-Heinemann, pp. 3–20.

**This page intentionally left blank**

# 7

## Engineering a Muscle: An Approach to Artificial Muscle Based on Field-Activated Electroactive Polymers

Roy Kornbluh, Robert J. Full, Kenneth Meijer, Ron Pelrine, and Subramanian V. Shastri

Biological creatures are, in general, more coordinated, more powerful, more efficient, and ultimately better than most robots at completing motor tasks in the unstructured environments found outside the laboratory or factory. Extracting principles of motion that are employed in nature might help us to design robots and other devices that perform as well as living organisms. Applying such biomimetic principles to robots might help us build walking robots that can traverse uneven terrain as quickly and as effectively as a cockroach, develop small flapping-wing flyers that can hover and fly as efficiently and as nimbly as a hummingbird, or design artificial or prosthetic hands that can grip and manipulate objects with the dexterity of organic human hands.

Many design and operational features of a robot affect the quality of its motion. These features include higher-level sequencing of joint motions, modifications of motions based on learning, and sensory-driven joint motion. Ultimately, the ability to produce motion in a controlled fashion depends on the use of actuators that can supply and control the requisite mechanical energy for energetic interaction with the environment. In the case of animals, we are talking about muscle. Therefore, our discussion of actuators for biomimetic robots will focus on the development of artificial muscle.

In this chapter, we first consider the function of biological muscles and the features and performance an actuator must have to be considered an artificial muscle. We then focus on one promising new actuation technology based on dielectric elastomers, a type of electroactive polymer. Experimental biology techniques are used to study this candidate artificial muscle technology in detail and determine if its performance fits within the space of natural muscle. Finally, we present examples of the application of the dielectric elastomer artificial muscle to biomimetic robots.



## ACTUATOR CHARACTERISTICS AND METRICS: WHAT IS AN ARTIFICIAL MUSCLE?

Many recently developed actuator technologies are termed “artificial muscles” by their proponents. These technologies include shape memory alloys (Hunter and LaFontaine, 1992); flexible pneumatic chambers (Klute, Czerniecki, and Hannaford, 1999); electrostatic devices (Bobbio et al., 1993); and a large variety of mechanochemical polymers (Kuhn, 1949; Brock, 1991) or electroactive polymers (Smela, Ingnas, and Pei, 1993; DeRossi and Chiarelli, 1994; Shahinpoor, 1995; Otero et al., 1999; Wax and Sands, 1999). Clearly, actuators based on these technologies are not like muscles in all respects. Some technologies match the stroke capabilities of muscle, while others match the force capabilities. Some technologies use chemical energy like muscle, while others rely on electrical energy. It is fair to ask “What is an artificial muscle?” In the context of biomimetic robots: An artificial muscle is an actuator that embodies certain desirable characteristics of natural muscle.

Many characteristics of natural muscle may be critical to its ability to produce the type of locomotion or dexterous manipulation found in biological creatures. Biologists have identified some of these characteristics, but many are yet to be identified. It is therefore a complex task to select suitable metrics to determine whether or not a given actuator technology is an artificial muscle. Some less obvious characteristics of muscle that may help to emulate the desirable motion of biological creatures include the degradation of performance with fatigue, the ability to self repair, sensitivity to temperature, the ability to be consumed to provide energy, structures composed of a large number of fibers operating in parallel, and diverse performance of individual muscle fibers. Although future biomimetic robots might incorporate sophisticated actuators with some of these features (should they be deemed important), the role of artificial muscle in the next generation of biomimetic robots will be based on its ability to produce and control motion. In this chapter, we focus on that role.

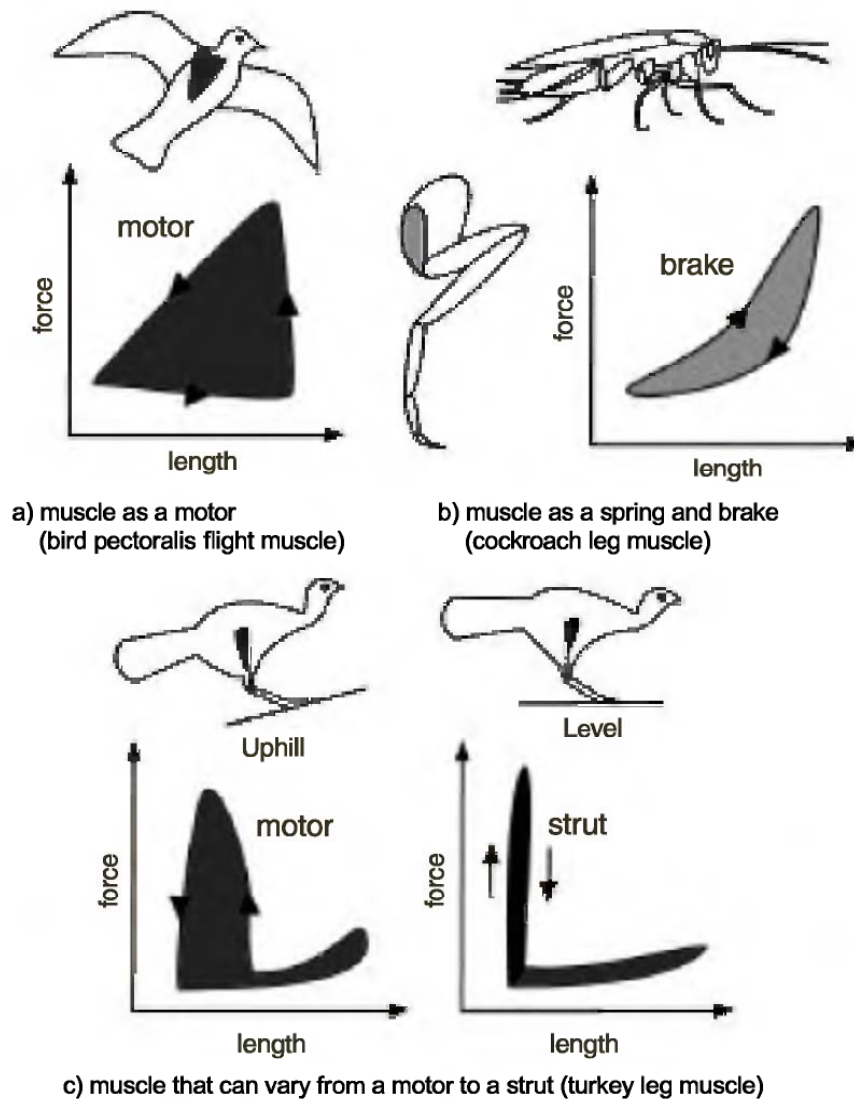
Even if we limit ourselves to the features of muscle that focus on motion, we must consider the diversity of function and performance found in nature. To establish whether or not a specific actuator embodies the desirable characteristics of natural muscle, we must know the function of biological muscles as well as the metrics that are used to describe that function. In the following section, we summarize some of the functions of biological muscles and then present a set of metrics that can be used to define artificial muscles and evaluate different actuator technologies.

## **Biological Muscles as Multifunctional Materials (Motor, Brake, Shock Absorber, Spring, and Strut)**

Muscles vary widely in their primary function. Conventionally, muscle is viewed as a linear motor, that is, an actuator that provides a motive force to an object. Indeed, many performance metrics focus on this role. Although this is certainly an important role of muscle, it is not the only role found in nature. Besides functioning as motors, muscles also play other important roles in the control of locomotion and manipulation. Some muscles also function primarily as brakes or shock absorbers, spring elements, or even relatively rigid struts. The functional role of an individual muscle may change with the task at hand, or even at certain points within a given task, depending on joint kinematics and external loads. It should also be noted that a muscle need not perform only one specific role at a time. For example, muscle can act simultaneously as a spring and as a motor or brake. Much current understanding of these different roles of muscle is relatively recent and can be attributed to analytical techniques such as work loops that more fully emulate the operating conditions found in nature (Dickinson et al., 2000). A true artificial muscle actuator should emulate all the functions that are important to the execution of motion. Some examples of the importance of the motor and nonmotor roles in biological systems are discussed below. Figure 7.1 graphically shows examples of the various roles of muscle found in nature.

**Muscle as a Motor** As previously noted, the conventional way of looking at muscles is to consider them as motors that generate the power to propel the body or body parts during locomotion. Indeed, some specialized muscles operate in a manner that maximizes their power output and optimizes the efficiency with which the power is generated. For example, in jumping frogs, the strain pattern and activation of their leg muscles are nicely tuned to enable the muscle to generate the maximal power needed for good jumping performance (Lutz and Rome, 1994). Furthermore, during cyclical activities like walking, the central nervous system tends to activate the muscles or parts of muscles that can operate most efficiently at the desired speed of locomotion (Rome et al., 1988).

**Muscle as a Brake or Shock Absorber** The motion of biological creatures can be more stable and robust in response to external perturbations if the muscles provide a certain amount of damping, acting much like a brake or shock absorber. The distinction between a brake and a shock absorber is that a brake absorbs energy over a significant portion of the cycle, while a shock absorber damps out energy due to short-duration perturbations.



**Figure 7.1** Examples of the varied roles of muscles in biological motion. Work loop diagrams help illustrate the behavior of specific muscles. (Adapted from Dickenson et al. 2000.) (a) Muscle as a motor (bird pectoralis flight muscle); (b) muscle as a spring and brake (cockroach leg muscle); (c) muscle that can vary from a motor to a strut (turkey leg muscle).

An example of a muscle acting as a brake is a certain leg muscle in the death's head cockroach that absorbs energy during normal walking gaits (Full et al., 1998).

The leg muscles of cockroaches also contribute to damping disturbances (Meijer and Full, 2000). Such a damped external disturbance response has been termed a *preflex* (Brown and Loeb, 1999). A preflex, in contrast to a reflex, does not use any sensory feedback or closed-loop control. Instead, it can act almost instantaneously on external perturbations, based on the open-loop dynamic response of the muscle and joints. There is evidence that the ability of a muscle to damp perturba-

tions while still operating as a motor or brake depends in part on the nonlinear mechanical behavior of muscles.

**Muscle as a Spring** The springlike behavior of muscles can also play an important role in making locomotion more efficient. The flight muscles of many insects demonstrate springlike behavior during operation at high speed (in other words, the passive springlike properties of the muscle dominate at high frequencies of operation).

In some instances in nature, the elasticity of muscles, in conjunction with elastic elements elsewhere in a creature, is used to produce a tuned dynamic system. In a tuned system, an interchange takes place between the kinetic energy in the inertia of body parts and potential energy stored in the elastic elements (Alexander, 1988). If this interchange has a resonant frequency equal to the normal cycle frequency of the creature's motion, then energy that might otherwise be used to accelerate or decelerate body parts, or to lift and drop the center of gravity, can be recovered for use elsewhere in the cycle, resulting in a net saving of energy per cycle. Although the importance of such resonant effects is still an active topic of research in biology, these tuned mechanical systems exist in the wing flapping of many insects and small birds, the hopping of kangaroos, the swimming of many fish, and the walking or running of many legged creatures including humans.

**Muscle as a Strut** In some cases, muscle acts as a rigid structural member (i.e., strut). Examples are certain muscles in hopping wallabies, running turkeys, or swimming fish, where the muscle transmits forces from a tendon or other muscle (Dickinson et al., 2000). This behavior enables the elastic tendons to store and return energy during the locomotion cycle, thereby reducing the metabolic cost of locomotion.

### **The Basis of Comparison: Choosing the Metrics**

Because we are focusing on the actuation aspects of muscle, it makes sense to build upon metrics already developed for artificial actuators.

Metrics based on performance and other characteristics that are commonly used to specify the requirements of actuators are given in table 7.1. Because muscles and artificial actuators exist in a wide range of sizes, it often makes sense to use scale-invariant metrics (normalized by mass or size). These parameters have some redundancy (e.g., power is energy times frequency) that reflects the differing manners in which actuator specifications are determined.

Unfortunately, these parameters alone may neither serve to accurately compare actuators to natural muscle nor determine whether or not a given actuator technology can fairly be called an artificial muscle.

**Table 7.1** Common actuator metrics and specifications


Parameter	Scale-Invariant Version	Comments
<b>Energy and Power</b>		
Energy output	Specific energy	Over full cycle
Power output	Specific power	Average or peak
Energy conversion efficiency	Same	Energy output over a full cycle/energy input over a full cycle (excluding recovered energy)
Response time, frequency bandwidth	Same (usually scale dependent)	For one direction or full cycle
<b>State Variables</b>		
Displacement, stroke	Strain	Peak over a cycle
Force	Stress, pressure	Peak over a cycle
Velocity	Strain rate	Peak over a cycle
<b>Impedance and Controllability</b>		
Stiffness	Elastic modulus	Usually nonlinear (not a constant)
Damping	Specific damping, loss factor, loss tangent	Usually nonlinear (not a constant)
Accuracy (displacement or force) (%)	Percentage of strain or stress	Usually percentage of maximums
Repeatability (%)	Same	Usually percentage of maximums
Linearity or sensitivity (%)	Same	Deviation from linear input-output relationship
<b>Operational Characteristics</b>		
Environmental tolerance	Same	Recommended ranges of temperature, humidity, etc., or effects of variations in temperature, humidity, etc., on the above parameters
Durability, reliability	Same	Number of cycles before degradation threshold or total failure, degradation per cycle or time
Input impedance (power supply requirements)	Specific impedance	Voltage and current requirements or pressure and flow, depending on power mode

The problem is twofold. First, the values of the parameters of different types of actuators and muscle may not have been determined under identical conditions. The performance of muscle, in particular, is highly dependent on the magnitude and timing of specific inputs, motions, and loading conditions. Thus, it is important to understand the definitions of the parameters in table 7.1 as well as the conditions under which the parameters are measured. The fact that measurements of muscle are typically performed by biologists with techniques and apparatus suitable to their needs, while measurements of most artificial actuators are made by engineers with their own set of techniques and apparatus, exacerbates the problem.

Second, the parameters in table 7.1 may not accurately or completely describe the behavior of muscle or an actuator in actual usage. For example, let us consider two of the most common parameters used to describe actuator behavior: force and stroke. Often, the force and strokes used in actuator specifications are the maximum blocked force (isometric) and maximum “free” stroke at a constant force loading (isotonic). However, these definitions do not provide enough information to determine the actual force or stroke output produced in a given implementation. Further, they do not directly indicate the amount of work that the actuators can perform. An actuator that moves freely may have large stroke but may exert no force and do no external work. An actuator whose motion is completely constrained may have a large force but zero stroke, and also may do no external work. Between these two extremes are forces and strokes that may not be the maximum forces or strokes that the actuator can produce but still allow external work to be done. Note, however, that the amount of work cannot necessarily be determined from the free stroke and blocked force.

The use of such parameters as maximum blocked force and maximum stroke at a constant force loading can give a quick indication of whether or not an actuator can approach the desired performance (e.g., if the free strain or blocked force are significantly less than required, the actuator technology is unlikely to be suitable). However, these parameters can give misleading indications of the performance actually produced in the particular application. Characteristics of muscle that make the use of common engineering metrics difficult are described in more detail below.

**Muscle Is Soft: The Importance of Impedance** The effect of loading conditions on an actuator’s output is particularly critical for measuring the performance of natural muscle. Muscle is a soft tissue: Thus, the stroke of a muscle is highly dependent on the external forces (loading conditions) that act upon it. Actuators, such as highly geared servomotors, whose output is little affected by loading conditions, cannot rightly be considered artificial muscles. The concept of impedance can clarify



**Figure 7.2** Simple lumped-parameter model of natural muscle, showing the importance of the inherent viscoelastic properties.


the importance of the response of muscle to the loading condition and why this response should be emulated in artificial muscle actuators.

Impedance is basically the relationship between the external force applied on an actuator and the resultant displacement (or sometimes the resultant velocity). In essence, muscle impedance describes the dynamic behavior of the actuator in response to external forces.

It is well known that the work output of an actuator is maximized when its impedance matches that of the external environment (Giurgiutui and Rogers, 1997). Of course, an exact match is not possible, especially because both the external load and the actuator are only approximately represented by simple lumped parameter models. Nonetheless, robot designers already consider impedance matching when they select actuators and transmission mechanisms. There is some evidence that natural selection has also favored biological creatures whose muscles exhibit a good level of impedance matching to the driven load (Alexander, 1997).

In simple linear systems where inertial effects can be ignored, impedance can be described by the stiffness and damping of the actuator. Although a simplification, a lumped-parameter model can be useful in understanding muscle behavior. This basic approach has been used by researchers attempting to develop a mechanical model of muscle behavior. Such a model was posited by Hill (1938). Over the years, many improvements and modifications have been made to the basic model. A common lumped-parameter model employed by biologists today is shown in figure 7.2. This model conveniently separates the active force-generating element from the passive properties of the muscle. Although, in general, impedance is a complicated nonlinear relationship that is affected by the state of the muscle and level of muscle stimulation, the separation of a muscle model into active and passive components helps to clarify muscle behavior.

The separation of the muscle model into passive and active components can also help explain the diversity of the roles that muscle plays.



**Figure 7.3** Results of common muscle characteristic tests. These curves provide a qualitative description of the nonlinear and time-dependent properties of muscle (Source: Full 1997).

In cases where muscle acts primarily as a spring, the spring element dominates. In cases where muscle acts primarily as a brake or shock absorber, the damping element dominates. The active component dominates when muscle acts primarily like a motor.

**Muscles are Nonlinear, Time-Variant, Multivariate Systems** Although some understanding of muscle behavior may be gained by using linear time-invariant elements to describe the passive behavior, the active behavior is time dependent and highly nonlinear. Thus, changes in measurement conditions can produce large and sometimes unpredictable changes in muscle performance. The situation is further complicated by the complex response of muscle to a control stimulus. Both the amplitude as well as frequency of a stimulus affect the response of the muscle. To fully map out the performance space of muscle performance is a daunting task. Simple parameters, such as those presented in table 7.1, can give misleading indications of muscle performance if complex behavior is not considered.

Biologists have developed several common tests that consider time-dependent nonlinear behavior that relates the force and stroke (or stroke velocity) as a function of stimulation level and time (Zajac, 1989). Figure 7.3 shows qualitatively the behavior that is often found in muscles when they are subjected to such testing at a particular level of stimulation. Note, for instance, that the effective stiffness of muscle (as seen from the slope of the force-length curve) can vary greatly. This illustrates how difficult it is to use a simple metric such as stiffness to describe the elastic response of a muscle.

Nonlinear, time-variant behavior not only may make it difficult to map out the performance space of muscle but may itself be responsible for some of the desirable features of biological motion (Full and Meijer, 2000). Artificial muscle actuators may need to emulate such behavior, so that the performance metrics will be capable of describing the behavior.

### Measuring Muscle: Work Loops and More

The tests that generate the curves in figure 7.3 are common tests used to characterize many unique properties of muscle behavior over a wide



range of performance. While the basic shape of the curves may be similar among muscles, the specific shapes differ greatly for different muscles. These tests alone do not, however, necessarily reproduce the loading conditions during natural motion and may not accurately represent the range of performance found in nature.

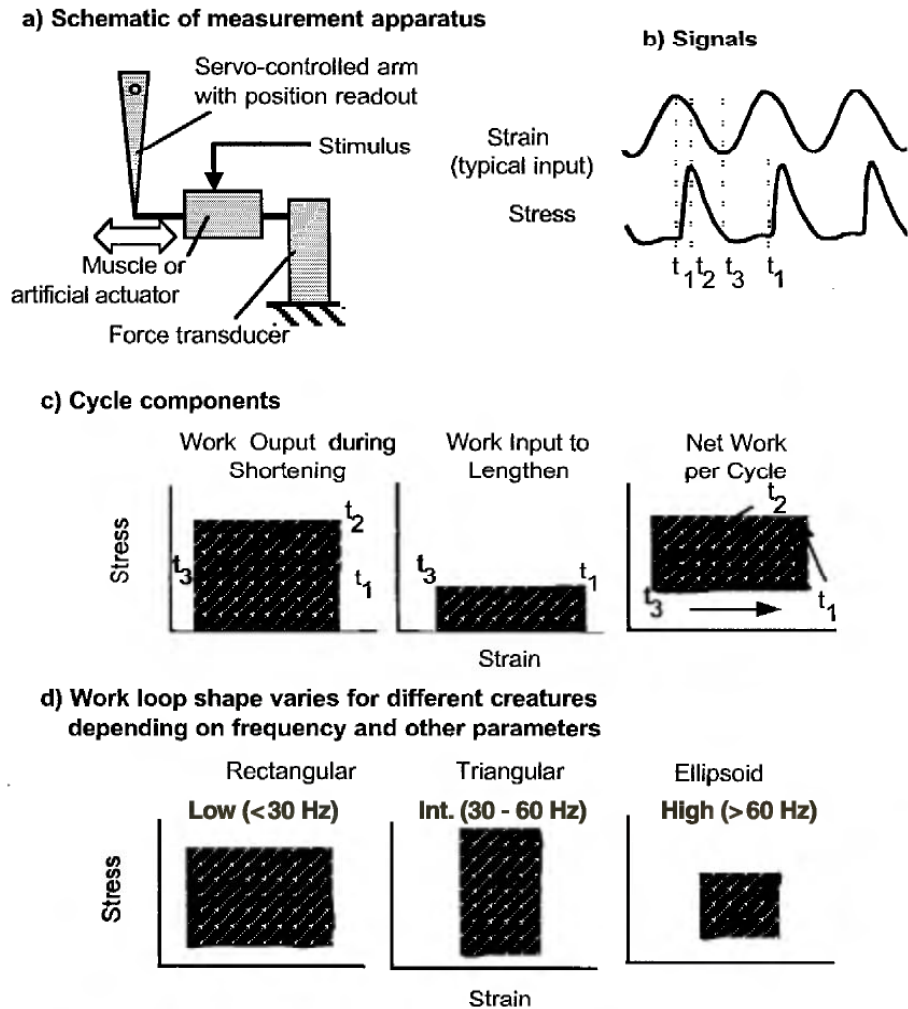
To more fully evaluate how the performance of actuators relates to the motion that they produce and control, it is important to accurately reproduce the load the actuators will experience in undergoing a given task. For example, the primary role of some muscles is actually to absorb rather than produce external energy. This role will be discussed in the following section.

Experimental biology measurement techniques have been developed to more accurately estimate the *in vivo* behavior of muscle. In order to conduct such measurements, the kinematics of the limb (or appendage or body segment) motion and the timing of the muscle stimulation are measured. The motion that individual muscles undergo during locomotion can be estimated from the kinematic data in combination with detailed muscle-skeletal models. A controlled experiment can then be set up that reproduces the motion and stimulation of the muscle. The muscle is attached to a transducer that measures the force it produces. Such an analysis of cyclic motion is termed a *work loop* (Josephson, 1985), because the area contained within the closed loop is equal to the work done by or on the muscle. Figure 7.4 describes such work loops.

By essentially reproducing the physical state of the muscle throughout its motion, work-loop techniques implicitly include all nonlinear, time-dependent effects as well as the effects due to loading conditions. A disadvantage of work loops is that they do not provide a general picture of the muscle behavior, which would enable performance to be determined if the motion or external loads change. For this reason, muscle performance will continue to be characterized by means of a variety of tests. It is up to the biologist to identify the features of muscle that are critical to its desirable performance. Similarly, it is up to the engineer to define the correct performance parameters and test protocols. We have identified some such features, parameters, and protocols in this chapter; however, the list is far from complete.

### **Muscle as Part of a System**

While this chapter focuses on muscle itself, in the study of biological motion we cannot consider muscle in isolation. The performance of muscle is closely dependent on the transmission mechanism and dynamics of the limb, body segment, or appendage. Depending on the size and speed of the motion, the inertial effects of the muscle and limbs also play significant roles. The end effector (the moving part that actually interacts with the environment) must also be considered. For ex-



**Figure 7.4** Work loop representation of muscle performance (Adapted from Josephson 1985.)

ample, opening and closing a lobster claw could not efficiently move air, nor could a bird's wing be used to crush shellfish.

It should be noted, however, that transmission mechanisms can effectively make trade-offs between the force and stroke output of an actuator, as well as change the effective impedance of the actuator. In some cases, it is most appropriate to consider the actuator and transmission mechanism together when comparing actuators with natural muscle. In the case of biomimetic robots, the artificial muscle is often attached to a skeleton-like linkage. Such linkages introduce significant nonlinearities in the forces transmitted to and from the muscle. These nonlinearities may be important to the proper functioning of the motive system.

Note also that this chapter focuses on muscles or actuators acting individually. In biological creatures, the motion of a limb, body segment, or appendage is often the result of the action of several muscles that act


directly upon it. In particular, muscles are often grouped in antagonistic pairs about a joint. Researchers on both the biological and engineering sides have recognized the importance of being able to control the impedance of a joint when it interacts with the environment (Houk, 1979; Hogan, 1985a, b, c; Winters, 1986). To understand the importance of impedance modulation in biological systems, we can consider some examples from our own experience. When opening a door, we want low effective stiffness in our arm joints because we let the door determine the path that the arm takes. On the other hand, when we use a plane to shave off excess wood on the edge of a door, we want higher stiffness because we want to dictate the path of the plane.

The modulation of the joint impedance can also be considered a preflex: that is, it need not include reflexive (closed-loop) control. An example of such modulation of impedance is that of the arms of a person riding a mountain bike down a bumpy trail: The rider can control the steering and support the body while damping out the higher frequency disturbances, even though the brain and even the central nervous system lack sufficient time to participate in this response. This example illustrates a more-general potential advantage of preflexes over reflexes for biomimetic robots (here, inertial effects may be important in addition to stiffness and damping). Preflexes use the inherent open-loop behavior of the muscle-limb system and do not require sensory feedback or closed-loop control. As such, preflexes can make biomimetic robots simpler and more robust. If the preflex is to be controllable, it makes sense that the primary damping mechanism reside in the artificial muscle. Thus, although the ability to set a preflex is a system-level issue, the ability to implement such a preflex in a biomimetic robot can depend on the characteristics of the actuators.


### **Range of Muscle Performance**

Even if we use consistent and well-defined performance metrics, natural muscle has a wide range of performance that makes a generalized comparison between an actuator technology and natural muscle difficult. An artificial muscle need not meet any single level of performance for any parameter. The criterion for determining whether an actuator is an artificial muscle therefore depends on the particular biological function that the robot is attempting to mimic. Simply listing a range of values for each performance parameter would give the false impression that an artificial muscle actuator need only meet the minimum performance level of each parameter.

It is, however, useful to place bounds on the performance of natural muscle as well as to observe some relationships between the values of different performance parameters. Much of this discussion derives from data presented by Full (1997). Figure 7.5 shows the work and power output of certain muscles for a wide range of biological



**a) Work and Power output per cycle (invertebrates)**




**b) Power output variation among species**

**Figure 7.5** Work and energy output as a function of frequency for a variety of species. Measured performance of dielectric elastomer (EAP) actuators is also shown. (From Full 1997.) (a) Work and power output per cycle (invertebrates); (b) power output variation among species.

creatures. These values were determined by many different researchers using work-loop techniques. Each data point represents the muscles operating at the frequency that gave the greatest power output.

The frequency at which the maximum power output is achieved varies significantly. The flapping of a small insect at a frequency of nearly 1,000 Hz is more than two orders of magnitude faster than the undulations of many swimming creatures. What is interesting, however, is that the average maximum power produced by muscles does not vary over as great a range. This power seems to be limited to values on the order of 100 W/kg. We can use this number as a rough basis of comparison between natural muscle and robot actuators.

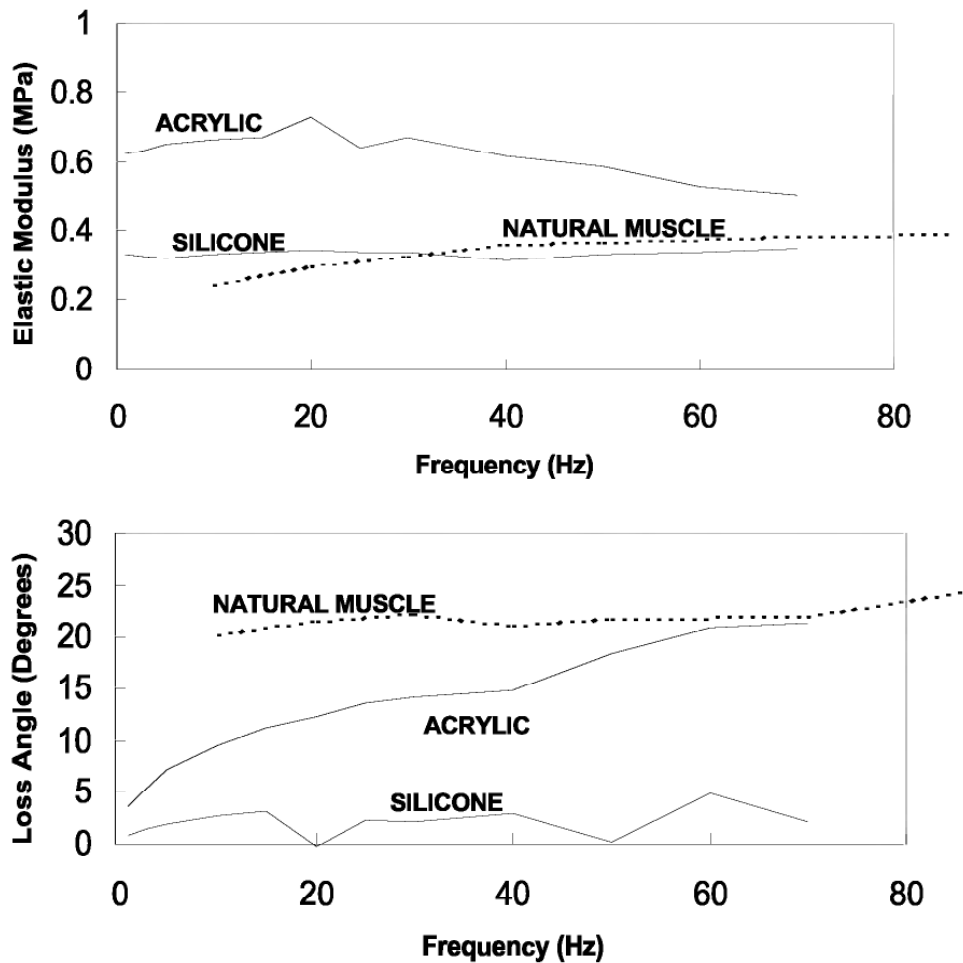


**Figure 7.6** Stress and strain of natural muscle for a variety of invertebrate species. Measured performance of dielectric elastomer actuators (EAPs) is also shown (From Full 1997.)

In some instances, the peak power (instantaneous maximum power generated during contraction) output is of more interest than the average power output. The peak power can be three times that of the average power, as determined by work-loop methods (Full and Meijer, 2000). The peak contraction velocity (usually measured isotonicly) varies greatly among muscles, ranging from 0.3 to 16 lengths/sec. Generally, the contraction velocity is less at greater force loadings.

While it is true that the highest response frequencies are encountered in small creatures, the response time of muscle can vary greatly even among creatures of similar size. The contraction time (the time to produce peak force after the application of the stimulus) of muscles varied from 0.004 to 0.79 sec, when isometric contraction measurements were used (Full, 1997).

Corresponding to the variations in work output, the stress and strain determined by observations of natural creatures and work-loop techniques also varies significantly among different species and is frequency dependent (figure 7.6). Work-loop techniques as well as force versus length measurements also show that muscles that operate at higher frequencies tend to generate stresses over a narrow range of strains. These muscles exhibit work loops that clearly show the inherent springlike behavior of the muscle (see figure 7.4). This result is not surprising. At high frequencies, there is less time for the muscle to generate work



**Figure 7.7** Comparison of the passive viscoelastic properties of natural muscle (cockroach leg muscle 177c) and dielectric elastomer EAP materials.

during a given cycle: Therefore, the passive properties (e.g., elasticity) are relatively more dominant.

Some testing has been done to measure the effective open-loop stiffness and damping characteristics of natural muscle over a range of frequencies (Meijer and Full, 2000). For example, figure 7.7 shows the elastic modulus (stiffness) and mechanical loss angle (damping) of an unstimulated cockroach leg muscle (here the assumed muscle model is simply a spring and damper operating in parallel). These passive measurements, performed at low amplitudes of motion, show that this muscle behaves much like a soft, highly damped polymer gel—which is exactly what muscle is. Because these passive measurements do not include the effects of the level of stimulation and state of stress or strain, they may not necessarily provide an accurate picture of in vivo muscle characteristics during tasks such as walking.

The efficiency of a muscle (defined as the mechanical power output divided by the chemical power input) is a particularly difficult param-


eter to accurately measure. Efficiency can be estimated by measuring the rate of oxygen consumption; although the estimates reported have varied from 8% to 25%, several measurements on a variety of species suggest that a value of 10% is typical (Stevenson and Josephson, 1990). This efficiency includes the efficiency of the conversion from the primary mode of energy storage. In the case of natural muscle, this efficiency includes basic metabolic energy costs. To achieve comparable efficiency, artificial actuators would need to include such conversion factors. The efficiency of electrical actuators must include the efficiency of power supplies and electrical driver circuits. The efficiency of fluid actuators (pneumatic and hydraulic) must include the efficiency of the compressor and transmission lines.

### **Form versus Function**

Our comparisons between natural and artificial muscle have thus far focused on function rather than form. That is, we have focused on the performance of the muscle rather than on its shape or internal structure. This approach makes sense as long as the shape or structure of the muscle allows for a robot with the desired kinematics. More sophisticated biomimetic robots might address issues related to form, such as mass distribution, redundancy, and diversity of muscle fibers, and pennate versus parallel arrangement of individual muscle elements.

### **DIELECTRIC ELASTOMERS: A CANDIDATE ELECTROACTIVE POLYMER ARTIFICIAL MUSCLE**

Softer elastomeric dielectric materials with compliant electrodes produce a significant electromechanical response due to the applied electric field. Such *dielectric elastomer* actuators have demonstrated characteristics such as high strain (exceeding that of many muscles), high efficiency, high energy density, fast response, good controllability, and the capabilities of impedance selection (based on material choice and design) and open-loop modulation. These characteristics make dielectric elastomer a candidate for consideration as artificial muscles for biomimetic robots (Kornbluh, Pelrine, and Joseph, 1995). The material cost is expected to be low. Because of this desirable combination of characteristics and the extensive experience of the authors in developing and testing these materials, we will focus on this technology as a representative artificial muscle technology in the remainder of this chapter. It should be recognized that the dielectric elastomer technology is not perfect for artificial muscle applications in all respects. Dielectric elastomer actuators need efficient high-voltage driver circuitry. Additionally, dielectric elastomer actuators are expected to be able to withstand a wide range of environmental conditions and to have good durability



**Figure 7.8** Principle of operation of the dielectric elastomer actuator.

(desirable features for artificial muscle actuators), but these features have yet to be proven in rigorous tests. As will be seen, the shape of such actuators and their mode of operation are different from those of natural muscle (which expands upon the application of a stimulus instead of contracting).

### Principle of Operation of Dielectric Elastomer Artificial Muscles

All actuators based on dielectric elastomer technology operate on the simple principle shown in figure 7.8. When a voltage is applied across the compliant electrodes, the polymer shrinks in thickness and expands in area.

The net volume change of the polymer materials that we investigated is small (Pelrine, Kornbluh, and Kofod, 2000). Therefore, the compliance of the electrodes is essential to allow the film to strain.


Typically, dielectric elastomer films are fabricated by conventional polymer fabrication techniques such as spin coating and casting. Electrodes can be deposited via spraying, screen printing, or photolithography (Kornbluh et al., 1999). The observed response of the film is caused primarily by the interaction between the electrostatic charges on the electrodes (Pelrine, Kornbluh, and Kofod, 2000). Simply put, the opposite charges on the two electrodes attract each other, while the like charges on each electrode repel each other.

Using this simple electrostatic model, we can derive the effective pressure produced by the electrodes on the film as a function of the applied voltage (Pelrine, Kornbluh, and Joseph, 1998). This pressure,  $p$ , is

$$p = \epsilon_r \epsilon_0 E^2 = \epsilon_r \epsilon_0 (V/t)^2 \quad (7.1)$$

where  $\epsilon_r$  and  $\epsilon_0$  are, respectively, the permittivity of free space and the relative permittivity (dielectric constant) of the polymer;  $E$  is the applied electric field;  $V$  is the applied voltage; and  $t$  is the film thickness. The





**Figure 7.9** Strain response of a circular area of an acrylic film with circular electrodes.

behavior of the polymer is functionally similar to that of electrostrictive polymers in that the response is directly related to the square of the applied electric field.

The resultant strains produced in the polymer are dependent on the boundary conditions and loads on the polymer. Further, the strains depend on the elastic modulus of the polymer, which for elastomers at large strains may be nonlinear. A further complication is that the polymers often are actuated with large initial prestrains. Such prestrains may be anisotropic and thus may cause the effective elastic modulus to be anisotropic as well. For these reasons, it is not possible to write a simple general equation for the resultant strain.

Fortunately, large strains can be easily measured. Because we typically operate with large prestrains, it is more convenient to use the relative strain, which is the strain difference from the prestrain conditions. Figure 7.9 shows the extremely large strain response that some actuator designs make possible. Strains of more than 100% have been observed in both acrylic and silicone elastomers.

The deformation of the polymer film can be used in many ways to produce linear musclelike actuation. For example, the film and electrodes can be formed into a tube shape, rolled into a scroll, stretched over a frame, or laminated to a flexible substrate to produce bending. The best configuration depends on the application and properties of the film. Several configurations that the Stanford Research Institute–based authors have made are shown in figure 7.10.

Many of these configurations are similar to those used for piezoelectric materials. This is not surprising, because piezoelectric materials also are actuated as dielectrics between relatively compliant electrodes. However, elastomeric polymer material is much softer and the amount of deformation is much greater than can be achieved with piezoelectric materials, allowing for a greater variety of actuator configurations.





Figure 7.10 Some configurations of dielectric elastomer actuators.

### Examples of Dielectric Elastomer Actuators

The bow-tie and trench actuators shown in figure 7.10, and variations on these designs, have been considered for artificial muscles. These actuators are based on simple two-dimensional fabrication procedures. The basic actuator configuration can be stacked in parallel to increase the force produced, or chained in series to increase the stroke. In this sense, the actuator is also like natural muscle.

Figure 7.11 shows a bow-tie actuator made from a single-layer silicone film about  $40\ \mu\text{m}$  thick. This actuator is 7 cm wide and 3 cm in length (including connectors). It is capable of a free stroke of 5 mm and a blocked force of roughly 1 N. The mass of such an actuator is 1 g. Of this 1 g, the active film mass is less than 0.1 g.



**Figure 7.11** Silicone film bow-tie actuator.



**Figure 7.12** Acrylic film double bow-tie actuator.

The linear actuator in figure 7.12 is composed of two layers of acrylic film sandwiched together. An advantage of the two-layer construction is that the high-voltage electrode can be sealed into the center of the actuator. It is then relatively simple to stack several actuators in parallel in order to produce the required force. In this case, two bow-tie units are connected in series to increase the stroke. This actuator has a blocked force output of about 2 N and a free stroke of up to 8 mm. Including the mechanical connections, the actuator measures about 3 cm on a side. The mass of the actuator is about 1.5 g; however, the mass of the active area is just 0.12 g.

A rolled artificial muscle is shown in figure 7.13. The advantage of the rolled actuator is that a larger cross-sectional area of film can be used with a correspondingly greater force output. The main disadvantage of the rolled actuator is that it is more difficult to fabricate. The rolled actuator illustrated in figure 7.13 could produce a blocked force of 250 g and a free stroke of 6 mm.

We have built and tested a great many other linear actuators. The ones mentioned above have proven to offer the best combination of reliability and performance with minimal added structure.




**Figure 7.13** Acrylic film rolled actuator.

## IS IT AN ARTIFICIAL MUSCLE? COMPARISON WITH NATURAL MUSCLE DATA


In order to determine whether or not a dielectric elastomer actuator can be considered to be an artificial muscle by the criteria set in the Actuator Characteristics and Metrics section, we conducted several tests on a sample actuator. The actuator was an acrylic bow tie, similar to the one shown in figure 7.12. These tests were performed by biologists using techniques and apparatus identical to those used for testing of natural muscle (Full and Meijer, 2000). These tests should be considered preliminary and incomplete. They did not include rigorous tests of all of the parameters, nor were the test protocols necessarily representative of actual loading conditions. Nonetheless, this testing illustrates how comparisons can be made between natural muscle and artificial actuators and provides a rough assessment of the musclelike features of the acrylic bow-tie actuator.

The leg muscles of a cockroach are used as a basis of comparison with the dielectric elastomer actuator. Because the acrylic elastomer actuators are intended for application to a hexapedal walker (see section entitled *Insect-Inspired Hexapedal Walking Robot*), such a comparison is valid. The hexapedal walker has fewer degrees of freedom than a cockroach. Because there is no direct mapping between the dielectric elastomer actuators and the leg muscles of a cockroach, several different cockroach leg muscles are used as a basis of comparison.

The basic setup for testing the artificial muscle is shown in figure 7.14a. A servo-controlled arm with force feedback is used to impart a motion to the artificial muscle actuator. This basic setup was used for several different tests. The muscle was made from an acrylic elastomer with carbon grease electrodes. The active area was 1.8 by 1.6 cm and the film thickness was 70  $\mu\text{m}$ . With this setup, we determined some of the parameters that define the functional workspace of the actuator, including maximum stress and strain, breaking stress, maximum power production, and elastic and viscous properties. Details of the experimental setup are given by Full and Meijer (2000).



a) Photo of an EAP actuator in the measurement apparatus




b) Sample work loop protocol consisting of sinusoidal motion with square wave activation of the muscle at 50% duty cycle

**Figure 7.14** Experimental setup for work loop testing of dielectric elastomer EAP actuator. (a) Photo of an EAP actuator in the measurement apparatus; (b) sample work loop protocol consisting of sinusoidal motion with square wave activation of the muscle at 50% duty cycle.

Four different sets of experimental tests were run. The first two tests were part of the basic set of tests shown in figure 7.3. In the first set, we studied the force output of the actuator as a function of time and stimulation level by applying different voltages (in the range 0–5 kV). During these experiments, the actuator was kept at a fixed length (i.e., the experiments were isometric). In the second set of experiments, the nonactivated actuator was stretched to different lengths to determine its passive stress-strain and viscoelastic properties (these were force-length experiments). In these experiments, the strain amplitude was varied from 2% to 20% and the strain rate was varied from 0.5 to 50 mm/sec. All experiments were performed at a pre-tension of 1 N.

The third set of experiments used a dynamic mechanical analyzer (from EnduraTEC Systems Corporation) to measure the mechanical impedance (viscoelastic properties) of the acrylic material used in the



**Figure 7.15** Sample work loops for a stretched film acrylic actuator. (a) Work loops; (b) actuator configuration.


actuators. In this test setup, a force feedback transducer is located between the test sample and ground. The other end of the test sample is driven by a linear voice-coil actuator. This set of tests indicates the passive mechanical impedance of acrylic artificial muscle actuators.

In the fourth set of experiments, the work-loop technique was used to determine the power output of the actuator. In these work-loop experiments, sinusoidal length oscillations were imposed on the actuator and the force response was recorded. The actuator was stimulated at different phases of the oscillatory motion cycle (see figure 7.14b). The duty cycle for stimulation was fixed at 50%; the imposed strain ranged from 2.5% to 20%, and the cycle frequency was 1 Hz. The oscillation experiments at 2.5% strain were repeated at higher frequencies for the stimulation phase that yielded the greatest power output. Unfortunately, due to limitations of the experimental setup, the actuator for the work-loop tests was different from the bow-tie actuator used in the other tests. The artificial muscle tested in this case was a stretched-film actuator with only a single active area (figure 7.15b). This actuator is made from the same material as the acrylic bow tie and has the same basic properties; thus the test results are meaningful. All of these experiments were performed at a pre-tension of 1 N.

It should be noted that the applied voltage was not the maximum for the actuator tested, nor was the actuator design ideal (the actuator was made to fit the needs of the test apparatus). Thus, the maximum output of the actuator reported here is substantially less than we would expect with actuators such as those reported in the previous sections. It should also be noted, however, that the work-loop protocol followed here was deliberately chosen for its simplicity and does not precisely represent the performance of the actuator in a robotic application (figure 7.15a). Nonetheless, the similarity between the motion of the lever arm and that of moving a robotic link are obvious, so that even this simple protocol can give a good idea of the utility of the artificial muscle in many robotic applications where power generation is required.

### Force-Time Response

The results of force versus time testing under isometric conditions is shown in figure 7.16 and compared to that of a cockroach muscle. As in the natural muscle, the force production of dielectric elastomer artificial muscle varies in a controllable manner with the level of stimulation.



**Figure 7.16** Comparison of the isometric stimulation response of natural muscle (cockroach leg muscle 177c-top) and dielectric elastomer artificial muscle (bottom).

The force response increased quadratically with voltage, as would be expected from the principle of operation of the dielectric elastomer technology. The maximum stress, calculated from change in force and the cross-sectional area ( $0.011 \text{ cm}^2$ ) of the actuator, was 0.53 MPa when the muscle was lengthened to a pre-tension of 1.9 N. The rate of force development and relaxation was rapid ( $\sim 10$  msec). The activation and relaxation kinetics depended largely on the electrical driver circuitry used; the response time could be shortened to less than 1 msec through optimization of the electrical circuit that charges and discharges the actuator. These preliminary results show that the maximum stress attained by the dielectric elastomer artificial muscle falls well within the range of values measured for natural muscle using similar measurement techniques. The maximum stress is similar to that of the cockroach muscle as well. The response time also appears very rapid, compared with that of most biological muscles.

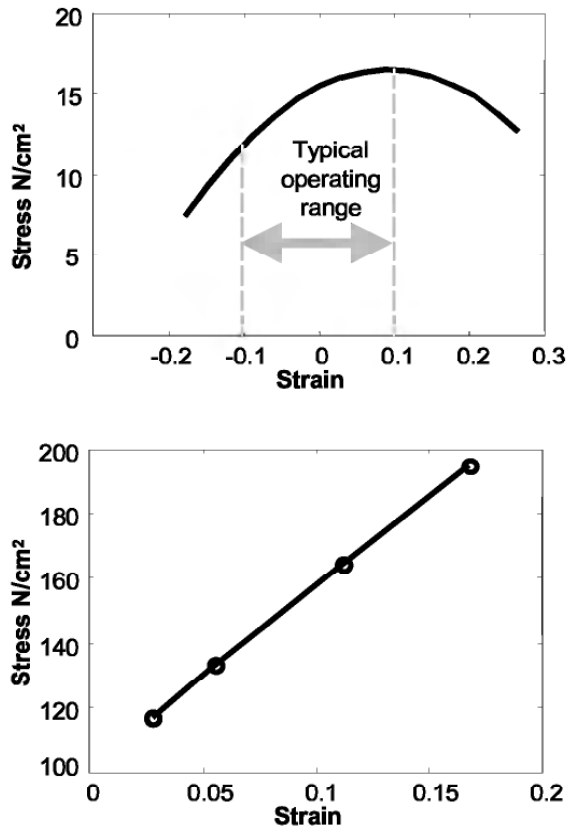
### **Force-Length Relationship**

The measured force-length function for the dielectric elastomer EAP actuator revealed a roughly linear stress-strain curve (figure 7.17). The average local Young's modulus was about 6.0 MPa. This value is about eight times greater than that of a natural muscle (in this case, a cockroach leg muscle). Other dielectric elastomer EAP materials (such as silicones) and actuator designs are less stiff, so we can say that EAP muscles can approximate the stiffness of comparable natural muscle. The EAP muscle does not show the parabolic force-length nonlinearity that is characteristic of active, natural muscle over a portion of the response. Studies of biological muscle show that animals in nature typically operate in the ascending or plateau region of the force-length relationship. The functional importance of this nonlinearity is unclear. More research is needed to determine if this feature of muscle has desirable effects on joint stability, impedance modulation, control, or other factors. If so, it may be possible to design artificial muscles with a preload such that they become slack and effectively exhibit the parabolic force-length relationship.

### **Mechanical Impedance**

The dynamic mechanical analyzer was used to measure the effective viscoelastic loss angle (a measure of the effective damping) of both the silicone and the acrylic polymer materials (see figure 7.7). These tests of the passive properties of the dielectric materials and natural muscle showed that the dielectric elastomer materials are more lightly damped than a natural (e.g., cockroach leg) muscle. However, it is easy to add materials that increase the damping. The tests also show that the elastic






**Figure 7.17** Comparison of strain-stress characteristics of muscle and dielectric elastomer EAP actuator. Top panel shows strain-stress function for maximally activated cockroach muscle (179). Arrow indicates range of strains where muscle functions under natural conditions. Bottom panel shows the strain-stress function for nonstimulated EAP actuator.

modulus of the materials themselves (as opposed to the entire actuator) are similar to that of natural muscle.

### Energy and Power Output: Work Loops

The dielectric elastomer EAP actuator produced work loops similar in shape to natural muscle (examples are shown in figure 7.15). Like natural muscle, the actuator could both generate and absorb energy. When the actuator was stimulated during the lengthening phase of the cycle, it could generate power.

Maximum power output of about 10 W/kg was obtained from the EAP actuator, using a cycle frequency of 10 Hz, a strain of 2.5%, and a stimulation voltage of 5 kV. This voltage is lower than the maximum of the muscle tested. Nonetheless, the power output falls within the lower boundary of values for natural muscle tested using work-loop techniques (see figure 7.5 above). The strains that maximized power output were low and the stress high relative to natural muscles operating at



**Figure 7.18** Self-contained hexapod robot powered by dielectric elastomer EAP artificial muscles.

the same frequency. However, both the stress and strain fell within the bounds of the performance of natural muscle.

### **APPLICATIONS OF EAP MUSCLES TO ROBOTS**

Now that the dielectric elastomer type of EAP actuators has been shown to fall within the parameter space of natural muscle in many respects, it is instructive to show how these actuators are being used in robotic applications and how these applications seek to benefit from the musclelike characteristics of these actuators.

#### **Insect-Inspired Hexapedal Walking Robot**

Given that most of our comparisons between natural muscle and dielectric elastomer artificial muscle are based on the leg muscles of the cockroach, it follows that the application of dielectric elastomer artificial muscle to insectlike walking robots is of interest. Figure 7.18 shows a completely self-contained, battery-powered six-legged robot that is

loosely inspired by a cockroach. Each leg has two degrees of freedom driven by a single actuator bundle that constitutes an artificial muscle. The robot weighs 650 g (including the battery) and is roughly 30 cm in length. In order to minimize the number of muscles required (until manufacturing procedures can be improved), we have opposed the muscles with springs located at the pivot points of the legs. Each muscle is a bundle of three acrylic double bow-tie actuators (see figure 7.12). Each muscle is driven by a small (4 g) DC-DC converter with a maximum output of 5 kV and 500 mW of power. The joint motion is controlled by a PIC.


At this point, the robot can only walk slowly over even terrain. Improvements are needed in the driving electronics and in the strength and durability of the muscles themselves for greater mobility. Nonetheless, this robot is significant because it is believed to be the first self-contained walking robot that is powered by dielectric elastomer actuators. The robot will also serve as a test bed to assess the advantages of musclelike actuation. In particular, we hope to exploit the viscoelastic behavior of the muscles, to enable robots to reject disturbances due to obstacles or uneven terrain, much as reflexes in the musculoskeletal system of a cockroach can help it run over uneven terrain in a stable manner and without large disturbances of its torso. Eventually, we hope to operate the muscles fast enough that their viscoelastic behavior can provide the correct amount of energy storage and absorption while the robot is walking or running.

### **Insect-Inspired Flapping-Wing Robot**

Figure 7.19 shows a flapping-wing mechanism whose design was inspired by the mechanics of the many flying insects whose wings are driven indirectly by muscles located in the thorax. The muscles flex the exoskeleton and move the wings, which are attached to the exoskeleton. In the same way, a bundle of artificial muscles can flex a plastic exoskeleton with wings attached.

Four silicone bow-tie actuators (see figure 7.11) drive the mechanism. The mechanism is designed so that its optimum flapping frequency coincides with the resonance of the muscles. Thus, the flapping amplitude can be increased and the power required to flap the wings minimized, because no energy is needed to counteract the inertia of the wings and mechanism. It has been shown that insect muscle does indeed operate in this manner (Alexander, 1988; Dickinson and Lighton, 1995). These artificial muscles show a behavior similar to that of the springlike behavior of many insect flight muscles.


The eventual goal of this effort is to make simple and robust flying platforms that can be used for reconnaissance in cluttered environments. Such missions will require efficient slow-speed and hovering



**a) schematic of the EAP-powered mechanism**



**b) demonstration mechanism with a six inch wingspan**



**c) Conceptual representation of muscle-powered flapping-wing flyer (Source: Dave Loewen, University of Toronto Institute for Aerospace Studies)**

**Figure 7.19** Insect-inspired artificial-muscle-powered flapping wing mechanism. (a) Schematic of EAP-powered mechanism; (b) demonstration mechanism with a six inch wingspan; (c) conceptual representation of muscle-powered flapping-wing flyer. (Courtesy of Dave Loewen, University of Toronto Institute for Aerospace Studies.)

flight capabilities. Inspired by nature, flapping-wing flight offers potential advantages in thrust-to-power ratio and stability over conventional flyers based on rotors. If the flyers are to be electrically powered, which is both quieter and logistically advantageous, then lightweight and powerful electric actuation is needed. Dielectric elastomer actuators, incorporated into a biomimetic flapping mechanism, can form the basis of such an electric flapping-wing propulsion system.

The mechanism has proven to be quite robust because it flaps wings at their spanwise bending resonant frequency of 18 Hz. However, it currently has only 4 muscles operating in parallel. As many as 25 muscles operating in parallel will be required to produce the necessary power output for hovering a 50 g vehicle, with the muscle bundle


resonating with the wings at 40 Hz (the optimal thrust-to-power point for the wings).

### Serpentine Robot

Figure 7.20 shows how dielectric elastomer artificial muscles can be used to make a serpentine manipulator or mobile robot with an extremely large number of joints. The large strain capability of the dielectric elastomer actuator is exploited to make the structure extremely simple. The robot would be composed of a large number of repeating



**a) Demonstration of modular design incorporating high-aspect-ratio linear actuators**



**b) CAD representation of a segment of a serpentine manipulator that is composed of a series of 1-DOF joints actuated by linear actuators**


**Figure 7.20** Serpentine manipulator joints actuated by dielectric elastomer EAP actuators (a) demonstration of modular design incorporating high-aspect-ratio linear actuators (b) CAD representation of a segment of a serpentine manipulator that is composed of a series of 1-DOF joints actuated by linear actuators.

units, each driven by a pair of trench-type dielectric elastomer actuators. However, these linear actuators have a very high aspect ratio, such that the robot cannot achieve bending in two axes in a single joint. Therefore, it would be necessary to build a serpentine manipulator consisting of a series of such joints, where every other joint would be rotated  $90^\circ$  with respect to the initial joint. Any two adjoining joints would form a spherical joint that can produce motion about an arbitrary axis.

In the demonstration joints, each actuator was capable of about 20% strain, which satisfied the requirements for a 1 m radius of curvature. The force produced by each trench was about 2 N. We estimate that a 1 m-long serpentine manipulator would need actuators that produce at least 26 N at the base, so we would need to stack about 13 layers in order to produce the necessary force. Our demonstration apparatus did not include more than one layer (or actuator in parallel) at each joint. However, the flat shape of each actuator would enable us to include 13 layers in a manipulator of this design. Shorter manipulators or mobile snakelike robots would not require as great an amount of force at any joint.

### Inchworm Robot

Figure 7.21 shows a small (16 mm long) robotic platform that crawls like an inchworm. The robot's "body" is a rolled actuator made from silicone, with an electrostatic clamp at each end. The clamps enable



**Figure 7.21** Inchworm-inspired propulsion system based on a silicone dielectric elastomer rolled actuator with electrostatic clamps.

the inchworm to travel on vertical as well as horizontal surfaces. In tests, the inchworm was able to travel at a maximum speed of about 10 cm/sec.

Robotic platforms like the inchworm could eventually be used to make small robots for tasks such as inspection in narrow pipes. The inchworm takes advantage of the large strain capability of the dielectric elastomer rolled actuator. It also shows how a musclelike actuator can function without a rigid supporting skeleton, as do many soft-bodied biological creatures such as worms.

### **Other Biomimetic Applications**

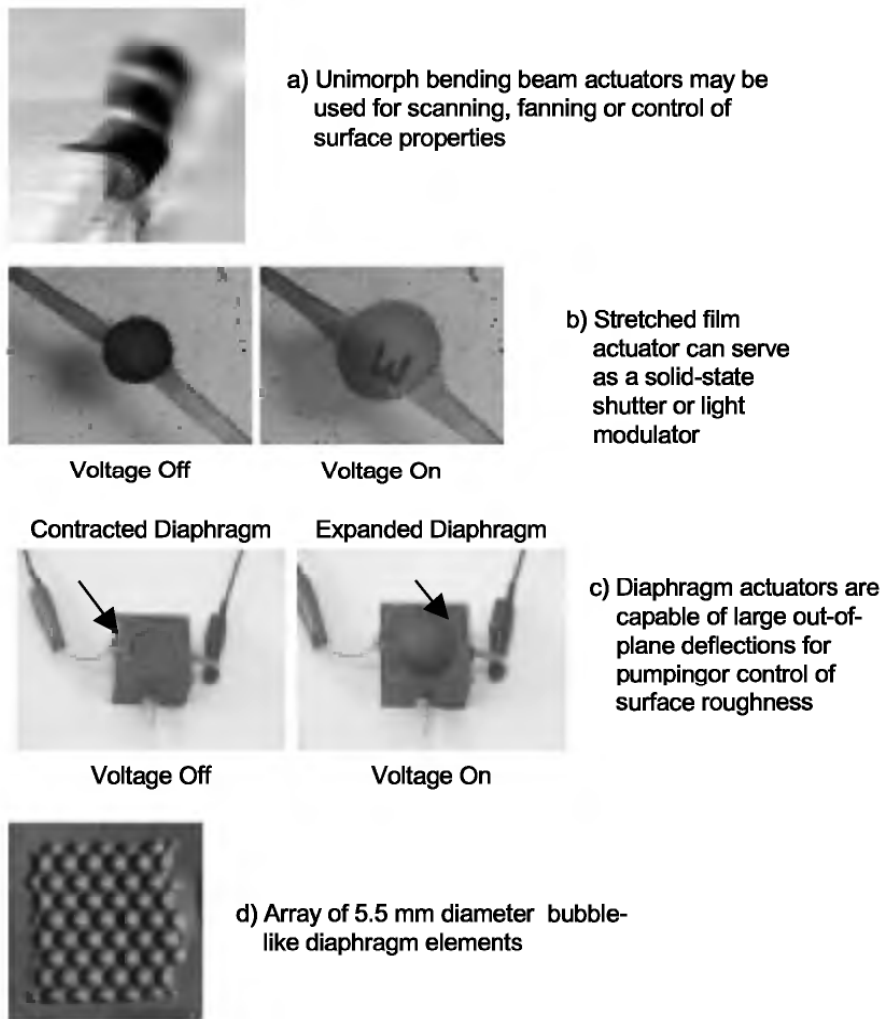
Thus far, we have focused on applications in which the muscle does mechanical work to move or modulate the impedance of an appendage or body segment. In many instances in nature, muscle has other functions. For example, muscle inflates and deflates the lungs to enable us to breathe. Muscle also acts as a valve in numerous places in the body. Muscle changes the shape of the lens of the eye and controls the amount of light that passes through the eye. Muscle changes the shape of our lips or eyebrows to allow us to express emotion. Muscle erects hairs when a creature is frightened or cold, to change its appearance or thermal properties.

Figure 7.22 shows a variety of dielectric elastomer EAP actuators in configurations that could relate to these other applications. The diaphragm actuator could act as a pump for airflow for a sniffing robot (analogous to the diaphragm of the lungs). An array of small diaphragms could change surface roughness (analogous to the ability to erect hairs in mammals). A stretched film actuator could act as a dilating pupil.

Muscle also produces sound in biological creatures. Usually, it does so indirectly by processes such as providing airflow to vibrate vocal chords or rubbing two surfaces together. While not necessarily biomimetic, the dielectric elastomer can produce sound directly by vibrating at sonic frequencies. Heydt and colleagues (2000) describe such an approach based on an array of diaphragm actuators.

### **SUMMARY AND CONCLUSIONS**

We have argued that musclelike actuators are critical to the development of robots with the dexterity and mobility of biological creatures. Although artificial muscle actuators cannot match all aspects of muscle behavior, they must be able to produce certain performances and must have operational characteristics that are similar to natural muscle used for similar tasks. We have also noted the importance of mechanical impedance similar to that of natural muscle.



**Figure 7.22** Other biomimetic applications of dielectric elastomer EAP actuators. (a) Unimorph bending beam actuators may be used for scanning, fanning or control of surface properties; (b) stretched film actuator can serve as a solid-state shutter or light moderator; (c) diaphragm actuators are capable of large out-of-plane deflections for pumping or control of surface roughness; (d) array of 5.5 mm diameter bubble-like diaphragm elements.

No single actuation technology exhibits all of the desirable qualities of natural muscle. Dielectric elastomers appear promising for artificial muscles and have been used to make several types of musclelike actuators. Preliminary testing with experimental biology techniques has demonstrated that such actuators can indeed exhibit musclelike behavior, based on several tests of performance parameters. The successful implementation of dielectric elastomer and other artificial muscle technologies in biomimetic robots will require improvements in electric driver circuits as well as further testing to determine the technologies' long-term durability. This research must be complemented by further development of control mechanisms and motion planning,



and by biological studies of natural muscle and biological motion, in order to realize robots with capabilities approaching those of biological systems.

## ACKNOWLEDGMENTS

The authors would like to thank the many individuals at SRI International who have contributed to the development of the dielectric elastomer actuation technology described above and to the preparation of this manuscript. The authors would also like to thank the individuals at the University of California, Berkeley who contributed to the measurements of natural muscle and evaluation of the dielectric elastomer artificial muscle. Much of the basic dielectric elastomer technology was developed under the management of the Micromachine Center of Japan as the Industrial Science and Technology Frontier Program, Research and Development of Micromachine Technology of MITI (Japan), supported by the New Energy and Industrial Technology Development Organization. The development of dielectric elastomer artificial muscle actuators and the application of the technology to large biomimetic robots was supported by the Defense Advanced Research Projects Agency, the Office of Naval Research, and the Naval Explosive Ordnance Technology Division. SRI visiting scientist Mr. Guggi Kofod of the Risø National Laboratory, Denmark, fabricated and tested the rolled artificial muscle actuator. Work performed at the University of California, Berkeley was supported by ONR MURI contract N00014-98-0747 and DARPA-ONR contract N00014-98-1-0669. The authors also thank Dr. Anna Ahn for the use of preliminary data on insect muscle. David Loewen, Dr. James DeLaurier, and Derek Bilyk of the University of Toronto Institute for Aerospace Studies provided inspiration and assistance in the development of the flapping-wing mechanism. The authors would also like to thank Naomi Campbell, Jean Stockett, and Peggy Nutsch for their assistance in the preparation of this manuscript.

## REFERENCES

- Alexander, R. (1988). *Elastic Mechanisms in Animal Movement*. Cambridge: Cambridge University Press.
- Alexander, R. (1997). Optimum muscle design for oscillatory movements. *J. Theor. Biol.* 184: 253–259.
- Bobbio, S., Kellam, M., Dudley, B., Johansson, S., Jones, S., Jacobson, J., Tranjan, F., and Dubois, T. (1993). Integrated force arrays. In *Proceedings of the IEEE Micro Electro Mechanical Systems Workshop*, Fort Lauderdale.
- Brock, D. (1991). Review of artificial muscle based on contractile polymers. Artificial Intelligence Laboratory Memo 1330, Massachusetts Institute of Technology, Cambridge.

- Brown, I. E., and Loeb, G. E. (1999). A reductionist approach to creating and using neuromusculoskeletal models. In J. M. Winters and P. E. Crago (eds.), *Biomechanics and Neural Control of Movement*. New York: Springer-Verlag (in press).
- De Rossi, D., and Chiarelli, P. (1994). Biomimetic macromolecular actuators. In *Macro-Ion Characterization*. American Chemical Society Symposium Series, vol. 548. pp. 517–530.
- Dickinson, M., and Lighton, J. (1995). Muscle efficiency and elastic storage in the flight motor of *drosophila*. *Science* 268 (7 April): 87–90.
- Dickinson, M. H., Farley, C. T., Full, R. J., Koehl, M. A. R., Kram, R., and Lehman, S. (2000). How animals move: An integrative view. *Science* 288 (7 April): 100–106.
- Full, R. J. (1997). Invertebrate locomotor systems. In W. Dantzer (ed.), *The Handbook of Comparative Physiology*. Oxford: Oxford University Press, pp. 853–930.
- Full, R. J., Stokes, D. R., Ahn, A. N., and Josephson, R. K. (1998). Energy absorption during running by leg muscles in a cockroach. *J. Exp. Biol.* 201: 997–1012.
- Full, R., and Meijer, K. (2000). Artificial muscle versus natural actuators from frogs to flies. In Y. Bar-Cohen (ed.), *Electroactive Polymer Actuators and Devices (EAPAD)*, Proceedings of SPIE vol. 3987, pp. 2–9.
- Giurgiutiu, V., and Rogers, C. (1997). Power and energy characteristics of solid-state induced-strain actuators for static and dynamic applications. *J. Intell. Mater. Syst. Struct.* 8 (September): 738–750.
- Heydt, R., Pelrine, R., Joseph, J., Eckerle, J., and Kornbluh, R. (2000). Acoustical performance of an electrostrictive polymer film loudspeaker. *J. Acoust. Soc. Am.*
- Hill, A. V. (1938). *Proc. R. Soc. Lond. B Biol. Sci.* 126: 136–195.
- Hogan, N. (1985a). Impedance control: An approach to manipulation: I. Theory. *Dynamic Syst. Meas. Control* vol. 107, no. 1: 1–7.
- Hogan, N. (1985b). Impedance control: An approach to manipulation: II. Implementation. *Dynamic Sys. Meas. Control* vol. 107, no. 1: 8–16.
- Hogan, N. (1985c). Impedance control: An approach to manipulation: III. Applications. *Dynamic Syst. Meas. Control* vol. 107, no. 1: 17–24.
- Houk, J. C. (1979). Regulation of stiffness by skeletomotor reflexes. *Annu. Rev. Physiol.* 41: 99–114.
- Hunter, I. W., and Lafontaine, S. (1992). A comparison of muscle with artificial actuators. *Technical Digest of the IEEE Solid-State Sensor and Actuator Workshop*, Hilton Head, S.C. pp. 178–185.
- Josephson, R. K. (1985). Mechanical power output from striated muscle during cyclic contraction. *J. Exp. Biol.* 114: 493–512.
- Klute, G. K., Czerniecki, J. M., and Hannaford, B. (1999). McKibben artificial muscles: Pneumatic actuators with biomechanical intelligence. In *Proceedings of the IEEE/ASME International Conference on Advanced Intelligent Mechatronics*, Atlanta, Ga.
- Kornbluh, R., Pelrine, R., and Joseph, J. (1995). Elastomeric dielectric artificial muscle actuators for small robots. In *Proceedings of the Third IASTED International Conference on Robotics and Manufacturing*, Cancun, Mexico.
- Kornbluh, R., Pelrine, R., Joseph, J., Heydt, R., Pei, Q., and Chiba, S. (1999). High-Field electrostriction of elastomeric polymer dielectrics for actuation. *Proceedings of SPIE International Symposium on Smart Structures and Materials: Electro-Active Polymer Actuators and Devices*, Newport Beach, Calif. pp. 149–161.

- Kuhn, W. (1949). Artificial muscle model. *Experientia* vol. 5, no. 8: 318.
- Lutz, G. J., and Rome, L. C. (1994). Built for jumping: The design of the frog muscular system. *Science* 263 (5145): 370–372.
- Meijer, K., and Full, R. J. (2000). Stabilizing properties of invertebrate skeletal muscles. *Am. Zoologist* 39(5): 104.
- Otero, T., Cantero, I., Sansiñena, J., and De Paoli, M. (1999). Artificial muscles working in both aqueous solutions or air. *Smart Structures and Materials 1999: Electroactive Polymer Actuators and Devices, Proceedings of SPIE*, vol. 3669. pp. 98–108.
- Pelrine, R., Kornbluh, R., and Joseph, J. (1998). Electrostriction of polymer dielectrics with compliant electrodes as a means of actuation. *Sensor and Actuators A: Physical* 64, pp. 77–85.
- Pelrine, R., Kornbluh, R., and Kofod, G. (2000). High-Strain actuator materials based on dielectric elastomers. *Adv. Mater.* 12: 1223–1225.
- Rome, L. C., Funke, R. P., Alexander, R. M., Lutz, G., Aldridge, H., Scott, F., and Freadman, M. (1988). Why animals have different muscle fiber types. *Nature* (London) vol. 335, no. 6193: 824–827.
- Shahinpoor, M. (1995). Micro-electro-mechanics of ionic polymer gels as electrically controllable artificial muscles. *J. Intell. Material Syst. Struct.* 6: 307–314.
- Smela, E., Inganas, O., and Pei, Q. (1993). Electrochemical muscles: Micromachining fingers and corkscrews. *Adv. Materials*, communications section 5: 630–632.
- Stevenson, R. D., and Josephson, R. (1990). Effects of operating frequency and temperature on mechanical power output from moth flight muscle. *J. Exp. Biol.* 149: 61–78.
- Wax, S. G., and Sands, R. R. (1999). Electroactive polymer actuators and devices. In Y. Bar-Cohen (ed.), *Proceedings of SPIE, Smart Structures and Materials: Electroactive Polymer Actuators and Devices*, vol. 3369. pp. 2–10.
- Winters, J. (1986). Muscle as an actuator for intelligent robots. Robotics International of SME, paper no. MS86-760. *Robotics Res. Trans.*, Scottsdale, Ariz. (18–21 August).
- Zajac, F. E. (1989). Muscle and tendon: Properties, models, scaling, and application to biomechanics and motor control. *Crit. Rev. Biomed. Eng.* 17: 359–411.

# 8

## Polymer-Based MEMS Actuators for Biomimetics

Robert Horning and Burgess Johnson

Actuators are crucial elements in biomimetic robots, finding uses in locomotion and stabilization (legs, wings, fins), and in pointing and manipulation of appendages (arms, fingers, antennae). The high level of adaptability that is sought in an autonomous robot requires an actuator with the adaptability of biological muscle. Muscles have a host of desirable properties. Among other things, they can generate large forces for their weight, they have high strain, they efficiently convert input (chemical) energy into mechanical motion, they are robust and reliable, and have built-in fault tolerance. The ideal engineered muscle should possess these properties as well. In this chapter, we describe our development work on polymer-based electrostatic actuator arrays that have many desirable musclelike properties and potentially can meet the requirements of many biomimetic robots.

For a typical robot, the locomotion actuators will require the largest forces and displacements. For example, in a several-kilogram robot, each locomotion actuator may need to generate up to  $\sim 1$  kg of force with 3–5% strain and a total throw of 2–3 mm. Power consumption during actuation must be held to a minimum, because batteries can become a significant fraction of the total robot weight. Preferably, the power consumption should be near zero for an actuator that is in an actuated position but not moving (i.e., a flexed muscle). A bandwidth of about 5 Hz should be sufficient for use in robots of this size. Manipulating actuators may require less force but more finesse, and a minimum of space and power.


Among the most widely used actuators for biomimetic robots are electromagnetic servos, pneumatic actuators, and shape memory alloys (SMAs). Servos are readily available in many sizes, but have a low output power-to-weight ratio, require high input power, and require some sort of “transmission” to convert rotary motion to linear motion. Pneumatic actuators also have low power-to-weight ratios, requiring an air pressure source, regulators, valves, and hoses. These add considerably to both the size and weight of the robot. SMAs are attractive because of their small size. They can be controlled easily and can

produce the desired high forces and  $\sim 5\%$  strain. Unfortunately, SMAs use a substantial amount of power because they are thermally activated. Electric current, flowing through the SMA wire, is typically used to heat and move the actuator. This current must be maintained to hold the actuator in an actuated position, even if it is not moving.

The polymer-based actuator, described in this chapter, has the potential to meet most of the above requirements for biomimetic robot applications. The polyMEMS actuator (PMA) is composed of many micro- to milliscale unit cells configured in a three-dimensional array. Related structures have been described by Minami and colleagues [1–2] and by Jacobson and colleagues [3]. Cabuz [4] developed an electrostatic air pump that operates on a principle very similar to the PMA. Unlike the earlier actuators, the PMA can generate macroscale forces and displacements sufficient for biomimetic robots, in an appropriately sized device. Unit cells are actuated electrostatically. This provides a strong force in each of the small unit cells, and multiplies to a macroscopic force when unit cells are joined in an array. Because it is electrostatic, the device has inherently low power consumption during actuation. Holding the device in its actuated position consumes almost no additional power.

## ACTUATOR CONCEPT

Figure 8.1 illustrates the unit cell and actuator concept. A unit cell is fabricated from two thin, flexible, polymer sheets, bonded together at localized points. When an external force is applied to the cell, it opens



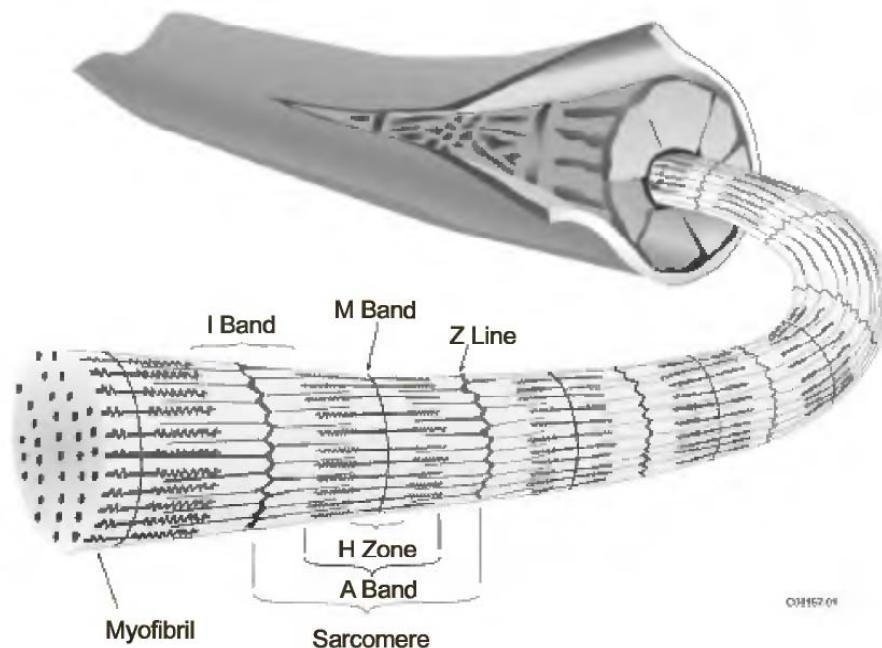
**Figure 8.1** A schematic representation of the PolyMEMS Actuator unit cells and 3-D array. An external force opens the unit cells like a spring. A voltage exceeding the pull-in voltage closes the cell. Layers of unit cells can be controlled individually or in groups.

like a spring. Thin metal and dielectric films are deposited and patterned on the inner surfaces of the cell, creating electrodes. When a sufficient voltage is applied between these electrodes, the electrostatic force causes the unit cell to completely collapse. Each unit cell supplies a certain force, and moves through a certain displacement or stroke. Additional unit cells in the plane of the sheet increase the force. Additional unit cells stacked on top of each other increase the displacement. Thus, a 3-D array of unit cells can generate large forces and large displacements simultaneously.

The modular nature of three-dimensional polyMEMS actuator arrays allows fabrication of a variety of actuator shapes and topologies to fit a variety of biomimetics applications. For example, cylindrical, prismatic, helical, or toroidal actuator shapes can be fabricated with only minor modifications to the basic device fabrication process. These actuators can be made to cover both large and small areas because the underlying plastic substrate can be made in almost any size or shape—the principles of the electrostatic actuation will be unchanged.

### COMPARISON WITH MUSCLE

The polyMEMS actuator has many features in common with biological muscles. Muscle tissue (figure 8.2) is composed of *unit cells* of myosin and actin protein filaments. External forces extend the filaments. When triggered by a signal from a nerve, chemical energy collapses the fila-



**Figure 8.2** Components of biological muscle tissue. The protein filaments form a 3-D array of unit cells that, acting together, are able to produce large forces and displacements. The PolyMEMS Actuator uses a similar 3-D array configuration.

ments. The protein unit cells are arranged in 3-D arrays—sarcomeres and myofibrils—to multiply both the force and displacement of the individual filaments. Although the unit cells in the polyMEMS actuator are structurally different than muscle cells, the multiplicative concept is the same.


Because there are many small units working together, the failure of one set of filaments in a muscle has only a small effect on the operation of the whole muscle. This is true also in the PMA, where one unit cell can fail without affecting the others. This fault tolerance is a key to biological survivability, and is expected to be a great advantage in biomimetic robots.

Muscles work in tension only, and antagonistic pairs are used for bidirectional motion and tuning the stiffness of a joint. The PMA also acts only in tension, because the electrostatic force is always attractive. Thus, two PMAs will be used together as antagonistic pairs. An advantage of this is that the stiffness of a joint can be tuned at will. If both of the antagonistic muscles/actuators are flexed or tightened simultaneously across a joint (e.g., at your elbow), the joint will not bend but will become stiffer. This ability is used frequently in holding objects, balancing, stabilization, and so on. The similarity between the PMA and muscles will allow joint stiffness to be tuned in biomimetic robots as well.

During locomotion, such as walking or running, muscles and tendons act as springs, alternately storing and releasing energy [5]. This greatly reduces the energy cost of locomotion. The PMA is also a spring. Thus, this feature could be exploited in biomimetic robots.

## UNIT CELL MODEL AND DESIGN

Figure 8.3 is a more detailed drawing of one unit cell, showing the base polymer sheets, the metal and dielectric electrodes, and the bond region. Initially, the sheets are flat and parallel to each other. In the absence of an applied voltage, an external force will pull the sheets apart at the center of the unit cell. The bending of the sheets results in a springlike behavior of the unit cell. When a voltage is applied between the electrodes, the electric field is largest at the edges of the unit cell, where the separation between the electrodes is smallest. The resulting attractive force causes the point of contact between the electrodes to move slightly in toward the center of the unit cell. Eventually, as the voltage is increased, a threshold is reached where the electrostatic force is sufficiently large that the contact point moves all the way to the center of the unit cell in a “zipping” motion. This threshold is called the *pull-in voltage*. The entire cell collapses, overcoming the externally applied force, and the two electrodes are in contact everywhere. The large



**Figure 8.3** A more detailed diagram of one unit cell, showing the substrate, electrodes, bonding sites, and the location of the strongest electric fields. The inset shows the parameters used in deriving the analytical model.

lateral zipping is converted to a vertical motion. Also, the strong electrostatic force is converted to a strong vertical force with the application of only a modest voltage.

The strongest force is achieved when the electrodes are at their minimum separation. Preferably, this would be intimate contact everywhere. Therefore, smooth surfaces are important to avoid unwanted and uncontrolled electrode separation. A bonding scheme that does not add to the gap at the unit cell edges is also needed.

The unit cell has been modeled both analytically and by finite element analysis (FEA). The analytical model uses an energy minimization approach. Consider the half cell in the inset of figure 8.3. The total energy,  $U_T$ , is the sum of the bending energy in the curved region,  $U_B$ , and the electrostatic energy in the flat, clamped region,  $U_E$ . Also, there is electrical energy in a very small region near the contact point, but FEA shows that this region is very small and can be ignored in the derivation of the model. The bending energy is nonzero only in the curved portion of the profile, from the centerline to the contact point,  $c$ . Using a standard profile [6], the bending energy is found to be

$$U_B = \frac{1}{2}kb \left(\frac{l}{c}\right)^3 \delta^2$$

where

$$k = \frac{Et^3}{2l^3}$$



The parameters  $l, c, b, \delta, E$ , and  $t$  are defined for the half unit cell in the inset of figure 8.3. Note that this is simply the energy of a spring with spring constant  $kb(l/c)^3$ .

The electrostatic energy is nonzero only in the flat, clamped portion of the profile, and is therefore just the energy stored in a parallel plate capacitor at voltage  $V$ :

$$U_E = \frac{\epsilon_0 \epsilon b(l-c)}{4d} V^2$$

where  $d$  is the dielectric thickness and  $\epsilon$  is the relative dielectric constant.

The total energy is minimized by varying the position of the contact point,  $c$ , while holding all other dimensions constant. This yields a simple expression for the pull-in voltage,  $V_{PI}$ , as a function of the material properties, unit cell dimensions, and the constant external force,  $F$ :


$$V_{PI} = \left( \frac{6d}{kl\epsilon_0\epsilon} \right)^{1/2} \frac{F}{b}$$

The strain is the ratio of the displacement,  $\delta$ , to the total thickness of the open cell,  $\delta + 2t$ . Strain is a function of the external force because the deflection is a function of the force. Therefore,

$$Strain = \frac{\delta}{\delta + 2t} = \frac{F}{F + 2kbt}$$

In the absence of a voltage,  $c = l$  and the unit cell is a spring with spring constant  $kb$ . An external force,  $F$ , will extend the actuator a distance  $\delta = F/kb$ . When a small voltage is applied, almost nothing happens. However, when the pull-in voltage is reached, the unit cell suddenly collapses. Figure 8.4 shows excellent agreement between the analytical and finite element models. The pressure ( $F/bl$ ) as a function of the displacement shows the linear, springlike behavior that occurs at large force or deflection and the pull-in at smaller deflections (for a given voltage). The intersection of the lines in the analytical model, and the minimum in the FEA curve, shows the maximum force that can be actuated at the indicated voltage.

This model is an extremely useful tool in adapting the actuator design to a specific application, such as biomimetic robots. For example, a desired actuator might be able to apply a force of 5–10 N ( $\sim 0.5$ –1 kg), have a strain  $\sim 5\%$ , and a total displacement  $\geq 2$  mm. The model shows that this is most easily accomplished by an appropriate choice of the unit cell length, the polymer thickness, and the dielectric thickness and constant. The unit cell length must be small, but assembly will be increasingly difficult for lengths under  $\sim 1$  mm. Therefore, this represents a lower practical limit using the current process. The dielectric must be thick enough to avoid breakdown, but thin enough to prevent all the




**Figure 8.4** A comparison of the finite element and analytical models. The analytical model yields two lines. The straight, positively sloped line is the purely mechanical response of the unit cell, with no applied voltage. The negatively sloping curve is the combined electrostatic/mechanical response at a nonzero voltage. Pull-in occurs at the intersection of these lines. The analytical model is in excellent agreement with the finite element results, shown as black dots.

field from dropping across the dielectric rather than the gap. Figure 8.5 shows the variation of force and strain as a function of Kapton [7] thickness for two different dielectrics and a bias of 200 V. The inset table lists design parameters that meet the desired specifications. Dielectrics such as polyimide or polymethylmethacrylate (PMMA) have  $\epsilon \sim 3.5$ , while polyvinylidene fluoride (PVDF) has  $\epsilon \sim 10$ . This actuator would consume only  $\sim 150$  mW of power while actuating continuously at 5 Hz operation, and power consumption would be even less at lower frequencies. If the actuator were held in its actuated (clamped) position, the only power consumption would be due to leakage current through the dielectric. Measured leakage currents are typically a few picoamps to  $\sim 10$  nA for a  $1 \text{ cm}^2$  unit cell area. Thus, the actuator described here would use, at most,  $\sim 1/3$  mW while holding in its actuated position.

## FABRICATION

The polyMEMS actuator is fabricated using a simple process adapted from flexible printed circuit manufacturing. As illustrated in figure 8.6, fabrication begins with a sheet of polymer such as Kapton [7] with a  $\sim 100 \text{ \AA}$ -thick aluminum film on one surface. Polyimide was chosen as the base polymer because of its excellent thermal and mechanical stability and because it is a widely used flexible circuit material. However, other polymer substrates can be substituted when desired, because the mechanical and thermal requirements often do not reach extreme



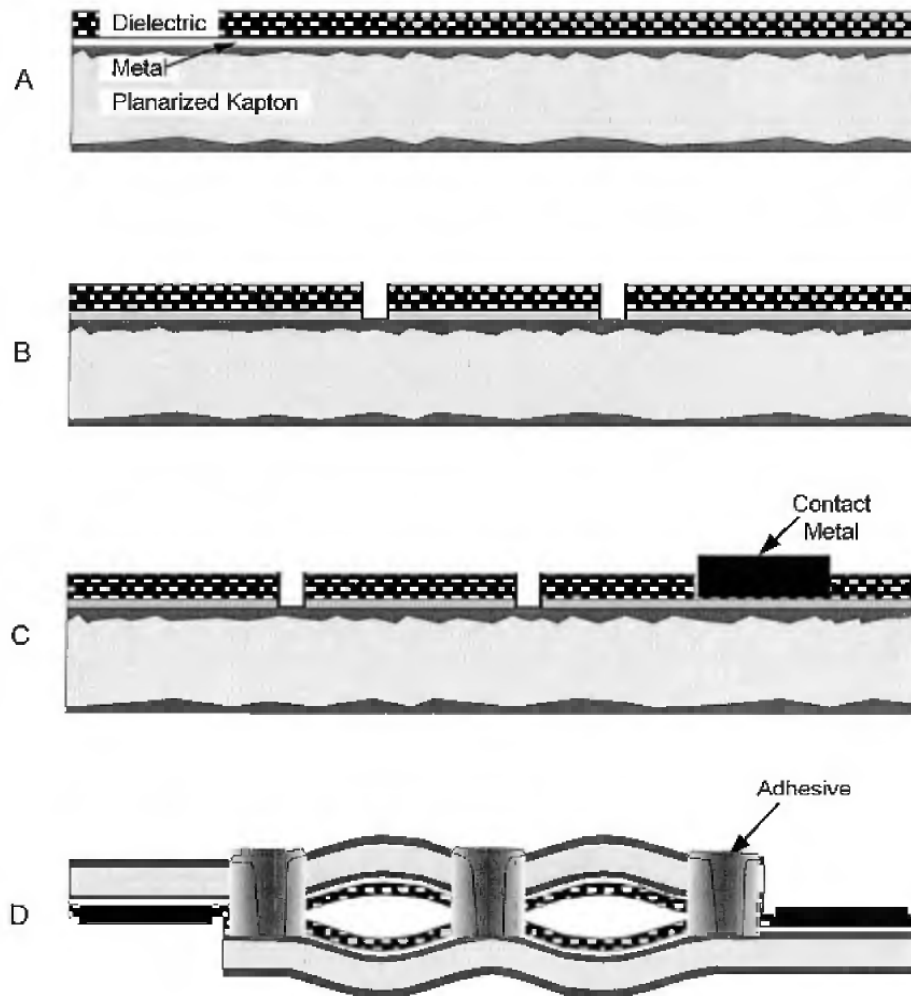
**Figure 8.5** Maximum load and strain as a function of the Kapton thickness and the dielectric constant. The strain is the value at the maximum load for that thickness. The target force and strain can be achieved with a 75  $\mu\text{m}$  thick Kapton substrate and dielectric constant of 10, or a  $\sim 100 \mu\text{m}$  thickness and dielectric constant of 3.5. This calculation assumed a 1 mm unit cell length, 10 mm width, 10 cells per layer (1  $\text{cm}^2$  area),  $0.45 \mu\text{m}$  dielectric thickness, and 200 V. This actuator would need 165 layers of unit cells to get 2 mm deflection at 1 kg ( $\sim 10 \text{ N}$ ) load. The fully collapsed length would be 2.5 cm.

values. For example, because they absorb moisture, polyimides may not be suitable for robots that work in or near water, and a polyester may be more suitable. The aluminum thickness is dictated by a self-healing effect that is discussed below.

The aluminum is coated with a dielectric film that prevents shorting between adjacent metal electrode films. The dielectric must be flexible, have a high breakdown strength, and preferably have a high dielectric constant. We have used polyimides ( $\epsilon \sim 3.5$ ) most frequently but are experimenting with materials with dielectric constants up to  $\sim 10$ . The dielectric thickness typically has ranged from  $\sim 0.3 \mu\text{m}$  to  $\sim 1 \mu\text{m}$ , with the thicker layers being preferred. A thicker dielectric requires a higher driving voltage for a comparable force, increasing as  $d^{1/2}$ . However, the maximum field within the dielectric decreases as  $1/d^{1/2}$ . Operation at this reduced field is further from the dielectric breakdown point and makes the dielectric less susceptible to charge trapping.

Next, the electrode areas, bonding areas, and actuator boundaries are defined by photolithography and oxygen plasma etching of the dielectric. Aluminum is removed from some areas. In other areas, metal is added over the aluminum to make more robust electrical contacts.

Finally, two patterned sheets are placed face to face and laminated together using one of several procedures. In figure 8.6, slits are laser cut into one sheet, aligned with the pattern on the other sheet, and adhesive is applied through the slits. It is critical that bonding takes place in




**Figure 8.6** The process flow for one layer of the PolyMEMS Actuator. (A) The polymer substrate is coated with a thin aluminum film and a dielectric film. (B) Unit cells are defined by photolithography and etching. (C) Contact pads are strengthened by adding more metal. (D) The lamination step can be accomplished in various ways. In the method shown, slits are cut in one sheet. That sheet is aligned to a second sheet, and the two are bonded together.

a way that does not increase the distance between electrodes, as this will reduce the available force. In addition, care must be taken to avoid deformation of the sheets near the bond, which also will degrade the actuator performance.

### ACTUATION RESULTS

The focus of our development to date has been on building and testing unit cells and small arrays. Unit cell lengths have ranged from 3 mm to 12.5 mm. The eventual cell length will be  $\sim 1$  mm or less. As can be seen from the model, larger unit cells have lower forces for a given voltage. Quantitative measurements of force and pull-in voltage have




**Figure 8.7** Photographs of a unit cell opened by application of an external load (a) and closed by applying a voltage (b).

been made on single layers of 1–3 unit cells. The actuators have been powered with both DC and square-wave AC inputs. Trapped charges can build up in the dielectric [8], and this is more prone to happen under DC conditions. However, equal amounts of mechanical work have been achieved under both DC and AC driving. Figure 8.7a is a photograph of a unit cell pulled open by an external force. Figure 8.7b shows the same cell closed after applying a voltage.

Figure 8.8 shows the measured pull-in voltage as a function of the external load for two devices. The upper data were taken from an early actuator that had some distortion due to assembly. Measured pull-in voltages are significantly above the model predictions. At the edges of these unit cells, the sheets were not parallel to each other as in the ideal case (figure 8.3). Instead, they diverged away from each other. The higher pull-in voltage was required to overcome this distortion. An improved assembly technique resulted in the lower data points that are much closer to the predictions. There is still a small amount of distortion visible at the bond points, and this distortion is believed to be the cause of the discrepancy between the prediction and the measurements.

Reliability of the PMA has yet to be studied. However, a number of features can be built into the actuator that point to the potential for high reliability. These include fault tolerance, thicker dielectrics, self-healing, and the ability to make an integral, hermetic package.

Multilayered actuators have been built with the elements electrically connected in parallel. These devices continued to function even if an electrical discharge caused one cell to be nonfunctional. Without this kind of fault tolerance, failure of a single cell would be catastrophic. Muscles have a similar characteristic: If one fiber is strained, there



**Figure 8.8** Pull-in voltage as a function of the applied load for two actuators. Much closer agreement with the theoretical predictions was achieved after making improvements in the assembly techniques.

are many others that continue working. Pain and weakness may be present, but total failure does not occur.

Most of the failures that have been observed involve high electric fields, which can result in trapped charges or breakdown in the thin dielectric films. These problems can be reduced substantially using thicker dielectric films. As shown above, a thicker dielectric raises the required pull-in voltage (which imposes a very minimal penalty on the system drive electronics as long as voltages remain below  $\sim 250$  V), but the maximum field at pull-in is reduced as  $d^{-1/2}$ . Initial results appear to verify this conclusion.

When breakdown does occur—for example, at a pinhole or defect—a “self-healing” phenomenon can eliminate the short. This effect has been exploited in commercial thin-film capacitors. When breakdown occurs, the discharge can completely vaporize and/or oxidize the aluminum electrode. The conductive path is burned away and the short is “healed.” Only a very small area around the defect is affected, so the capacitance change is negligible. For self-healing to occur, the metal must have a low melting point and it must be thin enough to prevent rapid dissipation of the heat from the discharge. Most metals have melting points too high for complete vaporization; aluminum is the preferred metal. The heat dissipation requirement limits the film thickness to less than  $\sim 100$  Å, but this is neither difficult nor degrading to the device performance. Small current spikes associated with self-healing are typically seen the first few times a new device is used, and disappear after that. These observations suggest that polyMEMS actuators can be made with good reliability.

A final reliability issue is moisture absorption in the plastic, which can cause swelling and film delamination. This can be especially problematic in polyimides. Barrier films and hermetic packaging should be able to control this in most terrestrial or aerial applications. Special packaging may be required for robots that are continually exposed to water or high moisture environments. A bellowslike integral hermetic package can be provided using the same polymer sheets and adhesive used in the basic actuator structure. Such a package would provide a controlled, low-humidity and dust-free atmosphere in the actuator regions where high electric fields are present.

## SUMMARY

The polyMEMS actuator is a 3-D array of micro- to milliscale electrostatic actuator unit cells that produces macroscale forces and displacements. The actuators are lightweight, consume small amounts of power during actuation, and consume almost no power when holding in the actuated position. Many structural and conceptual features of the actuator are analogous to features in muscle tissue, and these result in similar performance features such as a high strength-to-weight ratio and fault tolerance. An accurate analytical model has been derived that aids significantly in designing for specific specifications and features. In particular, it shows that these actuators can meet the requirements for many applications in biomimetic robots. Fabrication is fairly simple and fast. Tests on unit cells and small arrays show successful operation, with performance approaching the theoretical prediction.

## ACKNOWLEDGMENTS

The authors would like to acknowledge the excellent technical work of Terry Stark, Jerry Bolser, Son Thai Lu, Khanh Nguyen, Steve Swanson, Tom Rezachek, and Coline Amaury on this program. Dr. Cleopatra Cabuz, Eugen Cabuz, and Dr. Bill Herb have provided many useful comments and insights. This program was funded by DARPA under contract F30602-98-C-0217 and managed by Dr. William Tang and Mr. Duane Gilmour.

## REFERENCES

- [1] Minami, K., Kawamura, S., and Esashi, M. (1993). *J. Microelectromechanical Syst.* 2: 121.
- [2] Yamaguchi, M., Kawamura, S., Minami, K., and Esashi, M. (1993). *J. Micromech. Microeng.* 3: 90.
- [3] Jacobson, J. D., Goodwin-Johansson, S. H., Bobbio, S. M., Bartlett, C. A., and Yadon, L. N. (1995). *J. Microelectromechanical Syst.* 4: 139.

- [4] Cabuz, C., Cabuz, E. I., Herb, W. R., Rolfer, T., and Zook, D. (1999). *Transducers '99: Proceedings of the Tenth International Conference on Solid-State Sensors and Actuators*, Sendai, Japan, 7–10 June. p. 1890.
- [5] Alexander, R. McNeill. (1992). *Exploring Biomechanics: Animals in Motion*. New York: Scientific American Library, pp. 25–32.
- [6] Roark, R. J., and Young, W. C. (1975). *Formulas for Stress and Strain*, 5th ed. New York: McGraw-Hill, case 1b, table 3, p. 96.
- [7] Kapton is a registered trademark for polyimides manufactured by E. I. du Pont de Nemours and Company.
- [8] Cabuz, C., Cabuz, E. I., Ohnstein, T. R., Neus, J., and Maboudian, R. (1998). *Hilton Head '98: Proceedings of the 1998 Solid-State Sensor and Actuator Workshop*, Hilton Head Island, S.C. 8–11 June. p. 296.



**This page intentionally left blank**

# 9

## A Piezoelectrically Actuated Mesoscale Robot Quadruped

Michael Goldfarb, Michael Gogola, Gregory Fischer, Nikola Celanovic, Ivan Celanovic, and Ephraim Garcia

A mesoscale device is generally considered to be one that contains length scales on the order of millimeters to centimeters. Because volume scales with the cube of a geometric scaling, mesoscale machines contain significantly less volume than their conventional-scale counterparts. The disproportionate decrease in volume in a mesoscale machine results in two significant challenges. The first is limited design space, which in turn suggests the use of simple and minimalist designs. The second is that, because electrochemical battery energy densities are essentially constant, the amount of available energy will decrease with the cube of size. A mesoscale robot that is geometrically similar and smaller than a conventional-scale robot therefore has significantly less available energy to travel a given distance. For example, a battery-powered mesoscale robot that is designed as a one-tenth-scale version of a conventional-scale robot would ideally travel a distance ten times shorter than its larger counterpart. Assuming that the locomotive efficiency remains invariant, however, the smaller version will travel a distance one thousand times shorter. The smaller robot will therefore have a range that is one hundred times smaller than the larger, relative to its own dimensions. The development of mesoscale robots is therefore aided significantly by the implementation of simple design approaches that result in efficient locomotion. This chapter describes such a design approach. The approach, called *elastodynamic locomotion*, utilizes a lightly damped skeletal structure that generates efficient locomotion when vibrationally excited near a skeletal resonance.

### BACKGROUND

A considerable amount of work has been conducted in the design of autonomous conventional-scale robot quadrupeds. These designs generally incorporate independent control of the motion of each leg and multiple sensors for gathering information regarding leg motion and terrain. These features, along with onboard computational ability, afford these machines the ability to traverse various terrains and

obstacles. Such machines are generally not appropriate for implementation at the mesoscale because the required complexity (e.g., number of actuators and sensors) is mismatched to the previously mentioned volume constraints of a mesoscale device. Further, the energetically non-conservative nature of closed-loop control of leg motion is additionally mismatched to the previously described need for efficient locomotion.

Though the authors are not aware of other previously developed piezoelectrically actuated quadrupeds, other researchers have incorporated piezoelectric actuation into the development of small-scale mobile robots. De Ambroggi and colleagues have developed piezoelectrically actuated robots that have a footprint of approximately 2 cm by 2 cm [1]. These robots are actuated by two legs—each consisting of two serially mounted piezoelectric bimorphs—with a third passive leg used for stability. Because the piezoelectric actuators are utilized in a direct-drive sense (i.e., no displacement amplification), the robots are generally confined to operate on smooth surfaces. Martel and colleagues have reported on the development of a twelve degrees-of-freedom piezoelectrically actuated three-legged robot with a footprint of approximately 2 cm by 2 cm [2]. This robot also incorporates no displacement amplification of the piezoelectric actuators, and is therefore also confined to operate on smooth surfaces.

Kawahara and colleagues have developed a piezoelectrically actuated in-pipe wireless robot [3]. The robot is cylindrical in shape with a diameter of approximately 1.5 cm and a length of approximately 4 cm. This robot incorporates multilayered piezoelectric bimorphs as actuators and incorporates an on-board energy converter to convert microwave power traveling in the pipe into DC electric power. Other than the work presented herein, this is the first wireless (i.e., not power-tethered) piezoelectrically actuated robot of which the authors are aware.

Several researchers, such as Higuchi and colleagues [4] and Breguet and Renaud [5], have additionally developed stepping-type or inertial-drive-type piezoelectric motors, which can be considered similar to small-scale mobile robots. These types of motors are generally not designed to be self-powered, and because they incorporate no displacement amplification of the piezoelectric actuators, they are confined to operate on smooth surfaces.


The robot design described in this chapter is a piezoelectrically actuated quadruped that was designed to be capable of self-powered operation and to accommodate reasonably rough surfaces. The minimalist philosophy upon which the robot design is based was motivated in part by the work of Yasuda and colleagues, who have developed a microrobot that is actuated by a vibration energy field [6]. Specifically, they developed a passive silicon micromachined device that measures

approximately 1.5 mm by 0.7 mm. The silicon robot sits on a vibration table, which acts as the energy and control source. Vibrationally exciting the robot through the table excites each leg, which in turn generates contact between the robot and the table, which pushes the robot along. Each limb of the device is designed to exhibit resonance at a different frequency from the other limbs. The motion of the passive silicon device can therefore be controlled by superimposing sinusoidal excitation of the respective leg resonant frequencies onto the robot, and thus differentially controlling the amplitudes of leg motion. The work presented in this chapter is similar to the approach of Yasuda and colleagues [6], but instead utilizes two-degrees-of-freedom leg motion and incorporates the energy and actuation source on-board the robot.

## ELASTODYNAMIC DESIGN APPROACH

The design described incorporates a lightly damped skeletal structure that generates locomotion when vibrationally excited at a skeletal resonance by piezoelectric actuators. This primarily elastic approach to locomotion provides two significant advantages. First, because the motion is elastic, it is energetically conservative (i.e., the acceleration and deceleration of each limb requires no net power). The combination of a piezoelectric ceramic (which exhibits low energy dissipation) and the use of resonance (which is energetically conservative) results in a highly energetically conservative device. Second, utilizing a lightly damped structure enables control of limb amplitude by frequency shifting (as was demonstrated by Yasuda and colleagues [6]), which minimizes the number of independent actuators required for control of the robot. Specifically, by tuning the resonant frequencies of the legs on each side of the skeleton to slightly different frequencies, the leg amplitudes can be differentially controlled by slightly shifting the frequency of excitation. This concept is illustrated in the frequency response diagram shown in figure 9.1. A left or right turn can be achieved by shifting the excitation frequency to match the resonant frequency to the limbs on the side of the vehicle that is opposite the desired turning direction. Straight motion is achieved by exciting the system at a frequency in between the left and right side resonances, whereby both sets of legs (i.e., left and right sides) operate slightly out of resonance, but with an equal degree of amplification. Resonant (or near-resonant) operation therefore provides a mechanism for dynamically controlling a robot, which in turn minimizes the requirement for independently actuated degrees of freedom.

Finally, in addition to providing for efficient locomotion and a minimum actuator configuration, utilization of an elastic mechanism (i.e., a flexure-based metallic skeleton) eliminates the need for roller or journal



**Figure 9.1** Frequency response characteristics of quadruped legs for frequency-based directional control.

bearings and therefore results in a more environmentally robust system, providing immunity from environmental contaminants such as sand and dust.

## PIEZOELECTRIC CERAMIC ACTUATION

The use of piezoelectric (PZT) actuators provides several beneficial features. First, PZT is an elastic actuator, and therefore well suited for the elastodynamic approach. Second, PZT is a highly energetically conservative means of electrical-to-mechanical energy transduction. Specifically, the only significant mechanisms of dissipation are an electrical hysteresis, which dissipates approximately 10% of the input energy per cycle of actuation, and mechanical damping in the ceramic, which typically exhibits a damping ratio of approximately 0.5 [7]. Third, PZT has among the highest volume-specific power of any electrical-to-mechanical transduction mechanism [8]. Finally, PZT also has favorable force-scaling characteristics relative to other common transduction mechanisms [9].

Though PZT provides significant benefits as an actuator, it also has some significant drawbacks relative to mesoscale robotic applications. Perhaps the most limiting of these is that it typically provides relatively small displacements. Polycrystalline piezoelectric ceramic typically provides material strains on the order of 0.1%. In a direct-drive sense, a 2 cm-long PZT stack will provide only 20 microns of displacement. If a mesoscale robot is to traverse reasonably rough surfaces, the small displacements provided by the PZT ceramic must be amplified by significant mechanical amplification mechanisms. From a power perspective,


though, PZT actuation is volumetrically power dense—the power is provided in a high force and low displacement fashion, relative to that required by a mesoscale quadruped. Optimal use of the power provided by a PZT actuator therefore requires a large mechanical transmission to convert the high-force, low-displacement work into relatively low-force, high-displacement work. The amplification (or transmission) mechanisms utilized to match the actuator with the work output of the robot should additionally be energetically conservative in order to maintain the energetic benefits of the PZT actuator.

The primary mechanisms of PZT displacement amplification are structural amplification, kinematic amplification, and dynamic amplification. Specifically, structural amplification is that which results from utilizing a configuration of actuators that amplifies the strain output of the PZT ceramic and, in the design described herein, takes the form of a unimorph actuator. Kinematic amplification is achieved through use of compliant mechanism-based linkages. Finally, dynamic amplification utilizes the frequency response characteristics of the elastic transmission, and specifically operates the device at or near the primary resonance of the transmission. The authors incorporate all of these mechanisms into the design of the elastodynamic robot in order to maximize the displacement output of the PZT.

Another significant drawback of piezoelectric actuation is the relatively high voltage required for actuation. Just as the mechanical (output) power of the PZT actuator is in a different regime from that required by a mesoscale mobile robot, so is the electrical (input) power in a different regime from that provided by a typical electrochemical battery. Just as a large and efficient mechanical transmission is desired to convert the PZT output power to a useful regime, so is a large and efficient electrical transmission desired to convert the power from a battery into that required by the PZT. This issue is addressed in the Power Electronics section of this chapter.

## PROTOTYPE DESIGN

The elastodynamic robot quadruped is shown in figure 9.2. The robot incorporates two piezoelectric unimorph actuators (manufactured by Face, Incorporated) that drive four legs in a trot-type gait. Each leg contains two degrees of freedom and is configured as an inertially decoupled five-bar linkage. The two actuators are excited with a sinusoidal voltage of the same amplitude and frequency, but shifted in phase by approximately  $90^\circ$ . The robot measures approximately 9 cm long, 6.5 cm wide, and 5 cm tall, and weighs 51 grams without the on-board power electronics or battery. A description of the design morphology follows.




**Figure 9.2** Piezoelectrically actuated elastodynamic quadruped robot.

### **Drive System**

Though robot quadruped locomotion can be achieved with curvilinear foot motion (i.e., each foot moves forward and backward along the same path relative to the robot reference frame), use of a closed path (e.g., an elliptical path) generally provides greater mobility (e.g., ability to traverse rough surfaces) and greater efficiency (i.e., less frictional energy loss due to sliding). Due to these advantages, the objective of the robot quadruped design was to obtain closed-path foot motions for each leg, subject to the minimalist elastodynamic design philosophy previously described.

The minimum number of actuators for any active device is one. Obtaining a closed-path trajectory in the case of a nonelastic system that can utilize continuous rotation—such as with an electric motor—is trivial. The closed-path motion can either be obtained directly by using the center of rotation as the center of a circular path, or indirectly by incorporating a crank-rocker-type four-bar linkage. In an elastic system, however, obtaining a closed path requires, at minimum, the use of two degrees of freedom, which may or may not be coupled. More specifically, generation of a repeatable closed path in an elastic system requires at least two degrees of freedom that are excited at the same frequency and separated by some degree of phase lag. Obtaining a closed path in an elastic system with a single actuator therefore requires that the mechanism incorporate mechanical phase lag between two degrees of freedom. The authors initially attempted the development of a robot quadruped with a single actuator that incorporated passive mechanical phase lag to produce elliptical foot motion. Though theoretically such a system is possible, obtaining consistent foot trajectories proved difficult, due primarily to the fact that the phase in a



**Figure 9.3** Diagram of five-bar linkage-based two-actuator robot drive system.

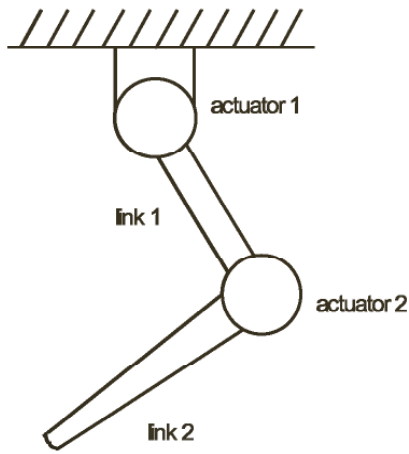
lightly damped system is highly sensitive to any variation in the input frequency. Such behavior is a particular impediment when purposefully utilizing variation in input frequency for directional control, as is the case with the elastodynamic approach.

Upon determining that a single-actuator elastodynamic quadruped would not provide sufficiently robust performance, the authors developed an elastodynamic quadruped design based on two actuators. The phase lag required for closed-path foot motion is therefore provided electronically rather than mechanically. The resulting two-actuator robot drive system, which is diagrammed in figure 9.3, incorporates an inertially decoupled five-bar linkage, as explained in the following section.

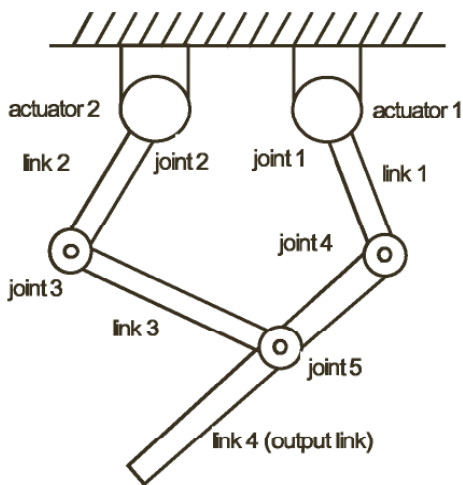
In designing a two-degrees-of-freedom revolute mechanism, one can utilize either a serial or a parallel configuration. In a serial configuration, each joint is actuated relative to the previous joint, as illustrated in figure 9.4. In a parallel configuration, each joint is actuated relative to a reference frame. The five-bar linkage illustrated in figure 9.5 is a common implementation of a two-degrees-of-freedom parallel revolute mechanism.

The legs of the elastodynamic robot quadruped incorporate a parallel configuration based on a five-bar linkage rather than a serial configuration for two primary reasons. First, the base of both actuators can be attached to the same reference frame, which in turn enables all four legs of the quadruped to be actuated by the same pair of unimorph actuators. Specifically, by crossing the connecting rods shown in figure 9.3, the linkage will be driven  $180^\circ$  out of phase with respect to the same set of actuators. A trot-type gait requires that diagonal pairs of legs





**Figure 9.4** Serial configuration of a two degrees-of-freedom revolute joint mechanism.



**Figure 9.5** Parallel configuration of a two degrees-of-freedom revolute joint mechanism (five-bar linkage).

operate  $180^\circ$  out of phase. With one connecting rod attached to each corner of each unimorph, all four legs can be driven from the same two actuators. The second significant benefit of a five-bar linkage configuration is the ability to inertially decouple the two degrees of freedom, which is of great importance when operating the mechanism near resonant frequencies.

### Leg Design

The configuration of five-bar linkage shown in figure 9.3 is a special configuration of a five-bar mechanism where the reference axis of rotation of links 1 and 2 are coaxially arranged. In the initial implementation of this leg configuration, elliptical foot motion was obtained quasi statically by driving the two actuators at the same amplitude and fre-

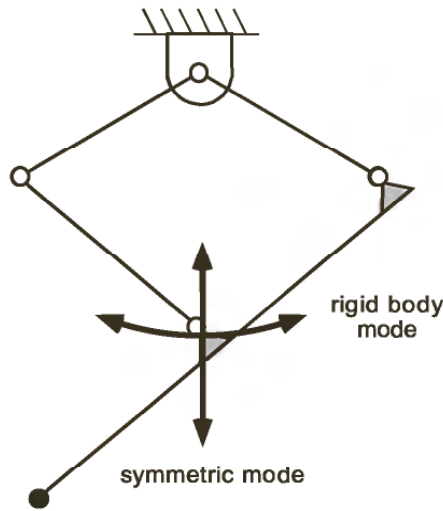


Figure 9.6 Illustration of vibrational modes in inertially coupled five-bar mechanism.

quency, but  $90^\circ$  out of phase. As the excitation frequency approached the dynamic regime of the mechanism, however, the linkage exhibited undesirable dynamically coupled behavior. Specifically, the mechanism demonstrated one of two fundamental modes of vibration. One of these was a symmetric mode in which the parallelogram structure would expand and contract, and one was a “rigid-body” type mode in which the entire parallelogram would maintain approximately the same geometry while swinging around the grounded revolute joint. These modes are illustrated in figure 9.6. Ideally, these modes would be tuned to the same frequency and occur simultaneously, but an effort to match the frequencies was unsuccessful, because even slight differences in the resonant frequencies of the two modes would introduce the undesired behavior.

This problem was addressed by utilizing the fact that a five-bar linkage can be inertially decoupled, so that the two input degrees of freedom will operate essentially independently of each other (i.e., the mechanism will behave dynamically essentially the same as it does quasi statically). Specifically, the inertial relationship between the torques about the ground reference joint of the two input links to the motion of those two links can be written as


$$\begin{bmatrix} \tau_1 \\ \tau_2 \end{bmatrix} = \begin{bmatrix} H_{11} & H_{12} \\ H_{12} & H_{22} \end{bmatrix} \begin{bmatrix} \ddot{\theta}_1 \\ \ddot{\theta}_2 \end{bmatrix} \quad (9.1)$$

The components of the inertia tensor,  $H$ , are given by

$$H_{11} = I_1 + m_1 g_1^2 + I_3 + m_3 g_3^2 + m_4 l_1^2 \quad (9.2)$$

$$H_{22} = I_4 + m_4 g_4^2 + I_2 + m_2 g_2^2 + m_3 l_2^2 \quad (9.3)$$

$$H_{12} = (m_3 l_2 g_3 - m_4 l_1 g_4) \cos(\theta_2 - \theta_1) \quad (9.4)$$



**Figure 9.7** Definition of geometric parameters of the five-bar mechanism.

where geometric parameters are as defined in figure 9.7 and the inertial parameters are the masses and rotational inertias of each respective link. Inertial decoupling between the two links requires that the inertia tensor be diagonal, which requires that the off-diagonal element,  $H_{12}$ , be zero, which requires

$$(m_3 l_2 g_3 - m_4 l_1 g_4) = 0 \quad (9.5)$$

or

$$\frac{m_4}{m_3} \cdot \frac{g_4}{g_3} = \frac{l_2}{l_1} \quad (9.6)$$

Inertial decoupling is therefore achieved by selecting geometrical and inertial properties of the links to satisfy equation 9.6. This procedure is commonly utilized in the design of parallel robot manipulators to simplify manipulator control [10]. Implementation of an inertially decoupled five-bar linkage as the leg structure of the quadruped eliminated the tendency of the leg to “choose” one of the two modes illustrated in figure 9.6. Instead, the inertially decoupled legs exhibited robust and repeatable elliptical foot trajectories when excited at or near resonant frequencies.

## ROBOT PERFORMANCE

Locomotive performance of the (power-tethered) robot was evaluated while carrying a 50 gram payload on a level surface. The payload was selected to be of approximately the same mass as the power electronics that are described in a section titled “Power Electronics.” The two uni-

**Table 9.1**


Operating Frequency (Hz)	Walking Speed (mm/sec)	Average Power (mW)
1	8	44
15	150	580
32	300	1400

morph actuators were driven with a sinusoidal voltage of  $\pm 240$  volts amplitude and at the same frequency, with one actuator lagging the other by a phase of  $90^\circ$ . Table 9.1 indicates the average speed of the robot quadruped operating at several different excitation frequencies, in addition to the total average power required by the actuator pair. As shown in the table, the maximum speed of the robot was 300 mm/sec at an excitation frequency of 32 Hz. The instantaneous power required by the actuators at each of the frequencies listed in table 9.1 is shown for several cycles of gait in the graphs of figures 9.8, 9.9, and 9.10. The fact that the average required power increases nonlinearly with speed is indicative of the increased slipping that occurs at higher frequencies between the feet and the ground surface. The authors have not yet addressed this problem, though utilizing footpads with a high coefficient of friction would most likely decrease the power required at the higher speeds.

As previously described, a significant component of the elastodynamic approach to mesoscale robot design is the dynamic control of the device (i.e., frequency-based turning capability). Dynamic control of the robot was verified by demonstrating repeatable control of both right and left turns. Specifically, within the frequency region of 14–24 Hz, right turning was repeatedly executed in the 14–18 Hz range, with a minimum turning radius of approximately 5 cm; straight walking was achieved in the 18–19 Hz range; and, finally, left turns repeatedly occurred in the 19–24 Hz range, also with a minimum turning radius of approximately 5 cm. With the removal of the 50 gram payload, the operating ranges shifted slightly to 15–17 Hz for right turns, 17–18 Hz for straight-line motion, and 18–21 Hz for left turns.


## POWER ELECTRONICS

As mentioned in the discussion on piezoelectric actuation, the use of PZT to actuate a small-scale mobile robot requires power electronics that convert the low-voltage and high-current power of an electrochemical battery to the relatively high-voltage and low-current power required by the PZT actuator. Specifically, the full range of output



**Figure 9.8** Instantaneous and average power for each PZT actuator at an excitation frequency of 1 Hz.


for the actuators on the elastodynamic quadruped robot require  $\pm 240$  volts, while a typical lithium-type battery supplies approximately 3 volts. It should be noted that Kawahara and colleagues [3] cleverly avoided the requirement of high voltages by developing multilayered PZT bimorph actuators, though the development of such actuators is nontrivial. In addition to requiring a large electrical transmission, the power electronics must also generate a sinusoidal output, because any other type of AC output would contain multiple frequencies and thus would excite multiple modes in the robot. This would significantly impair the ability to perform frequency-based directional control and would additionally sacrifice locomotive performance. Further, the elas-



**Figure 9.9** Instantaneous and average power for each PZT actuator at an excitation frequency of 15 Hz.

todynamic quadruped required two output channels, one of which is shifted in phase by approximately  $90^\circ$  from the other.

Figure 9.11 shows the power electronics circuit that was developed to drive the elastodynamic robot. This circuit incorporates a surface-mount microprocessor for pulse-width-modulated control of an H-bridge output stage, which efficiently generates a sinusoidal output. As required by the robot, the circuit generates two outputs, both sinusoids of the same amplitude and frequency, but different only in phase. The output amplitude can exceed the required 240 volts, depending on the required power, which is limited to approximately 1.25 watts per channel. The circuit is a one-layer circuit board that measures approximately 6 cm




**Figure 9.10** Instantaneous and average power for each PZT actuator at an excitation frequency of 32 Hz.


by 9 cm and weighs approximately 37 grams. Power is provided by a lithium-manganese dioxide (NEDA-5024LC) battery that weighs approximately 16 grams and carries 700 milliamp hours of energy at 3 volts. With this combination of power electronics and battery, the total weight imposed on the robot is approximately 53 grams.

### SELF-POWERED OPERATION

The previously described power electronics were mounted on the robot quadruped, as shown in figure 9.12, and self-powered crawling was demonstrated. The battery, which is not visible in the figure, is



**Figure 9.11** Power electronics circuit.



**Figure 9.12** Power electronics mounted on robot quadruped for self-powered operation.

mounted below the lower actuator in the center of the robot. The electronics presently do not incorporate a wireless link by which the excitation frequency can be controlled. In the current configuration, the frequency of excitation (and thus the direction of the robot) is set manually by adjusting a potentiometer on the board. A test was performed by adjusting the excitation frequency to approximately 15 Hz, which resulted in a right turn with a radius of approximately 25 cm. In this configuration, the robot operated for a duration of approximately 30 minutes, which, at an average speed of 150 mm/sec, corresponds to a distance of approximately 270 meters.



## DISCUSSION AND CONCLUSION

An elastodynamic design paradigm has been proposed, and the design of a mesoscale robot quadruped based on this paradigm presented. The authors have demonstrated frequency-based control of the robot, and have demonstrated self-powered operation. Future work includes the implementation of a wireless control link to remotely control the frequency of excitation, and thus the direction of the robot. Additionally, the authors wish to investigate and implement the effects of controlling the phase imposed between the two actuators. Specifically, the shape of the foot trajectory can be controlled by modulating the amount of phase between the two actuators, which can be utilized to decrease forward speed and increase foot height for rough surfaces or, conversely, to increase forward speed and decrease foot height for smooth surfaces. Finally, if the phase between the actuators is reversed completely, the direction of motion of the robot will also reverse, providing the complete range of forward, left, right, and backward operation of the quadruped.

## ACKNOWLEDGMENTS

Support for the work presented in this chapter was provided by DARPA DSO (contract DABT6397C0052).

## REFERENCES

- [1] De Ambroggi, F., Fortuna, L., and Muscato, G. (1997). PLIF: Piezo light intelligent flea—New microrobots controlled by self-learning techniques. In *Proceedings of the IEEE International Conference on Robotics and Automation*, vol. 2, 20–25 April. pp. 1767–1772.
- [2] Martel, S., Madden, P., Sosnowski, L., Hunter, I., and Lafontaine, S. (1999). Nano-Walker: A fully autonomous highly integrated miniature robot for nano-scale measurements. *Proceedings of SPIE* (15–16 June): 111–122.
- [3] Kawahara, N., Shibata, T., and Sasaya, T. (1999). In-pipe wireless micro robot. *Proceedings of SPIE* (21–22 September): 166–171.
- [4] Higuchi, T., Yamagata, Y., Furutani, K., and Kudoh, K. (1990). Precise positioning mechanism utilizing rapid deformations of piezoelectric elements. In *Proceedings of the IEEE Micro Electro Mechanical Systems: An Investigation of Micro Structures, Sensors, Actuators, Machines, and Robots*, 11–14 February. pp. 222–226.
- [5] Breguet, J. M., and Renaud, P. (1996). Four-degrees-of-freedom microrobot with nanometer resolution. *Robotica* vol. 14, pt. 2, (March–April): 199–203.
- [6] Yasuda, T., Shimoyama, I., and Miura, H. (1994). Microrobot actuated by a vibration energy field. *Sensors Actuators, [A] Phys.* vol. 43, nos. 1–3 (May): 366–370.
- [7] Goldfarb, M., and Celanovic, N. (1997). A lumped-parameter electromechanical model for describing the nonlinear behavior of piezoelectric actuators. *ASME J. Dynamic Syst. Meas. Control* vol. 119, no. 3: 478–485.

- [8] Kuribayashi, K. (1993). Criteria for the evaluation of new actuators as energy converters. *Adv. Robotics* 289–307.
- [9] Trimmer, W., and Jebens, R. (1989). Actuators for Micro Robots. In *Proceedings of the IEEE International Conference on Robotics and Automation*. pp. 1547–1552.
- [10] Asada, H., and Youcef-Toumi, K. (1987). *Direct-Drive Robots*. Cambridge: MIT Press, pp. 48–52, 129–132.

**This page intentionally left blank**

# *Biomimetic Control Architectures*

**This page intentionally left blank**

# 10 Optimization of Swimming Locomotion by Genetic Algorithm

David Barrett

MIT's RoboTuna is a prototype of an autonomous undersea vehicle that moves through the water by swimming (much like a biological fish). The seven characteristic parameters, defining the traveling body wave used by the RoboTuna for locomotion, create an enormous experimental space within which to search for optimum performance. For the sake of argument, let us say each variable's possible range will be divided into 10 rough quanta. For the seven-dimensional space defined, this produces:

$$(7 \text{ variables})^{10 \text{ quanta}} = 282,475,249 \text{ combinations of parameters}$$

to search through. The size of this experimental space is enormous. Consider that a single experimental run down the MIT testing basin (including swimming down the tank, processing the data, backing the tuna up to its starting position, and letting the basin water settle) takes approximately 5 minutes, then the time required to blanket search the parametric space would be:

$$282,475,249 \text{ runs} \left( \frac{5 \text{ min}}{\text{run}} \right) \left( \frac{1 \text{ day}}{1440 \text{ min}} \right) \left( \frac{1 \text{ year}}{365 \text{ days}} \right) = 2687 \text{ years}$$

Obviously, some more-efficient form of search technique must be used to allow optimum-performance body controllers to be found, before either the apparatus wears out or available tank time expires. There are many existing techniques for structured searches through unknown multidimensional spaces. An excellent summary of them can be found in *Search Methods for Artificial Intelligence* [1]. Most such techniques require certain continuity requirements on the part of the space being searched to both operate and/or achieve convergence. Fluid dynamic state-spaces, however, often have large, chaotic, discontinuous regions (consider the transition from laminar to turbulent flow, for example) and it is highly probable the search will have to pass through these discontinuities in order to find the intended maximal (or minimal) value.

One particular technique, which is demonstrably good at searching through a wide variety of fragmented, discontinuous, multidimensional

state-spaces, is the genetic algorithm (GA). Genetic algorithms were invented by Holland in 1975 [2], and are basically probabilistic search engines based on the mechanics of natural selection and biological genetics. They have been shown to be very robust and typically rapidly find optima in a wide range of difficult (almost pathological) state-spaces (see [3–5]). GAs have typically been used by researchers in the machine learning community for parametric search and prediction tasks (see [6–8]) and in the growing field of artificial life for direct work on the evolution of artificial creatures and ecosystems (see [9–11]). Based on the successful use of genetic algorithms in all these areas, as well as their ability to robustly find optimums in difficult state-spaces, it was chosen as the algorithmic method to be used to search for optimum sets of body-wave parameters. The target of this optimal search was to achieve maximum propulsive efficiency in a self-propelled state. As an additional incentive, the biologically derived nature of the genetic algorithm closely parallels the biologically derived design of the RoboTuna itself. This proved to be an extremely effective method to guide the design of the RoboTuna and, as will be shown shortly, also proved to be a remarkably efficient method to find optima in a large, complex, multidimensional parametric space.

This chapter contains a simplified description of the genetic algorithm used to search the experimental parameter space for those optimum swimming modes possible within the hardware limits presented by the RoboTuna. In simple terms, it will describe how the genetic search engine enabled the RoboTuna to “learn” to swim efficiently, based on real experimentally measured swimming performance.

## DESCRIPTION OF GENETIC SEARCH ALGORITHM

Genetic algorithms (GAs) are probabilistic search engines based on the mechanics of natural selection and natural genetics. GAs typically start out with an initial population of randomly selected possible problem solutions (in the RoboTuna’s case, a solution is one complete set of seven body-wave control parameters). They then use a combination of survival of the fittest among solutions and a structured randomized information exchange between solutions to evolve toward better and better solutions. This evolution continues until either a pragmatically acceptable level of performance is achieved or a preset termination condition (typically specified as something like no improvement for five generations) is reached.

Key to the function of a genetic algorithm is the way the state-space to search and the possible solutions to that search are represented. Typically, a chromosome-like string represents each possible solution in the overall state-space. Each location on the chromosome structure,

called a “gene,” represents one of the fundamental parameters of the state-space to be searched. The particular value stored in each gene in the chromosome represents the value of its corresponding parameter. In order to function, the GA requires this string-based encoding—a method of creating an initial population of solutions. A common objective function is to quantify the fitness of each solution (i.e., how well it performs in the environment) and a set of genetic transformation operators to alter the genetic structure of the population from one generation to the next. Each item will be described in detail in the sections that follow.

### String-Based Encoding

GAs work on an encoding of the solution sets and not on the solutions themselves. This encoding typically takes the form of a finite-length string, or chromosome, in which each element, or gene, comes from a finite alphabet of possibilities. The most common current encoding scheme uses a binary transformation of solution parameters such that the chromosomes end up looking like: 011011011001111011100. Wanting to tie the gene structure more closely to the fundamental kinematics of the RoboTuna system, I decided to break with convention here and choose a gene structure that directly incorporated the seven fundamental parameters of the body wave, in a form that was both easily and physically understandable, namely:

$$\text{Example Chromosome: } \{U, St\#, \alpha, \psi, TIP, \lambda, c_1\} \quad (10.1)$$

where:  $U$  = forward speed of RoboTuna,  $St$  = Strouhal number of Tuna’s tail,  $\alpha$  = angle of attack of tail,  $\psi$  = phase angle between tail pitch and heave,  $TIP$  = maximum excursion of tail’s tip,  $\lambda$  = wavelength of body wave, and  $c_1$  = amplitude envelope of body wave.

One way to look at this is that each combination of body-wave parameters, in the 7-D state-space of all such possible combinations, represents a single fish, which swims with a body wave defined by its unique set of parameters. These particular parameters can be expressed in terms of the fish’s “genetic code” and stored in a linear array of the form shown above. This fish’s genetic code stored in such a chromosome can now be operated on and processed by the remainder of the genetic algorithm.

### Generation of Initial Population

The second step in the genetic algorithm is to create a fixed-size initial population of  $N$  solutions and to somehow load the individual genes of each solution with appropriate values. Typically, for research purposes, these genes are loaded with totally random values from the possible



alphabet of each gene. The larger the population size (typically in the 100s), the greater the possibility that the appropriate genetic material for the optimal solution being sought lies within this initial population and the greater the probability of rapid convergence to it. This fixed population size of  $N$  solutions remains constant for all succeeding generations.

In the case of the RoboTuna, the alphabet of possible values for each gene was either dictated by the physical limits of the particular parameter involved or chosen to widely bracket pragmatically usable values based on previously collected biological and experimental data. In another break from traditional GAs, I chose a very small population size (typically only 10 fish) in order to be able to physically run an entire generation in about an hour's time. This was done explicitly to minimize wear and tear on the RoboTuna's hardware. Because in this case the GA is operating on real hardware and the maximum improvement in system performance comes during generational jumps, not within the members of a generation, it was desirable to get as many generations as possible within the minimum number of tests. Previous GA simulations indicated that the dramatic increase of information storage caused by encoding parameter values directly into genes (rather than by breaking this information down into a binary code) allowed much smaller populations of fish to converge in an acceptably small number of generations.

### **Definition of Objective (or Fitness) Function**

Genetic algorithms search through multidimensional parametric spaces by rewarding more "fit" solutions with relatively more descendants in future generations, and by punishing less-fit solutions by allowing them less descendants (i.e., replicating survival of the fittest in natural evolution). In order to carry out this allocation of offspring, it is essential to have a clearly defined, easily quantifiable measure of a solution's performance. In pure GA research work, this measure of performance is typically specified as an objective function (i.e., maximize this particular function). Inappropriate selection of an objective function can lead to premature convergence at a local, but not global, maxima or minima. Or, if too-loosely defined, can lead to excessive meandering and an intolerable delay in convergence. As such, proper selection of this "fitness" function is crucial.

Fortunately, in the case of the RoboTuna, the objective function was clearly defined (i.e., efficient propulsion), precisely experimentally measurable, and had a direct, useful, physical meaning. Consider that each set of parameters in the 7-D experimental state-space represents a single fish that swims with the body motion described by these parameters. By swimming these fish down the tank, a clean numerical

measure of their swimming performance can be determined directly from the experimental data recorded during the run. In some sense, each one of these combinations of parameters is a “virtual tuna,” and each virtual tuna takes a turn at inhabiting the RoboTuna’s body for its respective swim down the tank. At the end of each run, each virtual tuna’s performance is given a direct numerical score. And this score determines whether or not this particular set of control parameters will be allowed to propagate on to the next generation.

### **Genetic Transformation Operators**

Genetic algorithms use probabilistic (not deterministic) transformation operators to convert the genetic material stored in the population of solutions from one generation to the next. These operations typically fall into three classes; reproduction, crossover, and mutation. Each performs a biologically inspired, evolutionarily derived function on the information stored in the genes of the chromosomes of each solution in the population. The net effect of all three operations is the evolution of populations of solution sets, from less to more fit with each subsequent generation. Not only are superior performing sets of parameters passed in increasing shares from generation to generation, but also, these top-performing parameters receive exponentially increasing shares in the overall population’s gene pool—saturating the pool with top performers, and facilitating the rapid convergence to maximal values. In terms of the previous RoboTuna analogy, specific virtual fish that perform better than average for their current population have a much better chance of passing their genetic material on to the next generation of Tuna than fish that do not. Each transformation operation will now be discussed in detail.

The *reproduction operator* is effectively the artificial version of natural selection. It applies a common fitness (or objective) function to each solution set in the current population of solutions. Based on the performance rating generated by this operation, individual solutions with a higher fitness have a higher probability of contributing one or more offspring to the next generation than those that do not. Existing GAs typically use reproduction operators based on schemes like the “roulette wheel,” where, conceptually, the area of a roulette wheel is broken down into  $N$  pie segments—one segment for each member of the population. The area of each segment is proportional to its respective solution’s fitness, relative to the mean fitness of its generation. Typically, better performers get proportionately bigger slices of the roulette pie. The wheel is then spun  $N$  times, randomly selecting the  $N$  solutions that will be allowed to reproduce. Based on the larger amounts of wheel real estate claimed by the more-fit solutions, they will have a higher probability of being allowed to reproduce.

A similar operator was created for the RoboTuna, but was found to converge fairly slowly in simulation. It performs a more robust search through space, but tends to allow too many poor performers to squeeze by from generation to generation. So, in another departure from standard GAs, a harsher reproduction operator, one that only lets the top 50% of performers in each generation reproduce, was created. In keeping with previous decisions, it was felt that this step would lead to faster convergence, work better with the small population size, and minimize the number of experiments that needed to be run. In terms of the RoboTuna's GA then, a population of 10 virtual fish are subjected to a common swimming performance test and only the top 5 performers in each generation are allowed to reproduce. The top performers get two offspring apiece. This brings the number of fish in the next generation back up to 10, reforming the initial fixed population size.

The *crossover operator* is the artificial version of sex. It performs the mixing and joining of the genetic material from the two parent's chromosomes to produce the chromosomes of their children. In a typical genetic algorithm, the crossover operator is performed as follows. Survivors of the reproduction operator are randomly paired off. A random location (gene) along the chromosome is chosen. At this location, both parent's chromosomes are split and then all genes on one side of the split are swapped. This produces two offspring, both of which have some genetic material from one parent and some from the other. Depending on the chromosomes of the parents, the two offspring may be similar or wildly different. In some sense, the crossover operation is a way to more robustly search through state-space, while still retaining "high-performance" genetic material during the search.

The crossover operator used by the RoboTuna's GA is similar to that used in a standard GA. Mating pairs of virtual fish are chosen randomly from the survivors of the reproduction operation. But instead of a random break location, the location of the chromosome break, swap, and splice is fixed between the center genes. Based on the results of early simulations, the location of each gene in the Tuna's chromosomal structure was specifically chosen such that related body-wave control parameters would occupy adjacent positions. Because related parameters control real related physical actions (i.e., the angle of attack and the phase angle of the tail), it was thought that keeping them adjacent and intact during crossing would allow fish with clearly superior performance features (e.g., good tail motion) to cleanly pass this information onto their children without disruption. Furthermore, fixing the crossover location at mid-gene, during simulations with this particular chromosomal structure, cost little in the way of search robustness, but promoted a much more rapid convergence to optima values than when using a variable crossing location. In the ongoing effort to minimize the wear and tear on the tuna by reducing the number of experimental

runs required to find optimal performance, using the fixed location was a clear winner.

The *mutation operator* is the artificial equivalent of genetic mutation. It performs the occasional (with small probability) random alteration of one or more genes at randomly selected sites in the chromosomes of all the solutions in a population. Mutation is needed to compensate for two possibilities. First, the reproduction and crossover operators may occasionally lose potentially useful genetic material that is buried in the chromosome of a poorly performing solution (i.e., really good tail motion on an under-performing fish). Mutation provides a way to recover it. Second, the genetic material required for an optimum solution may not have been present in the initial population of solutions in the first place. In this case, a GA without mutation cannot find the optimum because the raw material required to build it is just not there to work with. By itself, mutation is just a random walk through the solution space. But when used in conjunction with reproduction and crossover, it acts as a form of insurance policy against premature convergence at nonoptimal solutions.

The mutation operator used by the RoboTuna's GA is slightly more sophisticated than that described above. Early work with simulations demonstrated that a high initial mutation rate dramatically helped make up for the small number of individuals in each population, but was counterproductive, and in fact interfered with convergence in the terminal stages of the search. Based on this, the mutation operator, finally developed, starts with a relatively high mutation rate and then trails off to a low rate as the number of generations increase (in a process akin to simulated annealing).

In terms of the virtual-fish analogy, the crossover operation leaves a population of  $N$  new fish, each inheriting half of their genes from their two randomly selected parents. The mutation operator then randomly selects a fixed number of these genes (from throughout the entire population of fish) to mutate, and randomly alters them within each gene parameter's pre-set boundaries. The number of genes selected starts out at a high percentage and trails off with each successive generation, as described above. At the completion of this operation, a fully formed population of offspring (the next generation of virtual tuna) has been created and is ready to be subjected to experimental performance trials in the tank.

#### **EXAMPLE OF ROBOTUNA GA**

The combination of all of the above features allows the RoboTuna's genetic algorithm to perform a structured, parallel, robust, probabilistically based, search through the tuna's 7-D parametric state-space. How this search is performed in practice will be presented next (by

**Table 10.1** Initial population

Tuna No.	U m/sec	St #	$\alpha^\circ$	$\psi^\circ$	TIP BL*	$\lambda$ BL	$c_1$	TPR
1	0.9	.18	15	90	.19	1.1	002	.96
2	0.6	.30	20	76	.30	1.5	001	.53
3	1.1	.28	14	83	.11	1.3	003	.92
4	0.3	.35	36	111	.29	.96	002	.17
5	0.7	.22	28	92	.25	1.1	002	.87
6	0.5	.10	11	100	.16	1.4	019	.64

Mean TPR = 0.68.

\*BL = w.r.t. Tuna's body length.

way of a direct example taken from actual experimental runs). Why the search works (the theoretical basis for genetic algorithms) will follow at the end this section.

In practical terms, the RoboTuna's GA is implemented in the following way: Each possible set of body-wave control parameters represents a single point in a seven-dimensional state-space consisting of all possible sets of such parameters. Consider each one of these points to be a chromosome made up from 7 genes, where each gene is simply one of the 7 body-control parameters ( $U, St \#, \alpha, \psi, TIP, \lambda, c_1$ ). In this light, each chromosome represents one possible virtual fish and the objective of the GA is to find the fish with the best swimming performance (the fit-test tuna), while searching through the least possible number of fish.

The genetic search engine accomplishes this using the following structure: A fixed-size starting population of  $n$  (in the following example,  $n = 6$ ) fish is created. The individual genes of each fish in this population are loaded with randomly selected (but physically possible) values of the appropriate parameter (table 10.1). Each fish is then subjected to a common "fitness test." In this case, the fitness test consists of a full experimental swim down the tow tank followed by the generation of the fish's performance rating in terms of its thrust power ratio, defined as:

$$\text{Thrust power ratio} = \frac{\text{power output as thrust}}{\text{power input to undulate Tuna's body}}$$

The thrust power ratio (TPR) is a quantitative measure of how well the power entering the Tuna is converted into usable output thrust power.

Based on the results of this fitness test, this initial population of fish is subjected to the aforementioned "survival of the fittest" selection process, where, in this particular instance, only the top 50% of the population are allowed to pass their genes on to the next generation of fish. This procreation of the next generation takes place in three basic operations derived from those present in natural systems. These three operations are called reproduction, crossover, and mutation, as noted above.

**Table 10.2** Reproduction operation

Tuna No.	U m/sec	St #	$\alpha^\circ$	$\psi^\circ$	TIP BL	$\lambda$ BL	$c_1$	TPR
1	0.9	.18	15	90	.19	1.1	002	.96
5	0.7	.22	28	92	.25	1.1	002	.87
3	1.1	.28	14	83	.11	1.3	003	.92
1	0.9	.18	15	90	.19	1.1	002	.96
5	0.7	.22	28	92	.25	1.1	002	.87
3	1.1	.28	14	83	.11	1.3	003	.92

Mean TPR = 0.91.

**Table 10.3** Gene-crossing operation

Tuna No.	U m/sec	St #	$\alpha^\circ$	$\psi^\circ$	TIP BL	$\lambda$ BL	$c_1$	TPR
1	0.9	.18	15	<b>92</b>	<b>.25</b>	<b>1.1</b>	<b>002</b>	
5	0.7	.22	28	<b>90</b>	<b>.19</b>	<b>1.1</b>	<b>002</b>	
3	1.1	.28	14	90	.19	1.1	002	
1	0.9	.18	15	83	.11	1.3	003	
5	0.7	.22	28	83	.11	1.3	003	
3	1.1	.28	14	92	.25	1.1	002	

In the reproduction operation, the fittest 50% of all fish are selected and allowed to have two offspring each, thus expanding the next generation's population size back to that of the original population's (table 10.2). Note that the average fitness of the offspring exceeds that of their parent's generation.

In the gene-crossing operation (table 10.3), adjacent pairs of offspring swap one half of their genes for one half of their partner's genes. Reproduction, by itself, quickly finds the best of the existing population but does nothing to explore new regions of state-space. Crossover allows a structured randomized search with a minimum of disruption to the allocation strategy dictated by reproduction. Careful inspection of the above gene map will reveal that the first and second fish crossed genes at  $\alpha, \psi$  interface, preserving their forward 3 genes and swapping the aft 4 genes. The remaining pairs (3 + 4, 5 + 6) did the same.

The remaining genetic operation, mutation, performs occasional (with small probability) random alteration of one or more genes at randomly selected sites within the chromosome of each fish in the population (table 10.4). It is needed because reproduction and crossover may occasionally lose potentially useful genetic material depending on initial population selection or on a too-narrowly defined fitness function. (In table 10.4, mutations are designated by numbers in bold-face type.)

**Table 10.4** Mutation operation

Tuna No.	U m/sec	St #	$\alpha^\circ$	$\psi^\circ$	TIP BL	$\lambda$ BL	$c_1$	TPR
1	<b>1.5</b>	.18	15	92	.25	1.1	002	
5	0.7	.22	<b>33</b>	90	.19	1.1	002	
3	1.1	.28	14	90	.19	<b>0.9</b>	002	
1	0.9	.18	15	83	.11	1.3	003	
5	0.7	<b>.35</b>	<b>29</b>	83	.11	1.3	003	
3	1.1	.28	14	92	.25	1.1	002	

At the end of the mutation operation, a new generation of fish is ready to swim down the tank and be subjected to the same common fitness test as their parents. Based on this second generation's swimming performance, the basic operations of reproduction, crossover, and mutation are again repeated, and a subsequent population of new offspring is generated. And the entire cycle is repeated over and over, with each generation performing better than the last, until a stable optimum is reached (SuperTuna) or a preselected level of performance is achieved (the search times-out). In some sense, the genetic algorithm only uses a few simple operations to perform its search mission, but it achieves its power by performing these simple operations on the full population of virtual Tuna in parallel, over and over again, until an optimum is achieved. As can be seen in table 10.4 above, optimal performance was achieved in only ten generations.

At this point, it is probably appropriate to stop and address the question: What exactly is being processed by the GA? Basically, GAs propagate short-length, highly fit groups of genes from generation to generation in exponentially increasing numbers until the gene pool is saturated with them. These groups of "fit" genes are actually the building blocks of the full problem solution we seek. Instead of trying to find this full solution by trying every conceivable combination of genes, we can build increasingly better partial solutions by concatenating together the exponentially improving building blocks found in each new generation. Because this process of building better solutions goes on in a massively parallel fashion, high state-space search efficiencies are achieved.

In more precise terms, it can be said that GAs process schemata. Schemata are defined as the set of all chromosomes that at certain pre-defined positions have similar gene values. Schemata are a notational device for rigorously discussing and classifying chromosome similarities. They also permit the theoretical analysis of the effect of genetic transformation operators (like reproduction, crossover and mutation) on gene structures in the population. For example, consider the binary

alphabet of 0, 1, \* where "\*" can represent either a "1" or a "0." In this case, the schema \*01\* has 0010, 0011, 1010, and 1011 as its members. For a fixed-length chromosome then, the number of possible schemata that exist is given by:

$$\text{Number of possible schemata} = (k + 1)^n \quad (10.2)$$

where  $k$  = size of gene alphabet (for binary  $k = 2$ ) and  $n$  = length of chromosome.

In the previous example,  $n = 4$  and  $k = 3$ , creating 64 possible schemata. For the RoboTuna, the space is of considerably higher complexity with  $n = 7$  and  $k \geq 300$ , yielding  $2.238 \times 10^{17}$  possible schemata. During this processing, the schemata present in each fish's chromosomes are operated on by the three basic operators described above. As will be described more rigorously in what follows, the reproduction operator gives an increasing proportion of the population to the fittest schemata. The crossover operation tends to pass on short-length, highly fit schemata and chops up longer ones. The mutation operator, when operating at normal (low) rates, does not really disrupt fit schemata. The sum total effect is that genetic algorithms propagate highly fit, short-length schemata from generation to generation in ever-increasing numbers until the overall population's gene pool is saturated with them. Consider next the effect of each transformation on the schemata in the gene pool.

The effect of the reproduction operator on the expected number of schemata in the population can be determined as follows; for a given generation  $t$  let the total number of a particular schema  $H$  contained in the population  $A(t)$  be given by:

$$m(H, t) = \text{total number of chromosomes of schema } H \text{ contained} \\ \text{in the population } \vec{A}, \text{ in generation } t \quad (10.3)$$

During reproduction, a chromosome gets copied according to its fitness, more precisely, the  $i$ th chromosome in generation  $t$  gets selected with the probability:

$$P(A_i) = \frac{Fit_i}{\sum_{j=1}^n Fit_j} \quad (10.4)$$

where:  $P(A_i)$  = probability of chromosome  $A_i$  being allowed to reproduce,  $Fit_i$  = fitness of solution described by chromosome  $A_i$ ,  $Fit_j$  = fitness of chromosome  $A_j$ , and  $n$  = population size.

In the next generation, the total number of chromosomes  $A_i$  contained in the population is given by:

$$m(A_i, t + 1) = m(A_i, t) \frac{Fit_i}{Fit} \quad (10.5)$$



where:  $m(A_i, t + 1)$  = number of chromosomes  $A_i$  in generation  $t + 1$  and

$$\overline{Fit} = \frac{\sum_{j=1}^n Fit_j}{n} = \text{mean generational fitness.}$$

By extension, for any particular schema  $H$ , with average fitness (taken over all chromosomes containing  $H$ ) of  $\overline{Fit}_H$ , the total number of schema  $H$  contained in the subsequent population is given by:

$$m(H, t + 1) = m(H, t) \frac{\overline{Fit}_H}{\overline{Fit}} \quad (10.6)$$

where:  $m(H, t + 1)$  = number of schema  $H$  in generation  $t + 1$ ,  $\overline{Fit}_H$  = the mean fitness of all chromosomes containing schema  $H$ , and

$$\overline{Fit} = \frac{\sum_{j=1}^n Fit_j}{n} = \text{mean generational fitness.}$$

This is known as the reproductive schema growth equation and it states that the number of a particular schema  $H$  in each generation grows at the ratio of the average fitness of the schema, to the average fitness of the generation. It is important to note that this behavior is carried out with every schema (i.e.,  $H, I, J, K, L \dots$ ) in the population in parallel. Schemata with average fitness values above the generation's mean receive more descendants, those below, less. How fast do the numbers of a particularly successful schema grow in a population? Consider the case of a particular schema  $H$  whose mean fitness remains above the average generational fitness,  $\overline{Fit}$ , by the amount,  $c\overline{Fit}$ , where  $c$  is a constant. Then:

$$m(H, t + 1) = m(H, t) \frac{(\overline{Fit} + c\overline{Fit})}{\overline{Fit}} = (1 + c)m(H, t)$$

then:

$$m(H, 1) = (1 + c)m(H, 0)$$

$$m(H, 2) = (1 + c)^2 m(H, 0)$$

⋮

$$m(H, f) = (1 + c)^f m(H, 0) \quad (10.7)$$

In this case, the reproduction operation allocates exponentially increasing (or decreasing) percentages of the overall population to above- (and below-) average performing schemata. And it is important to note that it does so in parallel, in each operation.

Clearly, the reproduction operator quickly finds the best schema (genes) in an existing population but does nothing to explore new

regions of solution space. The crossover operator allows a structured randomized search with a minimum of disruption to the schema allocation strategy used by the reproduction operation. The effect of crossover is illustrated in the following example:

Let:  $A_1 = \{011/1000\}$  where the “/” represents the crossover site.

Then:  $H_1 = \{*1*/**0\}$  schema  $H_1$  is destroyed.

But:  $H_2 = \{***/10**\}$  schema  $H_2$  is passed.

On average, long schema are more likely to be destroyed because the gene-slicing point is more likely to fall between extreme chromosome positions. The probability that any particular schema will survive crossover intact is given by:

$$P_{survival}(H) = 1 - \left( \frac{d(H)}{L_{chromosome}} \right) \quad (10.8)$$

where:  $P_{survival}(H)$  = the probability that schema  $H$  will survive the crossing,  $d(H)$  = the defining length of schema  $H$  (the distance between the first and last specific string positions), Ex:  $d\{011*1**\} = 5 - 1 = 4$ , and  $L_{chromosome}$  = length of chromosome (in the above example  $L_c = 7$ ).

If the crossover site is randomly selected with a probability of  $P_c$ , then the probability that any particular schema will survive crossover intact becomes:

$$P_{survival}(H) \geq 1 - P_c \left( \frac{d(H)}{L_{chromosome}} \right) \quad (10.9)$$

Combining the effects of both the reproduction and the crossover operations show that the total number of a particular schema  $H$  contained in the population in generation  $t$  is:

$$m(H, t + 1) \geq m(H, t) \frac{\overline{Fit}_h}{\overline{Fit}} 1 - P_c \left( \frac{d(H)}{L_{chromosome}} \right) \quad (10.10)$$

The net total effect being that schemata with both above-average performance and short defining lengths propagate at exponentially increasing rates.

Neither the reproduction nor the crossover operators can add “useful” material to the gene pool. In order to compensate for lost genetic material (schema) or to make up for genetic material that was never included in the initial population in the first place, the mutation operator must be employed. The mutation operator effects random alterations of one or more genes in a schema with small probability  $P_m$ . In order for a schema  $H$  to survive mutation, all of its fixed positions must survive. The probability that each gene will survive is  $(1 - P_m)$ . The probability that the entire schema will survive is given by:

$$P_{survival} = (1 - P_m)^{o(H)} \quad (10.11)$$

where:  $P_m$  = probability of mutation and  $o(H)$  = order of schema  $H$ , basically the number of fixed positions in the template. Ex:  $o(111*0*) = 4$ . For small values of  $P_m$  ( $P_m \ll 1$ ) this probability of survival can be approximated as:

$$P_{survival} = o(H)P_m \quad (10.12)$$

Combining all total effects of all three operators together leads to the following equation:

$$m(H, t + 1) \geq m(H, t) \frac{\overline{Fit}_h}{\overline{Fit}} \left\{ 1 - P_c \left( \frac{d(H)}{L_{chromosome}} \right) - o(H)P_m \right\} \quad (10.13)$$

Commonly called the *schemata theorem*, it is one of the fundamental cornerstones of genetic algorithm theory and it shows that the net effect of reproduction, crossover, and mutation is that schema that have a fitness above average, a small defining length, and a low order receive exponentially increasing numbers of offspring in each successive generation.

It is probably fair to ask again at this point: Why is this a good thing to do? Basically, these short, low-order, highly fit schemata are actually the building blocks of good solutions (i.e., they are partial solutions to the optimization problem). By passing ever-increasing numbers of them from generation to generation, the GA builds exponentially improving chromosomes (bigger and better partial solutions) from the best building blocks in each generation. This concatenation of ever-increasingly fit building blocks from generation to generation goes on in a massively paralleled fashion (building ever bigger and better solutions) until an optimal answer is found. This parallel processing of the information leads to both high computational efficiency and rapid convergence to optima (the full best solution).

It was this capability of robustly and rapidly converging on optima that made genetic algorithms a prime candidate for the optimization of performance within the chaotic, unknown space formed by the parameters of the RoboTuna's body-wave controller.

## REFERENCES

- [1] Bolc, L., and Cytowski, J. (1992). *Search Methods for Artificial Intelligence*. San Diego: Academic Press.
- [2] Holland, J. (1995). *Adaptation in Natural and Artificial Systems*. University of Michigan Press.
- [3] Goldberg, D. (1990). *Genetic Algorithms in Search, Optimization, and Machine Learning*. Reading, MA: Addison-Wesley.
- [4] Holland, J., and Goldberg, D. (1988). Genetic algorithms and machine learning. *Machine Learning* 3: 95.
- [5] Goldberg, D. (1986). The Genetic Algorithm Approach: Why, How, and What Next? In *Adaptive and Learning Systems*. New York: Plenum Press.

- [6] Rodgers, D. (1990). Weather prediction using genetic memory. NASA Tech. Rep. 90.6. Ames Research Center.
- [7] Koza, J. R. (1992). *Genetic Programming: On the Programming of Computers by Means of Natural Selection*. Cambridge: MIT Press.
- [8] DeJong, K. A. (1990). Genetic-algorithm-based learning. *Machine Learning* 3: 611–638.
- [9] Langton, C. (ed.). (1995). *Artificial Life, An Overview*. Cambridge: MIT Press.
- [10] Collins, R. J., and Jefferson, D. R. (1992). AntFarm: Toward simulated evolution. In *Artificial Life II*. Addison-Wesley, pp. 579–601.
- [11] Wilson, S. W. (1985). Knowledge growth in an artificial animal. In *Proceedings of the First International Conference on Genetic Algorithms*. Erlbaum.

**This page intentionally left blank**

# 11

## Biomechanical Primitives and Heterarchical Control of Limb Motion in Tetrapods

Simon F. Giszter

The aim of this chapter is to present a particular perspective and overview of our understanding of low-level motor control in legged vertebrates. The chapter focuses on spinal organization of movement, with a view to emulating this in a neuromimetic robot design. The perspective presented emphasizes the need for an extensible design that supports flexible planning, position control, and interaction force control and thus is intimately wedded to the mechanical plant to be controlled (see Full and Koditschek, 1999; Abbas and Full, 2000).

### NEURAL ORGANIZATION OF CONTROL: HIERARCHY OR HETERARCHY?

In considering the role of the vertebrate spinal cord in movement organization, a hierarchical view has often been favored historically. The spinal cord is often seen as a device that allows the cortex or other higher centers to address muscles and very little more. However, since the time of Sherrington (1910) or earlier, there has been a parallel view that many capabilities are embedded in the spinal cord. How these local autonomous capabilities and descending requirements are integrated into seamless adaptive action is an important topic of active investigation. The likelihood may be that the system acts as a heterarchy that is reconfigured for different tasks and conditions with the locus of control moving back and forth among centers arrayed throughout the central nervous system (CNS) neuraxis.

Marr's analysis of the computational problem may be extremely useful when considering this heterarchical organization of the CNS, and particularly the spinal organization, in order to build a biomimetic system. Task specifications at the highest level must be supported by solution algorithms at an intermediate level that can be implemented in numerous ways in the embodied agent. The algorithms chosen are constrained both by (1) task requirements, which are partly or wholly determined by the physical world (e.g., see McFarland and Houston, 1981; McFarland and Bossert, 1993), and by (2) possible

implementations given the hardware and technology available, which are also physically constrained. It is likely to be desirable to locate a simple extensible algorithmic structure that can enable solutions of many tasks in various hierarchies, and is capable of using differing implementation hardware in terms of both human design of artifacts, and biological evolutionary potential and survival. Perhaps biological motor control can provide insights into this algorithmic structure.

A central tenet of this chapter is that the appropriate computational or algorithmic structure to satisfy the concerns outlined is likely to be based around a modular design. A design based around an appropriate set of reconfigurable and recombined building blocks (Kiehn, Hounsgaard, and Sillar, 1997) or primitives (Giszter, Mussa-Ivaldi, and Bizzi, 1993) may support rapid local synthesis of simple movement and extension to complex or unusual domains by motor learning.

Here I will explore the primitives or holons that we believe are used in the spinal cord of lower vertebrates and mammals and relate these to a specific physical “mission” and to a strategic goal. The algorithmic framework should: (a) achieve the mission of a stable delegated execution of useful self-protective motions to defend the mechanical plant in a broader context of ongoing motions and (b) support the strategic goal of providing a language or basis composed of elements for useful synthesis of new motions by other components of the CNS hierarchy based on novel task demands.

## BACKGROUND

The appropriate units of movement or action in biological systems have been examined extensively but are not generally agreed upon. Kinematic units are one idea frequently explored (Berkinblitt, Feldman, and Fukson, 1986), and indeed, the notion of a kinematic plan or template is strongly supported by some investigators. However, it seems clear that animals must meet multiple task requirements and that a significant fraction of these are kinetic, dynamic, and energy based (McFarland and Houston, 1981; Raibert, 1986; Full and Koditschek, 1999; Abbas and Full, 2000). Depending on task or subtask, it is likely that position and force controls both must be used, and that these must be integrated and switched seamlessly in tasks such as locomotion. Investigation of neural organizations, which might shed light on these processes, has taken one of two approaches.

The first, which has unambiguously established the importance of the neural controller and a central representation of action, uses paralyzed or fictive preparations in which movements are abolished (see Grillner and Wallen, 1985). This is achieved either by excising the animal’s CNS and recording it in a dish, or by paralyzing the musculature. These approaches allow exceptionally fine and detailed examination of

neural activity and has led to the notion of central pattern generators (CPGs—to be distinguished from control plan generators). CPGs are circuits in the spinal cord able to organize complex adaptable behavior even when the rest of the brain is removed (see Grillner and Wallen, 1985; and Kiehn, Hounsgaard, and Sillar, 1997). CPGs are operationally defined neural systems that generate neural activity that approximates the neural patterns for various animal behaviors (locomotion, grooming, burrowing, struggling, righting, turning, escaping) in the absence of phasic feedback and in a largely or entirely feedforward fashion. These systems comprise oscillatory cells or circuits and pattern synthesizing and control circuitry. Flexible adaptable motor patterns presumably must be elaborated from spinal central pattern generators. This approach has led to several very important insights, but the method finesses the issues of what is controlled and represented (i.e., what are the controlled units) in the synthesis of action by the CNS/mechanical-plant combination.

The second approach examines the intact or semi-intact animal (Bizzi, Mussa-Ivaldi, and Giszter, 1991). However, this approach cannot currently provide the refinement of neural recording possible in fictive preparations. Many analyses of the second kind have remained at the level of an energetic or purely biomechanical description. Ideally, these two approaches and frameworks must be bridged both in terms of neural control and synthetic biomimetic approaches.

Over the last 10 years, I and colleagues at MCPHU, MIT, and Northwestern University have been investigating a framework that attempts to bridge these disparate but not incompatible bodies of work.

These experiments in unparalyzed spinalized animals (rat, Tresch and Bizzi, 1999; and frog, Giszter, Mussa-Ivaldi, and Bizzi, 1993, 2000a) have shown that there is a collection of functionally separate modular circuits or primitives that is likely to be recruited by CPGs in different behaviors and perhaps also used by descending systems. These primitives offer features that may have a considerable impact on the construction and organization of motion at spinal and higher levels. Primitives may form a core design feature in biological motor competences that may assist in designing biomimetic and humanoid systems (see also Giszter et al., 2000c). We will review now these primitives in more detail.

### **Spinal Motor Primitives**

Many functions are embedded in the spinal cord and the spinal cord itself can execute complex, graceful, and well-planned movements. Indeed, even some aspects of motor learning are believed to be directly embedded in the spinal cord (e.g., Wolpaw laboratory work). The spinal cord is now thought by us to contain a modular organization of



motor primitives (Bizzi, Mussa-Ivaldi, and Giszter, 1991; Bizzi et al., 1995; Giszter et al., 2000a, b; Kargo and Giszter, 2000). The perspective we favor views the CPG as one type of element of a modular system: The tasks of rhythm generation, pattern assembly, and output adjustment are separated into modular components. The output of the CPG is believed by us to comprise motor primitives (Bizzi, Mussa-Ivaldi, and Giszter, 1991; Giszter, Mussa-Ivaldi, and Bizzi, 1993). The motor primitives are seen as separate from the timing or rhythm generation system in this perspective. We speculate that motor primitives may be reused in different tasks (Berkinblitt, Feldman, and Fukson, 1986; Mussa-Ivaldi and Giszter, 1992; Bizzi et al., 1995). The precise composition of these motor output elements has been a point of discussion. At this point in the chapter it will be useful to define a motor or force-field primitive with more formal precision. Motor primitives are defined here as fixed groups of muscles, their motoneuron pools, and their afferent regulatory circuits, which together act as a unit to produce a fixed set of motor output properties in a limb.

At the simplest level, a motor primitive could be defined as a single muscle (the myotatic unit). On the one hand, the myotatic unit would allow the system complete flexibility to produce any motor output by combining activity across multiple muscles. However, the single muscle and its homonymous control exposes the controller to the full impact of the degrees-of-freedom problem described by Bernstein (1967). Because the combinations of muscle activity are constrained by the joints and limbs, we posit that a more-useful definition of motor primitives is a unit that combines muscles and their controls across joints and acts to represent a biomechanical force field. A force field, in this context, is a function that maps the forces generated in a limb to the limb's configuration. Ideally, these force fields would combine to synthesize behaviors.

In 1991, based on microstimulation, we suggested that spinal circuits are in part organized so as to support and control such force patterns in the limb. Remarkably, only a few types of force fields were observed in the spinal cord. It was discovered in dual stimulations that these force fields could be combined, in a simple linear fashion, by vector summation. It therefore seemed that a collection of spinal force field primitives could form a basis for the construction of arbitrary force fields or movements by recruiting them in combinations. It was hypothesized that such recruitment might be performed either by spinal pattern generators or by descending pathways. Subsequent studies by our laboratory (Giszter and Kargo, 2000a; Kargo and Giszter, 2000a, b) and that of Bizzi (Bizzi, Mussa-Ivaldi, and Giszter, 1991; Giszter, Mussa-Ivaldi, and Bizzi, 1993; d'Avella and Bizzi, 1998; Tresch and Bizzi, 1999) have tested the organization and control of spinal primitives and have demonstrated that reflex behaviors are likely to be built and adjusted using these primitives (Mussa-Ivaldi et al., 1995; Tresch et al., 1999;

Kargo and Giszter, 2000a). Further, Tresch has shown that closely comparable force-field primitives exist in mammals (Tresch, 1997; Tresch and Bizzi, 1999), and can be elicited from lumbar spinal cord in chronically transected adult rats. Modular muscle groups have also been observed in man, in cycling (Ting et al., 1998).

Ten years after the original publications, the idea of the force-field primitive in frog and rat has survived several sets of critical tests and corroborating experiments. We have shown (Kargo and Giszter, 2000a) that in the precisely targeted multiphase wiping reflex, both aimed-trajectory formation and rapid on-line corrections to circumvent obstacles can be described as summation of primitives. Alterations of patterns could be described as losses of primitives. Further, and perhaps most significant of all, in each of these studies we have found that every force-field primitive (and its underlying motor activity) that contributed to the behavior or to trajectory adjustments exhibited common activation dynamics in spinalized frogs. Taken together, our data strongly support a modular control framework for these behaviors based on force-field primitives. Moreover, the data suggest that the force-field primitive forms a very particular kind of building block for movement, and this leads naturally to several testable hypotheses and learning schemes, which may be of wide-reaching importance.

### **Force-Field Primitives as Building Blocks for Motor Learning**

Force-field primitives were defined as active force fields that exhibited invariant force vector directions and magnitude balances over time—that is, primitives had conserved structures through time (Giszter, Mussa-Ivaldi, and Bizzi, 1993). Force-field primitives are the effect of ensemble feed-forward and feedback circuits, which together act as a unit. These circuits generate a biomechanical multijoint force field. We will follow the definition used in Kargo and Giszter (2000a): A force-field primitive is observed as a structurally invariant force field over time. In its most general form, a force-field primitive would be a vector function of both position and velocity, modulated by a time-varying scalar. We use  $\mathbf{r}$  here as a configuration vector. The optimal coordinate system to express primitives (joint based, end-point, polar, Cartesian) is not addressed here directly. A primitive can thus be expressed in the form

$$F(\mathbf{r}, \dot{\mathbf{r}}, t) = a(t)\phi(\mathbf{r}, \dot{\mathbf{r}}) \quad (11.1)$$

$F(\mathbf{r}, \dot{\mathbf{r}}, t)$  is the observed field derived from a primitive, which can be expressed as a scaling  $a(t)$  of a fixed-field structure  $\phi(\mathbf{r}, \dot{\mathbf{r}})$ , a function only of position and velocity.

In the experiments described above, isometric testing was used (in isometric testing, the limb is held rigidly clamped and not allowed to

move), and in this case these descriptions can be written without velocity terms. Under such isometric conditions, at a single limb position, the force vectors generated by a primitive will increase and decrease in magnitude with activation-deactivation dynamics of the primitive. At the same time, force magnitude ratios among the sampled limb positions remain constant, that is, a primitive's time evolution can be expressed in the form  $a(t) \cdot f(\mathbf{r})$ . This is the description most often seen in neurobiological papers on primitives. Further, this is the framework in which linear superposition was tested experimentally (Mussa-Ivaldi, Giszter, and Bizzi, 1994).

Using elements of this form, Mussa-Ivaldi (1992) showed in a theoretical study that arbitrary smooth force fields could be approximated using combinations of appropriately chosen basis fields. In particular, he showed that an arbitrary smooth field  $F(\mathbf{r})$  could be generated by appropriate linear combinations of nonlinear conservative and rotational basis fields:

$$F(\mathbf{r}) = \sum_i A_i \phi_i(\mathbf{r}) + \sum_i B_i \psi_i(\mathbf{r}) \quad (11.2)$$

Where the field  $F(\mathbf{r})$  is constructed from conservative primitives  $\Phi_i(\mathbf{r})$  and rotational or circulating primitives  $\Psi_i(\mathbf{r})$  modulate by the individual scaling parameters  $A_i$  and  $B_i$ . This theoretical framework was particularly interesting for two reasons. First and foremost, these studies raised the interesting possibility that motion control by the brain and by spinal reflexes might synthesize the range of dynamic force patterns needed for real limb behaviors by recruiting, scaling, and combining the force-field primitives. Second, the theory showed arbitrary smooth fields could only be created by inclusion of nonconservative fields in the basis set (i.e., by inclusion of the second summation term  $\Psi_i[\mathbf{r}]$  in equation 11.2). Nonconservative fields could be used as energy sources; however, nonconservative fields might present a serious challenge to stability of the limb.

The repertoire of force patterns that could be constructed from primitives embedded in the spinal cord all appeared to be solely conservative. Was this a failing or a design feature? Colgate and Hogan have shown that the class of conservative and passive systems could be coupled to other passive systems without danger of instability. The set of isometric fields that could be constructed by combinations of frog's primitives belong to this class. More recently, the stability of classes of nonlinear systems created from linear combinations of stable nonlinear dynamical systems has also been demonstrated by Lohmiller and Slotine (1998). These results suggested it was feasible to design a simple principled motion control based around dynamic force-field primitives (Mussa-Ivaldi and Giszter, 1992). Primitives could be recruited and combined by sets of planning, patterning, and control systems. Further,

this design might offer specific stability guarantees during interactions with different environments. The laboratories of Mussa-Ivaldi (1997, 1999), Bizzi, and Shadmehr, among others, have examined biological motor learning and developed algorithms and descriptions of motor adaptation using this framework (see Shadmehr and Mussa-Ivaldi, 1994). The task for the motor system that is learned is the tracking of an arbitrary (usually straight) trajectory in different mechanical and visual environments. More-complex tasks, such as ball catching, ball bouncing, or pole balancing, could also be used.

Recently, we (Kargo and Giszter, 2000a; Giszter and Kargo, 2000a) observed in our experiments that not only field structure appeared conserved in spinal behaviors. The temporal dynamics of activation of primitives also seemed conserved in different spinal behaviors and corrections. Our data from the frog currently support the idea that all primitives used in a (spinal reflex) behavior can be described as having similar temporal dynamics—that is, they are driven by a common function  $a(t)$ . Similar results can be found in human motion (Burdet and Milner, 1998; Doreinger and Hogan, 1998; Krebs et al., 1999). We currently believe that during motion control in the spinal frog, any primitive  $\theta_i$  can be expressed in the form

$$\theta_i(\mathbf{r}, \dot{\mathbf{r}}, t) = A_i \cdot a(t) \phi_i(\mathbf{r}, \dot{\mathbf{r}}) \quad (11.3)$$

Where  $a(t)$  is similar in all primitives, (resembling a fixed half-cycle oscillation, or a fixed-impulse response, such as a cosine packet) and  $A_i$  is a scaling factor. This structure constrains the method of construction of a time-varying force field  $F(\mathbf{r}, \dot{\mathbf{r}}, t)$ . The free parameters for generation of a behavior are the selection of the scalings  $A_i$  and the phasings  $\tau_i$  of the component primitives. In this formulation, the motion control field  $F(\mathbf{r}, \dot{\mathbf{r}}, t)$  is constructed as:

$$F(\mathbf{r}, \dot{\mathbf{r}}, t) = \sum_i A_i a(t + \tau_i) \phi_i(\mathbf{r}, \dot{\mathbf{r}}) \quad (11.4)$$

### Data Supporting the Decomposition of Behavior into Primitives

It is worth reviewing the data that support the decomposition in equation 11.4 above. Giszter and Kargo have now amassed several sets of experiments supporting this framework. The strong predictions of the framework are (1) that new primitives added to a behavior should exhibit dynamics  $a(t)$ , (2) that primitives deleted from a behavior should cause loss of patterns of the same temporal form  $a(t)$ , and (3) that feedback modulation of primitives by natural or artificial stimuli should cause independent alterations of amplitude ( $A_i$ ) or phase ( $\tau_i$ ) but not alter  $a(t)$ . Each of these is observed in controlled experiments. In collision with obstacles, frogs make trajectory corrections by adding in an amplitude-modulated and externally triggered primitive with an

identical temporal pattern at different phases of the motion (Kargo and Giszter, 2000a). In immobilized wipes, deletions of phases occur (Giszter and Kargo, 2000; Kargo, 2000). Each deleted phase has similar temporal dynamics. In figure 11.1, we compare data from deletion of a hip extension primitive reported previously by Giszter and Kargo (2000) and data from deletion of knee extension reported by Kargo (2000). Kargo (2000) also describes effects of muscle and tendon vibration modifying proprioceptive feedback control signals. These manipulations alter the phase and amplitude of primitives while preserving temporal activation patterns. Moving beyond purely spinal motor patterns, D'Avella and Bizzi (1998) have demonstrated decomposition of more-complex brain-driven (vestibular system) behaviors into force-field elements. Use of the ideas of primitives to describe data in human motor learning has also been highly successful. Tresch and Bizzi (1999) in rat and in cat, Lemay and Grill (see Giszter et al., 2000b) have shown that force-field primitives exist in the mammalian spinal cord and are not a quirk of amphibian evolution. Perhaps the most interesting issue in this area as a whole is the relation among the various primitives in these several lines of research.

The framework expressed in equation 11.4 might be adequate for reflex behaviors, but seems limited. At the level of the set of execution primitives, how could motor learning be developed from such a starting point or bootstrap?

In motor-skill development and learning it seems likely that there should be the capacity to fully explore the degrees of freedom of the motor apparatus. Each primitive reduces the potential degrees of freedom implicit in the  $n$  muscles, constraining the primitive to a single field. Similarly, the activation of the primitive is constrained to a fixed pattern. It is likely that very small time increments  $\tau_i$  are not possible. The value of the primitive may be to place the motor learning system within an initial useful domain of the degrees of freedom, enabling initial adaptive survival skills to execute.

We speculate that from this initial constrained and adapted framework, motor-skill learning elaborates more-flexible controls. Our perspective on the spinal primitives described above in lower vertebrates and in rats and cats is thus that these form a bootstrap system for motor development and motor learning. The idea is that a collection of primitives defined by evolution "seed" subsequent motor development.

The natural question arising from this idea, for which there is currently no answer, is whether in mammals, the primitives that may constitute such a bootstrap survive or dissolve in the subsequent motor learning and plastic changes. This is an important issue from the point of view of design of a neuroprosthetic or robotic interface: Is a bootstrap still available in the spinal cord or expected by cortex? Does each adult cortex expect different "individualized" primitives as a result of

ontogenetic modifications? Does a neuromimetic robot of mammalian capability need a similar bootstrap set of primitives? How should the heterarchy of the robot controller recruit and be allowed to alter primitives? These are important questions; however, these issues have only a peripheral impact on a first generation of neuromimetic robots based on these ideas. Below, we outline several biologically plausible arrangements of control and methods for motor learning using primitives as we currently understand them.

## BUILDING A CONTROLLER USING PRIMITIVES

How are primitives controlled in the nervous system to provide a universal approximation system? How can this be implemented in a neuromimetic robotic controller?

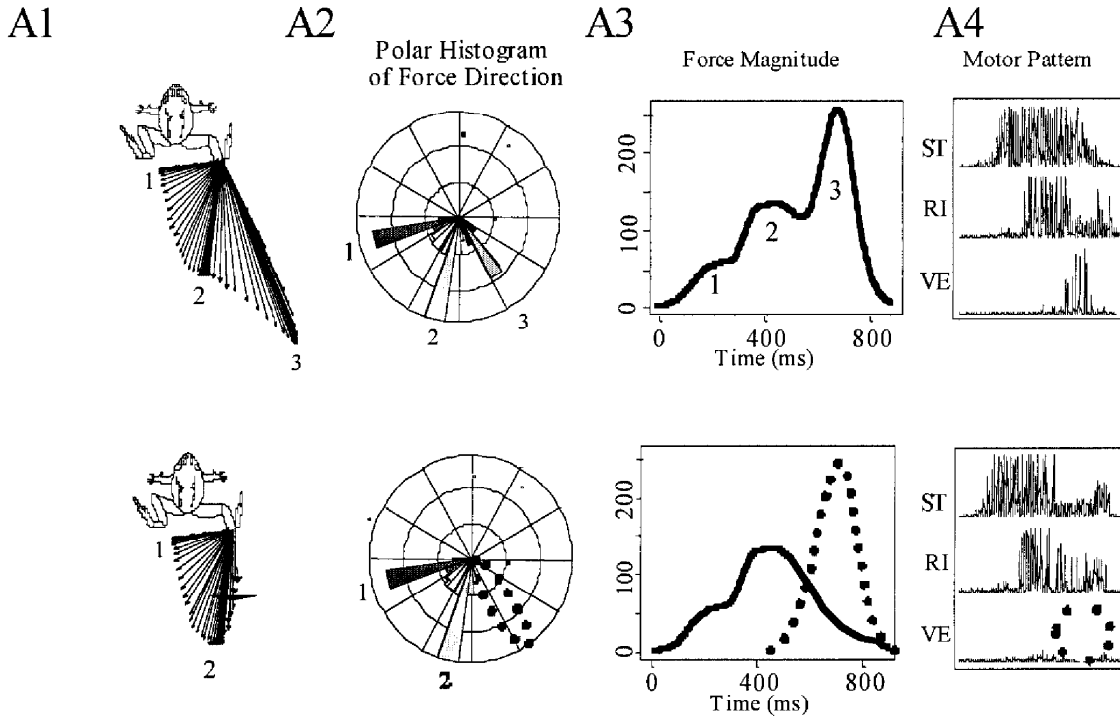
Thinking about this from the biological perspective tends to revolve around the relationship of the primitive to the pattern generator and the role of oscillatory processes in this relationship. The task in equation 11.4 is to sequence, combine, amplitude modulate, and perhaps time dilate the primitives. This could be related to present ideas of pattern generation in one of several ways, as outlined in figure 11.2. (1) The primitives could form a pattern-shaping system driven by the rhythm/pattern generator. This framework is presaged in the descriptions of Grillner and Jordan's burst generators, and has been given formal shape in an engineering application by Abbas. (2) The primitives could form a constitutive part or modules of the pattern generator and the timing function  $a(t)$  could represent intrinsic properties that contribute to the rhythm organization with perhaps different modules dominating control of rhythmicity in different contexts. This scheme is supported by Stein in his shared core notion. (3) Primitives could be an emergent and task-dependent functional organization of a more or less monolithic CPG. (4) Rhythmicity could emerge from a collection of modules interacting in a network to form series of action decisions, with an emphasis on the computation and planning of motion. This scheme has been suggested by Prochazka, and a model implementing such a system was used in my own work. (5) Rhythmicity and sequence, the organization of pattern formation, and execution strength could be separated so a scheme similar to (4) operates with a (set of) watchdog timer(s) and timers for information access.

The value of schemes 4 and 5 to the designer of a neuromorphic controller is the more intuitive, flexible, and modular design framework and the likelihood that it can be easily extended for new tasks and motor learning (e.g., backward walking in a cat could conceivably be achieved by loading of weights for a new set of action chains in the decision layer, while still using similar primitives and merely altering a few feedback pathways for the phase syntheses layer).

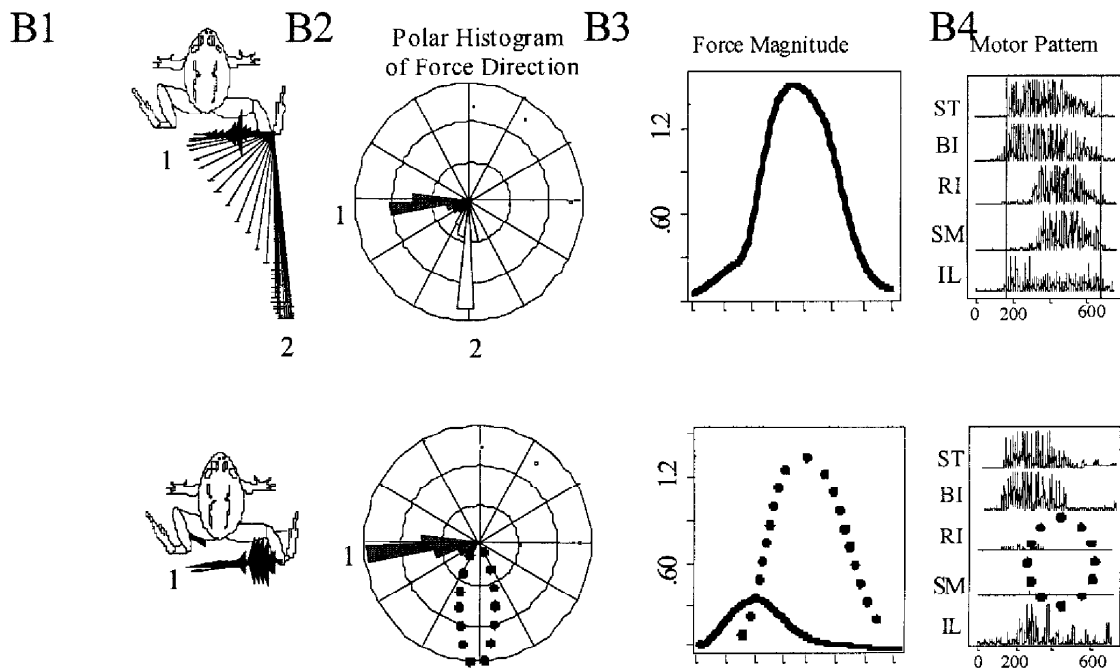
# Hindlimb to Hindlimb Wiping

Construction from force field primitives - deletion data

## A: Hindlimb to Hindlimb Wiping : knee extensor deletion



## B: Hindlimb to Hindlimb Wiping : hip extensor deletion



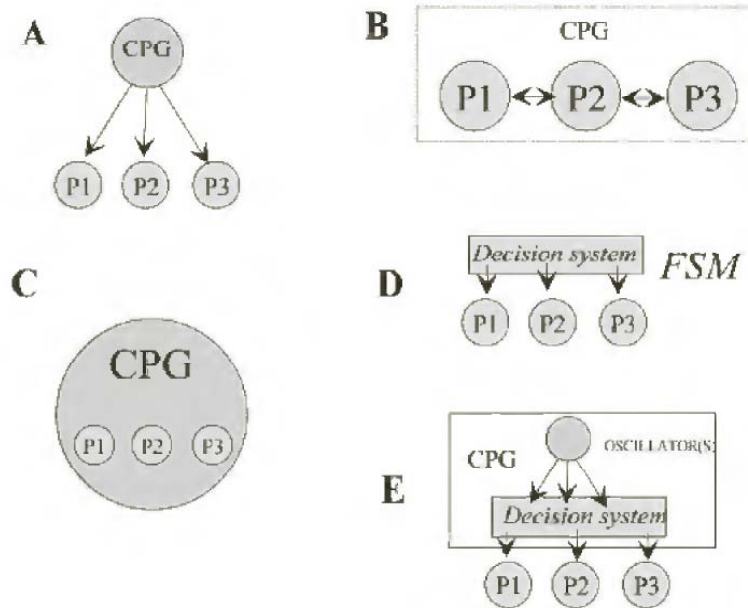
In comparing these schemes, which clearly vary in degree of explicit complexity, it is useful to consider the computational roles that oscillators might play in the task of organizing and controlling action and achieving perception for action. Oscillators offer precise phase locking and differing entrainment regimes and phase operations (e.g., see Ayers and Selverston, 1977, 1984). The oscillator need not be used simply as a driver system. Among other things, oscillators in engineering applications may be used for:

- Rhythm, waveform, pattern synthesis
- Phase locking and tracking
- Data synchronization
- Phase measurement, modulation, and demodulation
- Transponders and boosters
- Clocking computation cycles

In this context, it may be reasonable to consider these other roles of oscillators in relation to the operation of circuits that are operationally defined as a CPG by experiment in complex motor tasks. Several investigators, myself included, have argued about this relationship quite extensively. Taking this into account and with the caveat that various stability-of-oscillation issues must be addressed in any scheme, a modified version of the architecture in scheme 5 is likely to allow the most-flexible construction of motor actions, motor related perceptions, and the computations connecting the two and controlling and organizing them. This is outlined in figure 11.3. The use of a supervisory or driving rhythm generator above the decision and execution networks is particularly desirable for locomotion: It allows the decision cycles and execution elements to be driven at higher rates to match external mechanics and mission requirements in a classic CPG/limb phase-locking

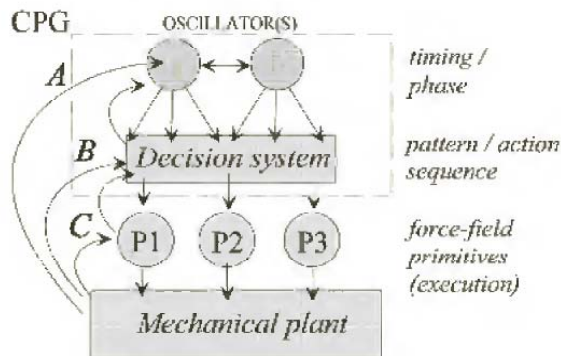
**Figure 11.1** Examples of deletion of primitives redrawn from Giszter and Kargo 2000 and Kargo 2000. (A) Deletion of the third knee extensor phase: A1 frog and forces recorded at ankle in relation to frog over time at a single limb posture (limb held immobile during force recording). A2 polar histogram of distribution of forces. Note three peaks reduced to two following deletion. A3 force magnitude over time for entire pattern (above) and after deletion (below). Shown dotted is the lost phase's magnitude. This was obtained by vector subtraction of deletion from full pattern. A4 electromyogram recordings of critical muscles: note deletion of vastus externus (VE) burst. (B) Deletion of the second hip extensor phase in a second frog in which knee extension was already deleted: B1 frog and forces recorded at ankle in relation to frog over time. B2 polar histogram of distribution of forces. Note two peaks reduced to one following deletion. B3 force magnitude over time for entire pattern (above) and after deletion (below). Shown dotted is lost phase's magnitude. This was obtained by vector subtraction of deletion from full pattern. B4 electromyogram recordings of critical muscles: note deletion of RI (rectus internus) and SM (semimembranosus) burst. ST, semitendinosus; BI, biceps (iliofibularis); II, iliopsoas; RI, rectus internus; SM, semimembranosus.





**Figure 11.2** Possible relationships between the CPG and the collections of biomechanical primitives discussed in the text. (A) The CPG recruits primitives in a feedforward fashion and rhythmicity and execution are separable. (B) The CPG is comprised of the interactions among primitives and little else. Phase and amplitude controls are integrated in the same network. (C) Primitives that appear functionally independent are an emergent property of a monolithic network that comprises the CPG. (D) The CPG comprises finite state machine sequencing primitives that are recruited in a feedforward pattern. Rhythmicity emerges from the state transitions. (E) A rhythm generating oscillator layer relays phase information and “suggestions” or “imperatives” to a decision and pattern formation network that decides on sequences and combinations of primitives based on task and execution space exigencies. For data in support of discrete decisions re execution of primitives and their combinations independent of rhythm construction see figure 11.1 and Kargo and Giszter 2000a and Giszter and Kargo 2000a.

process. Further, computation of decisions in the middle layer may incorporate both internal phase measurements, internal/external phase disparities, and kinematic and kinetic execution measures to control action sequences and combinations. One example of such a framework is an extension and augmentation of a scheme originally proposed for description of frog wiping behavior, based around Maes networks (Giszter, 1993, 1994; Giszter and Kargo, 2000b). The experimental data we now have in frogs (Kargo and Giszter, 2000b) forces us beyond the earlier simple scheme to a system that explicitly considers timing globally as well as local to a primitive. The spinal frog executes constant time trajectories from all areas of the work space in wiping. It does this using constant-duration primitives that are regulated in their phase and amplitude. Sensory information from both proprioception and contact sensing is used in on-line decisions that utilize phase information as part of their decision and control criteria. The new formulation of



**Figure 11.3** Proposed architecture for a locomotor/protective controller. A collection of oscillators (timing/phase layer) provide data clocking and synchronization, phase measurements and watchdog timing to a *decision system* that generates pattern action sequencing and interacts with and recruits oscillators (pattern/action sequence layer). This decision layer creates actions by combining and sequencing viscoelastic force-field primitives in the execution layer according to the task demands and effector system choices in place. Networks of the types proposed by Maes (1990) are suitable candidates for the decision layer. The recruited primitives interact with and sum at the mechanical plant to generate the driving and stabilizing force-fields needed to complete the task. Primitives have local feedback processes and may pass information local to their execution (C) back to the decision layer. The decision layer also receives information from other sensory modalities directly from in the plant itself (B). The decision layer receives and integrates information from the mechanical plant, the primitives and the timing, phase synthesis and phase tracking system. This layer receives information from the decision layer and from the mechanical plant in order to track phase and disparity in timing. Sensory information thus acts directly at the oscillator layer or is relayed through the decision layer (A). The system thus computes decisions both with phase information (layers 1 and 2), and with pattern/configuration/force information (layers 2 and 3). This system clearly allows flexible recombination of execution elements e.g. to support execution of forward and backward locomotion from similar gait patterns. It can be extended to different tasks and may be organized to allow separable controls of timing and execution needs based on load, and implement various task dependent optimizations of timing, execution and degree of precision and use of feedback.

the earlier model is powerful. However, it is important to ensure that the combined system coupling of the oscillators in the rhythmicity construction to the execution primitives does not violate the stability guarantees of either. The system can be scaled to provide simple low-level multijoint motor acts such as wiping, whole body multilimb acts such as locomotion, and a substrate for learning novel motor acts as new tasks are required. The temporal constraints on the pattern formation/decision layer can be severe, leading to a condition that is nearly feed forward (gallop or jump) or loose, allowing the decision layer to dominate timing unless some deadlock or ambiguity does not allow timely decisions to occur.

The use of a decision layer to pick a chain of actions defined by force-field combination patterns provides a kind of internal model and

credit assignment for the controller in a manner that may be harder in a pure oscillator system. It is easy to imagine using a Q-type learning reinforcement in the decision layer and optimizing the rhythm and primitive layers for minimal energy expenditure. In addition, it is clear that the patterns or decision chains embedded in the middle layer can be executed in different chunks and cadences in relation to the rhythm-generating system in different tasks and conditions.

The wiping-behavior simulation using this framework (Giszter, 1993, 1994) captured the flexibility of synthesis of switched, blended, and transitions of behaviors to multiple targets seen in real frogs' reflex scratches and could be tuned. However, it is now clear that the true performance of frogs embodies far-more-precise controls than the model. To date, we have not tried to emulate this with the architecture shown in figure 11.3.

## DISCUSSION AND CONCLUSION

The framework we propose is clearly a work in progress that requires both experimental and theoretical testing. In particular, from our perspective, the layered modularity satisfies a need for potentially separable controls of motion planning, of internal models implicit in the framework, of execution amplitude, and of timing and phasing. At the same time, these can be integrated and motion optimized while compromising among the competing requirements of each in a divide-and-conquer fashion. Any scheme must provide accounts of (1) efficient task completion and energetics with different loads, including coarticulation and blended motions; (2) provide adaptable adjustments of locomotory swing, stance, placement, and obstacles in locomotion; (3) provide for precise starting, stopping, turning, evasion, and pursuit; and (4) provide the capacity to generate and integrate defensive grooming, offensive reaching, aimed limb or body trajectories, and other integrated patterns of the biomimetic system. A system that can achieve these is likely to be able to complete the mission of autonomous self-preservation in a potentially hazardous and hostile environment. This is a minimal prerequisite to the more-cognitively defined mission goals such as foraging, reconnaissance, or payload transport to a specific target.

## ACKNOWLEDGMENTS

Discussions with Drs. S. Schaal, D. J. Sternad, T. D. Sanger, W. J. Kargo, J. Ayers, and several others and numerous discussions at the Neuro-mimetic Robotics meeting have proven useful in formulating many of the ideas presented and we thank all concerned for their generous interactions. Supported by NIH NS34640 to SFG, and NIH 5T32 HD07467.

## REFERENCES

- Abbas, J., and Full, R. J. (2000). Neuromechanical interaction in cyclic movements. In J. Winter and P. E. Crago (eds.), *Biomechanics and Neural Control of Posture and Movement*. New York: Springer Verlag, pp. 334–346.
- Ayers, J. L., and Selverston, A. I. (1977). Synaptic control of an endogenous pacemaker network. *J. Physiol. Paris* 73(4): 453–461.
- Ayers, J. L., and Selverston, A. I. (1984). Synaptic perturbation and entrainment of gastric mill rhythm of the spiny lobster. *J. Neurophysiol.* 51(1): 113–125.
- Berkinblit, M. B., Feldman, A. G., and Fukson, O. I. (1986). Adaptability of innate motor patterns and motor control mechanisms. *Behav. Brain Sci.* 9: 585–638.
- Bernstein, N. (1967). *The Coordination and Regulation of Movements*. Oxford: Pergamon.
- Bizzi, E., Giszter, S. F., Loeb, E., Mussa-Ivaldi, F. A., and Saltiel, P. (1995). Modular organization of motor behavior in the frog's spinal cord. *Trends Neurosci.* 18(10): 442–446.
- Bizzi, E., Mussa-Ivaldi, F. A., Giszter, S. F. (1991). Computations underlying the execution of movement: A novel biological perspective. *Science* 253: 287–291.
- Burdet, E., and Milner, T. E. (1998). Quantization of human motions and learning of accurate movements. *Biol. Cybern.* 78(4): 307–318.
- Cannon, M., and Slotine, J. J. E. (1995). Space frequency localized basis networks for nonlinear estimation and control. *Neurocomputing* 9(3): 293–342.
- Chapin, J. K., Moxon, K. A., Markowitz, R. S., and Nicolelis, M. A. L. (1999). Real-time control of a robot arm using simultaneously recorded neurons in the motor cortex. *Nat. Neurosci.* 2(7): 664–670.
- Cruse, H. (1996). *Neural Networks as Cybernetic Systems*. Stuttgart: Thieme Verlag.
- d'Avella, A., and Bizzi, E. (1998). Low dimensionality of supraspinally induced force fields. *Proc. Natl. Acad. Sci. U.S.A.* 95(13): 7711–7714.
- Dietz, V., Colombo, G., Jensen, D. M., and Baumgartner, L. (1995). Locomotor capacity of spinal cord in paraplegic patients. *Ann. Neurol.* 37: 574–582.
- Doeringer, J. A., and Hogan, N. (1998). Intermittancy in preplanned elbow movements persist in the absence of visual feedback. *J. Neurophysiol.* 80(4): 1787–1799.
- Full, R. J., and Koditschek, D. E. (1999). Templates and anchors: Neuromechanical hypotheses of legged locomotion on land. (Review.) *J. Exp. Biol.* 202, pt. 23 (December): 3325–3332.
- Giszter, S. F. (1993). Behavior networks and force fields for simulating spinal reflex behaviors of the frog. In J.-A. Meyer, H. L. Roitblat, and S. W. Wilson (eds.), *From Animals to Animats 2: Proceedings of the Second International Conference on Simulation of Adaptive Behavior*. Cambridge: MIT Press, 172–181.
- Giszter, S. F. (1994). Reinforcement tuning of action synthesis and selection in a virtual frog. In D. Cliff (ed.), *From Animals to Animats 3: Proceedings of the Third International Conference on Simulation of Adaptive Behavior*. Cambridge: MIT Press, 291–301.
- Giszter, S. F., Grill, W., Lemay, M., Mushahwr, V., and Prochazka, A. (2000b). Intraspinal microstimulation: Techniques, perspectives, and prospects for FES. In K. A. Moxon and J. K. Chapin (eds.), *Neural Prostheses for Restoration of Sensory and Motor Function*. CRC Press, pp. 101–138.

- Giszter, S. F., and Kargo, W. J. (2000a). Conserved temporal dynamics and vector superposition of primitives in frog wiping reflexes during spontaneous extensor deletions. *Neurocomputing* 32–33: 775–783.
- Giszter, S. F., and Kargo, W. J. (2000b). Movement organization in the frog spinal cord: Prerational intelligence? In *Prerational Intelligence: Adaptive Behavior and Intelligent Systems without Symbols and Logic*. Vol. 1. Springer-Verlag, pp. 323–342.
- Giszter, S. F., Kargo, W. J., and Davies, M. R. (1996). Motor Primitives in the spinal cord as a basis for motor learning and action. In M. Mataric (chair), *Embodied Cognition and Action: Papers of 1996 AAAI Fall Symposium*. AAAI Press, technical report FS-96-02.
- Giszter, S. F., Kargo, W. J., and Davies, M. R. (2000). Augmenting postural primitives in spinal cord: Dynamic force-field structures used in trajectory generation. In J. Winter and P. J. Crago (eds.), *Biomechanics and Neural Control of Movement*. Springer Verlag, pp. 334–346.
- Giszter, S. F., Loeb, E., Mussa-Ivaldi, F. A., and Bizzi, E. (2000a). Dense mapping of frog lumbar spinal cord: Organization of force and muscle use. *Human Movement Science* (in press).
- Giszter, S. F., Moxon, K. A., Rybak, I., and Chapin, J. K. (2000c). A neurobiological perspective on design of humanoid robots and their components. *IEEE Intelligent Syst.* 15(4): 64–69.
- Giszter, S. F., Mussa-Ivaldi, F. A., and Bizzi, E. (1993). Convergent force field organized in the frog's spinal cord. *J. Neurosci.* 13: 467–491.
- Grillner, S., and Wallen, P. (1985). Central pattern generators for locomotion, with special reference to vertebrates. *Annu. Rev. Neurosci.* 8: 233–261.
- Kargo, W. J. (2000). Afferent regulation of force-field primitives during real limb behaviors. Ph.D. thesis, Neurobiology and Anatomy MCPHU School of Medicine, Philadelphia.
- Kargo, W. J., and Giszter, S. F. (2000a). Rapid correction of aimed movements by summation of force-field primitives. *J. Neurosci.* 20(1): 409–426.
- Kargo, W. J., and Giszter, S. F. (2000b). Afferent roles in hindlimb wipe-reflex trajectories: Free-limb kinematics and motor patterns. *J. Neurophysiol.* 83: 1480–1501.
- Kawato, M. (1999). Internal models for motor control and trajectory planning. *Curr. Opin. Neurobiol.* 9: 718–727.
- Kiehn, O., Hounsgaard, J., and Sillar, K. (1997). Basic building blocks of vertebrate spinal central pattern generators. In P. Stein, S. Grillner, A. Selverston, and D. Stuart (eds.), *Neurons, Networks, and Motor Behavior*. Cambridge: MIT Press, pp. 47–60.
- Krebs, H. I., Aisen, M. L., Volpe, B. T., and Hogan, N. (1999). Quantization of continuous arm movements in humans with brain injury. *Proc. Natl. Acad. Sci. U.S.A.* 96(8): 4645–4649.
- Loeb, E., Giszter, S. F., Borghesani, P., and Bizzi, E. (1993). The role of afference in convergent force fields elicited in the frog spinal cord. *Somatosens. Mot. Res.* 10: 81–95.
- Loeb, E. P., Giszter, S. F., Saltiel, P., Mussa-Ivaldi, F. A., and Bizzi, E. (1999). Output units of motor behavior: An experimental and modeling study. *J. Cogn. Neurosci.* 12: 1–20.
- Lohmiller, W., and Slotine, J. J. (1998). On contraction analysis for non-linear systems. *Automatica* 34: 683–696.
- Maes, P. (1990). Situated agents can have goals. *Robotics Autonomous Syst.* 6: 49–70.

- Marr, D. (1982). *Vision*. San Francisco: Freeman.
- Massaquoi, S. G., and Slotine, J. E. (1996). The intermediate cerebellum may function as a wave-variable processor. *Neurosci. Lett.* 215: 60–64.
- Mataric, M. J., Williamson, M. M., Demiris, J., and Molan, A. (1998). Behavior-based primitives for articulated control. In *From Animals to Animats 5: Proceedings of the Fifth International Conference on Simulation of Adaptive Behavior*. Cambridge: MIT Press, pp. 160–170.
- McFarland, D. J., and Bossert, T. (1993). *Intelligent Behavior in Animals and Robots*. Cambridge: MIT Press.
- McFarland, D. J., and Houston, A. (1981). *Quantitative Ethology: The State Space Approach*. New York: Pitman.
- Mussa-Ivaldi, F. A. (1992). From basis functions to basis fields: Vector field approximation from sparse data. *Biol. Cybern.* 67(6): 479–489.
- Mussa-Ivaldi, F. A. (1997). Nonlinear force fields: A distributed system of control primitives for representing and learning movements. In *Proceedings of the IEEE International Symposium on Computational Intelligence in Robotics and Automation*. Computer Society Press, pp. 84–90.
- Mussa-Ivaldi, F. A. (1999). Modular features of motor control and learning. *Curr. Opin. Neurobiol.* 9: 713–717.
- Mussa-Ivaldi, F. A., and Giszter, S. F. (1992). Vector field approximation: A computational paradigm for motor control and learning. *Biol. Cybern.* 67(6): 491–500.
- Mussa-Ivaldi, F. A., Giszter, S. F., and Bizzi, E. (1994). Linear superposition of primitives in motor control. *Proc. Natl. Acad. Sci. U.S.A.* 91: 7534–7538.
- Raibert, M. H. (1986). *Legged Robots that Balance*. Cambridge: MIT Press.
- Sanner, R., and Slotine, J. J. E. (1995). Stable robotic learning using “neural networks.” *Neural Comput.* 7(4): 753–790.
- Shadmehr, R., and Mussa-Ivaldi, F. A. (1994). Adaptive representation of dynamics during learning of a motor task. *J. Neurosci.* 14: 3208–3224.
- Sherrington, C. S. (1910). Flexion-reflex of the limb, crossed extension reflex, and reflex stepping and standing. *J. Physiol.* 40: 28–121.
- Tresch, M. C., and Bizzi, E. (1999). Responses from spinal microstimulation in the chronically spinalized rat and their relationship to spinal systems activated by low threshold cutaneous stimulation. *Exp. Brain Res.* 129: 401–416.
- Vicario, D. S., and Ghez, C. (1984). The control of rapid limb movement in the cat: IV. Updating of ongoing isometric responses. *Exp. Brain Res.* 55: 134–144.
- Williamson, M. M. (1996). Postural primitives: interactive behavior for a humanoid robot arm. In *From Animals to Animats 4: Proceedings of the Fourth International Conference on Simulation of Adaptive Behavior*. Cambridge: MIT Press, pp. 124–135.

**This page intentionally left blank**

# 12

## A Conservative Biomimetic Control Architecture for Autonomous Underwater Robots


Joseph Ayers

We are adapting the command neuron, coordinating neuron, central-pattern-generator (CCCPG) model of the organization of motor systems to the control of autonomous robots. This model has been developed over the past 30 years to explain the generation of behavior by central neural networks (Kennedy and Davis, 1977; Selverston et al., 1980; Stein et al., 1997). According to this model, the behavior of the segmental motor apparatus is generated by central-pattern-generating networks composed of neuronal oscillators (Selverston, 1980; Pinsker and Ayers, 1983; Jacklet, 1987), coordinating neurons (Stein, 1974), and command neurons (Larimer, 1978; Kupferman and Weiss, 1978).

Command neurons are the organizers of behavioral acts and in some cases may be considered to be the locus at which the decision to execute a behavior is made (Kovac and Davis, 1980). During operation, commands may trigger a segmental oscillator to generate a behavior that outlasts the stimulus or *gate* on a behavior during its execution through continuous operation throughout (Stein, 1978). Although some commands may be unitary at the higher levels of the central nervous system (CNS) (Bowerman and Larimer, 1974a, b), there is strong evidence that they are employed in groups (Davis and Kennedy, 1972; Larimer, 1988). The segmental oscillators may be divided into a neuronal oscillator level that controls the timing of the rhythm and a pattern-generating level that specifies details of the final output (e.g., forward vs. lateral walking, etc.). Coordinating systems provide a governed oscillator with information about the timing of its operation relative to a governing oscillator (Ayers and Selverston, 1979). These signals reset the timing of the oscillation to maintain gait (Stein, 1978). Commands may be the locus at which exteroceptive feedback operates to mediate orientational reflexes. Other sensory feedback operates at the motor neuron level to modulate the amplitude of behavioral acts or at the oscillator level to mediate the timing of the output (Stein, 1978).

We have selected two model systems, the lobster and lamprey, which share this underlying architecture and feature reactive autonomy as their underlying behavioral control architecture (figure 12.1). We hold






**Figure 12.1** Biomimetic robot physical plants. (Top) Lobster-based robot. The electronics interfaces are housed in a Plexiglas electronics bay. Which forms the hull of the vehicle. Eight modularized walking leg subsystems are attached to the hull. Anterior claw-like hydrodynamic control surfaces are controlled by a dc motor and jackscrew system. The posterior abdomen-like hydrodynamic control surface is also controlled by a motorized jackscrew system. Paired MEMS antennae are controlled in the yaw plane by a motor/crown and pinion system. A dorsally mounted sonar transponder allows bi-directional acoustic communications. (Bottom) Lamprey-based robot. A water tight electronics bay houses the same circuit boards used in the lobster. Propulsion is mediated by a polyurethane “notocord” divided into segments by Teflon™ vertebrae and which support axial nitinol muscle modules.

that the conservative control architecture of animal motor systems lends itself to a simple modularized robotic architecture that can be adapted to a variety of habitats.

We employ a biomimesis process to implement robots based on animal models. This process allows us to share elements between different vehicles. The process involves five different organizational levels. First, the robots are based on a *biomorphic plant* that captures the biomechanical advantages of the model organism (splayed posture, etc.). Movements of the vehicles are mediated by *myomorphic actuators* that produce linear graded force in response to trains of current pulses. The *neural-circuit-based controller* is based on a simulation of a synthetic model of the operation of the animal CPGs (Ayers and Crisman, 1992).

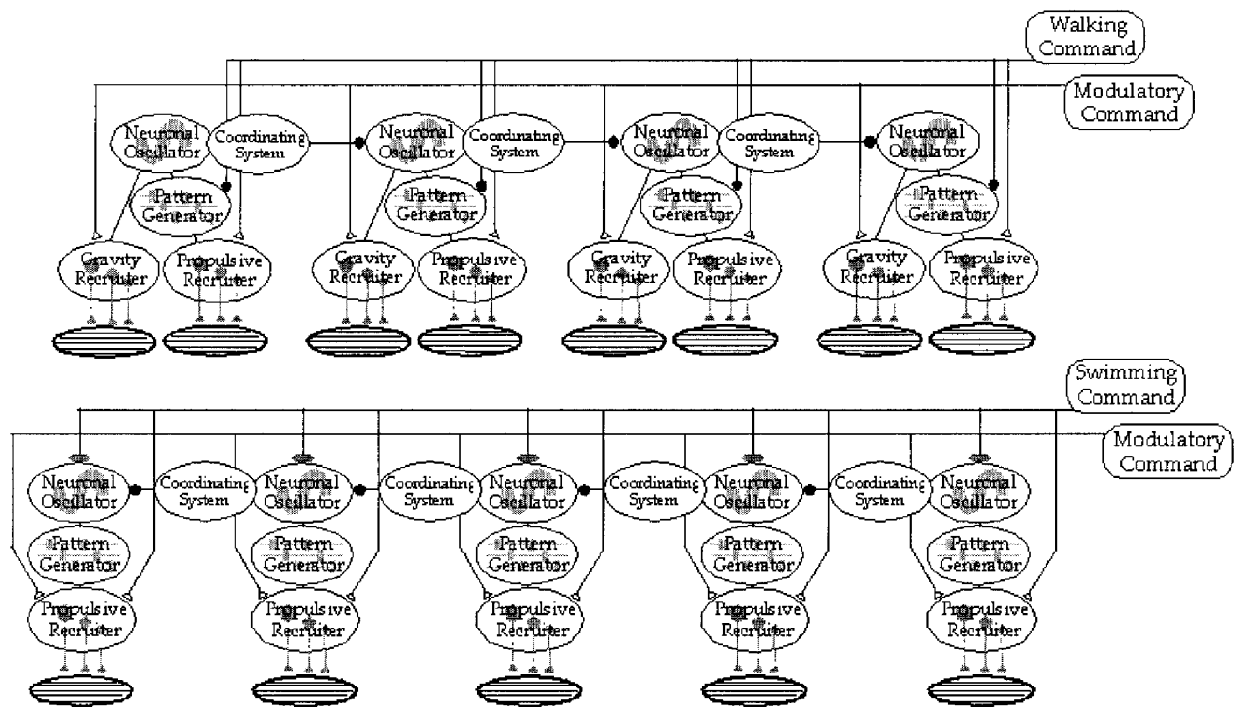


**Figure 12.2** Myomorphic actuators. (Left) Size-based recruitment of motor unit synergies. (Right) Pulse width modulation based force recruitment process. Control signals from the finite state machine gate on pulse trains with different duty cycles.

The processes that trigger or modulate behavior from environmental input or *releasers* are embodied in a labeled line code (Bullock, 1978) from *neuromorphic sensors* (McGruer et al., chapter 1 of this volume). Here, the sensory modalities we consider are orientation relative to gravity, direction, and contact. A reactive behavioral sequencer, which executes a behavioral library based on reverse-engineered behavior of the animal model, in turn controls these subsystems. The sequences specified from this behavioral library modulate the internal state variables that specify the ongoing behavior. This queue-based behavioral sequencer will reactively select behavioral sequences in response to sensor-mediated releasers. The following is a detailed description of how we derive a digital architecture from these neural functional elements.

## MYOMORPHIC ACTUATORS

The motor organization of our controller is based on central pattern generators activating motor unit synergies (Kennedy and Davis, 1977; Atwood, 1976). These systems grade force in the muscles by recruiting elements of a synergistic pool of motor neurons. In both invertebrates (Davis, 1971) and vertebrates (Henneman and Carpenter, 1965) this recruitment is on the basis of motor neuron size from smaller to larger, with the larger motor neurons having the greatest force-generating capability. In the lobster system, we distinguish two functionally different recruiter systems (figures 12.2–12.3). Propulsive recruiters act on synergies that produce translational propulsive force during locomotion.



**Figure 12.3** Functional components of CPGs. (Top) CPG for ambulatory robot. (Bottom) CPG for undulatory robot. Closed circles indicate inhibitory connections. Open triangles indicate excitatory connections.

These systems control the speed and thrust during walking. Anti-gravity recruiters, in contrast, act on the systems that support the body against gravity.

We realize this form of motor control by using artificial muscle formed from the shape memory alloy (SMA) nitinol (Dureing et al., 1990). The engineering aspects of these actuators are described in chapter 6 of this volume. From the controller perspective, nitinol can be activated approximately in the same fashion as muscle. We use trains of current pulses to activate nitinol to contract in the same way as trains of action potentials activate muscle. For each of the synergies, we provide three discrete levels of recruitment duty cycles analogous to the small, medium, and large motor neurons that characterize lobster neuromuscular systems (Kennedy and Davis, 1977). We grade the force of contraction in our artificial muscle by dividing the pulse-width duty cycle into three discrete levels as is found in lobster neuromuscular systems (Davis, 1971). During ongoing operation, the recruiter objects send messages to activate a given muscle at a given recruitment level. The kinematics of the shape memory activation cycle require that we divide the muscle activation cycle into a preheat phase, when the muscle is activated for a brief (ca. 150 msec) epoch to initially activate the muscle; then a contraction maintenance epoch, when the muscle is activated at the appropriate recruitment level for the ongoing activity;

and then a cooling period, when the muscle is cooled to the point where it can be mechanically elongated into the deformable state by the antagonist muscle.

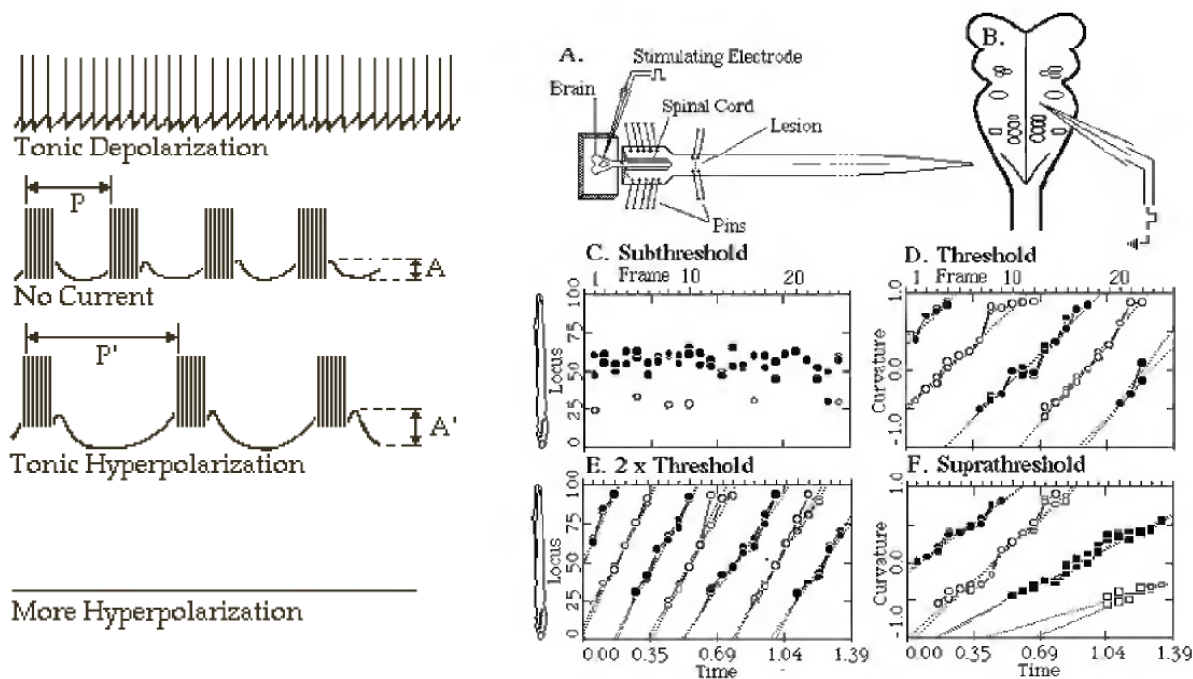
## NEURAL CIRCUIT-BASED CONTROLLER

The controller for our robots is organized at two levels. At the lowest level, a pattern-generating circuit generates the current pulse pattern that controls individual legs in the ambulator or segments of the body axis in the undulator (figure 12.3). At the highest level, the controller is organized around the behavioral libraries and a reactive sequencer. The sequencer maintains the state of the vehicle by switching between different command states in a temporal sequence specified by both orientation and magnitude of sensor input and the state sequences specified by the behavioral libraries.

The pattern-generating circuit we employ in our robots is indicated in figure 12.3. The segmental pattern generators include software objects for neuronal oscillation, pattern generation, and recruitment. These functional elements are implemented in the controller as software objects that pass messages contingent on the internal state variables of the controller. As indicated in figure 12.3, the components utilized for the ambulation controller can be reorganized to create an undulation controller. The primary differences are the antigravity recruiter and the polarity of the coordinating systems. The coordinating function can occur in both directions—that is, forward versus backward walking in the lobster or swimming versus backward crawling in the lamprey. Descending controls include commands that specify the direction of walking and modulatory commands that act upon the antigravity synergies to control height, pitch, and roll.

### Command Systems

Crustacean motor systems include single descending interneurons that can be stimulated to evoke recognizable behavioral acts (Bowerman and Larimer, 1974a, b). The action patterns evoked by these neurons therefore represent fundamental atoms of behavior (Barlow, 1968). Command systems are also present in the lamprey, where the mesencephalic locomotor system can be stimulated to initiate swimming (Ayers, 1989). Command systems can be either postural (Bowerman and Larimer, 1974a) or phasic (Bowerman and Larimer, 1974b), causing repetitive behaviors such as walking and swimmeret beating. In our robotic architecture, command systems exhibit several forms of control including parametric modulation, pattern selection, postural control, and superposition to generate hybrid behavioral acts. A detailed discussion of our implementation of these forms of control follows.

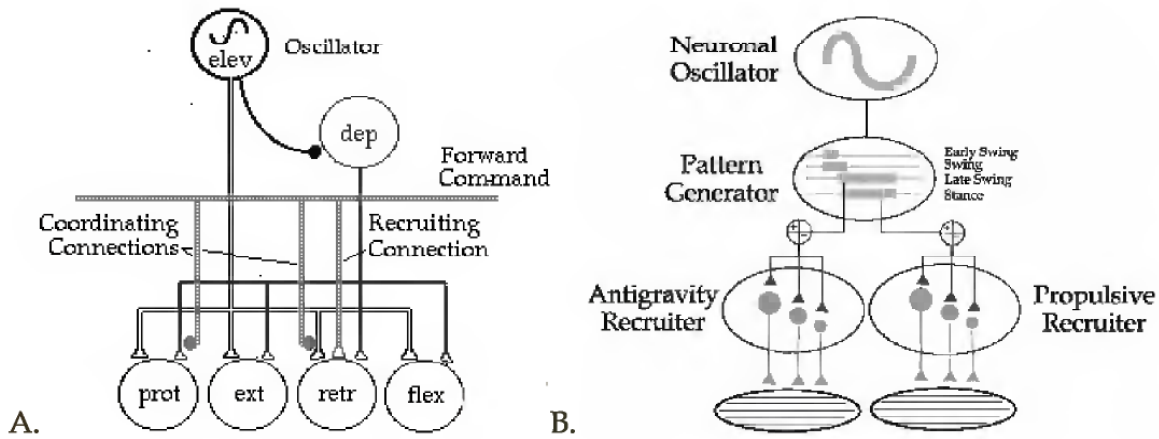


**A.** **B.**

**Figure 12.4** Parametric modulation. (A) Responses of an endogenous burster neuron to different levels on applied current. (After Pinsker and Ayers, 1983.) (B) Parametric control of lamprey swimming by the focal microstimulation of the mesencephalic locomotor region. (After Ayers, 1989.)

**Parametric Modulation** Parametric control by command systems initiates behavior acts mediated by central pattern generators (Pinsker and Ayers, 1983). Parametric modulation can control both the frequency and the amplitude of a rhythm. In *in vitro* models, parametric modulation can be mediated by current injected into a single neuron (figure 12.4A), electrical microstimulation of command systems (figure 12.4B), or bath application of neuromodulators (Harris-Warrick, 1990). In general, where parametric modulation occurs, there is a direct relationship between the frequency of a rhythm and its amplitude (Kennedy and Davis, 1977). In our robotic implementations, parametric control is mediated by control of an internal-state variable designated as *speed*. Speed determines the period of both stepping in the ambulator and axial undulations in the undulator. Speed also determines the recruitment level associated with the propulsive force actuators. In general, we quantized speed into three or four different magnitudes.

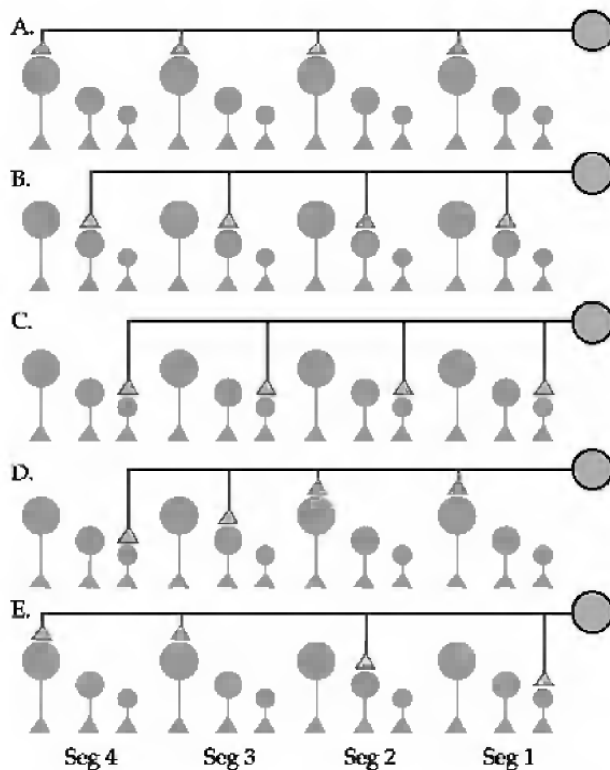
**Pattern Specification** One of the unique features of the lobster walking system is the ability to change its direction of walking on a step-by-step basis. This capability implies that there is a pattern-specifying level in the pattern generator that determines the direction of walking. We



**Figure 12.5** Network and functional models of the limb controllers. (A) Neuronal circuit model of the command neuron control of forward walking. The command elements include coordinating connections that abolish the connections between inappropriate synergies by presynaptic inhibition. The recruiting connection acts on the synergy that produces translational propulsive force. Elev, elevator synergy; dep, depressor synergy; prot, protractor synergy; ext, extensor synergy; retr, retractor synergy; flex, flexor synergy. (B) Functional model of the walking system indicating the different functional objects including neuronal oscillator, pattern generator and recruiters.

published a network model (Ayers and Davis, 1977) based on presynaptic inhibition of the connectivity linking the elevator and depressor synergies with the bifunctional muscles that shift their phase of discharge relative to the elevator during walking in the different directions. (figure 12.5A). According to this model, commands for different walking directions specifically inhibit two inappropriate connections as well as mediate a recruiting connection (figure 12.5A).

We have generalized this neuronal circuit model to a functional model that is the basis of our finite-state machine-based limb controller (figure 12.5B). The neuronal oscillator divides the step into an early swing phase where the limb is elevated and a late swing phase where the limb is depressed and applies propulsive force. The late swing phase is manifest as a delay between the initiation of discharge of the depressors and the synergies that generate propulsive force as observed in walking lobsters (Ayers and Davis, 1977). A pattern generator object specifies the walking direction and embodies the presynaptic inhibitory logic (figure 12.5A). The pattern generators select the specific synergies that operate together during walking in the different directions. The elevators are active only during the early swing phase while the depressors are active during both late swing and stance. The bifunctional muscles that control protraction and retraction of the thoracocoxal joint and extension and flexion of the merocarpoalite joint are active in a somewhat asynchronous pattern. For example, during the swing phase of forward walking, the protractors are active during both the early and late swing phases while the retractors are active only during the stance phase.




**Figure 12.6** Postural interneuronal objects. The gray circles represent the small medium and large objects controlling the antigravity synergies in each segment. Larger circles indicate larger units that produce a greater amplitude of depression in that segment. The gray circles with black outlines indicate postural interneuronal objects that act on the depressor synergies to control pitch roll and height. Triangles indicate excitatory connectivity. Rostral is to the right. (A) High posture. (B) Normal posture. (C) Low posture. (D) Rostrum up posture. (E) Rostrum down posture.

Thus there are at least three timing synergies—early swing, late swing, and stance—and all motor synergies are active during one or two of these epochs.

**Postural Control** The limb controller includes recruiters for the antigravity (limb depressor) synergies as well as the propulsive synergy appropriate to the direction of walking (see below). These antigravity recruiters control the height, pitch, and roll of the vehicle by regulating the degree to which the limb is depressed in the different segments. Height, pitch, and roll are controlled by postural interneuronal objects (figure 12.6) that activate the antigravity recruiters in the different body segments. These antigravity recruiters act on different elements of the antigravity synergy, and can mediate bilateral asymmetries in the height commanded for the two sides. Only one of the five postural elements can be active at a time on each side.

**Superposition** One of the features of this parallel command architecture is the possibility of superimposing commands to generate hybrid





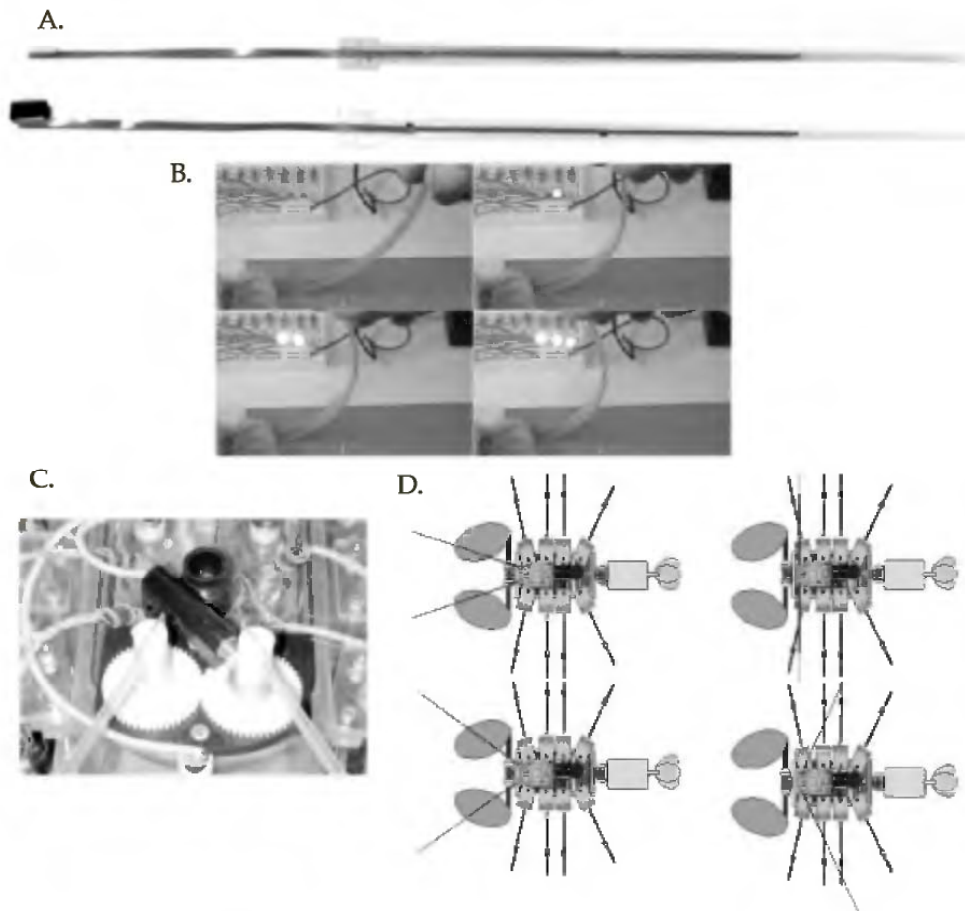
**Figure 12.7** Specification of the motor patterns for walking in the different directions. (A) Functional connectivity between command systems (F, forward; B, backward; L, lateral leading; T, lateral trailing). (B) Superposition of backward and lateral commands. In the left part of this panel only the backward commands are active. At the midpoint of the panel the leading and trailing commands are superimposed on the backward commands.

behavioral acts. For example, the only rule governing the interactions between walking commands is that the commands for walking in the opposite direction cannot be active at the same time (figure 12.7A). Command neuron objects can superimpose their effects. Figure 12.7B demonstrates the superposition of lateral walking commands on backward walking to generate diagonal walking. As described above, the posture of the hull can be controlled independently of the walking direction, as can the movements of the antennae, claws, and abdomen. Due to the capability for superposition, the vehicle has similar degrees of freedom to the animal models.

## NEUROMORPHIC SENSORS

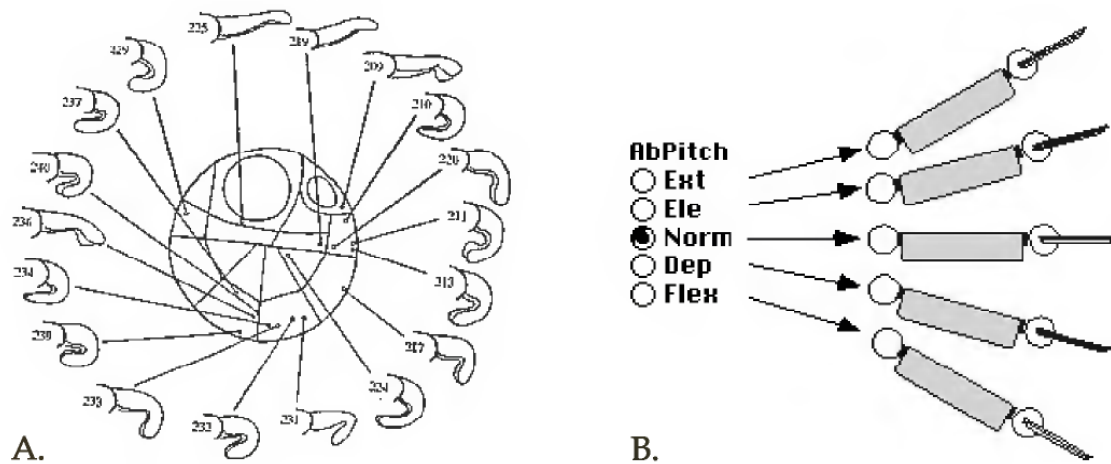
Crustacean sensors code environmental parameters in terms of labeled lines (Bullock, 1978). Each sensor is represented by an array of labeled line elements, each of which codes for a particular sensory modality (gravity, water current, etc.), as well as receptive fields (i.e., orientation relative to horizontal water currents from the front, rear, or sides, etc.). In the controller implementation, sensor objects pass messages to command objects based on the actual input. Thus the message contains the modality and orientation and evokes different methods from the command system, based on these parameters. During locomotion in currents and surge, lobsters predominately rely on three sensors, the antennae, water current receptors, and the statocysts or vestibular receptors. Although the actual biological sensors are complex, mechanical transducers can readily code environmental information in the same fashion as the lobster nervous system (figure 12.5).





**Figure 12.8** Operational aspects of MEMS antenna. (A) Two MEMS antennae demonstrating flex connectors. (B) Range fractionation of antenna bend. Note that as the antenna is bent to a greater degree LEDs powered by the MEMS switches indicate that the switches are progressively closed. (C) Motor drive used to move the antennae. (D) The four positions that the antennae move between.

Lobsters use their antennae extensively both as tactile receptors for collision avoidance and as water current sensors. In a robotic implementation, they require both active motor control as well as a set of labeled lines that code for contact as well as current forces. In our current implementation (see chapter 1 of this volume), the antennae are fabricated from a tapered beam of moderate restoring force (figure 12.8A). The antennal sensor is controlled in the yaw plane by a motor that either moves the antennae to one of four orientations relative to the hull axis or oscillates between two of four positions (figure 12.8B). All the switches code for both contact and flow. The end switches have very low flow sensitivity and maximum contact sensitivity, while the sensors near the base exhibit a combination (figure 12.8C–D). These switches implement a range fractionated labeled line code (Bullock, 1978) due to their having different cantilever lengths. Thus each switch constitutes a labeled line for the magnitude of bending of the antenna.



**Figure 12.9** Command-driven postures and abdomen behavioral quantizer. (A) Abdominal postures evoked by stimulation of different postural command neurons in the crayfish. This diagram indicates the locations of the command neuron axons in the brain connectives and the postural states evoked by each. (After Bowerman and Larimer, 1974a.) (B) Postural states used for quantizing abdominal posture in the lobster/robot.

### Variable Attention

We have implemented two analog sensors, a Densmore compass and an Analog Devices MEMS accelerometer. These sensors are housed on a logic board with a PIC16 microcontroller that converts the analog signals to quantized bit codes for the serial/sensor interface. The compass signal is converted to 8 bits that indicate one of 8 compass sectors and the accelerometers are converted to three bits that indicate one of 8 orientations in the roll and pitch plane. The controller has a polling interface that can request input from these sensors with a variable sample rate. In other words, the sample rate is different for different sensors and can be varied on the fly. Thus, for example, during an evoked turn, the sample rate can be increased to increase the preciseness of the turn.

### REACTIVE BEHAVIORAL LIBRARIES

The behavior of our robots is based on a library of action patterns derived from quantization of behavior observed in the animal models—the lobster and the lamprey. Sequences of the behavior of animals operating in different environments are maintained as digital movies and are subjected to stop-frame analysis. Our analysis resolves the command states from films of behavior.

### Finite-State Analysis of Behavior

Finite-state analysis assumes that behavior can be quantized into command states (figure 12.9). Further, it assumes that ongoing behavior is

based on dynamically changing sets of commands being active. The basic assumption of this analysis is that the posture and action of the different task groups is specified by a set of command neurons that instruct the task group to generate a different state (figure 12.9A). Thus the task of this analysis is to specify the state of the task groups in each of the frames of a movie. For example, we have found that 11 task groups that can exist in 2–5 different states can define the state of a lobster. During analysis of video sequences, a panel of buttons is used to specify the states of the different groups in each of the frames of a movie (figure 12.10).

We utilize the following states in our analysis

*Thorax pitch*: rostrum up, level, rostrum down.

*Thorax roll*: left up, level, right up

*Thorax yaw*: hard left, easy left, straight, easy right, hard right

*Thorax height*: high, normal, low

*Walking direction*: standing, forward, backward, lateral leading, lateral trailing

*Walking speed*: slow, medium, fast, stop

*Chelae pitch*: up, normal, down

*Chelae yaw* (left and right): extended, normal, meral spread, lateral spread

*Antennae yaw* (left and right): protracted, normal, lateral, retracted

*Uropod posture*: flared, normal, adducted

*Abdominal pitch*: extended, elevated, normal, depressed, flexed

To derive robotic control sequences from video data, we perform finite-state analysis task groups that mediate locomotion and searching individually to determine which synergistic sets are operant during different behavioral acts. Our analysis of the sequencing of these task groups borrows from a technique utilized by astronomers to detect motion of objects such as comets. As the analysis proceeds through each frame of the digital movie, the program flashes between temporally adjacent frames of the movie with a brief pause after each cycle. Appendages that are moving the most will flash in these projections. A panel of buttons that represent different states of the task groups (e.g., elevation vs. depression of the chelipeds, etc.) are available to the investigator to specify which groups are active. By clicking on the appropriate buttons for each frame, it is possible to efficiently quantify the activity of all task groups at high temporal resolution from video tapes of specimens behaving in a variety of situations. These state-change sequences constitute the behavioral library of the vehicles.

The controller can control all state variables including the posture of the claws, tail, and antennae using the same state-variable panel used

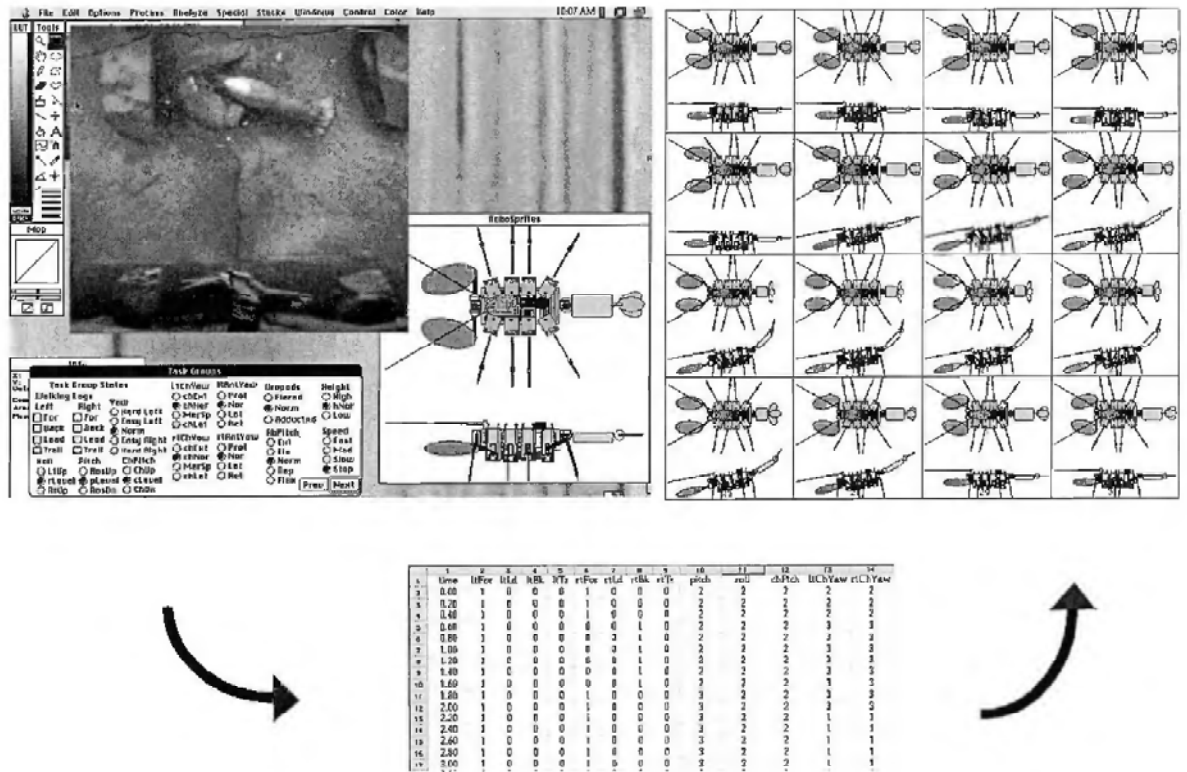


Figure 12.10 Behavioral quantization. (Top left) Screenshot of ColorImage in the process of decomposing a behavioral sequence into underlying commands. The output of ColorImage is a state matrix table (Bottom Panel) where each row is a frame of the movie and each column represents the state of one of the internal state variables. (Top right) Montage of the state changes underlying rheotactic behavior, reverse engineered from a movie.

for the behavioral reverse-engineering analysis. We have added the ability to adjust the pulse with modulation duty cycle “on the fly” through a modal dialog. This is essential for tuning the different postures and achieving the appropriate power levels for graded contractions of the SMA actuators.

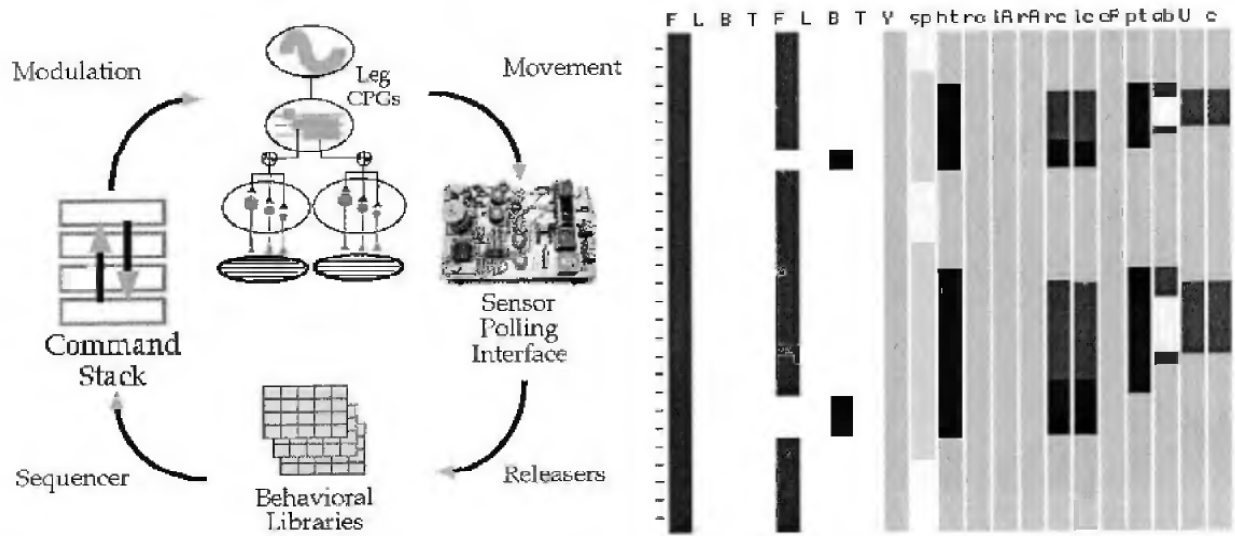
### **Behavioral Sequencer**

The overall control scheme we employ is supervised reactive autonomy. The behavioral set of the vehicles is based on a library of action patterns that are released by sensor input or sonar-based supervisory commands. During operation, the vehicle is given an underlying motivation (i.e., home in on a sonar beacon or initiate a search on an arbitrary compass heading). During execution of this motivated task, the controller will respond to environmental contingencies specified by its sensor suite by evoking preprogrammed sequences from its behavioral library.

The sequence of commands that underlie ongoing behavior is maintained on an event stack. Action components are pushed on this stack in order of time by releaser objects. Releaser objects are triggered by different methods of sensor objects that indicate orientational and magnitude components. Sensor objects receive input from an array of MEMS sensor switches and activate different releasers depending on the character and receptive field of the sensors activated by exteroceptive input. The sensor inputs are evaluated by bit masking to detect switch activation patterns resulting from an adequate stimulus.

**Releasers** A releaser is a software object generated by specific sensor input that pushes action component sequences on the command stack. Once initiated, a releaser (1) saves the current state of the system, (2) queues the times of parametric modulation events, and (3) queues the times of commands and transitions and then returns to the event loop to pop transitions off the queue. A releaser can thus activate multiple command objects in parallel, orchestrating a complete behavioral act.

**Choice** The command stack is the locus of implementation of behavioral choice. When presented with the releasers for two incompatible behavioral acts, the model animals typically decide to perform one act over the other through hardwired connections. In our controller, suppressor objects, triggered by releasers, embody lateral inhibitory connections between commands and are the locus of implementation of behavioral choice. Suppressor objects for a particular behavioral act maintain a look-up table of action components (which the act suppresses) and then clear incompatible action components from the event stack (figure 12.11A).



**Figure 12.11** (A) Event loop of the behavioral controller. See text for further explanation. (B) Waterfall display of internal state variables executing positive rheotaxis at two different temporal compressions. In this display, time flows from top to bottom.

**Goal-Achieving Objects** Goal-achieving objects mediate adaptations to exteroceptive feedback such as deviations of primary orientation relative to gravity, compass heading, or yaw sensor bias. Goal-achieving objects maintain ongoing states until a goal is achieved (e.g., turning to a particular compass heading).

**Linked Sequences** Linked sequences involve lists of both fixed and goal-achieving objects. Links consist of transitions between fixed and taxic components. Positive rheotaxis involves an initial fixed postural adjustment, then a goal-achieving yaw component, which persists until symmetric head-on flow is detected, then an achievement of the goal triggers the next subsequence in the list.

**Modulation** In addition to a temporal component, our sequencer mediates both taxic and parametric modulation. The yaw components of navigation are mediated reactively by taxes and kineses (Loeb, 1918; Braitenberg, 1984). Positive yaw taxes (or attraction) occur when sensor bias directs locomotion toward a source and are generally mediated by contralateral causality between sensors and effectors. Negative yaw taxes (or avoidance) occur when sensor bias directs locomotion away from a source and are generally mediated by ipsilateral causality between sensors and effectors. In the ambulation controller, such inputs send messages to both the parametric command and the recruiting component. In attractive reflexes, the messages are sent to the contralateral command objects, while during avoidance reflexes (negative taxis) the messages are sent to the ipsilateral command objects. As a

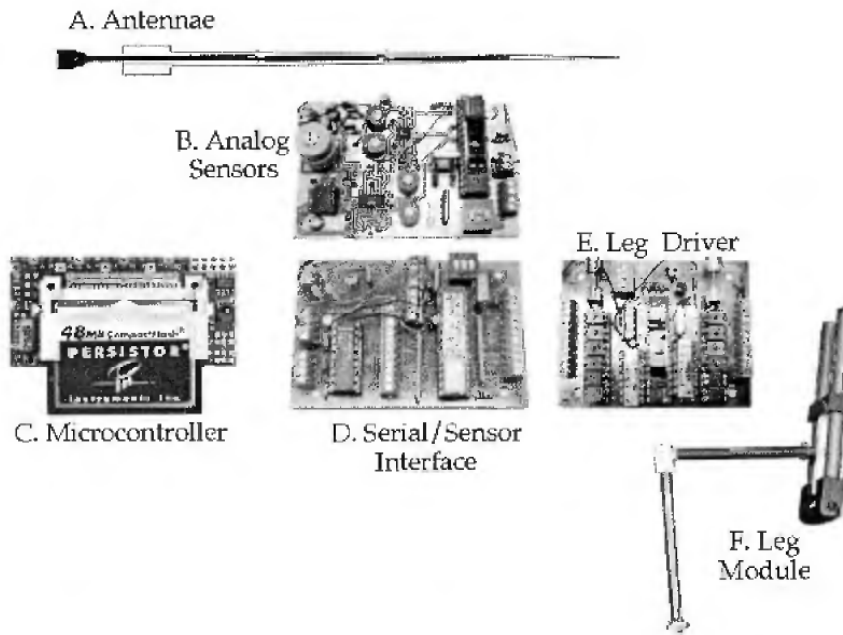
result, the controller biases both the period on the two sides (walking faster on the side turned away from) as well as the recruiters on the two sides to generate more-propulsive force on the faster side.

One of the features of animal behavior that has eluded most robotic implementations is the ubiquitous parametric modulation due to neuromodulators. The sequences of state changes specified in the behavior libraries can be modulated in intensity. In animal models, this modulation would be mediated by circulating neuromodulators. Our state matrix tables are derived from reverse-engineered movies and when played back by the controller, they can be placed on the event stack with variable temporal compression (figure 12.11B). Thus, increasing the temporal compression as the behavior is evoked by sensor input can increase the temporal intensity of the evoked behavior. Similarly, by systematically varying the pulse widths corresponding to low, medium, and high recruitment levels, the intensity can be varied in amplitude. Acting in concert, these two interacting forms of modulation allow parametric modulation of the intensity of evoked behavior.

**Supervised Autonomy** An operator can communicate with the system at two levels. First, using an attached float and radio antennae, one can operate the vehicles through a serial connection from a laptop. This allows the operator to monitor sensor input on the fly (for analysis of releasers, etc.). The off-board controller program supports a waterfall display that indicates the status of the different internal-state variables and allows one to control all internal-state variables on the fly. Second, the vehicles can be sent supervisory commands using sonar. The system includes a sonar transponder element (see figure 12.1A), through which supervisory commands can be delivered to the vehicle. We are developing a command grammar of the sonar teleoperation to allow the operator to select different motivations for the system (such as direction of walking; propensity for investigation of contacts of rocks, mine candidates, etc.; proportion of time spent statically sensing vs. moving, etc.). These commands are now issued as two byte codes, where the first byte indicates the nature of the data (compass heading, walking speed, etc.) and the second byte contains the magnitude.

## IMPLEMENTATION

The architecture described above leads to a simple board-level architecture. Both the ambulatory and the undulatory robot share the same modular organization (figure 12.12). MEMS sensors and analog sensors feature quantized outputs that are mapped on the logic pins of a PIC16 microcontroller. These logic pins are mapped on the bits of one or two bytes to be returned to the controller. We employ a polling interface wherein the controller requests specific data from individual sensors.



**Figure 12.12** Ambulator board level architecture. (A) MEMS antennae. (B) Analog compass, inclinometer and quantizer. (C) Persistor microcontroller. (D) Serial and sensor interface. (E) Leg current driver board. (F) Leg module. Not shown: sonar transponder interface.

The motor effectors are actuated by bits mapped on the output pins of several PIC16s, each of which controls the current drivers of a pair of legs. Thus individual muscles in individual legs can be addressed at one of three programmable pulse widths. The motors that control the pitch of the claw and tail as well as the yaw of the antennae are controlled by two bits from the same interface, absent the nitinol current drivers. These bits specify 1 of 3–4 positions through another PIC16 controller. All of these boards communicate to a host computer through a serial interface board that supports the sensor inputs. During operation, the controller sends a continuous byte stream to the serial interface that actuates the nitinol actuators and the motors and polls for sensor input.

## CONCLUSION

It is feasible to base the design of an autonomous underwater robot on biological principles. Sensors, controlling circuits, and actuators can be readily designed that operate on the same principles as their living analogs. Nature is conservative in the neuronal control strategies throughout the animal kingdom, relying on the command system, coordinating system, central pattern generator model (Stein, 1978). As we have shown, the same robotics controls can be adapted to both ambulating and undulating vehicles.



## REFERENCES

- Atwood, H. L. (1976). Organization and synaptic physiology of crustacean neuromuscular systems. *Prog. Neurobiol.* 7: 291–391.
- Ayers, J. (1989). Recovery of oscillator function following spinal regeneration in the sea lamprey. In J. W. Jacklet (ed.), *Neuronal and Cellular Oscillators*. New York: Marcel Dekker, pp. 371–406.
- Ayers, J., and Clarac, F. (1978). Neuromuscular strategies underlying different behavioral acts in a multifunctional crustacean leg joint. *J. Comp. Physiol. [A]* 128: 81–94.
- Ayers, J., and Crisman, J. (1992). The lobster as a model for an omnidirectional robotic ambulation control architecture. In R. Beer, R. Ritzmann, and T. McKenna (eds.), *Biological Neural Networks in Invertebrate Neuroethology and Robots*. Boston: Academic Press, 287–316.
- Ayers, J., and Davis, W. J. (1977). Neuronal control of locomotion in the lobster, *Homarus americanus*: I. Motor programs for forward and backward walking. *J. Comp. Physiol. [A]* 115: 1–27.
- Ayers, J., and Fletcher, G. (1990). Color segmentation and motion analysis of biological image data on the Macintosh II. *Adv. Imaging* 5: 39–42.
- Ayers, J. L., Jr., and Selverston, A. I. (1979). Monosynaptic entrainment of an endogenous pacemaker network: a cellular mechanism for von Holst's magnet effect. *J. Comp. Physiol.* 129: 5–17.
- Ayers, J., Zavracky, P., McGruer, N., Massa, D., Vorus, V., Mukherjee, R., and Currie, S. (1998). A modular behavioral-based architecture for biomimetic autonomous underwater robots. In *Proceedings of the Autonomous Vehicles in Mine Countermeasures Symposium*. Naval Postgraduate School. In press.
- Barlow, G. W. (1968). 1. Ethological Units of Behavior. In D. Ingle (ed.), *The Central Nervous System and Fish Behavior*. Chicago: The University of Chicago Press. pp. 217–232.
- Barlow, G. W. (1977). Modal action patterns. In T. A. Sebook (ed.), *How Animals Communicate*. Bloomington: Indiana University Press.
- Bowerman, R. F., and Larimer, J. L. (1974a). Command fibres in the circumoesophageal connectives of crayfish: I. Tonic fibres. *J. Exp. Biol.* 60: 95–117.
- Bowerman, R. F., and Larimer, J. L. (1974b). Command fibres in the circumoesophageal connectives of crayfish: II. Phasic fibres. *J. Exp. Biol.* 60: 119–134.
- Bullock, T. H. (1978). *An Introduction to Neuroscience*. San Francisco: Freeman.
- Davis, W. J. (1971). Functional significance of motor neuron size and soma position in the swimmeret system of the lobster. *J. Neurophysiol.* 34: 274–288.
- Davis, W. J., and Kennedy, D. (1972). Command interneurons controlling swimmeret movements in the lobster: II. Interaction of effects on motor neurons. *J. Neurophysiol.* 35: 13–19.
- Duerig, T. W., Metlon, K. N., Stöckel, D., and Wayman, C. M. (eds.) (1990). *Engineering Aspects of Shape Memory Alloys*. London, Butterworth-Heinemann.
- Duerig, T., Melton, K. N., Stockel, D., and Wayman, C. M. (1990). *Engineering Aspects of Shape Memory Alloys*. London: Butterworth Heinemann.
- Edwards, D. H. (1991). Mutual inhibition among neural command systems as a possible mechanism of behavioral choice in crayfish. *J. Neurosci.* 11: 1210–1223.

- Evoy, W., and Ayers, J. (1982). Locomotion and control of limb movements. In D. C. Sandeman and H. Atwood, (eds.), *Neural Integration and Behavior*, vol. 4 of *The Biology of Crustacea*. New York: Academic Press, pp. 62–106.
- Harris-Warrick, R. M. (1988). Chemical modulation of central pattern generators. In A. H. Cohen and S. Grillner (eds.), *Neural Control of Rhythmic Movements*. New York: J. Wiley. p. 285–331.
- Harris-Warrick, R. M. (1990). Mechanisms for neuromodulation of biological neural networks. *Advances in Neural Network Information and Processing Systems 2*: in press.
- Henneman, E. S. G., and Carpenter, D. O. (1965). Functional significance of cell size in spinal motoneurons. *J. Neurophysiol.* 28: 599–620.
- Kennedy, D., and Davis, W. J. (1977). Organization of invertebrate motor systems. In R. Geiger, E. Kandel, and J. M. Brookhart (eds.), *Handbook of Physiology*. Bethesda, Md.: American Physiology Society, pp. 1023–1087.
- Kovac, M. P., and Davis, W. J. (1980). Neural mechanism underlying behavioral choice in Pleurobranchaea. *J. Neurophysiol.* 43: 469–487.
- Kupferman, I., and Weiss, K. R. (1978). The command neuron concept. *Behav. Brain Sci.* 1: 3–39.
- Larimer, J. L. (1988). The command hypothesis: A new view using an old example. 11: 506–510.
- Pinsker, H. M., and Ayers, J. (1983). Neuronal oscillators. In W. D. Willis (ed.), *The Clinical Neurosciences*. Churchill Livingstone Publishers, pp. 203–266.
- Selverston, A. I. (1980). Are central pattern generators understandable? *Behav. Brain Sci.* 3: 535–571.
- Stein, P. S. G. (1978). Motor systems, with specific reference to the control of locomotion. *Annu. Rev. Neurosci.* 1: 61–81.
- Stein, P. G. S., Grillner, S., Selverston, A. I., and D. G. Stuart (1997). *Neurons, Networks, and Motor Behavior*. Cambridge: MIT Press.
- von Holst, E. (1973). *The Behavioral Physiology of Animals and Man*. Coral Gables: University of Miami Press.

**This page intentionally left blank**

# **II**    **Vehicles**

**This page intentionally left blank**

# Vehicles in Neurobiology

Polly K. Pook

Humans invented the wheel—the rotary joint. It is employed for locomotion in most of our vehicles: as a wheel, in tracks, or as a rotary propeller in air or underwater. The rotary joint offers large torque and efficient movement with simple actuation.

The rotary joint doesn't exist in nature (although some animals form wheels—e.g., rolling themselves into circles to wheel along the ground with speed). Instead, there are flapping, writhing, and legged gaits. What animal locomotion offers over the wheel is extraordinary maneuverability. Lampreys and snakes can wend their way into seemingly impassable crevices, then zip away with a few movements. The tuna and the dolphin, known for their speed and endurance in long-distance swimming, are able to accelerate and turn with surprising rapidity. Cockroaches scurry over broken surfaces with amazing stability and use their forelegs for precise manipulation. Moths glide and bank in dynamic air currents to find an efficient trace along a pheromone plume. Flies dart in space in a series of saccade-like motions, transitioning smoothly from flight to crawl.

These animals serve as design targets for the work presented here. In no case is the goal to mimic the animal systems precisely. To target the specific results of evolution, without considering developmental context (and chance), misses the boat. Moreover, we do not have skin, muscle, or metabolic energy sources at our disposal (although significant progress is being made to advance some of these technologies). Available materials dictate the design, as will manufacturing concerns. Instead, a synergistic analysis of salient biological and engineering principles guide design and development, often in a chicken-and-egg circle.

Both scientist and engineer seek an understanding of key functions that support the desired features of the animal. For example, vortices cast by a swimming tail or a twisting wing reduce drag and enhance lift for swimming and flight (cf. chapters 15, 22, 23, and 24). Such knowledge focuses the design model. The engineers' response may or may not be to duplicate the mechanism, but it will attend to the effect.

A necessary goal, driven by available components and computational limits in both fabricated and biological systems, is to reduce complexity where possible. Again, this approach informs both mechanical design and biological understanding. Thus, the degrees of freedom are drastically reduced in all of the engineered systems, with the resulting demonstration of the limitations such reduction places on functionality (often less than imagined, cf. chapters 13, 16, 20, and 21). The biologist, in turn, discovers surprising couplings that serve a similar effect—for example, that the tobacco hawkmoth couples/decouples its forewing and hindwing selectively to change the wing surface area during banking turns (cf. chapter 22). Other simplifications, such as the exploitation of passive compliance and natural resonant frequencies, which make such mechanical designs possible, provide hindsight into the ubiquitous use of such mechanisms in nature.

Control of subsystems and their interaction must be simple and scalable. The use of local central pattern generators, distributed subsystem interaction, and behavioral control to achieve sophisticated functionality with computational economy is continually being reinvestigated and addressed in both reverse-engineered models and robot designs (cf. chapters 14, 18, 19, and 20).

The challenges faced by these researchers are numerous: Powered complex actuation with sensing and control, all packaged in small, sometimes tiny, forms, that must operate in uncertain, dynamic, and often harsh environments. Their successes, thus far, are remarkable. All of these systems are dynamic, so static words and images don't do them justice. Readers are urged to visit the web sites to see video clips of the robotic animals in action.

## UNDULATORY AND SWIMMING MOTION

Gavin Miller (chapter 13) targets the snake in the first chapter of part II. Snakes have incredible mobility—they can slither into crevices, yet move and strike with astonishing rapidity. Snakes exhibit a number of gaits, ranging from a fast sidewinding to a creeping earthwormlike series of contractions. Snakes can climb and even “fly” from tree to tree. Miller presents five evolutions of a snake prototype, with a design for a sixth, that tackle a subset of gaits. The number of actuators required (37 in his most advanced prototype) demands thoughtful packaging to embed motors, processors, and a power source within the strict form. The results are elegant.

In chapter 14 Wilbur and colleagues look at undulatory motion underwater, inspired by the sea lamprey. While Miller takes a bottom-up approach, building successive prototypes to evolve initially simple motion with ever more complexity, Wilbur and colleagues start top down. They segment videos of actual lamprey motion in a variety of

habitats to identify a taxonomy of behaviors based on visual features. Coupled with an understanding of underlying kinematics and hydrodynamics, their first prototype robot mimics the underwater, undulating gait of the lamprey. The ultimate realization of this approach will allow high-level behavioral control that adapts motion parameters to changes in the environment.

For open water, few animals succeed as well as the tuna and the dolphin. They are capable of high speeds (30 km/hr or about 8.3 m/sec) with long endurance, and their maneuverability far exceeds typical unmanned underwater vehicles (UUVs). At the Draper Laboratory, Anderson and Kerrebrock (chapter 15) present a robot modeled on the large yellowfish tuna. Fish are able to perform surprisingly sudden turns, rapid accelerations, and enduring velocities due to vortices shed from their skin and cast into their wake that reduce the effects of drag. The eight-foot Draper robot takes advantage of this action via relatively simple control of its articulated tail. The result is a robot that can attain speeds of 1.2 m/sec and maneuver in a small area, exhibiting suggestively lifelike motion. Nakashima and Ono (chapter 16) go on to examine the propulsive efficiency of a robot dolphin. Two prototype robots were built to scale but with simplified actuation (and therefore limited maneuverability). They were able to achieve a very fast 1.9 m/sec propulsion speed (just over four times slower than the real dolphin) with an efficiency of 0.65.

The work on tuna and dolphin robots has focused on the tail, using the fins predominantly as stabilizers. Naomi Kato (chapter 17) looks at the motion of fish fins, with the black bass as his model. Fins allow the fish to make subtle and controlled movements, unlike the strong and initially undamped effects of typical UUV screwlike thrusters. His Bass fins flap in both the vertical and horizontal planes as well as twist, permitting controlled movement in all three dimensions of the underwater world.

## CLOSED-LOOP CONTROL IN AMBULATORY ROBOTS

Underwater robots must deal with the complexities of three-dimensional motion through a dense medium that is harsh on mechanical components. Ambulatory robots instead grapple with the problem of controlling legs over terrain that is uncertain and uneven, particularly from an insect's perspective.

Klavins and colleagues (chapter 18) provide an analysis of methods for closing the loop about the springy walking robot RHex. RHex takes inspiration from sprawled-posture hexapods, such as the cockroach. The running cockroach generates large lateral forces, unlike the upright-postured quadruped (e.g., the dog), which restricts its movement to the fore-aft and up-down planes. Cockroaches demonstrate remarkable



stability under large natural perturbations generated as they run over terrain that, to them, is very uneven. Kubow and Full (2001) hypothesize that the springlike legs of a sprawled-posture insect mechanically absorb the load generated by these lateral forces, in a form of reflex that stabilizes the dynamic system. RHex is an excellent example of such mechanical economy. Its six legs have only a single actuator each but are composed of a bowed compliant spring that accepts lateral loads. With open-loop control, under a central pattern generator, RHex scrambles reasonably well over an impressive obstacle course. Klavins and colleagues take the next step here and present simulations to predict the effects of closing the loop under a set of reflex controllers that are designed about the energy exchange of the reflex.

The work of Klavins and colleagues (chapter 18) recasts the canonical gait controller formalized by Cruse and colleagues (1995) into a mechanical realm of coupled oscillators that capture the occurrence of energy transfer. The model developed by Cruse is based on a reverse engineering of the stick insect (and later extended to the crawfish). This model identifies six “influences” among neighboring legs that modulate tripod or tetrapod gaits to accommodate natural perturbations. A key feature of this closed-loop model is that it is fully distributed; there is no centralized controller. By simply modulating the velocity of the leg as it moves through its swing or power strokes (or alternatively, adjusting the length of the power stroke) based on phase signals from neighboring legs, the insect is able to converge on a limit cycle following perturbation. In chapter 19, Cruse and colleagues investigate higher-level control mechanisms that superimpose on the stick-insect controller to respond appropriately to more-global parameters, such as obstacle avoidance, desired speed, and direction. They present the “walknet,” a neural network that controls key features in a simulated walking insect.

At Case Western Reserve University, Roger Quinn and Gabriel Nelson work with neurobiologist Roy Ritzmann to understand and build another cockroach robot (chapter 20). One of their earlier prototypes, Robot II, served as a proof of concept for Cruse and Dean’s stick-insect controller. In the work presented here, like Cruse, Quinn and colleagues superimpose additional controllers on the stick insect gait control. In their case, the added control addresses large external forces generated by the environment and by the addition of payload. Robot III, the latest prototype, incorporates three new controllers: posture, swing, and stance. Posturing serves to relocate the center of mass as the vehicle moves to climb, resist external lateral forces, and carry up to its own weight in payload. The swing controller transforms the four- and five-jointed leg motion (middle and front, respectively) into Cartesian coordinates and assists the gait controller in overall leg coordination. The stance controller responds to strain to reinforce the posture controller. Robot III is also a much closer model of the kinematics and

biomorphology of the cockroach than was Robot II, displaying differentiated front, middle, and rear legs, twenty-four degrees of freedom, pneumatic actuators, and skeletal design. The Cricket, another robot briefly mentioned, starts to address the problems of incorporating the mechanical, power, and computational requirements onboard the robot—in this case, a cricket-size robot, measuring just two inches in length.

Kirchner, Spenceberg, and Linnemann (chapter 21) have designed a novel hybrid open-loop/closed-loop controller for their eight-legged Scorpion robot. The controller relies on a central pattern generator to produce any one or a combination of open-loop motions belonging to a pattern library. A closed-loop reflex controller handles asynchronous environmental forces, such as sudden contact with an obstacle. Notably, the transition from one pattern to another can cause instabilities, just as large external perturbations might. The system deals appropriately with known transitions.

## FLYING ROBOTS

Creating and controlling a flying microrobot can be much more exacting than designing ambulatory or even swimming robots. Flight requires great attention to mechanical form, size, weight, and aerodynamics. The payoff is huge, however. Such a vehicle may attain unprecedented access to desirable vantage points and do so efficiently and quietly.

The chapters by Wendler and Michelson (22 and 24, respectively) each present an analysis of flight based on the tobacco hawkmoth, *Manduca sexta*. *Manduca* is a neurogenic flier (having a low wing-beat frequency) with a 10 cm wingspan that exhibits particularly accomplished flying capabilities. Wendler's work is to reverse engineer the moth. He presents first an analysis of muscle-activity patterns, wing movements, and aerodynamic resultants derived from a tethered moth under optical flight stimulation. To better correlate these parameters, Wendler goes on to stimulate specific muscles in a de-efferented moth, and reports preliminary findings. In the meantime, Michelson presents a far-reaching design of a flying vehicle, just twice the size of a *Manduca*. The success of the project will primarily depend on the realization of a chemically fueled muscle, the prototype of which is now being made at the appropriate scale. The very imaginative design incorporates a host of innovations, from a unique nontwisting wing that produces lift on the downstroke as well as the upstroke, to a routing of the muscle exhaust for functional reuse six times.

Chapter 23 presents one of the most ambitious design targets: a blowfly, to scale. The blowfly is a fast-beating myogenic flier with a wing-beat frequency of 150 Hz. Its mass is just 100 mg and wing length 11 mm. The issues of material strength, compliance, actuation, and

power delivery are critical at this scale. Fearing and colleagues present a highly instructive progress report toward the promising realization of their design. A 5-times-scale prototype thorax made of four flexural hollow bars controls the leading and trailing edge of the wing. This innovative structure magnifies the twisting motion of the folding wing so that only simple piezoelectric actuation is required. The design is necessarily keyed to durability. Fearing and colleagues have tested more than one million beat cycles at 1.3 times scale successfully.

## SUMMARY

If you were to build a small robotic vehicle, how would you go about it? The researchers in part II look interchangeably at biological and engineering principles in elegant detail. The results are informative with regard to the control, the necessary complexity, the possible simplifications, and the use of state-of-the-art materials.

The potential benefits of this research are enormous. Crawling snake and insect robots can search through rubble in the aftermath of a bomb or earthquake and enter buildings through small cracks and ductwork covertly. Underwater vehicles can explore the ocean, moving quickly and covertly near submarines, hot thermal vents, and underwater mines. They can gather data on ocean life and support underwater oil drilling. Eel-like swimmers can search murky tidal pools or swim up estuaries and sewage pipes to seek out the sources of chemical pollutants. They can search in shallow water for mines or perform maintenance inspections on bridge pilings. Microflyers can scout areas and provide a bird's-eye view, trace chemical or biological agents, enter upper-story building ventilation systems, and hover unobtrusively above eyesight. *All* of these robots could lead to fantastic toys and entertainment devices. Enjoy.

## REFERENCES

- Cruse, H. (1995). Walking: A complex behavior controlled by simple networks. *Adaptive Behav.* vol. 3, no. 4: 385–418.
- Kubow, T. M., and Full, R. J. (2001). The role of mechanical system in control: A hypothesis of self-stabilization in hexapedal runners. *Philos. Trans. R. Soc. Lond. B. Biol. Sci.* 354: 849–862.

# *Undulatory and Swimming Robots*

**This page intentionally left blank**

# 13

## Snake Robots for Search and Rescue

Gavin S. P. Miller

This chapter considers the applicability of current and future snake robots to a search-and-rescue mission. This imaginary scenario has been devised to act as a framework within which to discuss the advantages and challenges of using robots, inspired by biological snakes, to do useful work. The first section describes the mission, a search-and-rescue task following a major earthquake. The section entitled Prior Work explores precedents in this field, and the Snake Gaits section introduces a number of primary snake gaits for locomotion and discusses their impact on snake robot design. Snake Robot Prototypes describes the snake robots constructed by the author and reveals lessons learned. The chapter concludes with a description of ongoing work for a next-generation snake robot, a summary section, and references.

The question naturally arises, why snake robots? Biological snakes occupy a wide variety of ecological niches, ranging from arid desert to tropical jungle as well as swimming in rivers and oceans. Abandoning limbs and developing elongated spines has proved an effective survival strategy, allowing snakes to hunt underground in confined tunnels, above ground in grassy fields, and up in the treetops—even falling in a controlled glide from one tree to the next (Mattison, 1999). By attempting to build robots that emulate and perhaps match the capabilities of their biological counterparts, it is possible that we will create useful tools capable of carrying sensors, taking samples, and effecting physical changes in a wide variety of environments.

### THE SEARCH-AND-RESCUE MISSION

This scenario is purely imaginary and is presented to help provide a goal for future developments rather than stating current capabilities. It may be that for all practical tasks, some other form of robot, such as one with wheels or tracks, will prove superior. However, that is an open research question, worth asking every few years as computation, actuation, and scientific insight continue to develop.

Following a major earthquake in the San Francisco Bay Area, old Victorian houses lie in ruins, the ground floors having collapsed as upper stories remain standing but unstable, threatening to fall at the next aftershock. Willing rescuers stand ready to climb and dig, but the question remains of whether the risk of injury to them is worth the chance of finding someone buried under the rubble. Fortunately, a new generation of search-and-rescue snake robots is now available to answer that question.

Pulling out a 2-meter-long tube from under the stretcher in the paramedic ambulance, the operator opens one end and powers up the system. She removes a remote control unit that has a video screen built into it—fed by the camera from the head of the robot—with a graphical overlay allowing a range of optional behaviors to be activated. There are also two joysticks and a variety of other buttons and knobs for controlling subsidiary parameters of the system. The would-be rescuers carry the tube as close as possible to the wreckage, and slide the snake robot out of the end of it. A thin composite cable goes from the inside of the tube to the tail of the snake, supplying power and carrying data in both directions. The operator activates the snake's locomotion system and moves it forward. It automatically adapts to the uneven ground and rises up over small obstacles while maneuvering around the larger ones.

Pausing every few feet to listen for signs of survivors, the snake robot's head relays binaural stereo sound back to the operator. Hearing no cries for help, the snake is directed toward a 20-cm-wide gap where one house has collapsed and is leaning against another. Approaching the aperture, the snake transitions to rectilinear motion, and uses infrared distance measuring devices and flex-sensor whiskers to center it between the walls. As it moves further inside, the operator switches to the infrared-sensitive camera and illuminates the scene with high-power LEDs. As the snake progresses, it sweeps the area up ahead with a pyroelectric device to look for body heat. Its underbody (ventral) scales pull the robot along the ground like a small conveyor belt, even pushing it through the tangled heaps of cables left by the collapsed building. Then it reaches a region of shattered plaster and broken rubble, which provides insufficient grip for forward locomotion. The tether is also pulling on the snake robot's tail, caught on previous obstructions. At the system's on-screen suggestion, the operator switches the snake robot to internal power and detaches the tip of the tail by remote control. This now becomes a base station communicating wirelessly with the snake robot and relaying information back to the operator.

She instructs the snake to begin concertina motion, in which the snake robot coils into an S-shaped curve until it can feel the walls of the fallen houses pushing in on it from both sides. Small scales on the

skin of the snake robot grip the walls and allow it to push forward by changing the amplitude of its coils in one region while gripping with another. Once through this difficult area, the snake comes to a region of the original floor of the house. The snake swishes its head from side to side to sweep the area clean of rubble. Using a downward-sensing ultrasonic device in the chin of the robot, the operator determines that it is possible to make a hole leading directly to the basement of the building. The snake robot is instructed to detach, from under its head, a small explosive charge, which it leaves on the floorboards. The snake slithers back slightly before detonating the charge, which makes a 10-cm hole in the floor. Curling its head downward through the opening, it pushes forward a meter or so and then uses its neck to point its head in a variety of directions, looking for survivors. There is a peak detected by the pyroelectric sensor, and the snake freezes, going silent. The microphones pick up the faint sounds of breathing and the infrared camera indicates a blob in approximately the direction from where the sounds originated. The operator pinpoints on a map where the survivor is most likely to be found. The location is shown in relation to a reconstructed 3-D model of the path taken by the snake robot, along with the surfaces it sensed. Other rescuers are given the go-ahead to carefully approach the building as the operator talks to the survivor over a loudspeaker carried by the snake robot, letting him know that help is on its way and trying to discover the extent of his injuries.

## **PRIOR WORK**

Biological snakes have inspired a variety of robotic designs since the 1920s. A good overview of the field is given in Dowling (1997). Shigeo Hirose, in the 1970s (Hirose, 1993), attained the first successful locomotion by a tethered snake robot. Since then, there have been a variety of projects examining path planning and control (e.g., Chirikjian and Burdick, 1995). Innovative mechanical designs have flourished (such as that in Dowling, 1997), with one of the most sophisticated being GMD-SNAKE2 (Klaasen, 1999), which has actuated joints between each segment, along with powered wheels all around the circumference. This enables an approximation of rectilinear progression (see the following section), but such wheels may snag on fibrous obstacles. (An excellent list of current snake robot projects may be found in Worst, 2000.)

## **SNAKE GAITS**

Studies of biological snakes have identified a variety of fundamentally different strategies for snake locomotion (Miller, 1991). It is worth considering each of these in turn because they place different requirements



on the actuation and surface properties of biologically inspired snake robots.

### **Sidewinding**

Sidewinding is observed in a variety of snakes and is especially popular on crumbly terrain such as hot sand. It may serve the dual purposes of reducing contact area with the ground, to minimize heat transfer, while allowing a nonsliding means of locomotion on compliant surfaces.

Sidewinding may be achieved by having horizontal and vertical waves of undulation traveling down the length of the snake, one trailing the other by  $90^\circ$  of phase. Typically, the horizontal undulation has at least three times the amplitude of the vertical undulation. (It is worth noting that if the vertical amplitude is too large, the snake robot will corkscrew in place flipping upside down repeatedly. This may be useful as a separate skill, if the snake robot was accidentally flipped over and needs to be inverted.)

Sidewinding may be visualized as a single oval tank track that has been cut and stretched into a helical form to maintain lateral stability. It has been created successfully in computer simulations and using snake robots with two degrees of freedom at each joint. (See Miller, 1988; Burdick, Radford, and Chirikjian, 1995; and Dowling, 1997 as well as the Snake Robot Prototypes section.) Sidewinding requires relatively more energy than some other gaits, and also does not provide very smooth or stable motion for the head of the snake, leading to sensor and path-planning challenges.

### **Concertina Progression**

Concertina progression is achieved by using a nontraveling curve along the length of the snake with modulation of the curve amplitude to attain forward motion.

The gait takes two forms. The first is more suited to open ground and consists of a straight head and tail with a single S curve in the center of the snake. Changing the amplitude of this curve changes the distance between the head and tail. On an appropriate surface, the directional friction due to the ventral scales or "scutes" under the snake leads to locomotion. This gait is used occasionally during hunting because it allows for very accurate steering of the head (such as in Miller, 1991). The second version of the concertina gait works well in narrow passageways in which the floor of the tunnel is too smooth or crumbly to provide an adequate grip. The snake has a sinusoidal wave all the way along its length, but can control the amplitude of different regions independently. Expanding the amplitude of the tail section pushes that por-

tion of the snake against the walls of the tunnel. Reducing the amplitude of the middle portion of the snake then pushes the head forward. The head section amplitude is then increased so that it grips the walls. The tail is then relaxed and the amplitude of the center section expanded to pull the tail forward.

### **Horizontal Undulatory Progression**

Horizontal undulatory progression works by a different mechanism again. A traveling wave along the length of the snake attains forward locomotion either by directional friction with the ground (Hirose, 1993), by pushing against vertical obstacles, or by using vertical undulations in the ground.


### **Rectilinear Motion**

Rectilinear motion requires the most elaborate mechanism of all snake gaits, requiring the skin of the snake to be pulled forward and backward over the ribs as downward pressure is modulated on the scutes by a bulging and releasing of muscles. The gait has the advantage of smooth and steady progress for the head and works on rocky ground as well as when climbing over such surfaces as the boughs of trees. Rectilinear motion is often observed in larger snakes, such as pythons, and is not particularly fast.

## **SNAKE ROBOT PROTOTYPES**

The author of this paper has developed a number of snake robots of increasing complexity. Each prototype provided valuable lessons in different physical principals. It was decided to make the robots self contained in terms of power and computation, to eliminate the drag effects of external tethers, and radio control was used to allow intuitive interaction with the robot.

A first prototype (not illustrated), called S0, embodied a planar snake robot design with only horizontal flexing between adjacent segments. It employed 12 radio-controlled servos and a custom controller to allow a traveling square-wave of undulation to move down the length of the snake. The battery for the snake was contained in the tail, the controller in the head. Under each scale was an angled blade approximating a scute. Performance of the robot was disappointing, with only very limited locomotion from concertina progression and almost no success creating horizontal undulatory progression. A major flaw in the design was in the weight distribution. The battery in the tail concentrated all of the friction in one area, rendering the rest of the scutes largely ineffective. A second problem was due to the nature of the fixed scutes




**Figure 13.1** A central section of S1. The batteries are placed next to each servo.

themselves. On carpet and smooth cement the angled blades provided surprisingly uniform friction forward, backward, and laterally, only being useful on macroscopic features such as the edges of rugs. Based on the lessons learned, it was decided to produce S1.

### Planar Snake Robots with Wheels

The prototype called S1 (figures 13.1 and 13.2) improved on the design of S0 by having rechargeable cells distributed along the central sections of the snake. A second innovation, new for snake robots, was the use of tapering segments toward the head and tail. The torque on the motors, due to flexing the snake, is greatest in the central portion of the snake and diminishes toward each end, although less so for the neck because of the mass of the head. The use of progressively smaller servomotors allowed a more optimal use of weight for a given length of snake. S1 was used to test a variety of materials for the underside of the snake to determine suitability for locomotion. “Snakeskin” material, created for cross-country skis, provides directional friction on snow using angled fibers. This material slides well in one direction while gripping in the other, but, unfortunately, it generates large amounts of lateral friction. With S1, it proved ineffective on carpet and asphalt—the lateral forces overwhelming the strength of the motors. On dirt, the material was rapidly fouled with particles matting the fibers.




**Figure 13.2** S1 seen from above. Notice the tapering toward the tail and head.

Adding free-spinning wheels to the middle of each segment, however, proved very effective, allowing successful locomotion on smooth surfaces. On thick-pile carpet and grass, the wheels proved less effective. This design is a somewhat modified version of the “active cord mechanism” of Hirose (1993). An improved design, which hid the wiring and larger wheels within the structure, is shown in figures 13.3 and 13.4. S2 had a caster wheel under the head and moved more successfully when the wheel from the second segment was removed (figure 13.5).


The control mechanism of S1 and S2 were as follows: A single micro-computer, a Basic Stamp II, was used to decode the signals from the radio receiver and then drive each servo directly. This design is hard to scale up to more-complex designs, because the bundle of wires grows linearly with the number of servos and the computational load soon overwhelms the processor.

### **Nonplanar Snake Robots with Wheels**

The next prototype, S3, was developed to experiment with sidewinding gaits and to try transitions from one gait to another. The main feature of the new design (figure 13.6) is the use of two servos per segment arranged in an opposed configuration with a universal joint at the top (figure 13.7). Each servo arm is connected to the case of the other servo by a ball-jointed rod. The symmetry of the arrangement means that rotating each servo in the same direction flexes the joint laterally, while rotating them in opposing directions flexes the joint vertically. This



**Figure 13.3** A section of S2. The wiring and wheels are enclosed by the structure.




**Figure 13.4** S2 seen from above. The cylindrical structures act as hinges.

arrangement doubles the lateral torque available from a given servo and also allows a choice of mechanical advantage based on the placement of the ball joint on the servo arm relative to the servo's axis of rotation.

A more elaborate controller was required than for S2, due to the larger number of servomotors (37 in all, including a single-servo joint in the tail). One Basic Stamp II microcomputer was used to interface with the receiver and to drive two servo controllers (each of which




**Figure 13.5** S2 seen from below. Note the caster wheel on the head and the lack of wheel on the next section back.



**Figure 13.6** A section of S3. The servos are mounted parallel to the cross-section, with a universal joint at the top.



**Figure 13.7** S3 seen from above. The tail and neck taper, and the universal joints form a spine along the length of the robot.




**Figure 13.8** A section of S5. Refinements over S3 lead to a smaller cross-section for a given size of servo.

could control eight servomotors). A second Basic Stamp II acted as a slave processor and controlled three other servo controllers. Suitable programming of S3 allowed experiments with sidewinding as well as wheel-based horizontal undulatory progression. By changing the amplitude of the vertical component of the sidewinding gait, the robot could make a smooth transition from sidewinding to horizontal undulatory progression and back again.

An enlarged version of S3, called S5, was developed with a smaller cross section (figure 13.8) and 32 joints. It was successful at horizontal undulatory progression, forming a more sinuous curve than S3 due to the larger number of segments per wavelength (figure 13.9). The control system for S5 employed a Basic Stamp II, for high-level control, along with a Scenix microcontroller for interfacing to eight servo



**Figure 13.9** S5 seen from above. Two wavelengths of undulation fit within the length of the snake robot.



**Figure 13.10** The head of S5. A radio receiver can be seen below the main circuit board. At the front of the head-section is an infrared proximity sensor.

controllers over four unidirectional serial lines. Figure 13.10 shows a side view of the head, with the radio-control receiver underneath the main circuit board. The section at the front is an infrared proximity sensor.

Unfortunately, S5 proved unable to sidewind due to the scaling effects of larger length, and possibly because of a heavier head. Determining the exact cause of this limitation is an area for future work.



## **Toward a Snake Robot without Wheels**

Biological snakes prove that it is possible to have successful locomotion without wheels. Sidewinding has also been demonstrated in snake robots without wheels (Burdick, Radford, and Chirikjian, 1995; Dowling, 1997). S3 could also sidewind successfully when placed on the ground upside down.

Other gaits pose greater challenges. Horizontal undulatory progression, without wheels, requires sliding contact with vertical obstacles and uneven terrain. To achieve this will require a smooth continuous covering, rather than the segmented design of previous prototypes.

Successful locomotion for larger snake robots, without wheels, will probably involve the need for rectilinear motion. As we know from biological snakes, this can be achieved by pulling the external covering forward and backward over the skeletal structure while modifying the internal shape. (Design of a snake robot to achieve this successfully is an area for future work.) Adding a sense of touch to at least the underside of the snake robot is probably vital to allow it to locomote efficiently over undulating ground.

These considerations are being incorporated into a new prototype, S7, which is currently under construction. Adding sensors to each segment creates a need for bidirectional communication with each servo controller. Sensor-based locomotion highlights the need for a more powerful microcomputer in the head. These are both nearing completion. The off-the-shelf radio-control system is being replaced with a custom packet-based bidirectional link, with a liquid crystal display (LCD) in the control unit. This will enable messages to be reported back to the operator as well as allowing the initiation of different discrete behaviors, such as changing gait, or stretching out straight for stowage in the transportation tube.

## **SNAKE ROBOT COMPONENTS**

This section discusses individual components of the snake robots, both present and future, and helps to suggest additional directions of research.

### **Sensors for Snake Robots**

A sense of touch may prove vital for versatile snake robots, involving measurements of internal joint angles and of external pressures on the scales. It is possible that the force sensors for each surface scale could be binary, but the wiring requirements will still be formidable. In addition, the snake robots will require visual and other sensors, packed into a relatively small and light head.

## **Skin Requirements for Snake Robots**

One difference between man-made and biological snakes is the length of each segment compared to its diameter. Biological snakes have up to 300 joints along their length, whereas S5, the largest to date in the series, has only 32. This places a greater burden on the exterior covering, requiring it to act as a spline, smoothing out the relatively limited number of segments.

Kevin Dowling's (1997) approach to a snakeskin used a meshed material with a covering of Lycra. It did not prove completely effective for large deflections of the snake, due to the diagonal fibers of the mesh conserving their length causing buckling in the skin. Even real snakes have buckling under extreme deflections, but uniform materials are probably unable to smooth out the underlying irregularities of the mechanisms described in this paper. One possibility is to use a composite structure of flexible foam rubber with embedded stiffeners.

## **Actuators for Snake Robots**

The snake robots in the current series have all used off-the-shelf radio-control servomotors. They are cheap and compact, and certain models provide high levels of torque (up to 9.8 kg/cm at 6 volts) with metal gears. In early experiments, plastic-gear servos lasted less than a few hours, stripping their threads with expensive results.

Certain robots have employed memory metal alloys for making swimming eels (Jalpert, Kashin, and Ayers, 1995), but on land, the heat dissipation, time constants, and power requirements of these materials render them unsuitable, at least at present, for snake robots. Other materials hold promise, such as the electrostatic linear actuators, and electroconstrictive polymers described elsewhere in this book. Their power-to-weight ratios can exceed those of servomotors, and the actuators might be made numerous and small—more effectively approximating the high number of degrees of freedom of biological snakes—thus posing simpler requirements on the covering material. One potential drawback is the use of high voltages, but the binary control of larger numbers of actuators may make this system practical.

## **CONCLUSIONS**

Although no snake robot exists with all the capabilities described in the scenario at the beginning of this chapter, progress in that direction remains steady. With the advent of more-sophisticated electronics and advanced mechanical designs, future generations of snake robots may begin to approach the locomotion capabilities of biological snakes. More-

advanced actuators may yield additional improvements in lightness, agility, and endurance.

## REFERENCES

- Burdick, J. W., Radford, J., and Chirikjian, G. S. (1995). A sidewinding locomotion gait for hyper-redundant robots. *Adv. Robotics*, special issue on highly redundant manipulators, vol. 9: 195–216.
- Chirikjian, G. S., and Burdick, J. W. (1995). The kinematics of hyper-redundant robotic locomotion. *IEEE Trans. Robotics Automation* vol. 11, no. 6 (December): 781–793.
- Dowling, K. (1997). Limbless locomotion: Learning to crawl with a snake robot. Ph.D. diss., technical report CMU-RI-TR-97-48, Robotics Institute, Carnegie Mellon University.
- Hirose, S. (1993). *Biologically Inspired Robots (Snake-like Locomotor and Manipulator)*. Oxford: Oxford University Press.
- Jalbert, J., Kashin, S., and Ayers, J. (1995). Design considerations and experiments on a biologically-based undulatory lamprey-like AUV. In *Proceedings of the International Symposium on Unmanned Untethered Submersible Technology, Autonomous Undersea Systems Institute, Portsmouth, N.H.*
- Klaassen, B., and Paap, K. L. (1999). GMD-SNAKE2: A snake-like robot driven by wheels and a method for motion control. In *IEEE International Conference on Robotics and Automation*, IEEE Robotics and Automation Society. pp. 3014–3019.
- Mattison, C. (1999). *Snake*, 1<sup>st</sup> American ed. New York: DK Publishing.
- Miller, G. S. P. (1988). The motion dynamics of snakes and worms. *Computer Graphics (Proceedings of SIGGRAPH '88)*, Atlanta, Georgia, J. Dill (ed.), vol. 22, no. 4 (August): 169–178.
- Miller, G. S. P. (1991). Goal-directed animation of tubular articulated figures or how snakes play golf. In N. I. Badler, B. A. Barsky and D. Zeltzer (eds.), *Making Them Move: Mechanics, Control, and Animation of Articulated Figures*. San Mateo, Calif.: Morgan Kaufmann, pp. 233–241.
- Worst, R. (2000). Snake-like robot list on the world-wide web. Institute for Autonomous Intelligent Systems. On-line: <http://ais.gmd.de/%7Eworst/snake-collection.html>.

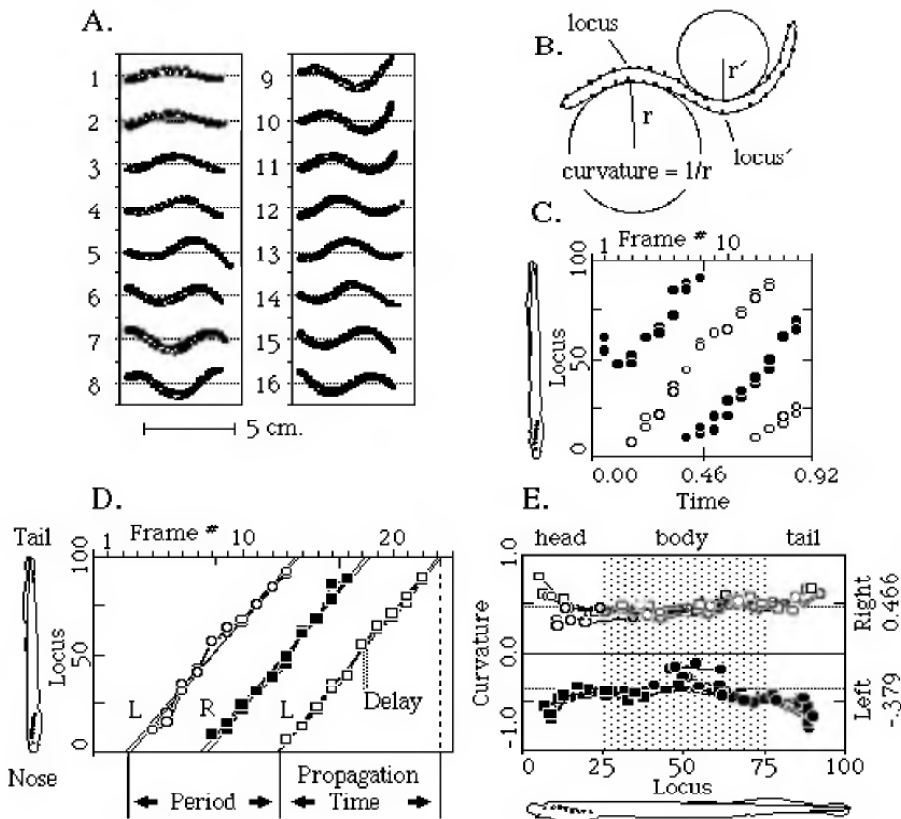
Cricket Wilbur, William Vorus, Yusong Cao, and  
Scott N. Currie

Sea lampreys (*Petromyzon marinus*) are among the most primitive vertebrates (Hardisty, 1981), having an elongated and flexible eel-like body (subcylindrical forward, laterally compressed dorsal area) and as mature adults, averaging 2–3 feet in length. Although their swimming movements are often considered eel-like or anguilliform, they are characterized by the propagation time of flexion waves being equal to the period of the movement, such that the body axis typically appears S-shaped (Ayers, 1989). We term this more-constrained form of swimming *lampriform*. Despite this relatively simple organization, lampreys are capable of a rich repertoire of undulatory behavioral acts (Ayers, 1989). During migration to and from the sea, lampreys encounter a variety of obstacles and hydrodynamic conditions. Different habitats through which the lamprey must travel could include streams, rivers, falls, rapids, lakes, estuaries, intertidal zones, and the open ocean.

In addition to its high adaptability, lampriform swimming affords an additional advantage of energetic efficiency. Unlike carangiform swimmers, which transfer considerable energy to the water in the generation of vortices (Triantafyllou and Triantafyllou, 1995), lampriform swimming features a reduced wake signature (Vorus, 1995). Thus the lamprey provides two major advantages of adaptability and efficiency as a model organism for a biomimetic underwater vehicle (Ayers, 2000).

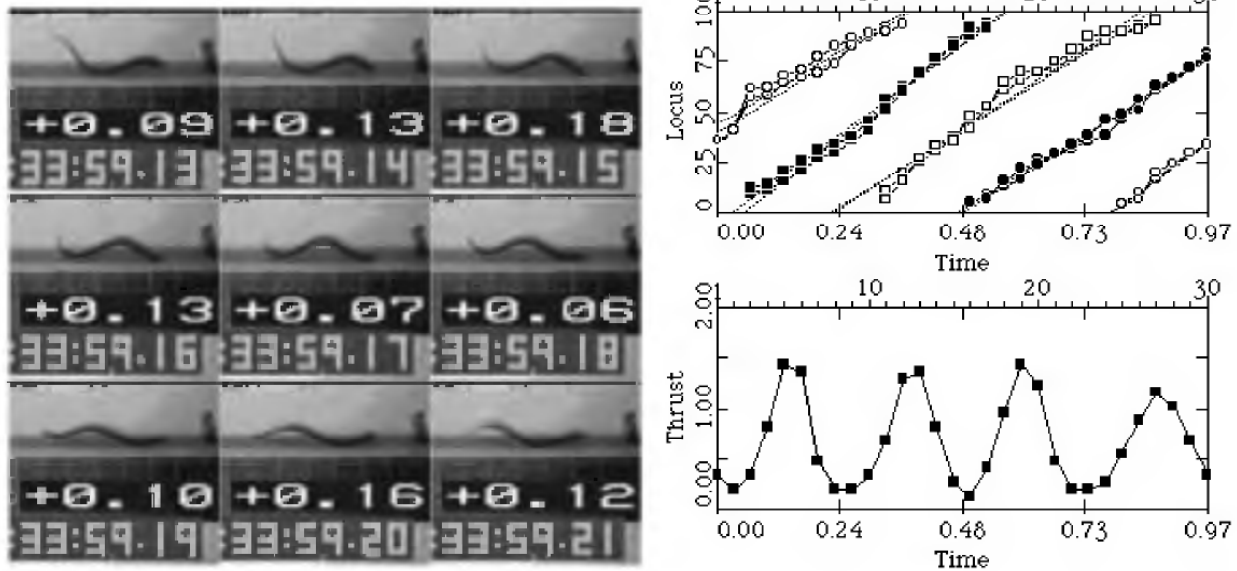
#### KINEMATICS OF SWIMMING

Lampreys swim by rhythmic lateral undulations of the body axis. The lamprey swims forward by propagating lateral axial undulations that increase in amplitude from nose to tail. Similar waves traveling from tail to nose can propel the lamprey backward. Rhythmic alteration of muscle activity on either side of the body axis produces a propulsive wave. Lamprey swimming is uncomplicated by pectoral, pelvic, or anal fins. Endurance, rather than high speed, is characteristic of pure anguilliform mode (Gosline, 1973).



**Figure 14.1** Analysis of lamprey swimming. (A) Graph of the normalized digitized images of swimming in a larval lamprey. A least squares regression line is fitted to the set of points outlining the body; the line is rotated to become the x-axis. The number to the left of each image is the movie frame number. The interframe interval is 18 msec. Analysis of the images presented in (A). (C) is a plot of the locus of flexion waves versus time, going from nose (0) to tail (100). (D) is similar to (C), only a least-squares regression line is fitted to the points for each side of the specimen. (E) shows curvature of lateral flexions ( $1/\text{radius}$  of the best-fitted circle to curvature maximum [B]) plotted as a function of locus.

To quantitatively analyze the undulatory behavior of the lamprey, we utilized a computer algorithm (Ayers et al., 1983) that characterizes behavior via stop-frame analysis of video images. Filmed sequences of swimming were digitized frame by frame, and the data points describing the shape of the body axis were smoothed, fit with a fifth-order polynomial, and the curvature determined along the body axis (figure 14.1A). Undulatory movements are quantified in terms of regions of maximal curvature, defined as lateral flexions. Kinematic analysis of the propagation of flexions, which demonstrates swimming behavior, is organized into lateral flexion waves that propagate either rostro-caudally or caudorostrally (figure 14.1B, C, D). The time between initiation of flexion waves to the same side define the period of swimming movements; the time for the wave to propagate from nose to tail is the propagation time of the wave. During swimming, the propagation time



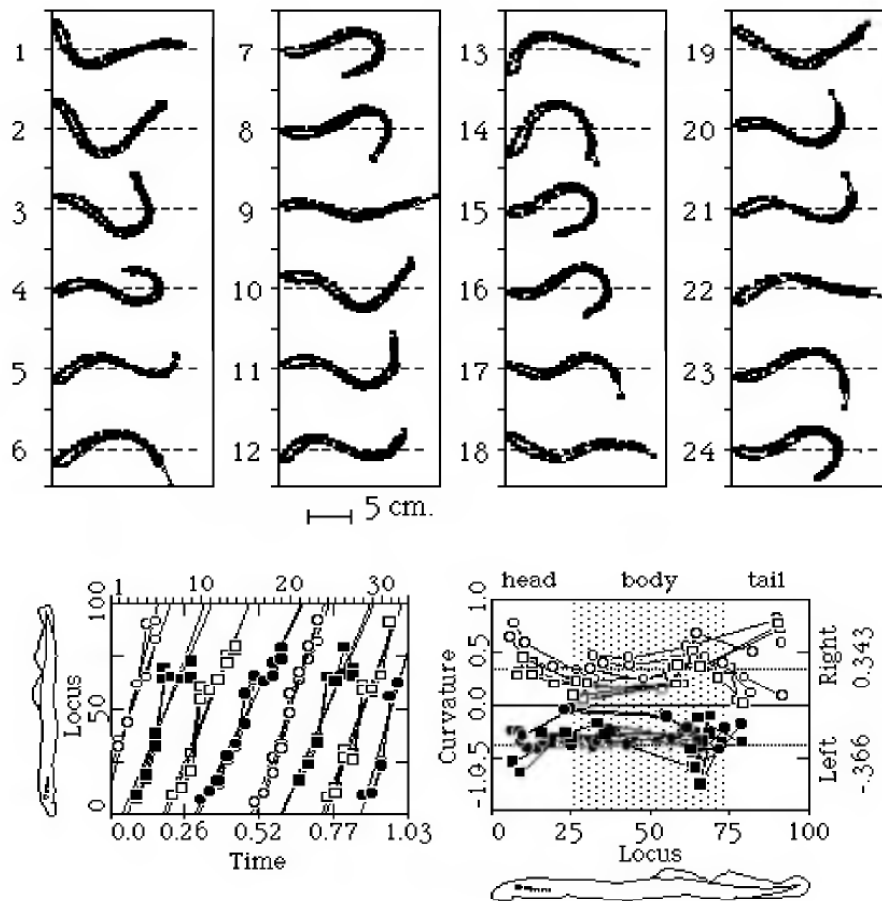
**Figure 14.2** (B) Plot of the locus of flexion waves versus time. Least-squares regression line is fitted to the points for each side of the specimen. (C) Plot of thrust versus time.

of flexion waves down the body axis is equal to the period, so specimens always maintain an S-shape during swimming (figure 14.1E).

Thrust is generated when a propulsive wave travels backward over the body at a velocity greater than the speed of forward swimming. Kinematic analysis from thrust experiments on larval lampreys (figure 14.2A) demonstrates that lampreys generate a peak of thrust as each flexion wave passes the cloaca (figure 14.2B). The thrust peaks at approximately 60% of the body and there are two peaks of thrust for every swimming cycle (figure 14.2C).

To evaluate how lamprey swimming movements are modulated by hydrodynamic load, we analyzed high frequency forward swimming of the lamprey in a shallow, wet pan sans water. The analysis of the swimming movements (figure 14.3) demonstrate that the active segments during swimming are located in the middle two thirds of the body in the curvature versus locus plot (figure 14.3C). The fact that the peak of thrust occurs at this same locus indicates that the mid-region of the body constitutes the propulsor for locomotion, and the tail responds passively to hydrodynamic resistance.

A behavioral library of the lamprey has been developed by reverse-engineering sequences of spontaneous lamprey action on video. This library consists of a table of timing parameters derived from the quantitative analysis of the various forms of behavior such as swimming, crawling, burrowing, withdrawal, and turning (figure 14.4, table 14.1). A look-up table generated from the behavioral library forms the basis for the controller of the vehicle.

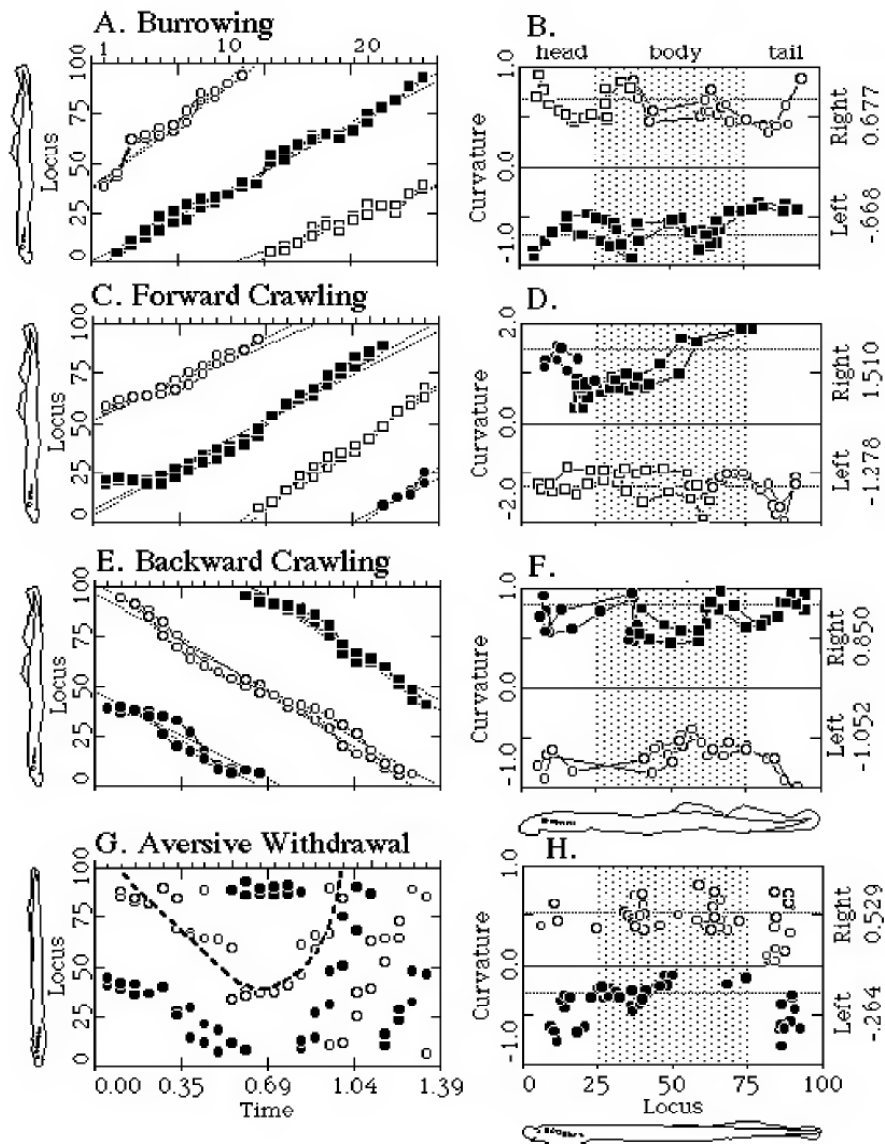


**Figure 14.3** (A) Graph of the normalized digitized images of swimming in a larval lamprey. (B) Plot of locus of flexion waves versus time. Least squares regression line is fitted to the points for each side of the specimen. (C) Plot of curvature versus locus.

## SYSTEM DESIGN

We adapted a conservative command neuron, coordinating neuron, central pattern generator (CPG) architecture to realize a finite-state machine that generates an alternating bilateral pattern that propagates along the body axis. This pattern actuates segmental muscles fabricated from shape memory alloy (SMA) arranged on a five-segment undulator formed from a polyurethane “notochord” to generate lateral axial undulations. The actuators are recruited to produce graded force through a pulse-width modulation system that recruits proportionally greater martensite state nitinol to the shorter austenite state (chapter 6 of this volume).

Control and power electronics are housed in a watertight hull (figure 14.5A). The system maintains trim and static roll stability using close-cell foam buoyancy elements and a lead shot “keel” covered by a Lycra skin (figure 14.5B). We are incorporating a sensor suite including a com-



**Figure 14.4** Analysis of normalized digitized images of different lamprey behavior: (A) burrowing, (B) forward crawling, (C) backward crawling, and (D) aversive withdrawal.


**Table 14.1** Timing and curvature parameters of different lamprey behavioral acts

Behavior	Period (sec)	Phase	Curvature
Slow swim	3.17	0.013	0.4
Average swim	1.90	0.013	0.37
Escape swim	1.52	0.013	0.43
Forward crawl	5.89	0.012	1.27
Backward crawl	10.61	-0.011	0.96
Burrowing	5.66	0.015	0.65

*Note:* Period was normalized to a 1 m vehicle based on extrapolation from measurements of period as a function of length for smaller lamprey. Phase is the intersegmental delay/period based on a 110 segment lamprey.

Adapted from Ayers, 1989.





**Figure 14.5** (A) The undulatory robot from nose to tail (average length 36 inches), watertight electronics bay, notochord of polyurethane, tail of fiberglass shimstock. (B) Detailed depiction of notochord and tail. Notochord comprised of teflon vertebrae (white), close-cell foam buoyancy elements between vertebrae, Lycra skin, and lead shot encased in latex.

pass, pitch and roll inclinometers, accelerometers, and both look-forward and look-down sonar range finders.

The undulatory vehicle design is a three-part system consisting of a rigid hull (~15% of the total body length), a flexible polyurethane notochord (~60% of the total body length), and a thin, flexible, passive tail (~25% of the total body length) (figure 14.5). This proportionality is based on our observations of the function of the lamprey body axis in the absence of hydrodynamic resistance (figure 14.4). Ten nitinol actuators for locomotion are located along the notochord (figure 14.5).

### Hull

The watertight cylindrical hull of plexiglass houses the electronics (a driver board, a controller board, two constant-current boards, a microprocessor [neural-circuit-based software]) and batteries. In the lamprey, mass distribution is biased toward the head—50% of body mass is located in the anterior 39% of the body (Currie, 1999). Watertight through-hull fittings allow for communication between the actuators and the processors.

### Axial Beam

The rectangular notochord of polyurethane consists of six Teflon vertebrae that attach and pin the SMA actuators along the midline. This arrangement transfers force to the flexible backbone causing it to bend. Each actuator is attached at three points on the notochord to minimize any offset distance from the polyurethane and to control shape changes (Chaudhry and Rodgers, 1991). Buoyancy elements of close-cell foam

are placed between the vertebrae and offset from the notochord so as not to impede bending of the polyurethane. These units also help maintain a uniform cylindrical body form. Lead shot weights encased in latex are placed parallel to and below the polyurethane but within the skin to keep the vehicle submerged and to maintain a dorsal-side-up orientation. Lycra covers the notochord, acting as a skin attached posterior to the hull, continuing the length of the notochord, and terminating between the fins of the tail. Lycra allows water to penetrate, effectively cooling the actuators but protecting them from any flotsam and jetsam in the water column.

## **Tail**

The tail, contiguous from the notochord, is tapered polyurethane overlain by fiberglass shimstock. The shimstock is shaped to resemble two dorsal fins, the rostral fin being much larger than the caudal. This arrangement allows for an increase in axial stiffness (shimstock has a higher Young's modulus) while maintaining the same bending stiffness, thus reducing energy loss in the tail region and inducing the undulatory wave to increase in amplitude as it progresses from head to tail.

## **Actuators**

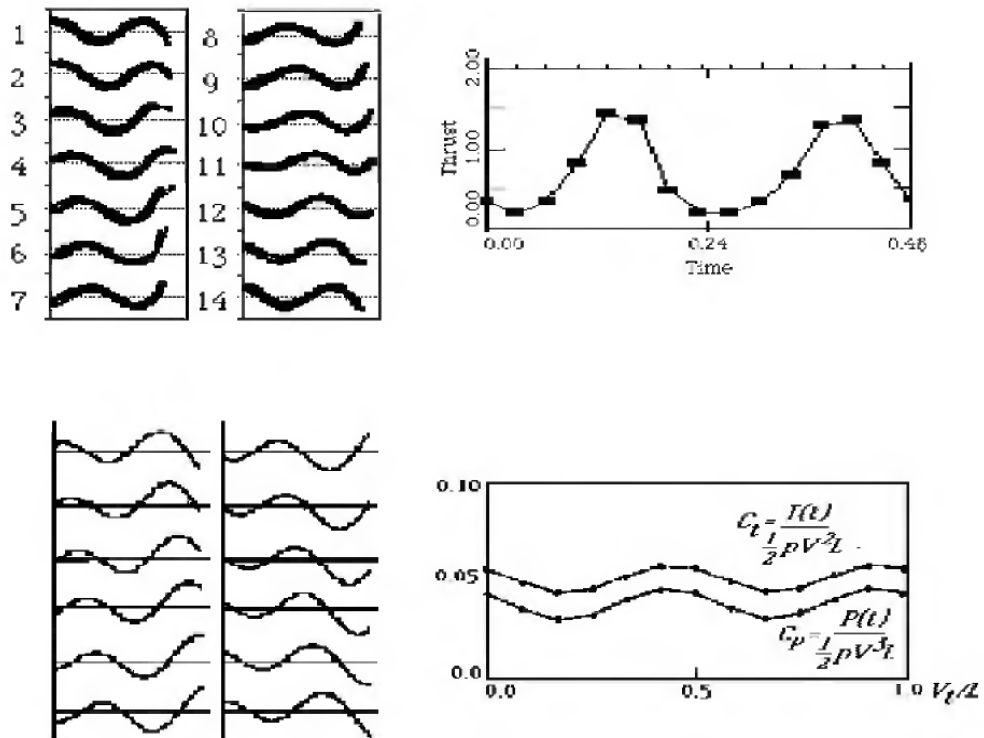
The vehicle is propelled by ten nitinol actuators, which consist of a watertight assembly of 0.01 in. nitinol, etched polytetrafluoroethylene (PTFE) tubing, stainless steel crimps, 32 AWG lead wire, Kevlar line, and Aquaseal, a waterproof sealant (chapter 6 of this volume).

This assembly is attached externally to the notochord to induce shape changes in the flexible polyurethane. Operating the actuators in a liquid (H<sub>2</sub>O) environment allows fast convective cooling necessary for shape control. Maximum cyclic frequency of the actuators is also important to generate the necessary propulsive forces.

The SMA actuators are controlled by neural-based circuit software that activates the actuators via pulse width modulation (PWM). Direct-current pulses of specific frequencies, amplitude, and duration are used, leading to the definition of a *duty ratio*, defined as the ratio of "on" time to the time between pulses (Furuya and Shimada, 1990).

## **HYDRODYNAMICS**

We modeled the hydrodynamics of the robot and optimized the hydrodynamics for wakeless propulsion by first comparing lamprey swimming to a hydrodynamic model of a wakeless propulsor (Vorus, 1995, 1996). Comparisons showed correlation between period and




**Figure 14.6** (A) Graph of normalized digitized images of swimming in a larval lamprey. (B) Analysis of images in (A), a plot of thrust versus time. (C) Graph of a theoretical wakeless propulsor. (D) Plot of thrust versus time.

propagation time as well as number of peaks of thrust per swim cycle (figure 14.6).

A beam model (figure 14.7) was then developed for the undulatory robot. This model assumes that the neural axis of a section of the composite beam structure (the notochord plus the nitinol [SMA] wires, one on each side of the notochord) is on the side of the notochord where the nitinol wire is not electrically active. The neutral axis switches back and forth from one side to the other side as the beam undulates. Given a displacement of the notochord (e.g., the wakeless motion) and external load, this model can provide the strain and stress of the wire, as well as the stress in the notochord, as functions of time for each section. The beam model was coupled with the 3-D slender body hydrodynamic model for external hydrodynamic load. This coupled model can be used for the prediction and/or evaluation of the external hydrodynamic loads, the internal load (stress), the input power, the propulsion efficiency, and the motion of the undulator.

## SYSTEM INTEGRATION

System integration is a simple input/output (I/O) interface from the processor to the controller. Each of the sensor elements is a switch and



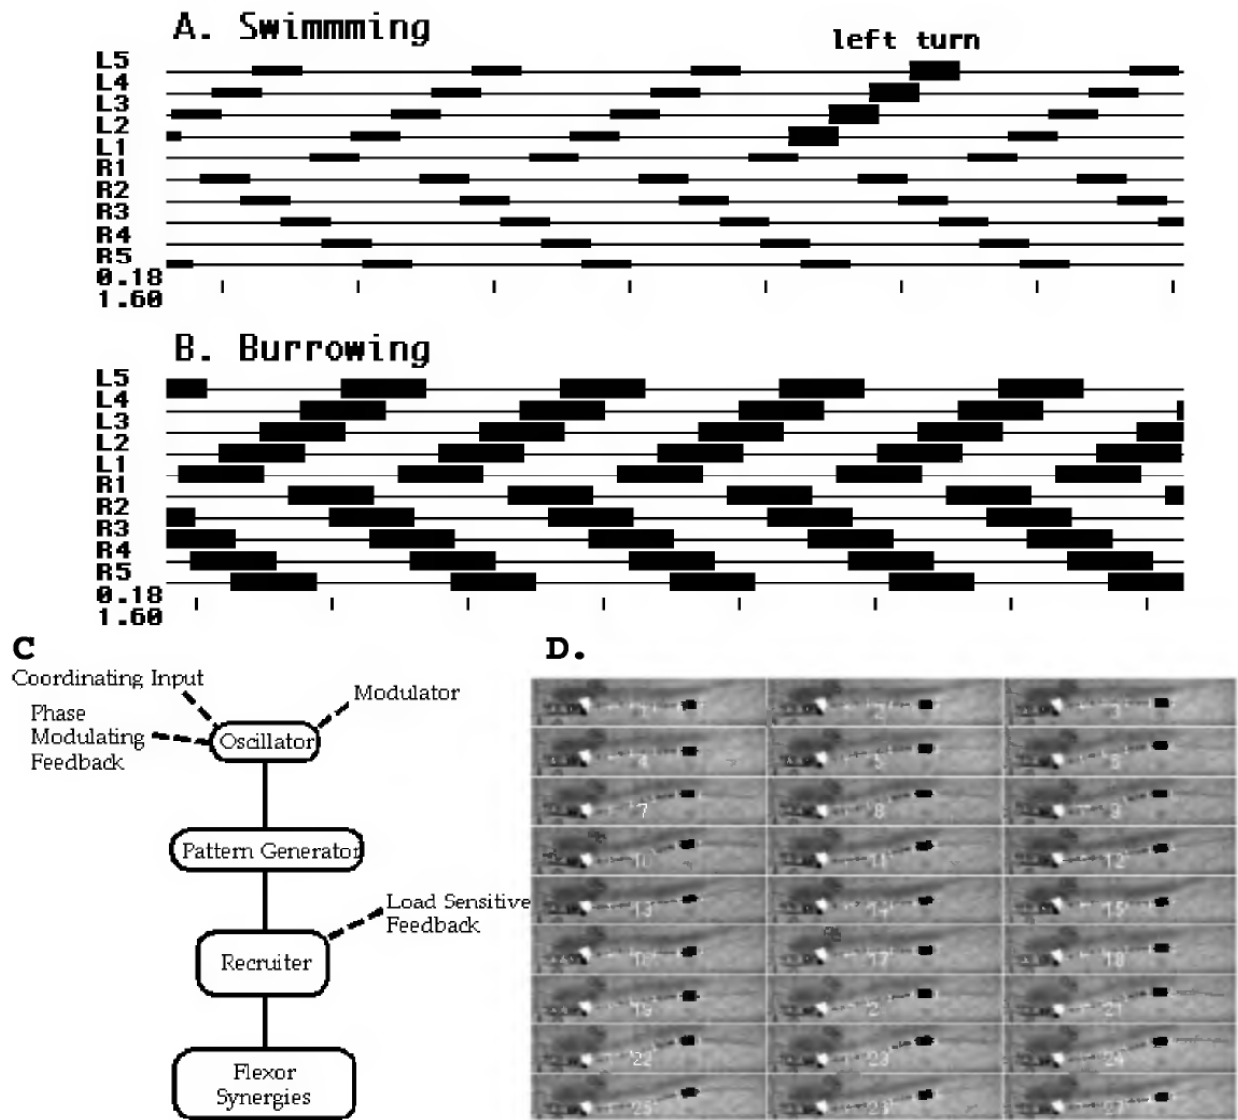
**Figure 14.7** A plot of the undulation predicted as needed to overcome a drag coefficient of 0.001 at a steady forward speed ratio of 0.9. This is the 3-D ideal displacement shape corresponding to “wakeless propulsion.”

each switch is mapped to the bit of a shift register, which the computer reads via a serial line. When a state change occurs in the motor pattern generating circuit, it writes, via a serial line, to a series of shift registers, each bit of which is mapped onto the current driver and then to a given actuator. The vehicle can be used autonomously with a small, embedded controller, or be supervised by a laptop for testing.

For our first prototype, we concentrated on reproducing the dynamic behavior of lampiform swimming. We developed a multisegment, headless vehicle comprising ten SMA actuators. The vehicle, when suspended in the water column and tethered from above, reproduced undulatory swimming motions (figure 14.8). Using the same design, only untethered and suspended in the water column with foam blocks, it produced straight-line forward swimming.

### SUPERVISED AUTONOMY

Interactions between sensors and commands will be established as releasers to increase the reactive capabilities of the vehicle. In addition, we plan to further formalize categories of modulation for different command systems—both in terms of parameteric or “ramp” modulation as well as taxic components. This releaser library associated with



**Figure 14.8** (A) Motor pattern for swimming behavior. Amplitude modulation characteristic of turning to the left is indicated about one third way through the trace. (B) Motor pattern for burrowing behavior. (C) Message passing hierarchy of the segmental state machine that generates the motor pattern. (D) Filmed sequences of the undulatory vehicle swimming.

segmental reflexes will form the basis of the reactive navigational capabilities of the vehicle and will be based on the behavioral library. These different adaptive behaviors will be implemented sequentially. Many behavioral acts involve basic sequences of more-fundamental action patterns and by expanding the library of sequences, we can include compensation for current and surge, collisions, disturbances of primary orientation, navigation, and searching.

The system will face numerous situations where it is presented simultaneously with the sensory releasers for incompatible behavior

(e.g., beacon tracking vs. surge compensation). We have established modulation of commands as the basis of choice by elaborating algorithms proposed by Edwards (1991) for crayfish. By establishing a behavior choice hierarchy based on a truth table of releaser combinations, we will mediate the behavioral hierarchy. This hierarchy will mediate interactions of a set of modal action patterns and goal-achieving behavior that will form the behavior set of the vehicle.

## CONCLUSIONS

This autonomous agent demonstrates the robust locomotion of its biological model. The robotic lamprey emulates the movement of a real lamprey by sequentially contracting its axial musculature. Sensory feedback is used to elicit behavioral changes due to environmental perturbations. For example, the robot will need to increase its swim speed when it encounters strong currents, turn away from or into obstacles and/or mine candidates, and dive and climb according to static pressure changes, therefore maintaining a constant depth in the water column. These are all examples of natural behavioral acts of the sea lamprey, *Petromyzon marinus*.

## ACKNOWLEDGMENTS

This work is supported by DARPA/DSO through ONR grant N00014-98-1-0381.

## REFERENCES

- Ayers, J. (1989). Recovery of oscillator function following spinal regeneration in the sea lamprey. In J. Jacklet (ed.), *Cellular and Neuronal Oscillators*. New York: Marcel Dekker, pp. 349–383.
- Ayers, J. (1992). Desktop motion video for scientific image analysis. *Adv. Imaging* 7: 52–55.
- Ayers, J. (2000). Finite state analysis of behavior and the development of underwater robots. In D. McFarland and O. Holland (eds.), *Artificial Ethology*. Oxford University Press (in press).
- Ayers, J., Carpenter, G. A., Currie, S., and Kinch, J. (1983). Which behavior does the lamprey central motor program mediate? *Science* 221: 1312–1314.
- Chaudhry, Z., and Rodgers, C. A. (1991). Bending and shape control of beams using SMA actuators. *J. Intell. Mater. Syst. Struct.* vol. 2: 581–602.
- Furuya, Y., and Shimada, H. (1990). Shape memory actuators for robotic applications. In T. W. Duerig et al. (eds.), *Engineering Aspects of Shape Memory Alloys*. London: Butterworth-Heinemann, pp. 338–355.
- Gosline, W. (1973). *Functional Morphology and Classification of Teleostean Fishes*. Honolulu: University Press of Hawaii.

- Hardisty, M. W. (1981). *The Biology of Lampreys*. London: Academic Press.
- Triantafyllou, G. S., and Triantafyllou, M. S. (1995). An efficient swimming machine. *Sci. Am.* 272: 64–70.
- Vorus, W. S. (1995). The concept of a traveling-wave propulsor for high efficiency and low wake signature. Twenty-fourth American Tow Tank Conference, College Station, Tex.
- Vorus, W. S. (1996). A reduced-wake marine propulsor concept. Annual report to Office of Naval Research, grant no. N00014-96-0124.

## Biomimetics in Action: Design and Performance of an Autonomous Robotic Fish

Jamie M. Anderson and Peter A. Kerrebrock

Unmanned undersea vehicle (UUV) technologies have evolved in recent years to produce highly functional and capable platforms for a wide variety of undersea missions. As the supporting technologies have progressed, so have the mission requirements. Today's UUV missions require a variety of capabilities that in some cases can be mutually exclusive: high transit speed, long range and duration, maneuverability, and station-keeping ability. Fish and marine mammals have captured the interest of vehicle designers, as they are able to cruise great distances at significant speed, maneuver in tight spaces, and accelerate and decelerate quickly from rest or low speed with the same integrated propulsion and steering system.

Recent research in the flow mechanisms used by fish and marine mammals for propulsion and maneuvering has demonstrated the utility of biopropulsion for undersea vehicles. Despite advances in UUV technology, little progress has been made in improving propulsive efficiency and maneuverability. Conventional UUVs employ a long slender hull with a propeller as the main propulsor and lifting surfaces that provide maneuvering control. Although several advanced demonstrations have been made in recent years with conventional designs, these types of vehicles are fundamentally limited in their maneuvering performance. They typically require several body lengths to execute a turn and can have fatally poor performance at very low speeds. Attempts to improve low-speed performance by using cross-axis thrusters have been effective, but the net result is generally loss of useful hull volume and degraded performance at higher speeds.

We present here the results of the first engineering demonstrations of the Draper Laboratory prototype flexible-hull UUV, which propels and maneuvers like a tuna. Named after the vorticity flow control mechanisms employed by fishes to propel and maneuver, the vorticity control UUV (VCUUV) mimics the form and kinematics of a large yellowfin tuna. Across the broad spectrum of fish form and movement, tunas are most desirable as a vehicle platform because they are very streamlined, relatively rigid in the forebody, and propel with low amplitude



movements in conjunction with a high-performance hydrofoil (caudal fin). Vorticity control propulsion is well suited to submarine-based UUVs, as it enables similar or greater transit distances and speeds as compared to conventional vehicles and order-of-magnitude improvements in maneuverability. The enhanced maneuvering performance is particularly desirable during launch and recovery and in missions that require operation in perturbed environments or in close proximity to the free surface and underwater objects.

A common misconception of the biopropulsion concept is that improved efficiency and maneuvering capability are not worthwhile when compared to the required displacement for the propulsion system. Past attempts to achieve both long-range and highly maneuverable rigid-body UUVs have met with disappointment, largely due to the need for both propulsion and maneuvering systems that rarely operate together effectively. Cross-axis thrusters, for instance, are both heavy and detract from useful vehicle displacement, and can only be used at zero speed without concern of control nonlinearities (that can even include control reversal). For high-speed maneuvers, for which thrusters are useless, a rigid body UUV must also have a conventional propulsor and control surfaces that cannot produce turning diameters less than several vehicle lengths.

By comparison, biopropulsion can provide both efficient propulsion and powerful maneuvering over a wide range of speeds with a modest investment in vehicle displacement. For example, the VCUUV propulsion system occupies only 23% of the total displacement (32% of the envelope displacement due to some free-flood volume). *Thunniform* (tuna) morphology and kinematics allow this modest propulsion system displacement with excellent steady swimming and maneuvering performance. Traditional propulsion and maneuvering systems typically occupy 5 to 10% of the envelope displacement.

The thunniform model possesses several advantages for use on an underwater vehicle system. As with other *carangiform* fishes, the propulsive movements are localized to the last 30–40% of the vehicle length and are moderate in amplitude. Tuna caudal-fin peak-to-peak excursions rarely exceed 15–20% of the body length. The localized tail motion allows the forward body to be used as a rigid housing for energy, intelligence, payload, and so on, without the complexity of flexible pressure hulls or the power required to move them. The tuna's fusiform forward body has nearly elliptical sections within which internal arrangements are more easily made than with cylindrical sections. The body is low drag but contains significant packaging space.

The utility of vorticity-control propulsion and maneuvering is apparent when one considers the current UUV missions of interest. Today's missions are challenging in that they often require long transit, long duration on site, loitering without loss of power, movement in close

proximity to objects for docking, tagging, and so on, and operation in dynamic environments such as shallow waters near the beach zone. Vorticity-control propulsion and maneuvering may prove to be essential in realizing these missions in that fishlike maneuverability may not be as energetically taxing as conventional means of generating large side forces (such as thrusters) in varying conditions. For example, fish can loiter at zero speed and rapidly accelerate in nearly any direction. Vorticity-control propulsion and maneuvering provides the range and speed of conventional low-drag hulls driven by propellers, with the maneuvering performance of ship-deployed remotely operated vehicles that can maneuver precisely at low speed.

## BACKGROUND

In recent years, researchers at the Massachusetts Institute of Technology have investigated rigid flapping foils (combined translation with simultaneous heave and pitch) as a means of generating large propulsion forces at very high efficiencies. Anderson and colleagues [1] have demonstrated propulsive efficiencies in excess of 85% with proper selection of the Strouhal number, angle of attack, heave amplitude ratio, and the phasing of the combined motions. Flow visualization studies with live fish have shown that the propulsion and maneuvering performance of fishes is related to their ability to control their wake vorticity [2]. Contrary to bluff bodies towed through the water, fish generate an oscillatory wake consisting of alternating vortices arranged in a jet pattern. Boundary layer relaminarization also affects the propulsive performance of fish [3].

The MIT research has culminated in the development of RoboTuna, a 1.2 m biologically inspired tow-tank model built to study propulsive efficiency and how it relates to body movement [4–5]. RoboTuna has been exercised in the MIT Ocean Engineering Testing Tank by prescribing a set of kinematic parameters (angular deflections of each joint, tow speed, phase relationships) and measuring the net power transmitted to the linkages and the reaction force between the tuna and carriage. Optimal motion was defined as the set of motion parameters that produced no net force on the carriage (self-propelled status) with minimum energy input for a given speed.

Hundreds of combinations of swimming parameters have been studied with promising results. The RoboTuna project has demonstrated that significant drag reduction can be achieved even without careful tuning of the drive mechanism. The RoboTuna system performed with thrust power ratios in excess of one and apparent drag reduction up to 70% compared to the “dead fish” drag [3]. For certain swimming parameters outside of the optimal range, drag amplification up to 300% was observed, which could be used as a very effective brake.

## VCUUV SYSTEM DESCRIPTION

At the Charles Stark Draper Laboratory, the next generation of flexible-hull robots has been developed and tested. The vorticity control unmanned undersea vehicle (VCUUV) (figure 15.1) is a freely swimming, tuna-shaped vehicle built to demonstrate tunalike swimming and maneuvering performance. Constructed as a proof-of-concept demonstration of vorticity-control propulsion and maneuvering, the VCUUV serves as a mission-scale exercise in design and packaging as well as a research platform with which swimming energetics and maneuvering performance are evaluated.

Although the VCUUV prototype was not intended as an ocean-going vehicle, the vehicle does contain all of the necessary components



**Figure 15.1** The Draper Laboratory VCUUV.

of autonomy required to execute fully functional autonomous missions: onboard energy, actuation, and control. The vehicle has been operated in shallow, controlled environments where the propulsion system performance has been studied. Vehicle design and fabrication issues addressed in the VCUUV project—including sizing of actuators, design of the articulated body, pressure hull design, and on-board intelligence, sensors, and power—are described in detail by Anderson and Kerrebrock [6–7]. The VCUUV was designed to operate in swimming pool and test-tank environments at depths less than 10 meters at typical live-animal cruise speeds (1 body length per second).


Twenty days of testing and approximately 100 autonomous experiments have been completed with the VCUUV, beginning with the vehicle's first swim in April 1998 in a local swimming pool. Preliminary trials were completed at the University of New Hampshire Ocean Engineering Laboratory in a 60 foot by 40 foot by 20 foot-deep freshwater tank. Autonomous swimming and simple maneuvers were achieved early in the experiments with no physical or electrical connections to the ground station.

In the summer and fall of 1998, the VCUUV was tested in two local freshwater lakes in Hopkinton State Park (Hopkinton, Mass.) and Nickerson State Park (Brewster, Mass.). The VCUUV lake experiments took advantage of the larger bodies of water by swimming longer and more-aggressive missions. Experiments in heading control were conducted yielding straight, circular, square, and zigzag trajectories. Swimming kinematic parameters were varied for straight swimming and maneuvering to identify the key parameters and how they affect swimming performance.

### **System Characteristics**

The need to demonstrate a fieldable vehicle that can serve as a research tool influenced both the mechanical and electrical system designs. The mechanical system was required to allow independent control of tail-element oscillation (phase and amplitude of each link) at variable frequencies. The articulated tail portion of the vehicle required nearly neutral buoyancy so as not to introduce undesirable rolling moments during oscillation. The hydraulic actuation system was designed to propel the vehicle at 2.5 knots for up to 3 hours of testing. The tail articulation range was based on the requirement for a desired turning circle diameter of 2.5 lengths or less. Although designed for modest depths (<10 m), the hull was robustly designed to accommodate the loads of handling and collisions.

The resulting design is 8 feet (2.4 m) long and displaces approximately 300 lb (1334 N). Figure 15.2 illustrates the system architecture and major elements. The forward half of the vehicle is a rigid pressure



**Figure 15.2** VCUUV system layout.

hull, while the aft portion is articulated and freely flooded. The pressure hull houses all of the dry components—including batteries, electronics (computer, sensors, etc.), hydraulic power plant, and payload.

The articulated tail structure consists of a planar, four-degrees-of-freedom robot arm that acts upon a tuna-shaped flexible exostructure. The entire tail area is freely flooded with water and covered with an impermeable skin that prevents exchange of water across the envelope boundary. Three equal-length links and the caudal fin are independently actuated with hydraulic cylinders integrated into the links themselves.

The tail exostructure converts the discrete movements of the underlying robot linkage to the smooth, fluid movements of a fish tail. The key elements of the exostructure are rigid buoyant foam “ribs” that are shaped like hollowed out tuna steaks and are connected to one another by flexible splines. The splines produce smooth curvature along the tail envelope, which simulates a linkage of much higher resolution (as though there are a greater number of links). At three locations along the tail, follower rod mechanisms transmit the hydrodynamic loads on the exostructure to the actuation linkage.

The structure-fluid interface is a dermal layer of overlapping scales that prevent transverse folds between the ribs, similar to their function in real fish. The outermost layer consists of neoprene rubber that is bonded to Lycra on both sides. This “skin” prevents water movement across the boundary while maintaining an extremely smooth flexible surface.

The VCUUV is outfitted with an array of sensors that serve as the science payload. The science mission in the field trials has been the evaluation of the design performance over a range of kinematic inputs:

frequency of tail oscillation, amplitude of tail lateral excursion, propulsive wavelength, and caudal fin angle of attack and phase. The sensor suite includes a six-axis inertial measurement unit (IMU), including three accelerometers and three gyros, a compass, tail linkage position sensors (4), tail linkage force sensors (4), a depth sensor, and other health-monitoring sensors (internal temperature, hydraulic pressure, leak detection, etc.). All sensors except for those dedicated to the tail linkage are located in the electronics assembly inside the pressure hull. In addition, an external speed sensor was utilized in some experiments. By measuring the linkage force and displacement as functions of time, the instantaneous power absorbed by the tail assembly can be calculated as the product of force times velocity for each link.

The requirements for control of the vehicle were precise and robust control of the tail linkage, such that its movements adhere to the prescribed linkage trajectory. Closed-loop control of the vehicle position was not implemented, as the goal of the experiments was to identify open-loop characteristics that could later be used to design a closed-loop trajectory controller. Heading control was explored through compass feedback on a variety of trajectories. Depth control was tested but not implemented, as the vehicle was passively stable in depth due to careful ballasting.

The key parameters in thunniform steady swimming are the vehicle speed, oscillation frequency and amplitude, propulsive wavelength, shape parameters (which describe the envelope of transverse motion), tail phasing, and angle of attack [4]. For steady straight swimming, a traveling propulsive wave is required that moves from head to tail, with increasing amplitude near the tail. The kinematics are described in detail by Barrett [4], Dewar and Graham [8], and Anderson and Kerrebrock [6–7]. Turning kinematics are less understood, and consequently, the VCUUV design has most of its design margin in the turning requirements (amplitudes and loading conditions).

As in actual tunas, dive-plane control is provided by means of two pectoral fins located near the mid-body. Real tunas, which are negatively buoyant, use the pectoral fins for speed-dependent lift to control vertical position in the water column. At low speeds, the fins are fully splayed for maximum lift; as the speed increases, the fin sweepback angle is increased until the fins are fully adducted, effectively reducing the lifting area [8].

The VCUUV pectoral fins are positioned in the same location as real tuna pectoral fins, slightly forward of the center of gravity. Thus, the pectoral fins act as canards. Unlike actual tunas, which can articulate their pectoral fins in pitch and sweep angle, the VCUUV pectoral fins only pitch, controlled by motors located inside the pressure hull. For simplicity, the pectoral fin shapes are swept and tapered NACA 0015 sections roughly matching the chord and span of actual tuna fins.

Because of the large metacentric height of the vehicle, pectoral fin lift causes little change in the vehicle while controlling depth.

## VEHICLE PERFORMANCE

Early field trials focused on engineering diagnostics and adjustments—especially in the tuning of the hydraulic system and control system performance. Hydraulic pressure, control gains, and the desired tail kinematics were adjusted to give good tracking performance for tail oscillation frequencies up to 1 Hz. Tail oscillations above 1 Hz were not pursued due to larger-than-expected forces that led to poor control performance. Simple maneuvers were explored including coasting turns and biased swimming movement resulting in circular trajectories. Because of the short distances available in the tank trials, steady-state swimming at terminal velocity could not be achieved. Thus, the field trials were moved to local freshwater lakes for longer duration and higher-speed tests.


In the summer and fall of 1998, the VCUUV was tested in freshwater lakes in Hopkinton State Park and Nickerson State Park. These sites were chosen for their proximity to Draper Laboratory and their clean, undisturbed fresh water. A systematic study of tail kinematic parameter variations was completed as well as several basic open- and closed-loop horizontal plane maneuvers.

A typical experimental run began with the vehicle at rest and submerged while baseline inertial data were collected. The vehicle then began the swimming tail motion that accelerated the vehicle to steady-state velocity. The typical acceleration period was 20–30 seconds for 1 Hz tail oscillation. After swimming a total of 1 minute, the vehicle coasted to a stop while the tail was held in a straight position. The tail movement for swimming adhered as well as possible to the kinematics reported by Barrett [4] for optimal performance.

### **Straight Swimming Performance**

The maximum speed attained in the field trials was 1.25 m/sec (2.4 knots) at 1 Hz tail oscillation while under heading control (closed-loop control using compass feedback). Without heading control, the maximum speed was 1.19 m/sec (2.3 knots). Figure 15.3 shows the steady-state speed as a function of frequency as well as the speed predicted by extrapolating the RoboTuna results to the VCUUV vehicle size using a body-undulation-wave-speed-based scaling law. The published RoboTuna result is at a single speed that is also marked in figure 15.3. The VCUUV speed was measured by both speed sensor and time trials with good agreement between the two measurements.

The speed increases with frequency up to 1 Hz at which point the tail kinematic performance suffers and the speed decreases due to hard-



**Figure 15.3** VCUUV steady state speed as a function of tail oscillation frequency.


ware limitations of the hydraulic system. The actual speed of the vehicle was less than the predicted speed by 16%, which may be due to differences in the kinematics between the VCUUV tail motion and that of the RoboTuna. Using the same length reference (prepeduncular length), the VCUUV achieved 0.61 lengths per second (L/sec) whereas the RoboTuna achieved 0.65 L/sec. Unfortunately, degraded performance above 1 Hz did not allow the vehicle to achieve the design speed of 1.0 L/sec.

### Maneuvering Performance

The VCUUV possesses excellent maneuverability that matches the performance reported for live animals [9]. Several maneuvering trajectories were explored, including abrupt discrete turns and continuous turning, which yielded circular and spiral trajectories. Small stationary turning circles (8–10 ft diameter) were easily achieved by simple biasing of the straight swimming tail motion.

Figure 15.4 illustrates the heading angle measured during a coasting turn with full body deflection executed at full speed (1.2 m/sec). Two measurements are indicated in the figure: the integral of the heading gyro yaw rate and the absolute heading angle measured by the compass. The data from both sensors are in agreement and indicate that the vehicle turned through approximately 180° in less than 10 seconds. The maximum measured heading rate during this maneuver was 75°/sec. As a point of comparison, a conventional UUV with control






**Figure 15.4** Heading angle as a function of time during a hard turn.

surfaces turns at approximately 3 to 5°/sec, often requiring several body lengths and as long as a minute to complete a 180° turn. This drastic improvement in maneuverability may enable mission elements previously considered beyond the capabilities of UUVs, including close inspection tasks, rapid course changes, and operation in highly perturbed environments such as launch and recovery to a moving submarine.

Figure 15.5 illustrates the heading angle achieved in an aggressive zigzag trajectory. The vehicle swam straight for 20 seconds to attain steady-state velocity and then swam zigzag legs with 90° heading changes every 10 seconds. The vehicle maintained forward velocity during the maneuvers at 2.4 knots. The maximum yaw rate for this maneuver was approximately 30°/sec, which also outperforms conventional systems by several factors. The zigzag maneuver indicates the level of performance that can be achieved in simple search patterns or in close inspection missions. The vehicle can make aggressive course changes in-stride while preserving forward speed.

### SUMMARY

The Draper Laboratory VCUUV is the first autonomously operated fishlike UUV that is realistically sized for real-world missions and pay-



**Figure 15.5** Heading angle as a function of time during zigzag maneuvers.

loads. It is also the first vehicle of this type able to move with arbitrary thunniform kinematics for scientific exploration of straight swimming and maneuvering. The fish propulsion paradigm offers order of magnitude improvement in maneuvering capability that is required for today's challenging UUV missions in dynamic, cluttered environments. The VCUUV prototype proves the concept while serving as a research test bed. Lessons learned from the VCUUV can be applied to new vehicles mechanically optimized to achieve fishlike capabilities in engineered vehicles.

Further work will demonstrate that this novel propulsion mechanism can be incorporated into vehicles capable of practical missions that are sufficiently compelling to justify the risk of their development. This effort will include demonstrating the scalability of vorticity-control propulsion (both up and down) and developing more-stealthy actuators, three-axis maneuvering capability, and sensors and control systems to support fully closed loop propulsion and maneuvering control.

#### **ACKNOWLEDGMENTS**

This project has been funded internally by the Charles Stark Draper Laboratory.

## REFERENCES

- [1] Anderson, J. M., Streitlien, K., Barrett, D. S., and Triantafyllou, M. S. (1998). Oscillating foils of high propulsive efficiency. *J. Fluid Mechanics* 360: 41–72.
- [2] Anderson, J. M. (1996). Vorticity control for efficient propulsion. *Ph.D. thesis*, Massachusetts Institute of Technology/Woods Hole Oceanographic Institution Joint Program.
- [3] Barrett, D. S., Triantafyllou, M. S., Yue, D. K. P., Grosenbaugh, M. A., and Wolfgang, M. J. (1999). Drag reduction in fishlike locomotion. *J. Fluid Mechanics* 392: 183–212.
- [4] Barrett, D. S. (1996). Propulsive efficiency of a flexible hull underwater vehicle. *Ph.D. thesis*, Massachusetts Institute of Technology.
- [5] Triantafyllou, M. S., and Triantafyllou, G. S. (1995). An efficient swimming machine. *Sci. Am.* (March).
- [6] Anderson, J. M., and Kerrebrock, P. A. (1997). The vorticity control unmanned undersea vehicle—An autonomous vehicle employing fish swimming propulsion and maneuvering. In *Proceedings of the Tenth International Symposium on Unmanned Untethered Submersible Technology*, Durham, N.H., 7–10 September. pp. 189–195.
- [7] Anderson, J. M., Kerrebrock, P. A., and Triantafyllou, M. S. (1997). Concept design of a flexible-hull unmanned undersea vehicle. In *Proceedings of the Seventh International Off-shore and Polar Engineering Conference*, vol. 2, Honolulu, 25–30 May. pp. 82–88.
- [8] Dewar, H., and Graham, J. B. (1994). Studies of tropical tuna swimming performance in a large water tunnel: III. Kinematics. *J. Exp. Biol.* 192: 45–59.
- [9] Blake, R. W., Chatters, L. M., and Domenici, P. (1995). Turning radius of yellowfin tuna (*Thunnus albacares*) in unsteady swimming manoeuvres. *J. Fish Biol.* 46: 536–538.

# 16

## Development of a Two-Joint Dolphin Robot

Motomu Nakashima and Kyosuke Ono

The swimming speeds of tuna and dolphins are greater than 30km/hr, and the swimming motion of these animals is thought to be highly efficient; therefore, the swimming movements of these animals have been studied extensively. Lighthill [1] applied the two-dimensional airfoil theory of aerodynamics to the caudal fin of fish for a high aspect ratio and analyzed the performance of a flat-wing model of the caudal fin with respect to the ratio of amplitudes of the heaving and pitching motions and with respect to the location of the pitching axis. Chopra [2–3] added to these findings by studying three-dimensional rectangular wings and two-dimensional wings having large amplitude. In addition, Chopra and Kambe [4], Cheng and Murillo [5], and Karpouzian, Spedding, and Cheng [6] analyzed the performance of a three-dimensional lunate tail. G. S. Triantafyllou, M. S. Triantafyllou, and Grosenbaugh [7] showed experimentally that the optimal efficiency of an oscillating foil is obtained at the frequency of maximum amplification of wake vortices.

On the other hand, few studies have examined experimentally the consumed power and propulsive efficiency of a large (approximately the same size as an actual tuna or dolphin) self-propelled swimming robot. Triantafyllou and colleagues [8] made a full-scale tuna robot (RoboTuna) and performed experiments. Although the experiments used a tow carriage, they did consist of “self-propelled” tests by driving the system to have zero net force on the tow mast. Anderson [9] developed the VCUUV (vorticity control unmanned undersea vehicle), which is the mission-scale fishlike UUV, and demonstrated its propulsive speed and maneuverability. However, the experimental data of propulsive characteristics on this type of swimming robot are not yet sufficient. Therefore, the present authors developed two two-joint dolphin robots as simplified models of high-speed swimming animals. Both robots have the following features:

- Designed to study forward swimming at constant speed.
- Self-propelled and self-contained.

- Have two joints (minimum, to actuate the caudal fin). The first joint is actuated by an active motor and the second joint is moved passively by springs.
- No buoyancy control (swims near water surface).

The propulsive characteristics of these robots, such as the propulsive speed and the propulsive efficiency, are experimentally investigated.

This type of robot is suitable for the mission to research the fragile ocean nature, such as coral seas, because the local velocity of the caudal fin is basically an equal level to the propulsive speed of the body. However, the local velocity of the conventional screw propeller is basically much higher than the propulsive speed of the body, resulting in a large wake and harm to the fragile ocean nature.

### EXPERIMENTAL APPARATUS OF THE FIRST ROBOT

Figure 16.1 shows the first two-joint dolphin robot constructed for the present study, and figure 16.2 shows the schematic diagram of the robot, which consists of two sections: a streamlined body section, which has a circular cross section along the entire length, and the caudal fin section connected to the rear of the body section. The total length of the body is 1.75 m. The two joints are actuated vertically in the same direction as marine mammals'. The first joint axis is located at a position three-quarters of the body length from the front of the robot. The second joint is located at the connection point of the body and the caudal fin (also the leading edge point of the fin). The first joint is actuated by an air motor, whose maximum power is 230 W. The robot contains a compressed air tank, thus eliminating the need for an umbilical cable.



**Figure 16.1** Photograph of the first two-joint dolphin robot.

Figure 16.3 shows the air flow circuit for actuation of the first joint. First, compressed air from a separate air compressor is used to fill the air tank. The air in the tank then flows to the mechanical valve, which can be opened and closed by remote control. Next, the air flows into the air motor through the flow speed controller. The drain from the air motor flows into the drain tank and the exhaust flows out from the body through the check valve, which prevents water from flowing back into the body.

The rotational motion of the air motor is transformed into reciprocal motion through a system of gears and a crank, as illustrated in figure 16.4. First, because the rotation axis of the motor is parallel to the body axis, a pair of gears is used to make the rotation axis perpendicular. This rotation becomes the rotation of the first link (see figure 16.4) and is then transformed into the reciprocal motion of the second link. Finally, this reciprocal motion is magnified twice through a pair of gears and is used to actuate the caudal fin. The amplitude of the motion of the first joint can be altered easily by adjustment of the length of the first link. For example, in figure 16.4, three amplitudes are obtained.

The second joint is moved passively by springs. As shown in figure 16.5, two wires are fixed to a pulley, and a spring is connected to each of the wires. When the fluid force acts on the caudal fin, the caudal fin moves passively, reducing the fluid force. The motion of the second joint becomes similar to the ideal motion of a propulsor. Figure 16.6 shows the motion of the caudal fin. When the body moves up, the fin is




Figure 16.2 Schematic diagram of the two-joint dolphin robot.

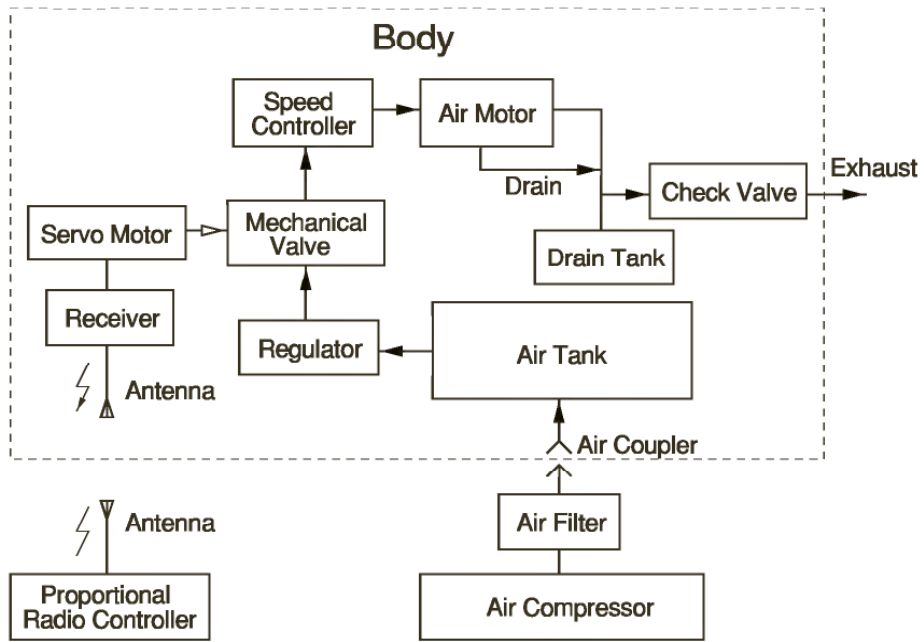


Figure 16.3 Air flow circuit for the first joint.




Figure 16.4 Crank mechanism for the first joint.




Figure 16.5 Spring mechanism for the second joint.




Figure 16.6 Motion of the caudal fin.

rotated in the clockwise direction due to the fluid force. Therefore, the fin motion becomes an appropriately blended motion that consists of both heaving and pitching, as suggested by Lighthill [1].

The aspect ratio of the caudal fin is 5. The section of the caudal fin was constructed according to NACA 65<sub>3</sub>-018. The planform of the caudal fin is rectangular to allow for ease of production.

The body shape was constructed according to NACA 65<sub>4</sub>-021. The body consists of a hexagonal aluminum frame and a urethane foam skin to provide a streamlined shape, and the surface was coated using plastic dip (rubber coating). A tubal rubber sheet was glued around the first joint so that it was free to move over a wide angle. The posterior part of the body was not circular but elliptical in order to avoid unnecessary drag that occurs in reciprocal motion.

## EXPERIMENTS WITH THE FIRST ROBOT

Figure 16.7 shows the system used to measure the angle and torque of the first joint. A variable rotational condenser, connected coaxially to the first joint, is used to detect rotation. Variation of the joint angle produces variation in the capacitance of the condenser, and because the condenser is placed in an oscillating circuit, the variation in capacitance is transformed into the variation of the oscillating frequency of the circuit. This oscillating signal is then transmitted to a radio receiver, and finally sent to a notebook computer through the f/V (frequency to voltage) converter. A strain gauge was placed on the shaft of the first joint axis in order to detect twist of the joint axis. The signal from the strain gauge is sent to a receiver by a telemeter transmitter, and then sent to a notebook computer through the telemeter receiver.

Figure 16.8 shows the robot swimming in water. At present, the robot can control neither its depth nor its pitching angle, and therefore swims near the surface, causing surface waves to occur. The distance traveled by the robot is recorded using an 8 mm video camera. Figure 16.9 shows the propulsive speeds  $U$  attained by the robot when the amplitude of the first joint  $\theta_{max}$  is  $20^\circ$  and the spring constant for the second joint is 2.016 nm/rad. Because  $U$  is calculated as a reciprocal



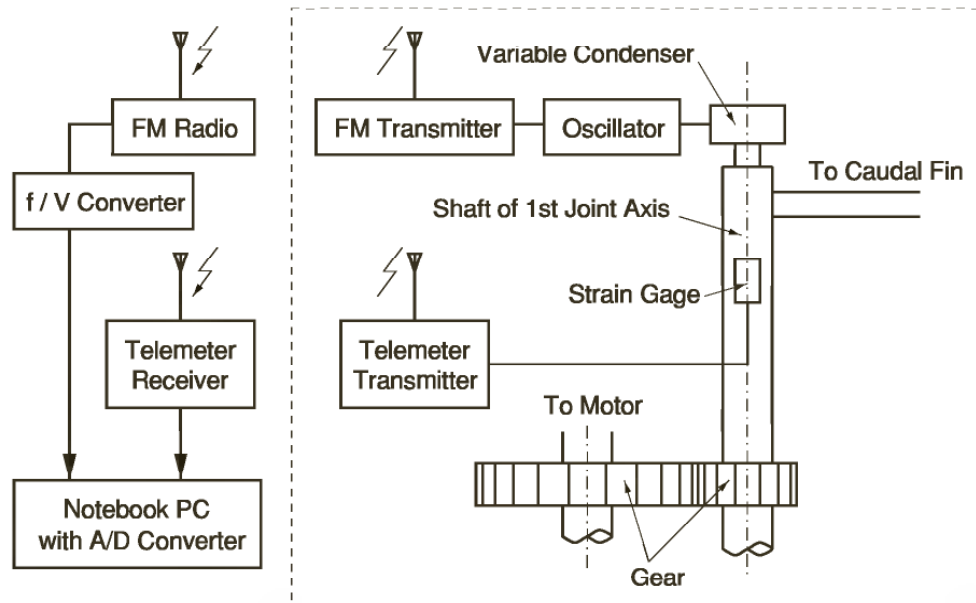



Figure 16.7 Measurement system.




Figure 16.8 Swimming dolphin robot.

number of the period for a one-meter advance of the robot, the graph is somewhat spiky. The dashed-dotted line represents the average propulsive speed. Figure 16.10 shows the variation of joint angle  $\theta$  and joint torque  $T$  with respect to time. Using joint angle  $\theta$  and joint torque  $T$ , the consumed power  $P$  is calculated as:

$$P = T \times \frac{d\theta}{dt} \quad (16.1)$$



**Figure 16.9** Propulsive speed of the dolphin robot. (The amplitude of the first joint  $\theta_{max}$  is 20 deg and the spring constant for the second joint is 2.016 Nm/rad.)



**Figure 16.10** Variation of the joint angle  $\theta$  and the joint torque  $T$ . (The amplitude of the first joint  $\theta_{max}$  is 20 deg and the spring constant for the second joint is 2.016 Nm/rad.)

Figure 16.11 shows the consumed power  $P$  obtained using equation (16.1) and the results shown in figure 16.10. The graph of the consumed power  $P$  becomes somewhat spiky. This is because the spikes of the joint angle  $\theta$  shown in figure 16.10 are emphasized when  $d\theta/dt$  is obtained from differentiating  $\theta$  numerically.

The valid work done by the robot is considered to be a product of the towing drag  $D$  for the robot at rest and the average propulsive speed  $U$ . Therefore, the propulsive efficiency  $\eta$  is defined as:

$$\eta = \frac{DU}{\bar{P}} \quad (16.2)$$

where  $\bar{P}$  represents the time average of  $P$ .

In order to obtain the towing drag  $D$ , two measurement methods were conducted. Figure 16.12 shows the towing method in a swimming pool. The body of the robot is towed near the water surface in this method. Figure 16.13 shows another towing method in a circulating




Figure 16.11 Consumed power with respect to time. (The amplitude of the first joint  $\theta_{max}$  is 20 deg and the spring constant for the second joint is 2.016 Nm/rad.)




Figure 16.12 Towing method at pool.




Figure 16.13 Towing method at circulating tank.

tank. In this measurement, the two values of drag—that is, one at 200 mm depth (near the water surface) and another at 500 mm depth—can be measured. Figure 16.14 shows the experimental results of body drag. The values of the pool and those of the tank at 200 mm depth become very close. On the other hand, the values of the tank at 500 mm depth become considerably small. This may be because the effect of the wave-induced drag is large when the body is near the water surface.

Figure 16.15 shows the relationship between the motion frequency  $\nu$  and the propulsive speed  $U$  when the amplitude of the first joint  $\theta_{max}$  is




Figure 16.14 Experimental results of body drag  $D$ .





Figure 16.15 Relationship between motion frequency  $\nu$  and propulsive speed  $U$  ( $\theta_{max} = 20$  deg).

Table 16.1 Spring constants for four springs of the second joint (first robot)

	Spring Constant (Nm/rad)
Spring 1	0.240
Spring 2	0.978
Spring 3	2.016
Spring 4	3.060

20°. In this figure, the output power of the air motor and the spring of the second joint are varied as parameters. The spring constants of these springs are listed in table 16.1. Generally, higher motion frequency was found to produce higher propulsive speed, and a stiffer spring in the second joint was found to produce higher propulsive speed except for the case of a fixed joint (meaning the spring constant is infinity). A maximum propulsive speed of approximately 1.15 m/sec was attained by



**Figure 16.16** Relationship between reduced frequency  $\sigma$  and propulsive efficiency  $\eta$  ( $\theta_{max} = 20$  deg).

the robot. The propulsive efficiency was calculated using the measured  $U$  and  $\bar{P}$ .

Figure 16.16 shows the reduced frequency  $\sigma$  and the propulsive efficiency  $\eta$  obtained using these calculations. Lighthill [1] defined the reduced frequency  $\sigma$  as:

$$\sigma = \frac{\omega c}{U} \tag{16.3}$$

where  $c$  represents the chord length of the caudal fin. A lower reduced frequency  $\sigma$  was found to produce a higher propulsive efficiency  $\eta$  except for the case of a fixed second joint. This tendency may occur because a lower reduced frequency  $\sigma$  produces a higher efficiency of the caudal fin—proven theoretically by Lighthill [1]. A stiffer spring produces a lower propulsive efficiency  $\eta$  for a constant reduced frequency  $\sigma$ . The maximum propulsive efficiency  $\eta$  reaches around 0.7 and does not vary much with respect to the spring constant except for the case of a fixed second joint.

### EXPERIMENTAL APPARATUS OF THE SECOND ROBOT

Figures 16.17 and 16.18 show the outline and a photograph of the second dolphin robot, respectively. The total length of the body is 1887 mm. The first and second joints are located at the 68% position of the body length and at the end of the body, respectively. The body shape was produced with reference to NACA 65<sub>4</sub>-021. A basket-shaped aluminum frame is used for the anterior part of the body. In addition, a carbon fiber reinforced plastic (CFRP) case wraps the aluminum frame. For the posterior part of the body, several FRP (fiber reinforced plastic) rings are connected so that the body bends freely and does not collapse due to fluid pressure. The whole body is surrounded by a special

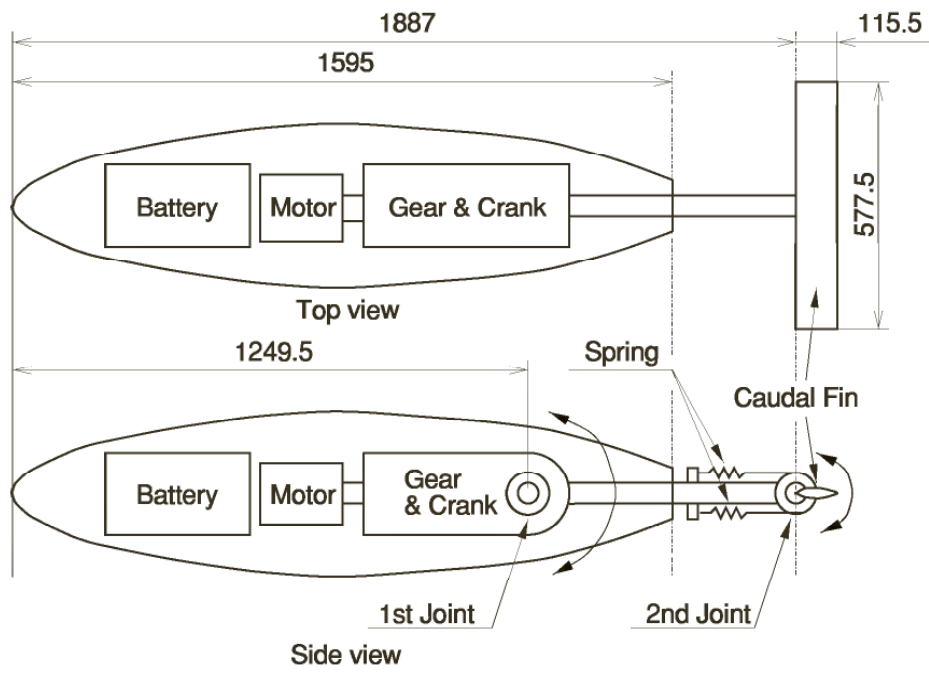


Figure 16.17 Outline of the second two-joint dolphin robot.



Figure 16.18 Photograph of the second dolphin robot.




Figure 16.19 Crank mechanism for the first joint.

Table 16.2 Spring constants for springs of the second joint (second robot)

	( $\times 10^4$ N/m)
Spring 1	14.056
Spring 2	11.471
Spring 3	8.070
Spring 4	5.544
Spring 5	3.216
Spring 6	1.484

waterproof swimsuit, made of the same material as dry suits for humans. A rectangular wing, which has a NACA 65<sub>3</sub>-018-shaped section and an aspect ratio of 5, is used for the caudal fin.

The first joint is driven by a car starter motor (its maximum power is 7.4 kW), which produces high power for a short period. The rotational motion of the motor is transformed into an up-and-down motion in the shaft to the caudal fin by a crank mechanism through reduction gears. Figure 16.19 illustrates the crank mechanism of the first joint. Figure 16.19 only shows the view from the left side of the body. The right side of the body has a similar mechanism, although there is no pin from the rod and the phase of the crank mechanism has a 90° lag in order to obtain smooth rotation. The rotation of the first link and the sublink is transformed into the up-and-down motion of the second link and the pin of the rod from the caudal fin. The up-and-down motion of the pin produces a corresponding motion in the rod. The amplitude of this motion is varied by changing the length of the first joint and the sublink. The second joint is supported by springs similar to the first robot's. Table 16.2 shows the spring constants of the six springs that are used for the second joint in the experiment.

The first dolphin robot did not have a direction control system in the yawing direction. But the second dolphin robot has a structure in which the posterior part of the body (which includes the starter motor, the crank mechanism, and the caudal fin) can perform the yawing motion. In addition, the yawing angle can be remotely controlled.

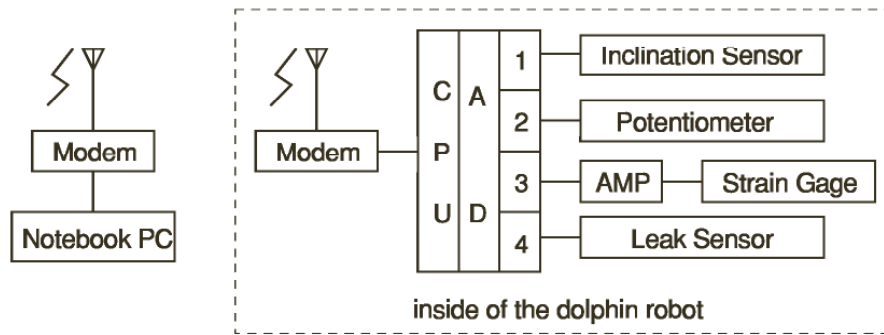


Figure 16.20 Measurement system.




Figure 16.21 Relationship between motion frequency and propulsive speed.

## SWIMMING EXPERIMENTS

The swimming experiment was performed using the second dolphin robot. Figure 16.20 shows the setup of the measurement system. The angle of the first joint, that of the second joint, and the torque of the second joint can be measured by this system. A potentiometer attached to the side of the body detects the relative angle between the body and the rod to the caudal fin in order to measure the angle of the first joint. For the torque of the first joint, a strain gauge attached to the base of the rod detects the bending of the rod. These analog signals are transformed into digital signals by A/D converters in the microcomputers in the dolphin robot, and the digital signals are sent to a wireless modem, which communicates with a notebook PC.

Figure 16.21 shows the relationship between the motion frequency and the propulsive speed. It was found that faster propulsive speed is achieved by a stiffer spring and that the gradient of the curve becomes more gentle as the spring becomes weaker. It was also found that the maximum propulsive speed reaches around 1.9 m/sec.






Figure 16.22 Relationship between towing speed and body drag.




Figure 16.23 Relationship between motion frequency and reduced frequency.

Figure 16.22 shows the results of the measurement of the body drag. All the points have little deviation from the fitting curve. This fitting curve is used for the calculation of the propulsive efficiency.

In order to investigate the relationship between the reduced frequency  $\sigma$  and the propulsive efficiency  $\eta$ , we first investigate the relationship between the motion frequency  $f$  and the reduced frequency  $\sigma$ . Figure 16.23 shows these results. It can be seen that, when the spring is weak (such as spring 6), the reduced frequency  $\sigma$  increases as the motion frequency  $f$  increases. However, when the spring becomes stronger, the reduced frequency  $\sigma$  has the minimum value—that is, there are two values of the motion frequency  $f$  for the one value of the reduced fre-




Figure 16.24 Relationship between reduced frequency and propulsive efficiency.

quency  $\sigma$ . Figure 16.24 shows the results of the relationship between the reduced frequency  $\sigma$  and the propulsive efficiency  $\eta$ . It was found that the propulsive efficiency  $\eta$  generally decreases as the reduced frequency  $\sigma$  increases. We note that the experiment at lower reduced frequency ( $\sigma < 0.9$ ) is impossible because the reduced frequency  $\sigma$  is determined as a calculated value from the propulsive speed  $U$  and the motion frequency  $f$ . The maximum values of the propulsive efficiency reach around 0.65 when the spring is relatively weak.

## CONCLUSIONS

Two self-propelled two-joint dolphin robots were constructed to study the propulsive mechanisms of high-speed swimming animals. The characteristics of the propulsive speed and the efficiency with respect to the spring stiffness of the second joint and the reduced frequency were investigated. Both robots attained a propulsive efficiency of around 0.7 and a propulsive speed of 1.15 m/sec for the first robot and 1.9 m/sec for the second robot.

## REFERENCES

- [1] Lighthill, M. J. (1970). Aquatic animal propulsion of high hydromechanical efficiency. *J. Fluid Mechanics* 44(2): 265–301.
- [2] Chopra, M. G. (1974). Hydromechanics of lunata-tail swimming propulsion. *J. Fluid Mechanics* 64(2): 375–391.
- [3] Chopra, M. G. (1976). Large amplitude lunata-tail theory of fish locomotion. *J. Fluid Mechanics* 74(1): 161–182.
- [4] Chopra, M. G., and Kambe, T. (1977). Hydromechanics of lunata-tail swimming propulsion: II. *J. Fluid Mechanics* 79(1): 49–69.

- [5] Cheng, H. K., and Murillo, L. (1984). Lunate-tail swimming propulsion as a problem of curved lifting line in unsteady flow: I. Asymptotic theory. *J. Fluid Mechanics* 143: 327–350.
- [6] Karpouzian, G., Spedding, G., and Cheng, H. K. (1990). Lunate-tail swimming propulsion: II. Performance analysis. *J. Fluid Mechanics* vol. 210, no. 4: 329–351.
- [7] Triantafyllou, G. S., Triantafyllou, M. S., and Grosenbaugh, M. A. (1993). Optimal thrust development in oscillating foils with application to fish propulsion. *J. Fluids Structures* 7: 205–224.
- [8] Triantafyllou, M. S., Barret, D. S., Yue, D. K. P., Anderson, J. M., and Grosenbaugh, M. A. (1996). A new paradigm of propulsion and maneuvering. *Trans. Soc. Naval Architects Marine Engineers* 104.
- [9] Anderson, J. M., and Kerrebrock, P. A. (1999). The vorticity control unmanned undersea vehicle (VCUUV) performance results. *Eleventh International Symposium on Unmanned Untethered Submersible Technology*, pp. 360–369.

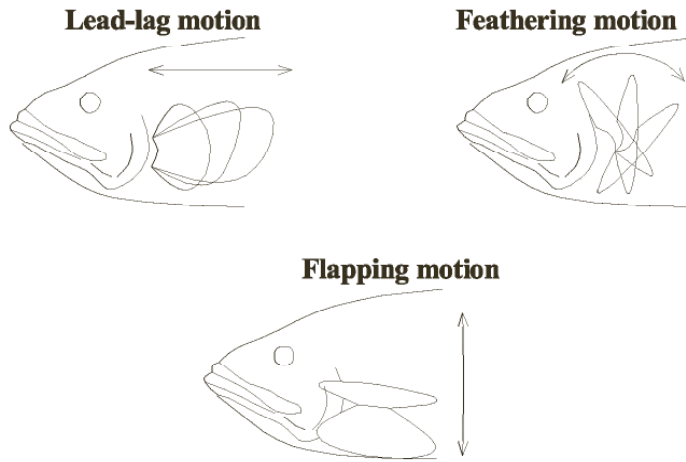
# 17 Pectoral Fin Controllers

Naomi Kato


The existing underwater robots, the motion of which is controlled by thrusters and wings, exhibit poor performance—not only in hovering and turning in a vertical or horizontal plane over intricate seabed terrain in strong currents, but also in dexterous manipulation under floating conditions [1]. The study of locomotion of aquatic animals can provide new insight into maneuverability and stabilization for underwater robots. This paper focuses on biomimesis of aquatic animals as a source for developing a new device for maneuvering underwater robots.

Many previous studies have focused on fish fins and their motion as an oscillating foil for propulsion [2–5]. However, due to the prevalent use of screw-type thrusters, application of the oscillating foil as the main thruster for underwater vehicles is not straightforward. A major drawback of these screw-type thrusters is the sudden generation of a thrust force when the underwater vehicle is hovering, which leads to imprecise control of the position and attitude of the vehicle and of its manipulators. Few researchers have considered the utilization of fish-fin motion to maneuver and stabilize underwater vehicles at low speeds [6]. Given that many types of fish use oscillating pectoral fins for their maneuvers at low swimming speeds [7], utilization of the pectoral fin motion could possibly overcome the drawbacks of the screw-type thrusters.

Kinematic studies indicate that pectoral fin movements are variable among species and across a range of swimming speeds with complex fin ray patterns [8–11]. We experimentally analyzed the pectoral fin motion of a black bass (*Micropterus salmoides*) in order to develop a set of mechanical pectoral fins [12]. Webb characterized the bass as a generalist blessed with several fish gaits, including cruising/sprinting, accelerating, and maneuvering [13–14]. Fin motion like the beating of a bird's wings will generally consist of four fundamental motions: (1) flapping motion in vertical plane; (2) lead-lag motion, which denotes a posterior and anterior motion in the horizontal plane; (3) feathering motion, which denotes a twisting motion of the fin pitch; and (4) spanning motion, which denotes an alternatively extending and contracting



**Figure 17.1** Illustrations of lead-lag motion, feathering motion and flapping motion of a pectoral fin.



**Figure 17.2** Coordinate system.

motion of the fin span [15] (figure 17.1). In a previous study [12], the lead-lag motion in the  $X$ - $Y$  plane, and the feathering motion in the  $X'$ - $Z'$  plane, were employed in the analysis of pectoral fin motion (figure 17.2). The flapping motion, which produces vertical motion of fish, and the spanning motion were not analyzed because the swimming of tested fish was restricted to the horizontal plane, and because the pectoral fin was treated as a rigid plate. The study revealed that (1) the combination of bilateral lead-lag motion and feathering motion is the predominant generator of the forward swimming, backward swimming, hovering, and turning in the horizontal plane; (2) a two-motor-driven mechanical pectoral fin (2MDMPF) making the feathering motion and the lead-lag motion generated thrust force in a certain range of phase difference between both motions; and (3) the fish robot composed of a rigid fish body and bilateral two-motor-driven mechanical pectoral fins can perform forward swimming and backward swimming and turn in a horizontal plane.

The author developed three-motor-driven mechanical pectoral fins (3MDMPFs), which are capable of making not only the feathering motion and the lead-lag motion, but also the flapping motion, and incorporated these fins into an autonomous underwater vehicle capable of high maneuverability at low speeds in both the horizontal and the vertical planes.

This chapter summarizes the ongoing project on pectoral fin controllers from the following four aspects:

1. Maneuver control of the fish robot with the pair of 2MDMPFs in the horizontal plane [16].
2. Control by the pair of 2MDMPFs in the horizontal plane of interaction of motion between an underwater vehicle in a free-swimming condition and its manipulator.
3. Hydrodynamic characteristics of 3MDMPF.
4. Swimming performance of an autonomous underwater vehicle equipped with a pair of 3MDMPFs.

## MANEUVER CONTROL OF THE FISH ROBOT

This section discusses the performance of guidance and control in the horizontal plane of the fish robot with the pair of 2MDMPFs for rendezvous with and docking at an underwater post in water currents.

### Two-Motor-Driven Mechanical Pectoral Fin

Figure 17.3 shows a pair of two-motor-driven mechanical pectoral fins. Each fin system has two servomotors, the angles of which are sensed by potentiometers and controlled by a personal computer by use of proportional-derivative (PD) control. The servomotors generate the lead-lag motion and the feathering motion of the pectoral fin. Although the fin system can generate any kind of programmed motion within the performance of the motors, sinusoidal motion was used. The pectoral fin model is a flat plate made of stainless steel (0.6 mm) with a chord length of 0.155 m.

We define the lead-lag angle  $\phi_L$ , the feathering angle  $\phi_{FE}$ , and the flapping angle  $\phi_{FL}$  on the right-hand side as shown in figure 17.2. We set  $\phi_{FL}$  as zero and the X-axis parallel to the longitudinal axis of the fish body, the positive direction of which points to the head. We make Y'-axis lie in the plane of the X-Y axis.

An open-water test of the 2MDMPF system was carried out in a water-circulating tank measuring 2.2 m long, 1.4 m wide and 0.9 m deep to investigate the hydrodynamic characteristics of the pectoral fin alone, in terms of the phase difference between the lead-lag motion and the feathering motion and the nondimensional frequency  $K$ , as defined




Figure 17.3 Two-motor-driven mechanical pectoral fin.

below:

$$K = c \cdot \omega / U \quad (17.1)$$

where  $c$  denotes the chord length of the fin,  $\omega$  the angular velocity of the fin motion, and  $U$  the uniform velocity of the water. Here,  $c$  is 0.155 m, and  $U$  is set as 0.163 m/sec. The lead-lag angle and the feathering angle are controlled to produce sinusoidal motion in the following formulas:

$$\phi_L = 35 - 35 \cdot \cos(2 \cdot \pi / T \cdot t) \quad (17.2)$$

$$\phi_{FE} = -35 - 35 \cdot \cos(2 \cdot \pi / T \cdot t + \Delta\phi) \quad (17.3)$$

Here  $\Delta\phi$  represents the phase difference between the lead-lag motion and the feathering motion,  $T$  denotes the motion period, and  $t$  denotes time.

A three-component block gauge was used to measure the hydrodynamic forces. Net hydrodynamic forces on the fin alone were estimated by subtracting the hydrodynamic forces on the mechanical system without the fin from those with the fin, because the support setups were submerged in water.

Figure 17.4 shows the coefficients of mean thrust during one period,  $Cx^*$ , for  $K$  of 2, 3, and 4 against the phase difference  $\Delta\phi$ . Here we define  $Cx^*$  as follows:

$$Cx^* = \frac{1}{T} \int_0^T Fx / (0.5 \cdot \rho \cdot U^2 \cdot S) dt \quad (17.4)$$




Figure 17.4 Coefficients of the mean thrust of a 2MDMPF.




Figure 17.5 Photograph of fish robot “Bass II.”

Here  $\rho$  denotes the fluid density and  $S$  denotes a projected area of the fins of  $0.029 \text{ m}^2$ .

We can see that the maximum value of  $Cx^*$  increases as  $K$  increases, that  $Cx^*$  reaches its maximum value in a region of  $\Delta\phi$  from  $90^\circ$  to  $120^\circ$ , and that  $Cx^*$  reaches its minimum value in a region of  $\Delta\phi$  from  $-60^\circ$  to  $-90^\circ$ . These hydrodynamic characteristics are then used in the guidance and control of a fish robot equipped with a pair of 2MDMPFs.

### Fish Robot with a Pair of 2MDMPFs

Figure 17.5 shows a photograph of the fish robot Bass II with a pair of 2MDMPFs. The fish body is made of fiber reinforced plastic (FRP). Its dimensions are as follows: 1.08 m by 0.28 m by 0.39 m (length, width,



height). A pair of 2MDMPFs were installed at a quarter of the fish body length from the head.

### **Experimental Setups**

The carriage for the measurement of motion of the fish robot in the horizontal plane was mounted on the water circulation tank. The fish robot was submerged to half the depth of the tank. The fish robot and the carriage were rigidly connected by a support rod located a quarter of the fish body length from the head. The X0-axis component and Y0-axis component of the position of the fish robot were measured by two rotary encoders, and the yaw angle  $\varphi$  was measured by a potentiometer. The mass of a pair of 2MDMPFs is 8.42 kg, the volume of the fish body is 0.0494 m<sup>3</sup>, the mass of the fish body is 15.74 kg, the mass of the carriage is 7.0 kg, and the friction force between the carriage and the rails on the tank is 0.49 N. A pipe, 0.1 m in length and 0.02 m in diameter, was attached to the front of the fish robot. The rendezvous and docking test was performed using this pipe in conjunction with an underwater post with a hole of 0.08 m diameter. The weight of the fish robot including the support rod in water was set as almost zero. The control frequency was set as 10 Hz.

### **Control Laws for Rendezvous with and Docking at an Underwater Post in Water Currents**

The variables for the motion control of the fish robot in the horizontal plane are basically (1) each phase difference between the lead-lag motion and the feathering motion on each side, (2) each angle of the central axis of the lead-lag motion on each side, (3) each angle of the central axis of the feathering motion on each side, (4) phase difference between the feathering motions on both sides, (5) each amplitude of the lead-lag motion on each side, (6) each amplitude of the feathering motion on each side, and (7) motion frequency; in total, 11. The input variables are the X-axis component of the position, the Y-axis component of the position, and the yaw angle.

Because the motion of the fish robot is highly nonlinear with regard to the control variables and because it is not possible to express the equations of the motion explicitly in terms of the control variables, we employ a fuzzy control algorithm, which is applicable to this sort of problem.

The prearranged course from point O to point C was given as shown in figure 17.6. The underwater post was set at point C. The control laws consist of (1) a fuzzy control law for forward swimming and hovering from point O to point A and from point B to point C, (2) lateral swimming from point A to point B, and (3) a switch of the control laws.





Figure 17.6 Loci and yaw angle of fish robot in water current of 0.05 m/sec.

**Control Law of Forward Swimming and Hovering** We set the lead-lag angle  $\phi_{L\_right}$  (degrees) and the feathering angle  $\phi_{F\_right}$  (degrees) on the right side as follows:

$$\phi_{L\_right} = 30 - 30 \cdot \cos(2 \cdot \pi/T \cdot t) \quad (17.5)$$

$$\phi_{F\_right} = -30 - 30 \cdot \cos(2 \cdot \pi/T \cdot t + \delta p(t) + \delta x(t)) \quad (17.6)$$

Here, the phase angle  $\delta p$  is obtained using outputs  $\delta E y_0$  and  $\delta E \varphi_0$  from fuzzy rules on transverse deviation from the prearranged course and deviation from the set yaw angle, respectively, as follows:

$$\delta p(t) = \delta p(t - \Delta t) + W1_{\delta E y_0} + W2_{\delta E \varphi_0} \quad (17.7)$$

Here we take  $W1$  and  $W2$  to be 0.75 and 0.25, respectively. The phase angle  $\delta x$  is obtained using output  $\delta E x_0$  from fuzzy rules on longitudinal deviation from set points as follows:

$$\delta x(t) = \delta x(t - \Delta t) + \delta E x_0 \quad (17.8)$$

Similarly, we set the lead-lag angle  $\phi_{L\_left}$  (degrees) and the feathering angle  $\phi_{F\_left}$  (degrees) on the left side as follows:

$$\phi_{L\_left} = 30 - 30 \cdot \cos(2 \cdot \pi/T \cdot t) \quad (17.9)$$

$$\phi_{F\_left} = -30 - 30 \cdot \cos(2 \cdot \pi/T \cdot t - \delta p(t) + \delta x(t)) \quad (17.10)$$

**Control Law of Lateral Swimming to the Right-Hand Side** We set  $\phi_{L\_right}$  (degrees) and  $\phi_{F\_right}$  (degrees) as follows:

$$\phi_{L\_right} = -20 \cdot \cos(2 \cdot \pi/T \cdot t) \quad (17.11)$$

$$\phi_{F\_right} = -40 - 20 \cdot \cos(2 \cdot \pi/T \cdot t) \quad (17.12)$$

Using the phase angle  $\delta s$ , we set  $\phi_{L\_left}$  and  $\phi_{F\_left}$  (degrees) as follows:

$$\phi_{L\_left} = -45 - \delta s(t) - 30 \cdot \cos(2 \cdot \pi/T \cdot t) \quad (17.13)$$

$$\phi_{F\_left} = -30 \cdot \cos(2 \cdot \pi/T \cdot t) \quad (17.14)$$

The phase angle  $\delta s$  is obtained using output  $\delta E x_0$  and  $\delta E \varphi_0$  from fuzzy rules on longitudinal deviation from the set points and deviation from set yaw angle, respectively, as follows:

$$\delta s(t) = \delta s(t - \Delta t) + W1_{\delta E x_0} + W2_{\delta E \varphi_0} \quad (17.15)$$

## Experimental Results

Figure 17.6 shows the loci of the fish robot and its yaw angle in a water current of 0.05 m/sec in the negative direction of the  $X_0$  axis, respectively. Although some degree of offset from the prearranged course at the start point is given, the fish robot swims forward, stops swimming forward around point A, swims laterally toward point B with a small yaw angle, stops swimming laterally to swim forward, and finally docks with the underwater post. We found that the amount of drift in the direction of the water current between point A and point B becomes larger and the yaw angle becomes more oscillatory as the velocity of water current increases; that is, it takes more time for the vehicle to reach the underwater post.

## COORDINATED MOTION CONTROL FOR FREE SWIMMING IN A HORIZONTAL PLANE BY A PAIR OF 2MDMPFs

The interaction of motion between the vehicle and its manipulator(s) forms a complex, time-varying, nonlinear, multivariable dynamic system, with a large number of degrees of freedom (DOF); the control of the system is not straightforward. However, the additional DOF made available by the manipulator(s) can be put to efficient use in the design of the vehicle at an early stage to achieve more-precise motion control. For example, if only one arm is required to execute the mission, other arms can be used to assist in stabilizing the vehicle—either by attaching to a structure or by moving to dynamically counterbalance the reaction forces and movements from the mission arm on the vehicle. Thus, the redundant degrees of freedom are efficiently used for controlling the robot.


Thus far, the majority of research on such use of underwater manipulators has been limited to the condition where a manipulator is attached to a structure [17]. However, several reports have focused on the problem with the vehicle in a free swimming condition [1, 18–21].

In a previous paper [1], I reported on analysis and simulation studies for the coordinated control by multiple manipulators of a free-swimming underwater robot. The control objective was to perform an open-loop control of the end-effector position of one of the arms. A formulation was derived for both the inverse dynamics and the inverse kinematics of an underwater robot with multiple manipulators, having a redundancy in degrees of freedom of motion, in the presence of external generalized forces. The method was applied to a model of an underwater robot comprising a 6-DOF mission arm and an additional pair of 2-DOF paddling arms used for stabilization in the free-swimming condition. The numerical simulation results showed the effectiveness of the coordinated control method with the assistance of the stabilizing arms.

This section shows the experimental results in horizontal plane on the stabilization of motion of an underwater robot model equipped with a manipulator and a pair of 2MDMPFs instead of the paddling arms used in the simulation.

### Model of the Underwater Robot with a Manipulator

Figure 17.7 shows the model of the underwater robot with a manipulator. The dimensions of the model are 0.9 m by 0.6 m by 0.5 m (length, width, height). The mass of the model is 25.95 kg. The manipulator consists of 3 links where 3 DC motors and 3 potentiometers are installed. The dimensions of the manipulator are 0.651 m by 0.040 m by



**Figure 17.7** Model of underwater robot with a manipulator.

0.0.095 m (length, width, height). The distance between links is 0.2 m. The mass of the manipulator is 2.55 kg. The pair of 2MDMPFs were installed 0.34 m behind the base of the manipulator.

### Experimental Setups


The underwater robot model and the carriage for the measurement of the motion in the horizontal plane were rigidly connected by a support rod located 0.34 m behind from the base of the manipulator. The fish robot was submerged to half the depth of the water-circulating tank.

### Control Algorithm

A fuzzy algorithm was used for the control of the yaw motion of the underwater robot model with a manipulator. The control variables were one pair of phase differences between the lead-lag motion and the feathering motion on both sides. The amplitudes of both motions were fixed.

### Experimental Results

Figure 17.8 shows an example of the experimental results of yaw motion of the model during a period (11.1 sec) of oscillatory motion of the manipulator. The position and the posture of the model with the manipulator are illustrated each 0.6 sec. We found that a strong interaction of motion occurs between the model and the manipulator.



**Figure 17.8** Motion of an underwater robot model induced by an oscillatory motion of its manipulator.

Figure 17.9 shows the ratio of the root-mean-square of yaw motion of the model without control to that with control during 5 periods of sinusoidal motion of the manipulator; these periods ranged from 5.5 sec to 11.0 sec and those of the pectoral fin from 1.0 sec to 2.0 sec. We found that slower motion of the manipulator produces better control performance, which is only slightly dependent on the period of motion of fin, and that the control performance increases as the fin moves faster with the faster motion of the manipulator.

## HYDRODYNAMIC CHARACTERISTICS OF A 3MDMPF

### Mechanical Structure

We have developed a pair of three-motor-driven mechanical pectoral fins (3MDMPFs) in which gimbal structures are incorporated to allow three DC servomotors to independently generate the lead-lag motion, the feathering motion, and the flapping motion, respectively, and to produce vertical motion of the fin. All of these components are housed in a compact pressure vessel (figure 17.10). Its dimensions are 0.375 m by 0.212 m by 0.40 m (length, width, height). Its mass is 11 kg. It takes 0.5 sec to achieve a particular type of motion in the range of  $\pm 45^\circ$ . Three potentiometers are installed in each 3MDMPF to measure the

rms\_yaw ( with control ) / rms\_yaw ( without control )





Figure 17.9 Ratio of yaw motion of model with control to that without control.



Figure 17.10 Three-motor-driven mechanical pectoral fin.



**Figure 17.11** Configuration of a body equipped with 3MDMPF.


angles of the three DC motors. The addition of the flapping motion to the 2MDMPF produces vertical force. It enables a fish robot equipped with a pair of 3MDMPFs to swim not only in the horizontal plane, but also in the vertical plane.

### Hydrodynamic Characteristics

**Experimental Setups** An open-water test of the 3MDMPF was carried out in the water-circulating tank. The 3MDMPF was installed in the body as shown in figure 17.11, with care taken to make the flow around it smooth. The same pectoral fin model as in figure 17.3 was used in this experiment. The hydrodynamic characteristics were investigated in terms of the phase difference between the lead-lag motion and the feathering motion and in terms of that between the lead-lag motion and the flapping motion. The mean inflow velocity was 0.175 m/sec at the location of the pectoral fin model. The nondimensional frequency  $K$  was 3.76.

We used a three-component strain gauge to measure the hydrodynamic forces, which involved three types of support setup for the gauge: one for measurement of  $F_x$ ,  $F_y$ , and  $M_z$ ; one for  $F_x$ ,  $F_z$ , and  $M_y$ ; and one for  $F_y$ ,  $F_z$ , and  $M_x$ . Here  $F_x$ ,  $F_y$ , and  $F_z$  denote hydrodynamic forces in the direction of X axis, Y axis, and Z axis, respectively.  $M_x$ ,  $M_y$ , and  $M_z$  denote hydrodynamic moments around the X axis, Y axis, and Z axis, respectively. The three types of hydrodynamic force were measured separately. For each type, the hydrodynamic forces on the apparatus without the fin and those with the fin were measured separately, due to the fact that the support setups were submerged in water.





**Figure 17.12** Configuration of a pectoral fin for computation and the division of its surface.

To estimate the net hydrodynamic forces on the blade alone, the net value of hydrodynamic forces on the apparatus without the fin was subtracted from the net value of those on the apparatus with the fin.

The lead-lag angle  $\phi_L$ , the feathering angle  $\phi_{FE}$ , and the flapping angle  $\phi_{FL}$ , all of which are defined in figure 17.2, were varied sinusoidally using the following definitions.


$$\begin{aligned}\phi_L &= \phi_{LA} \cdot (1 - \cos(\omega_{fin} \cdot t)) \\ \phi_{FE} &= -\phi_{FEA} \cdot (1 + \cos(\omega_{fin} \cdot t + \Delta\phi_{FE})) \\ \phi_{FL} &= \phi_{FLA} \cdot \sin(\omega_{fin} \cdot t + \Delta\phi_{FL})\end{aligned}\tag{17.16}$$

We set  $\phi_{LA}$ ,  $\phi_{FEA}$  and  $\phi_{FLA}$  in equation (17.16) as  $35^\circ$ ,  $35^\circ$ , and  $20^\circ$ , respectively.

### Theoretical Analysis of Hydrodynamic Forces on the Pectoral Fin

Separated vortical flow is dominant around a pectoral fin operating in drag-based labriform mode because the pectoral fin forms a high angle in relation to the horizontal axis of the fish body in this mode. The author has developed an unsteady vortex lattice method (UVLM) including the effect of vortex diffusion for the theoretical analysis of separated vortical flows [22].

Figure 17.12 shows the configuration of the pectoral fin for computation with 130 segmented panels. The uniform flow  $U$  flowing in the




**Figure 17.13** Comparison of  $Cx^*$  without flapping motion between computation and experiment.

positive direction of the X-axis and time step  $\Delta t$  were set as 1.0 m/sec and 0.026 sec, respectively. The computation was done for 1.5 periods for the nondimensional frequency  $K$  of 3.76.

Here, hydrodynamic force coefficients are defined as follows:

$$\begin{aligned}
 C_X &= F_X / (1/2 \cdot \rho \cdot U^2 \cdot S), \\
 C_Y &= F_Y / (1/2 \cdot \rho \cdot U^2 \cdot S), \\
 C_Z &= F_Z / (1/2 \cdot \rho \cdot U^2 \cdot S), \\
 Cm_X &= M_X / (1/2 \cdot \rho \cdot U^2 \cdot S \cdot c), \\
 Cm_Y &= M_Y / (1/2 \cdot \rho \cdot U^2 \cdot S \cdot c), \\
 Cm_Z &= M_Z / (1/2 \cdot \rho \cdot U^2 \cdot S \cdot c)
 \end{aligned}
 \tag{17.17}$$

**Comparisons between Computation and Experiment** Figure 17.13 shows the comparison of mean thrust-force coefficient  $Cx^*$  during one period against phase difference  $\Delta\phi_{FE}$  without flapping motion between computation and experiment. Figure 17.14 shows the comparison of  $Cx^*$  against phase difference  $\Delta\phi_{FL}$  with flapping motion where  $\Delta\phi_{FE}$  was set as  $110^\circ$ . We can see that the mean thrust-force coefficient is largely dependent on  $\Delta\phi_{FE}$ , while it is not largely dependent on  $\Delta\phi_{FL}$ .



**Figure 17.14** Comparison of  $Cx^*$  with flapping motion between computation and experiment.


The maximum mean thrust-force coefficient with flapping motion is slightly larger than without flapping motion. The computed results agree fairly well with the experimentally measured values.

Figure 17.15 shows the comparison of mean vertical force coefficient  $Cz^*$  during one period against phase difference  $\Delta\phi_{FE}$  without flapping motion between computation and experiment. Figure 17.16 shows the comparison of  $Cz^*$  against phase difference  $\Delta\phi_{FL}$  with flapping motion where  $\Delta\phi_{FE}$  was set as  $110^\circ$ . We can see that  $Cz^*$  is also largely dependent on  $\Delta\phi_{FE}$ , and that  $Cz^*$  becomes minimum near  $\Delta\phi_{FE}$  where  $Cx^*$  becomes maximum.


## AUTONOMOUS UNDERWATER VEHICLE BASS III

### Structure

An autonomous underwater vehicle equipped with a pair of 3MDMPFs was constructed and its maneuverability was tested in both the horizontal plane and the vertical plane. Figure 17.17 shows a photograph of Bass III. Figure 17.18 shows its configuration and arrangement. The mass is 104.1 kg. The maximum swimming depth is 10 m. The pressure vessel for the battery is arranged outside the cover to bring the center of gravity close to the location of the pectoral fin's root, which enables the lateral swimming. Bass III has one thruster to use in water currents.



**Figure 17.15** Comparison of  $Cz^*$  without flapping motion between computation and experiment.



**Figure 17.16** Comparison of  $Cz^*$  with flapping motion between computation and experiment.



Figure 17.17 Photograph of BASS III.




Figure 17.18 Configuration of BASS III and its arrangement.

It has a 3-axes gyro, a 3-axes inclinometer, a depth meter, an altimeter, and 3 CCD cameras as visual sensors. One downward-facing CCD camera will be used to locate the position in the horizontal plane and two forward-facing CCD cameras will be used to localize the relative distance to an underwater structure. During the initial stage of testing, AC power was supplied to the vehicle through a power cable, and a LAN cable was connected between the vehicle and the ground computer to enable their communication.

**Table 17.1** Setups of angles

	$\phi_{LA}$	$\phi_{FEA}$	$\phi_{FLA}$	$\Delta\phi_{FE}$
Ascent in forward swimming	25°	20°	0°	90°
Descent in forward swimming	25°	-20°	0°	90°
Ascent in backward swimming	25°	-20°	0°	-105°
Descent in backward swimming	25°	20°	0°	-105°
Left turning (right)	25°	20°	0°	90°
Left turning (left)	25°	20°	0°	-105°

**Table 17.2** Setups of angles for lateral swimming to the right-hand side

	$\phi_L$	$\phi_{FE}$	$\phi_{FL}$
Lateral swimming (right)	$20 \cdot \sin(\omega_{fin} \cdot t)$	$-20 \cdot \sin(\omega_{fin} \cdot t)$	0°
Lateral swimming (left)	$30 \cdot \sin(\omega_{fin} \cdot t)$	$20 \cdot \sin(\omega_{fin} \cdot t)$	0°

## Swimming Tests in a Water Tank

**Experimental Setup** We carried out swimming tests in a JAMSTEC water tank 40 m long, 4 m wide, and 2 m deep. To measure the location of the vehicle in the horizontal plane, we set up ropes equipped with white panels (5 cm square), set at every 50 cm in the longitudinal direction. The ropes were set up at 50 cm intervals in the transverse direction. We analyzed the pictures taken by a video camera, which was set up on the carriage, to localize the position of the vehicle.

We used pectoral fins made of plastic and 1.5 times as large as those shown in figure 17.3. We set the frequency of oscillation of the pectoral fin at 0.5 Hz.

We carried out six types of swimming tests actuating the pair of 3MDMPFs, including (1) ascending in a forward direction, (2) descending in a forward direction, (3) ascending in a backward swim, (4) descending in a backward swim, (5) left turning, and (6) lateral swimming to the right-hand side. Table 17.1 shows setups of angles for ascending in the forward direction, descending in the forward direction, ascending in the backward swim, descending in the backward swim, and left turning. Table 17.2 shows setups of angles for lateral swimming to the right-hand side.

**Experimental Results** Figure 17.19 shows the experimental results of depth, distance from the start point, yaw angle, pitch angle, and roll angle against time in the ascending in the forward direction. We can see that the vehicle ascends with small changes in the attitude. The path angle reaches about 38°. The period of the oscillatory pitch angle corresponds with the natural period of the vehicle itself.

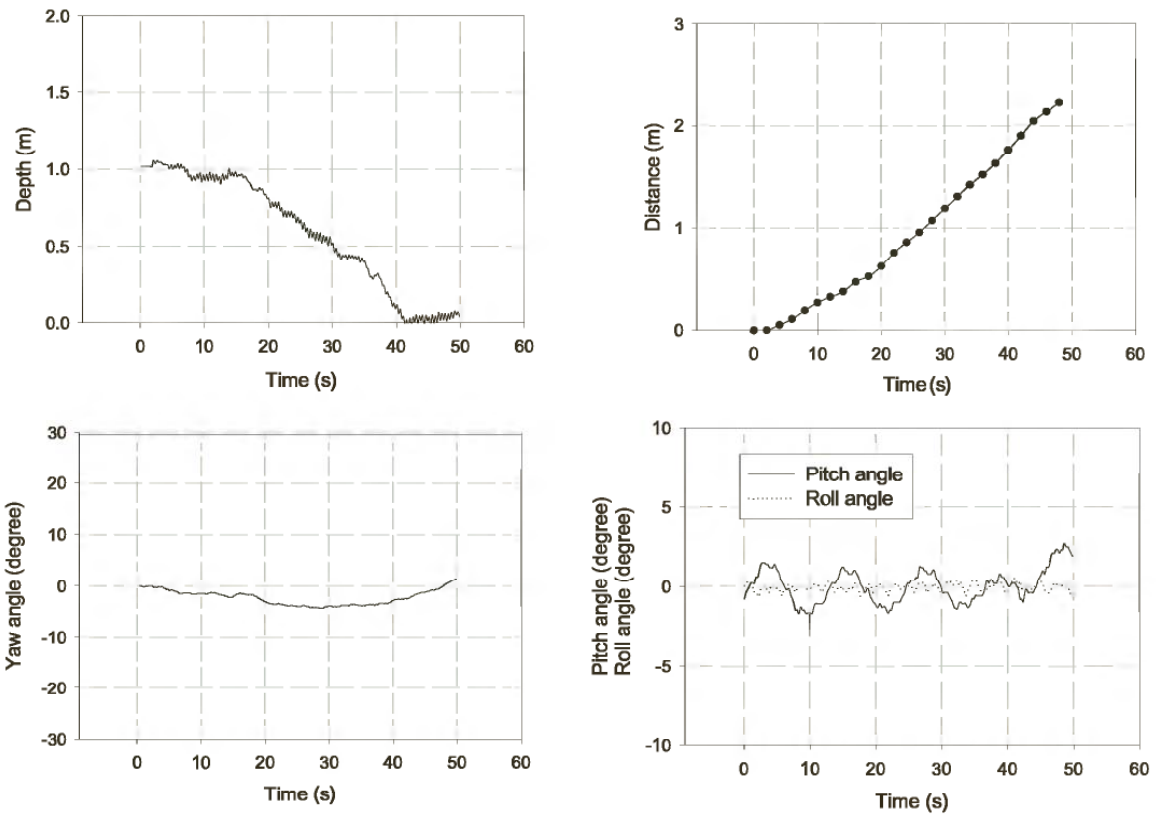


Figure 17.19 Ascending in forward swimming.

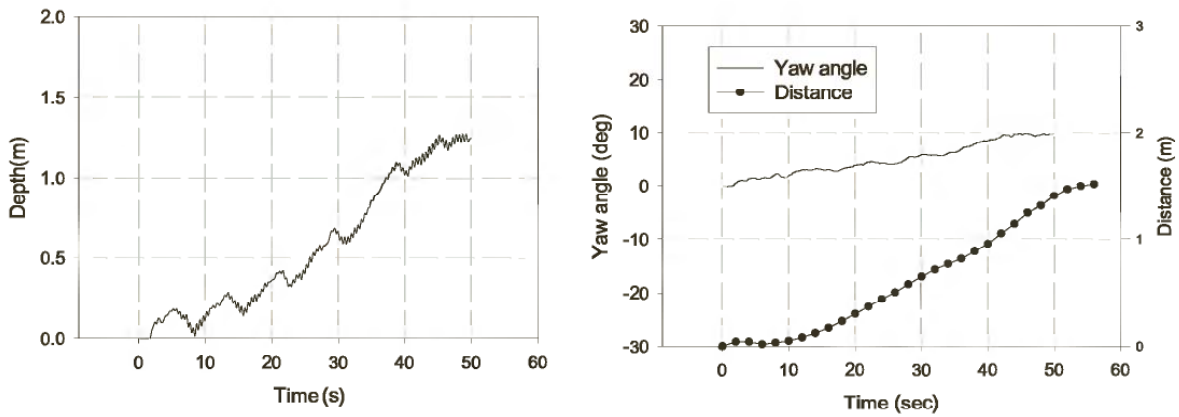


Figure 17.20 Descending in backward swimming.

Figure 17.20 shows the experimental results of depth and of distance from the start point against time in the descending in the backward direction. Even during backward swimming, the vehicle is stable. The path angle reaches about  $53^\circ$ . We found that  $\phi_{FEA}$  is closely related to vertical movement and  $\Delta\phi_{FE}$  is related to back-and-forth movement.

Figure 17.21 shows the experimental results of yaw angle and of location in the horizontal plane against time for left turning. Here we




Figure 17.21 Left turning.






Figure 17.22 Lateral swimming to the right-hand side.

define X coordinate as the longitudinal direction of the water tank. The vehicle turns at almost the same position.

Figure 17.22 shows the experimental results of yaw angle and of location in the horizontal plane against time in the lateral swimming to the right-hand side. Although the yaw angle slightly increases, the vehicle moves laterally to the right-hand side.

## CONCLUSIONS

This ongoing project on mechanical pectoral fins has been presented from the viewpoint of its application to underwater vehicles conducting various types of maneuvers. This chapter has demonstrated that the fish robot equipped with a pair of two-motor-driven mechanical pectoral fins has a high maneuverability in the horizontal plane; that the induced yaw motion of an underwater robot model in the horizontal plane caused by a forced motion of a manipulator can be largely decreased with the assistance of a pair of two-motor-driven mechanical



pectoral fins; that the three-motor-driven mechanical pectoral fin has a slightly larger mean thrust than that of the two-motor-driven mechanical pectoral fin; and that the Bass III equipped with a pair of three-motor-driven mechanical pectoral fins has a high maneuverability—not only in the horizontal plane, but also in the vertical plane.

## REFERENCES

- [1] Kato, N., and Lane, D. M. (1996). Co-ordinated Control of Multiple Manipulators in Underwater Robots. In *Proceedings of the 1996 IEEE International Conference on Robotics and Automation*, Minneapolis, 22–28 April. pp. 2505–2510.
- [2] Hertel, H. (1966). *Structure-Form-Movement*. New York: Reinhold.
- [3] Isshiki, N., and Morikawa, H. (1982). Study on dolphin-style fin ship (in Japanese). *Bull. Soc. Nav. Architects Japan* 642: 2–9.
- [4] Barrett, D. S., and Triantafyllou, M. S. (1995). The design of flexible hull undersea vehicle propelled by an oscillating foil. In *Proceedings of the Ninth International Symposium of on Unmanned, Untethered Submersible Technology*. pp. 111–123.
- [5] Tanaka, I., and Nagai, M. (1996). *Hydrodynamics of Drag and Propulsion—Learning from High Speed Swimming Performance of Aquatic Animals* (in Japanese). Ship and Ocean Foundation.
- [6] Bandyopadhyay, P. B., Castano, J. M., Rice, J. Q., Philips, R. B., Nedderman, W. H., and Macy, W. K. (1997). Low-speed maneuvering hydrodynamics of fish and small underwater vehicles. *J. Fluids Eng. Trans. ASME* 119: 136–144.
- [7] Lindsey, C. C. (1978). Form, function, and locomotion habits in fish. In W. S. Hoar and D. J. Randall (eds.), *Locomotion Fish Physiology VII*. Academic Press, pp. 239–313.
- [8] Webb, P. W. (1973). Kinematics of pectoral fin propulsion in *cymatogaster aggregata*. *J. Exp. Biol.* 59: 697–710.
- [9] Geerlink, P. J. (1983). Pectoral fin kinematics of *Coris formosa* (Teleostei, Labridae). *Neth. J. Zool.* 39: 166–193.
- [10] Gibb, A. C., Jayne, B. C., Lauder, G. V. (1994). Kinematics of pectoral fin locomotion in the bluegill sunfish *Lepomis macrochirus*. *J. Exp. Biol.* 189: 133–161.
- [11] Walker, J. A., and Westneat, M. W. (1997). Labriform propulsion in fishes: Kinematics of flapping aquatic flight in the bird wrasse *Gomphosus varius* (Labridae). *J. Exp. Biol.* 200: 1549–1569.
- [12] Kato, N., and Furushima, M. (1996). Pectoral fin model for maneuver of underwater vehicles. In *Proceedings of 1996 IEEE Symposium on Autonomous Underwater Vehicle Technology*. pp. 49–56.
- [13] Webb, P. W. (1982). Locomotor patterns in the evolution of actinopterygian fishes. *Am. Zoologist.* 22: 329–342.
- [14] Webb, P. W. (1984). Body form, locomotion, and foraging in aquatic vertebrates. *Am. Zoologist.* 24: 107–120.
- [15] Azuma, A. (1992). *The Biokinetics of Flying and Swimming*. New York: Springer-Verlag.
- [16] Kato, N. (2000). Control performance in horizontal plane of a fish robot with mechanical pectoral fins. *IEEE J. Oceanic Eng.* vol. 25, no. 1: 121–129.

- [17] Ioi, K., and Itoh, K. (1990). Modelling and simulation of an underwater manipulator. *Adv. Robotics* vol. 4, no. 4: 303–317.
- [18] Mahesh, H., Yuh, J., and Kakshimi, R. (1991). A coordinated control of an underwater vehicle and robotic manipulator. *J. Robotic Syst.* vol. 8, no. 3: 339–370.
- [19] McMillan, S., Orin, D. E., and McGhee, R. B. (1995). DynaMechs, an object oriented software package for efficient dynamic simulation of URVs. In J. Yuh (ed.) *Underwater Robotic Vehicles Design and Control*. pp. 73–98.
- [20] Lapierre, L., Fraise, P., and M'Sirdi, N. K. (1998). Hybrid position/force control of a ROV with a manipulator. *Oceans '98*.
- [21] Diaz, E. O., Wit, C. C. d., and Perrier, M. (1998). A comparative study of neglected dynamics on an underwater vehicle/manipulator system under nonlinear robust control. *Oceans '98*.
- [22] Kato, N. (1999). Hydrodynamic characteristics of mechanical pectoral fin. *J. Fluids Eng.* vol. 121, no. 3: 605–613.

**This page intentionally left blank**

# *Ambulatory Robots*

**This page intentionally left blank**

# The Role of Reflexes versus Central Pattern Generators in Dynamical Legged Locomotion

Eric Klavins, Haldun Komsuoglu, Robert J. Full,  
and Daniel E. Koditschek

## INTRODUCTION

### A Research Agenda: Functional Biomimesis

Animals execute locomotor behaviors and more with ease. They have evolved these breathtaking abilities over millions of years. Cheetahs can run, dolphins can swim, and flies can fly like no artificial technology can. It is often argued that if human technology could mimic nature, then biological-like performance would follow. Unfortunately, the blind *copying* or mimicking of a part of nature (Ritzmann et al., 2000) does not often lead to the best design, for a variety of reasons (Vogel, 1998). Evolution works on the “just good enough” principle. Optimal designs are not the necessary end product of evolution. Multiple satisfactory solutions can result in similar performances. Animals do bring to our attention amazing designs, but these designs carry with them the baggage of their history. Moreover, natural design is constrained by factors that may have no relationship to human-engineered designs. Animals must be able to grow over time, but still function along the way. Finally, animals are complex and their parts serve multiple functions, not simply the one we happen to examine. In short, in their daunting complexity and integrated function, understanding animal behaviors remains as intractable as their capabilities are tantalizing.

*Functional biomimesis* describes a design methodology embraced within our Computational Neuromechanics Project (CNM) (Koditschek, 1999) inspired by, but not slavishly yoked to, biological example. The trick in deriving practicable engineering design from biological understanding is to develop analogies at the appropriate level of abstraction. Biological materials are vastly different from those available to current engineering practice. Biological constraints on computation and signal processing are quite distinct from those facing contemporary control engineers. But the problems of performing work upon a dynamical environment are common, and even the simplest animals are masters at the kind of multifunctional capabilities yet absent from the horizon of



**Figure 18.1** The RHex robot running at a body length ( $\sim 50$  cm) per second over rough terrain.

contemporary robotics. These considerations have focused our attention on the *mechanical* end of the animal behavioral spectrum where the biology is most constrained by the physical world. By this approach, the engineering benefit should be maximized with the result of high performance machines. At the same time, we anticipate a substantial benefit to pure science arising from the ability to leverage necessary conditions resulting from these physical constraints.

### **Functional Biomimesis at Work: The RHex Robot**

RHex, our recently prototyped autonomous hexapod robot (Buehler, Saranli, and Koditschek, 2000), depicted in figure 18.1, represents a significant success for functional biomimesis. Weighing 7 kg and roughly the size of a shoebox (about 50 cm by 30 cm), this machine aims toward autonomous function in the rubble and debris caused by disasters such as earthquakes, fires, and bombings. Despite its distinctly unnatural appearance and apparent kinematic simplicity, RHex is designed to mimic certain functions exhibited by sprawled-posture, many-legged runners such as the cockroach species *Blaberus discoidalis*.

Decades of accumulating evidence from biomechanics suggest that running animals—whether bipeds, quadrupeds, hexapods, or myriapods—perform work on their center of mass so as to exhibit the dynamics of a pogo stick in the sagittal plane (Full and Farley, 1999). Their muscles are tuned and postures are selected so as to engage the ground in a compliant manner and bounce from stride to stride. Along

with this general strategy, cockroaches and other sprawled-posture animals—arguably among the most successful runners—combine an apparent commitment to compliant operation in the horizontal plane, simplifying significantly problems of pitch and yaw stabilization. Inspired by the cockroach example, we have designed RHex as a sprawled-posture mechanism with passively compliant legs.

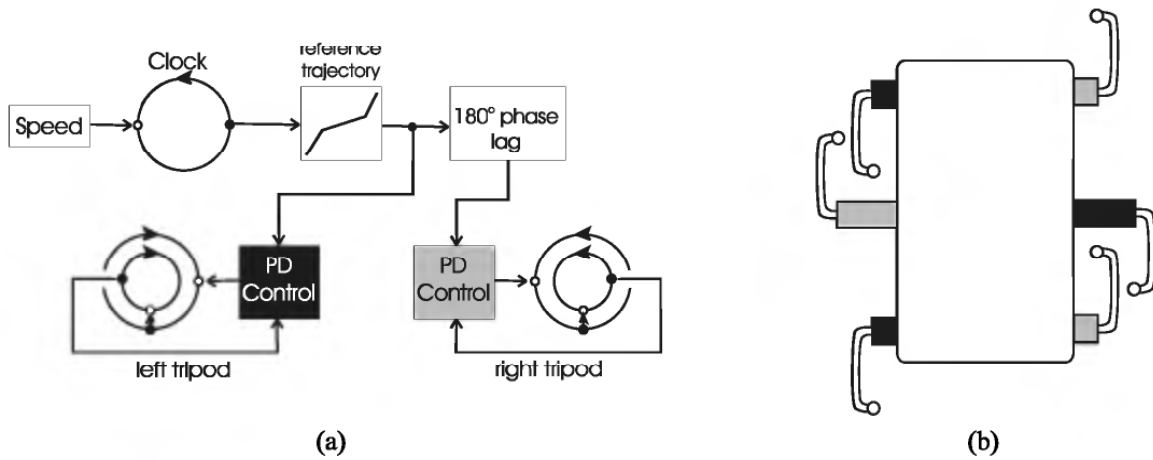
Recent studies (Kubow and Full, 1999) underscore dramatically the previously reported (Delcomyn, 1985) evidence for highly stereotyped cockroach limb motion, quantifying the very few degrees of freedom actually used in steady-state running (notwithstanding the potential for complex high-dimensional motions afforded by these animals' many muscle groups and joints). Because the power densities of commercial electromechanical actuators are so limited, RHex borrows from considerations of biological function on this count as well as replacing all but one degree of freedom per leg with a fixed (compliant) shape, locating just one actuator, a 20 W brushless DC servomotor at each of the machine's six "hips."

Accumulating empirical evidence confirms that the functional aspects of fast cockroach locomotion we have identified above do indeed appear to confer a high level of performance—morphological and even kinematic details notwithstanding. In a recently reported set of experiments (Buehler, Saranli, and Koditschek, 2000), RHex negotiates widely varied and dramatically adversarial terrain at speeds exceeding one body length per second. Moreover, initial runs over a force platform bear out the presumption that RHex's controller may be tuned so that leg coordination patterns supporting center-of-mass trajectories indistinguishable from a pogo stick in the sagittal plane (Full et al., 2000) result.

### **From Work to Tasks: Coordination Control as Phase Regulation**

In running insects, the leg patterns that support the pogo-stick template take the form of an alternating tripod gait wherein the first and third ipsilateral legs move in phase with the contralateral middle leg—all three antiphase to the opposite tripod. In this fashion, a tripod acts as a *virtual* single leg, and the tripod pair is coordinated in the manner of a virtual biped (Full and Tu, 1990). The simplest biologically plausible model of how the cockroach generates the alternating tripod running gait takes the form of a CPG (central pattern generator) hypothesis, wherein the legs simply track a properly shaped feed-forward periodic reference signal. In fact, animals do, in general, match this model with the very regular stereotyped pattern of their steady-state limb motions. They tend to have the shortest swing period possible and as speed increases, swing period changes little, and stance period decreases. In this view of the animal's coordination control, the pogo stick template (Full





**Figure 18.2** (a) Identical copies of the phase and anti-phase version of the two-stroke clock signal are used as reference trajectories by decoupled PD controllers at each of the hip motors comprising the respective tripods. (b) A schematic top view of RHex showing in black the left tripod and in gray, the right tripod.

and Koditschek, 1999) emerges as an immediate mechanical response, or “preflex,” before any classical reflexes contribute (Loeb and Brown, 2000).

Naive though it may seem, such a clock-driven preflex-stabilized mechanism is exactly the coordination controller presently employed in RHex. As depicted in figure 18.2, a two-phase periodic signal describes the conceptual operation of a wheel borne on a single spoke. The two phases are determined by a speed profile (elaborated later in figure 18.6) intended to match the physical state of a leg. A slow-speed phase corresponds to the contact (“retraction”) state of a leg, during which it acts like the spoke of a wheel to propel the axle forward. The much faster phase corresponds to the recovery (“protraction”) state of a leg, during which the spoke must cover a much greater arc in a comparable amount of time as it swings around over the axle to return for its next contact. This idealized signal functions as a reference trajectory tracked by a proportional-derivative (PD) controller driving the hip motor at each leg. Leg coordination is simply imposed by delivering identical copies of the clock reference to the participating motors, respectively, in the phase and antiphase tripods, as depicted in figure 18.2 (We defer to the discussion in the section entitled Oscillatory Systems that Perform Work Can Only Regulate Phase by Adjusting Energy regarding oscillator models depicted in figure 18.5: the explanation for the single and double circle icons used in figure 18.2.)

Our initial performance studies of RHex (Buehler, Saranli, and Koditschek, 2000) document that clock-driven preflex-stabilized coordination works quite well when in a steady-state mode of operation. There is apparently significant “mechanical intelligence” exhibited by the variously coupled analogue computers (i.e., the individual leg-body oscil-


lators) that coordinate the intermittent application of ground reaction forces to the robot's body by virtue of its clock-excited leg spring compliance. Yet, however capable, this intelligence is surely very local in nature. It seems intuitively clear that some additional feedback from the leg to its controller and then on to its neighbors should confer substantial benefits in the face of significant perturbations, rapid maneuvers, and, in general, operation away from steady state. Of course, the imperatives of autonomous operation in highly unpredictable environments demand exactly this sort of general operational capability. Thus, the question arises as to how to build a "reflex" layer over the existing reflex controller that takes advantage of and extends the mechanical intelligence rather than fighting it.

In the specific context of RHex's laboratory development, this issue emerges most sharply as we take the first steps toward improving and effectively utilizing the robot's sensory suite. The current implementation does not afford any task-level perception whatsoever. However, future versions will employ a broad array of sensors proposed to increase RHex's energy efficiency and afford autonomous locomotion without a human operator. Clarifying the role of feedback and feed forward in simulation studies and RHex experiments will lead to the development of sufficient mathematical conditions for the stable coordination of dynamically effective gaits. Such formal results should stimulate novel refutable hypotheses for biology.

## **Two Models for Comparing and Challenging Control Architectures**

At the present time, a significant gap in our understanding of animal and machine locomotion surrounds the issue of sensor recruitment for coordination. As the central goal of the simulation study presented in this chapter, we wish to understand controller capability along two architectural dimensions that appear to loom hugely for any complex system designed to perform work within and upon its environment. First, we seek to understand the manner and extent to which sensory data about the environment should be used to control actuators along an axis that might be termed "clock" versus "reflex." Second, along an axis that might be termed "central" versus "peripheral," we seek to understand the advantages and disadvantages of distributed coordination. This chapter explores the first of these axes in some depth, touching upon the second only in some speculative concluding remarks.


To this end, we introduce two simple mathematical models of locomotion that we have endowed, in simulation, with various coordination controllers. The first model, depicted in figure 18.3, is a bipedal bead on a rail (BBR) and the second, depicted in figure 18.4, is a bipedal spring loaded inverted pendulum (BSLIP) similar to a pogo stick, but with two alternating legs.



**Figure 18.3** The bipedal bead on a rail (BBR) model. (a) A bead is constrained to move on a rail (central horizontal line in the figure). The legs may grab the parallel outer rails to pull the body forward, which they do via their springs. Specifically, during protraction, a leg is not grasping the outer rail and is extending its spring well beyond its rest length. When it grasps the rail in retraction, the spring pulls the body toward the toe position, attempting to restore itself to the rest length. See the appendix for details of the BBR model. (b) A sequence showing the legs alternating between protraction and retraction. Open indicates the leg is not grasping the rail and closed indicates that it is.

The BBR model consists of a bead on a rail pulled by springy legs whose toes are alternatively fixed (“retraction” phase) or allowed to slide freely (“protraction” phase) on parallel rails. The mechanism propels its “body” directly through the work of its alternating leg springs when they are fixed to the parallel “ground” rails, as we detail mathematically in the Appendix. The only losses occur due to friction between the bead and the rail. This simple (essentially one-degree-of-freedom) model eliminates the issues of posture, and almost all other extraneous degrees of freedom, thereby, focusing, on the role of coordination for the underlying task of moving around a body. In contrast, the BSLIP model consists of a mass affixed to two spring-loaded legs that may rotate around the mass. The mass is constrained to lie in the sagittal plane and is acted on by gravity. It is an extension of the spring-loaded inverted pendulum model (SLIP) that has been proposed as a model for the motion of the center of mass of an animal (or robot) while running.

We will use the BBR model to illustrate comparisons—points of similarity and difference between various control architectures. We will use the more-realistic BSLIP model to illustrate challenges—aspects of the locomotion task that will defeat control architectures inadequately informed about the physical world.



**Figure 18.4** The bipedal spring-loaded inverted pendulum (BSLIP) model. (a) Two spring-loaded legs alternatively support the body and swing over the top of it. (b) A sequence showing the legs alternating in stance. Images 1 and 5 correspond to the touch-down of a leg, images 2 and 6 to maximal compression, and 3 and 7 to liftoff. Image 4 corresponds to the apex of the flight phase when both legs are in swing mode.

For both models, the ground reaction forces imparted to their bodies by their legs determine center-of-mass acceleration. In turn, the body's twice-integrated position (together with a leg coordination strategy) determines the degree to which the leg springs are compressed while in retraction. The BBR model has only one body degree of freedom: fore-aft progress along the uniform rail. Hence, absolute position has no consequence. Because the exchange of body kinetic energy with leg spring potential energy does not play an important role in body translation, this model operates effectively in the quasi-static regime. In contrast, adding a vertical degree of freedom in the BSLIP model introduces the key interchange of kinetic and gravitational potential energy: in the BSLIP model, the body, operating in a true dynamical regime, can be tripped by poorly executed leg behaviors. A coordination controller for BSLIP must therefore work in harmony with a body stabilization controller, which is not the case for the BBR model.

## REVIEW: GENERAL MODELS OF MECHANISM AND ENVIRONMENT

### Mechanism: Design Along the Feedforward/Feedback and Centralized/Decentralized Axes of Control Architecture

Two distinct schemes for coordination have generated significant discussion in the animal motion control literature. The first is coordination due to a central pattern generator (CPG) wherein the state of an actuator, independent of its environment, follows a pattern produced by some centralized, oscillatory mechanism. The second is reflex-driven coordination, wherein the state of an actuator is a function of its environment, which may include the surface of the ground or the states of other actuators. In the engineering control literature, these notions

correspond roughly to feedforward (or open-loop) and feedback (or closed-loop) operation, referring to the flow of information between the animal or robot and its environment.

The biological reality of pattern generators—units of identified neural tissue that exhibit fictive locomotion in isolation, excite true motor activity in more intact animals, but cannot sustain the oscillation upon further subdivision—has been established (Pearson, 1976; Orlovsky, Deliagina, and Grillner, 1999). Beginning with the seminal work of Cohen, Holmes, and Rand (1981/82) and continuing through more elaborated refinements (Cohen, Rossignol, and Grillner, 1988; Cohen et al., 1992; Kopell, 1995, 2000), a flourishing applied mathematical literature has employed the language and methods of dynamical systems theory to model the nature of such feedforward motor-control signals. In the robotics literature, purely feedforward-driven control schemes have been examined by a small, but growing, community as well (Swanson, Burridge, and Koditschek, 1995; Ringrose, 1997; Berkemeier and Desai, 1999; Komsuoglu and Koditschek, 2000).

In parallel, a long history of behavioral animal science has resulted in an equally large literature concerned with the role of reflexes in motor coordination. In his seminal analysis of insect leg coordination, Wilson (1966) proposed simple local rules capable of generating a great diversity of steady-state insect gaits. Some of the most compelling recent work within this tradition is reported by Cruse and colleagues (1998) respecting the disturbance recovery properties of stick insect locomotion. Cruse and his colleagues have discovered several mechanisms that, in a *kinematic* model of the stick insect *Carausius morosus*, reproduce the gaits and recovery characteristics exhibited by the animal. These models, emphasizing the role of reflex, represent a feedback-based approach to coordinated locomotion. Research by Chiel and colleagues (1992) has successfully demonstrated that such reflex-generated coordination mechanisms can elicit robust locomotion in hexapod robotic platforms operating in the quasi-static regime.

Feedback coordination schemes not necessarily inspired by the natural world have, of course, been employed in robot control as well. For example, the “virtual model controllers” of Pratt and colleagues (Pratt, Dilworth, and Pratt, 1997) impose joint torques on their walking machines’ limbs in such a fashion that the center of mass exhibits the compliance of some desired quasi-static template—typically, an overdamped spring. The juggling machines built by the last author and his former students (Rizzi, Whitcomb, and Koditschek, 1992; Buehler, Koditschek, and Kindlmann, 1994; Rizzi and Koditschek, 1996) constitute another such class of purely reflex-controlled robots. In the work on juggling, there is an explicitly dynamical (one-degree-of-freedom) template—a mass bouncing on an active vertical post—designed to self-excite stable periodic motion at a prescribed recurrent height. A

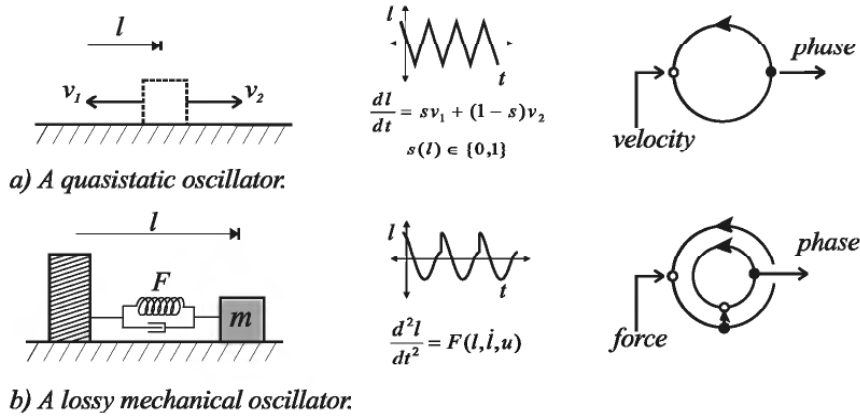
different one-degree-of-freedom dynamical template, the hanging pendulum, is used by Nakanishi, Fukuda, and Koditschek (2000) to elicit stable locomotion in a brachiating robot—again by self-excitation to a prescribed total (i.e., kinetic plus potential) energy.

The design axis we call *centralized versus decentralized* is less well understood and we only comment on it briefly, leaving further exploration as future work. On one hand, decentralization of control may be a matter of implementation convenience, and on the other, a matter of computational necessity. For example, the reflex control model for the stick insect, proposed by Cruse (1990), is completely decentralized. There is no central processing unit that collects the state of the constituent limbs and redirects their actuation relative to some global prescription. Individual actuators are allowed very limited information concerning the state of their neighbors, and the body's coordinated gait emerges from their mutual interaction via the body and ground. In contrast, the reflex schemes described in the previous paragraph presume the availability of a global view of limb and body state that can be accessed freely by the constituent actuators in making their contribution toward the global reference template. This same distinction between centralized and decentralized architectures can be made in the context of purely feed-forward control schemes. For example, RHex's feedforward controller presents a completely centralized architecture, as clearly depicted in figure 18.2. In contrast, lamprey locomotion models of the kind championed by Cohen and colleagues (1992) have the form of coupled peripheral pattern generators. Recent research (Cohen and Boothe, 2000; Lewis et al., 2000) suggests that similar completely local feedforward coordination can yield robust legged robot locomotion in the quasi-static regime. The challenge is to understand whether and how such an approach might be used more generally—for example, in the dynamical regime.

In this chapter, we concentrate on the feedback/feedforward axis, developing a framework for thinking about coordination control that permits an integration of the two extremes in both the robotics and the animal motion control literature. We will do so by reinterpreting the reflex model in a form that promotes its incorporation within the more traditional mathematical framework of coupled oscillators.

### **Environment: Phase Regulation Mechanisms for Dynamical versus Quasi-static Tasks**

**Oscillatory Systems that Perform Work Can Only Regulate Phase by Adjusting Energy** We depict in figure 18.5 two oscillator models commonly encountered in the literature on biological motor control but whose crucial differences represent a significant obstacle to the kind of unified perspective we hope to form. A mass whose position (length  $l$



**Figure 18.5** Two different types of oscillator with their corresponding equations and icons. (a) A first-order system cannot oscillate without some switching controller. Because velocity is the control input to the system, we denote a first order oscillator by a single circle. (b) A second-order oscillator has a phase mechanically determined by its energy. If it has damping, periodic control inputs (underscored in the figure by the discontinuities in the mass’s trajectory) are needed to reintroduce energy. Because the control input to the system is force, which only indirectly affects phase, we denote a second order oscillator by two circles.

in the figure) describes a periodic function of time may be considered an oscillator possessed of a “phase” variable that measures the percentage a given cycle has completed relative to some agreed-upon fixed reference, such as its point of maximal extension along the line. Continuous signals cannot be both periodic and nonreversing (i.e., exhibiting nonnegative speed) unless they lie on a circle. This motivates the terminology “clock,” as suggested by the icons appearing in the right-hand column of the figure. Quasi-static models are characterized by masses whose velocities can be altered directly by applied forces as depicted in figure 18.5a. First-order systems (i.e., their state is one dimensional), cannot be made to oscillate unless some regulatory switch is used to mediate the reversal of applied force as a function of extension relative to some reference lengths (that we simply denote by  $l_+$  and  $l_-$  for now). The resulting hybrid (i.e., mixed-continuous and discrete-state) system exhibits immediately an exactly periodic oscillation. Such models yield intrinsically stable clocks that may be instantaneously reset and which are tuned by position variables, the reference lengths,  $l_+$  or  $l_-$ .

In contrast, models suitable for studying mechanical work have state defined by their velocities as well as by the positions of their masses. Applied forces alter the acceleration of the body and, hence, can only affect position via a cumulative integrated effect. In the presence of positive stiffness, potential forces such as the Hooke’s law spring depicted in figure 18.5b, such systems oscillate with an intrinsically (neutrally) stable phase velocity determined by their energy level. When they

perform work upon a lossy environment (as in this study: recall that the leg springs of the BBR and BSLIP models are subject to viscous damping), then their energy gradually leaks out and the phase velocity decays to zero—effectively stopping the clock. There is no other way to maintain a lossy mechanical oscillator other than by pumping in energy—for example, as depicted here, by tuning the potential energy stiffness,  $k$ , or by adjusting the reference lengths  $l_+$  and/or  $l_-$ . However, because phase results from the integrated effects of these energy changes, the resulting lag can destabilize the clock if the regulatory policy is not carefully applied.

**Example: The Reflex as a Quasi-static Mechanism for Feedback-Based Phase Regulation** A major objective of this chapter is to recast a class of reflex models introduced by Cruse (1990) in a form amenable to regulating the exchange of total energy rather than merely prescribing kinematic relationships between limb positions and velocities. In Cruse’s model, each leg alternates between *protraction* (swinging the leg forward) to an anterior extreme position (*AEP*) and *retraction* to a posterior extreme position (*PEP*). (Notice that we deviate slightly from the literature and let *AEP* and *PEP* denote lengths and not angles, the former being a simple yet convenient kinematic transformation of the latter.) Each leg receives a combination of signals (which we describe in more detail in the next section) from certain neighboring legs, which alters the *PEP* of the leg. Changing *PEP* effectively changes the period of the leg, synchronizing it in or out of phase with its neighboring legs. For example, in one mechanism, a leg in protraction signals the next rostral, ipsilateral leg to decrease its *PEP*, thus causing the receiving leg to retract further. The combination of mechanisms operating in parallel according to an experimentally observed connection scheme results in a decentralized controller model.

Beer and colleagues (Chiel et al., 1992) have successfully used three of Cruse’s mechanisms to control the gait of a hexapod robot operating in the quasi-static regime. Furthermore, Calvitti and Beer (2000) have begun to show analytically how these mechanisms might operate in a model that is, roughly speaking, a logical abstraction of the kinematics presented by our BBR model. In particular, they examine a two-leg arrangement wherein one leg (the master) influences the other (the slave) via one of the three mechanisms. There is no reciprocal coupling. The legs are not attached to a body nor do they ever touch the ground, providing no means of mechanical or environmental coupling.

The authors are able to prove that each mechanism results in a neutrally stable, attracting region of phase relationships centered near  $180^\circ$  out of phase. That is, any phase relationship in a certain range is possible—phase relationships outside the range are *pushed* into the stable region, but no particular phase relationship within the stable



range is preferred. Such neutral stability is consistent with the hypothesis in Wilson (1966) that one gait may smoothly transfer to another, although other mechanisms may produce this behavior as well. It remains to analyze situations where legs are mutually coupled, coupled with more than one of the mechanisms Cruse describes, or where there are more than two legs. Furthermore, a model that includes the body of the insect coupled with its environment warrants analysis to relate the stability of body attitude and velocity to the stability of gait.

**Example: The CPG as a Dynamical Mechanism for Feedforward-Based Phase Regulation** For the purposes of this chapter, a CPG is an internally generated periodic signal used to produce a motor reference trajectory that is tracked by the actuators of the robot. Coordination results in a manner similar to that of a conductor leading an orchestra. In the control of RHex, we have chosen to realize this reference generator by passing a “clock”—a constant-speed periodic signal source—through a monotone output function that “shapes” the commanded speed profile, as depicted in the upper portion of figure 18.5. The speed profile presently used in RHex is graphed in figure 18.6c.

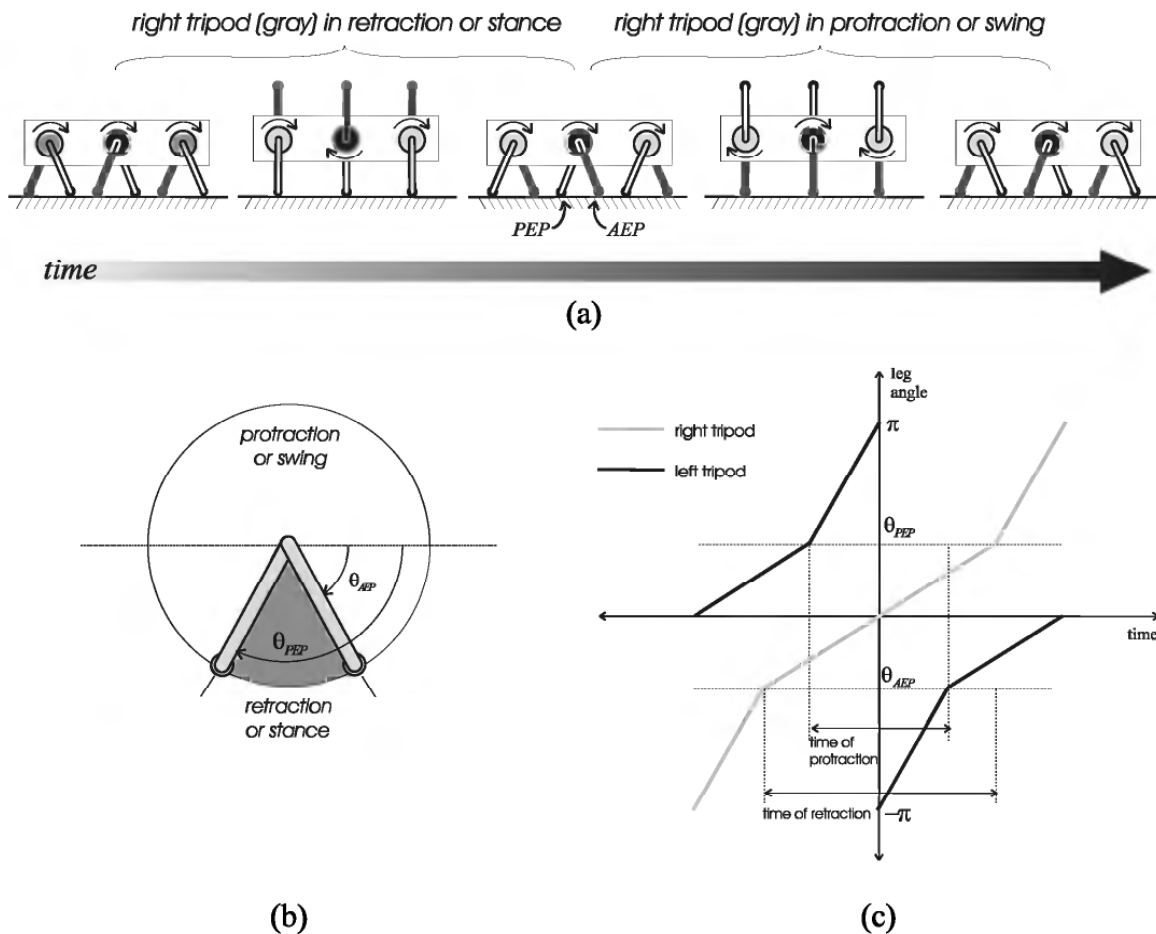
We model the clock as a harmonic oscillator, denoting its phase by  $\phi$ . The CPG thus effectively introduces an additional state whose dynamics are governed by the differential equation

$$\dot{\phi} = \frac{2\pi}{T_c}$$

where the parameter  $T_c$  represents the period of the CPG, and the “shape” of the speed profile is set by adjusting such constants as those depicted in figure 18.6 used in RHex.

Notice that the period of the harmonic oscillator essentially corresponds with its energy. Thus, we must explicitly manipulate the parameter  $T_c$  in order to change the phase velocity of the clock. In other words, clocks effectively show up as first-order (a single circle in the icon system introduced in figure 18.5) oscillators in the study we report here. Although second-order oscillators (e.g., Fitzhugh-Nagumo or other two-dimensional nonlinear approximations to the Hodgkin-Huxley equations; Holmes and Taylor, 1997) have been proposed as models of neural dynamics, we prefer to adopt the reasoning of Cohen, Holmes, and Rand (1981/82) and operate in the more neutral “phase coordinates” known to exist in the neighborhood of any limit-cycle oscillator.<sup>1</sup>

The CPG approach to control, avoiding any active measurements, is by its nature immune to sensor noise issues, which is a big advantage in practical implementations. On the other hand, steady-state behavior depends very strongly on the environmental parameters, and, therefore, external disturbances to the system cause it to vary its steady-state body speed wildly, thereby eliminating coordination between legs. We



**Figure 18.6** CPG control of RHex. (a) An animation sequence showing a full stride for the right tripod, shown in gray, of RHex. The counter motion of the left tripod is shown in black. (b) RHex's "clock"—the periodic feedforward reference signal presently used to drive an alternating tripod gait. Between the angles  $\theta_{AEP}$  and  $\theta_{PEP}$  the leg is in stance or retraction. The rest of the time it is in swing, or protraction. (c) The speed profile for the two tripods. During retraction, the leg moves slowly. While in protraction, it moves quickly, sweeping out a greater angle in the same amount of time.

have observed that the leg synchronization is maintained fairly well if the external changes are slow and do not drive the system into regions where no stable fixed operation exists. (Figures 18.13 and 18.14 illustrate the fragility of the CPG approach.)

In summary, coordination via a CPG does not require sensory information to generate control signals but instead drives the underlying cyclic mechanism in a purely feed-forward fashion.

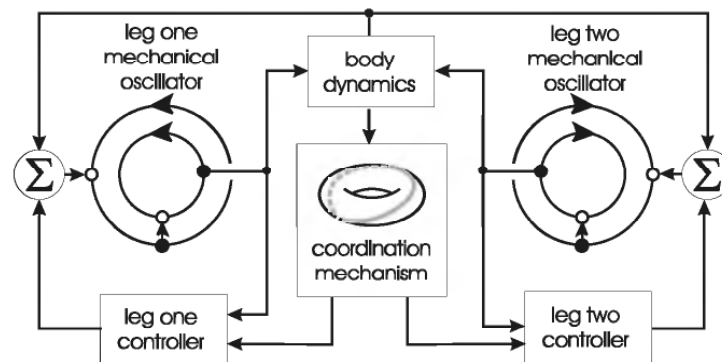
### SIMULATION STUDY: APPLICATION OF REFLEX AND CPG CONTROL TO TWO DYNAMICAL MODELS

RHex is presently controlled via the centralized pattern generator depicted in figure 18.2, which provides no mechanism for incorporating body-, leg-, or ground-state feedback. On the other hand, the feed-

back mechanisms discussed in the section entitled Example: The Reflex as a Quasi-Static Mechanism for Feedback-Based Phase Regulation target the quasi-static regime of operation. In this section, we use the “cartoon” BBR and BSLIP models to provide simulation evidence that the reflex coordination mechanism must indeed be reworked for successful operation in the dynamical regime. We then offer a modified representation of the reflex mechanism and simulation evidence suggesting that it works quite effectively in the dynamical regime. We also use the models to explore the feedback/feedforward axis of dynamical coordination control, presenting a means of interpolating between the two.

### The BBR and BSLIP Dynamical Models

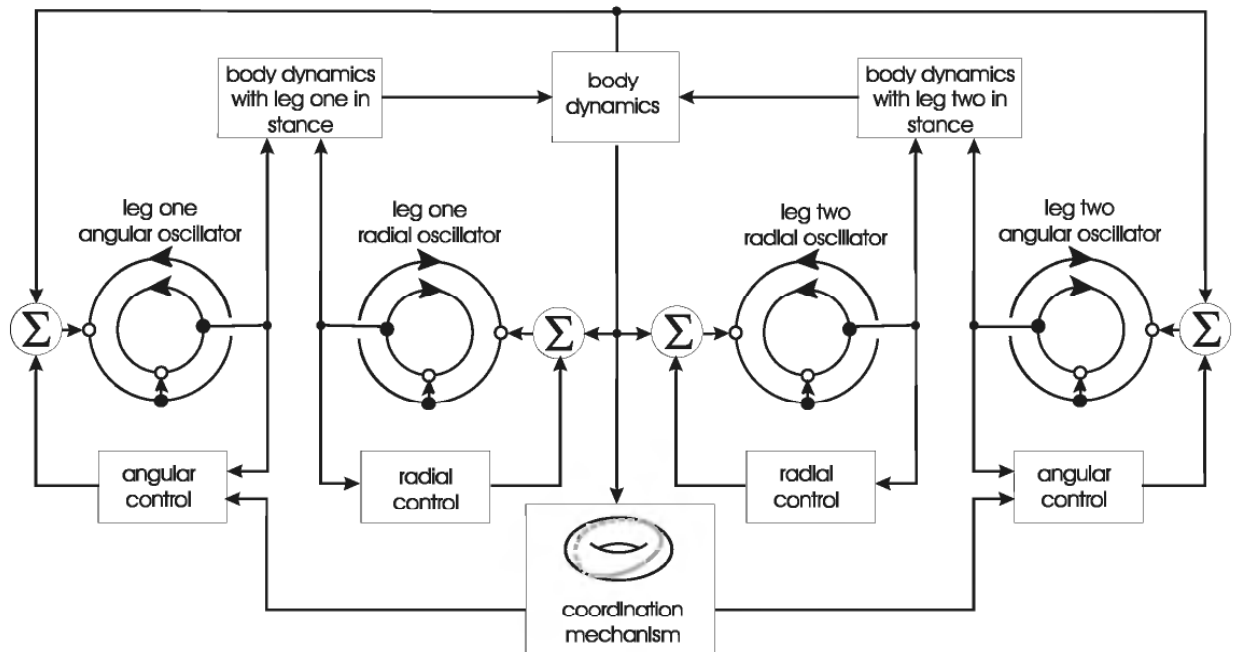
The BBR and BSLIP models are detailed mathematically in the appendix and it will suffice here merely to sketch the architecture of their dynamics and control blocks as in figure 18.7. Each leg periodically forces the body, defining second-order mechanical oscillators<sup>2</sup>—the double circles, according to the icon conventions established in figure 18.5. The dynamics are *hybrid* continuous/discrete—discrete modes being defined by whether one, two, or no legs are touching the ground. Because that condition may be read off the pair of leg phase angles, the block labeled “body dynamics” may be correctly imagined as rep-




**Figure 18.7** Block diagram suggesting the dynamics of the BBR model introduced in figure 18.3 (see the appendix for mathematical details). Each leg consists of an angular degree of freedom that periodically forces the body upon making ground contact, yielding, at each touchdown a second order mechanical oscillator. The resulting hybrid body dynamics (i.e., the evolution of the 1 degree of freedom center of mass position and velocity) together with the leg states determine whether one, two, or no legs are in contact. The coordination mechanism is allowed to set the posterior extreme position, *PEP*, of a leg. Reflex mechanisms adjust this input based on the state of the body and each leg; CPG mechanisms simply select the input in an unvaryingly periodic manner. In this figure and in figure 18.8, we have shown *two* oscillator blocks to illustrate the idea that each leg is treated as though it “owns” the body. The actual force on the body produced by each leg-oscillator is computed in the “body dynamics block.”

resented by an appropriate dynamical system on a torus (the cross product of two circular phases), as we have detailed in Klavins, Christ, and Koditschek (2000). Conceived of in this light, the job of the block-labeled coordination controller is to regulate a pair of angles.


In the case of the BBR model, the coordination controller acts to regulate that pair of phase angles by prescribing the *PEP* for each leg, determining in this manner the next angle of leg liftoff. In the case of the BSLIP model, each of the double circles should properly be repeated, as in figure 18.8, because there is a second degree of freedom (vertical as well as fore-aft or, equivalently, radial and angular) that must be coordinated as well. In the present study we have, following Raibert's lead in the control of hopping robots (Raibert, 1986), employed a feedback controller whose radial and angular components are decoupled.<sup>3</sup> The radial degree of freedom is controlled according to a PD (proportional derivative) law that seeks to regulate height by adjusting the spring stiffness independently of the leg angle. The angular degree of freedom is controlled by adjusting the touchdown angle (coincident with the




**Figure 18.8** Block diagram suggesting the dynamics of the BSLIP model introduced in figure 18.4 (see the Appendix for mathematical details). Each leg consists of a radial and an angular degree of freedom that periodically force the body upon making ground contact, yielding, each at touchdown a parallel composition of two second order mechanical oscillators. The resulting hybrid body dynamics (i.e., the evolution of the two degrees of freedom center of mass position and velocity) together with the leg states determine whether one, two, or no legs are in contact. The coordination mechanism is allowed to set either the posterior extreme position, *PEP*, or the swing velocity of the leg. Reflex mechanisms adjust this input based on the state of the body and each leg; clock mechanisms simply select the input in an unvaryingly periodic manner.



(a)



(b)



(c)

liftoff angle magnitude at steady state). In our study, we connect the output of the coordination block to the BSLIP angular controller by interpreting *PEP* as a phase coordinate prescription from the former and as a liftoff angle prescription to the latter block.

Throughout the remainder of this section, we will report on the results of simulating the dynamical system depicted in figures 18.7 and 18.8 with different coordination controllers substituted in the blocks marked “coordination control.”

## Dynamical Reflex Coordination Control

**Quasi-Static Reflexes Support “Walking” in the BBR Model** We first apply a reflex-coordination scheme similar to that described by Cruse and studied by Beer and Calvitti (see the section entitled Example: The Reflex as a Quasi-Static Mechanism for Feedback-Based Phase Regulation). Of the coupling mechanisms described by Cruse (1990), three are considered to be at work in regulating gait. In each, the *PEP* of one leg (the receiver) is offset based on the state of the other leg (the sender) as described in Beer and Calvitti’s analysis of Cruse’s work (Calvitti and Beer, 2000). In our model, each leg serves as both sender and receiver. In the first mechanism, the *PEP* of the receiving leg is decreased below its nominal value by some small amount,  $\alpha_1$ , when the sending leg is protracting, thereby delaying the start of protraction in the receiving leg. In the second mechanism, the *PEP* of the receiving leg is increased beyond its nominal value by a small amount,  $\alpha_2$ , when the sending leg is in retraction and its length is within some small distance of its *AEP*, thereby shortening the period of the receiving leg. In the third mechanism, the *PEP* of the receiving leg is increased from its nominal value gradually, throughout the retraction mode of the sending leg, by the value  $\alpha_3(AEP - l)$  where  $\alpha_3$  is a tunable parameter and  $l$  is the length of the sending leg.

As a kind of conceptual calibration to the previous literature—particularly the papers of Cruse and Beer mentioned above—we simulated the BBR model with each of the mechanisms in turn. With two-way

**Figure 18.9** Comparing the “return maps” generated by two styles of reflex coordination mechanism applied to the BBR model. Each figure exhibits the same BBR “Poincare section” wherein the horizontal axis represents body velocity at the instant leg one touches down while the vertical axis represents the position of leg two at that same instant. The dots represent samples from several hundred runs, each run having initiated at a distinct different state and sampled forward in time until steady state (depicted by a darker dot) is achieved. (a) Quasi-static style reflex coordination achieves a neutrally stable region (in dark dots). (b) Coupled oscillators style reflex coordination stabilizes at an exact, alternating gait—in other words, the stable region is a point. (c) The coupled oscillators style reflex coordination mechanism with a saturated energy function reproduces the effect in (a).

coupling, we found, as did Calvitti and Beer in their model with one way coupling, neutrally stable gaits for each mechanism. In addition, we found that the forward velocity of the body was regulated to a steady state. We found the third mechanism to be the most aggressive and we present simulation data for it in figure 18.9a. Note that the average velocity seems to stabilize after an initial acceleration and that the legs approach out-of-phase alternation but do not quite achieve it.



### **Quasi-Static Reflexes Do Not Support “Running” in the BSLIP Model**

We now replace the BBR model with the BSLIP and retain the reflex in the “coordination controller” block. It is not obvious that Cruse’s model of insect reflexes, concerned as it is with the behavior of decidedly quasi-static insect walking, should extend to the control of insect running. Specifically, because it does not address the exchange of potential and kinetic energy, that coordination mechanism must be supplemented with some additional scheme that can account for the second-order dynamics of running captured by the BSLIP model.

Initially, we attempted to supplement a Cruse-style reflex mechanism with a Raibert-style gait controller. This simulation coordinates the BSLIP legs by adjusting their liftoff length—the analog of *PEP* in this setting. Control of the remaining gait parameters (hopping height, forward velocity, and duty factor) is accomplished using mechanisms similar to Raibert’s, as discussed above. However, the domain of stability of this control scheme was quite small. Essentially, we found that changing the liftoff conditions interfered with the control of forward velocity and hopping height unless the synchronization gain,  $k_{sync}$ , was quite small, while a too-small  $k_{sync}$  resulted in a lack of synchronization—meaning that at some point the mass hit the ground because neither leg was “ready” to touch down. See the Appendix for an explanation of  $k_{sync}$ . Examples of these behaviors are shown in figure 18.10a and b.

**Dynamical Reflexes Use the Language of Coupled Oscillators** The quasi-static view of the world presumes that forces effect change in position. This is a useful fiction when operating in a highly damped regime, or in the context of behaviors that do not entail the exchange of kinetic energy. To understand intuitively the importance of managing explicitly the exchange of total energy intrinsic to mechanical systems that perform work in demanding regimes, consider the following analogy. When driving a car, there is a delay between the control input and the position output of the system. The driver of the car experiences this delay, for example, when he or she steps on the brakes: the car does not stop immediately but rather decelerates, gradually coming to a stop.

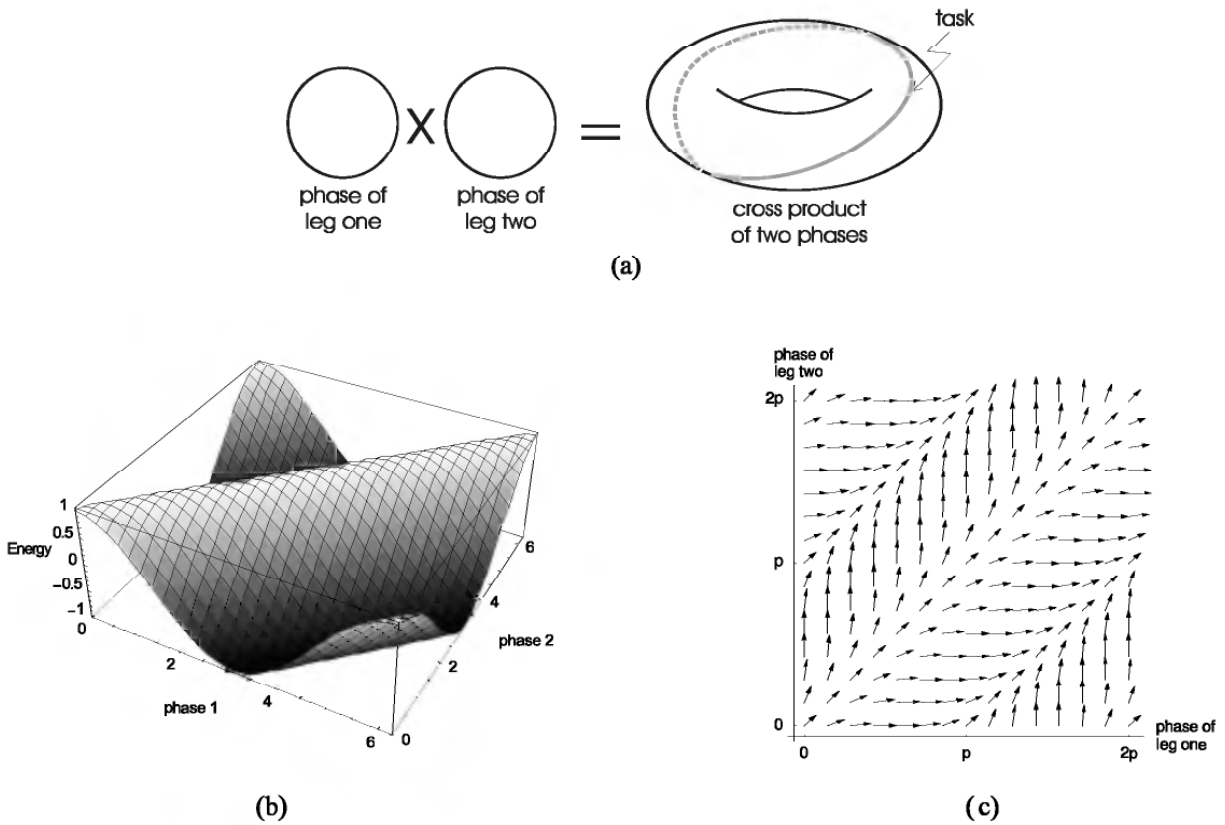
In this section we describe another, similar, coordination mechanism resulting from a formalization (Klavins, Christ, and Koditschek, 2000) of



**Figure 18.10** Two characteristic failure modes of the quasi-static style reflex control mechanism applied to the BSLIP model. (a) Using a higher synchronization gain and low swing velocity. The faltering trajectory of the center of mass height and the increasingly uncoordinated leg angles are plotted as functions of center of mass horizontal position. (b) Using a lower synchronization gain and higher swing velocity. The phase difference between the two legs, plotted as a function of time (solid line), fluctuates with increasing magnitude as each leg “waits” a longer time to touch down (because the swing velocity is set high) until, eventually, the body falls to the ground. Compare this to the neutrally stable phase difference trajectory (dashed line) resulting from the coupled oscillator style reflex control mechanism.

our phase-regulation algorithm for vertical juggling (Rizzi, Whitcomb, and Koditschek, 1992; Buehler, Koditschek, and Kindlmann, 1994) that presents (when viewed with the appropriate abstraction) an almost identical coordination problem to our present concern arising from locomotion. Implementation within the simple BBR and BSLIP models suggests that this representation of local reflexes is capable of imitating key properties of the quasi-static style reflex coordination mechanisms introduced by Cruse, described above. Expressed within the traditional framework of coupled oscillators, it enjoys the added benefit of straightforward analysis via standard tools from dynamical systems theory. Moreover, as we will soon show, for this reason it can be extended in a





**Figure 18.11** The coupled oscillator style reflex coordination mechanism results from correcting a measured pair of phase velocities at a measured pair of phases toward a desired pair at that measured phase. The desired pair of phase velocities is specified by a reference phase dynamical system constructed as follows. (a) The set of phase angle pairs (the “cross product” of two cyclic phase spaces) is a torus. The task of cycling out of phase is encoded as a circle embedded in the torus along with a desired instantaneous phase velocity for each point along the circle. (b) A “pseudo-energy” function  $V(\phi_1, \phi_2) = \cos(\phi_1 - \phi_2)$  yields a reference field derived by taking the gradient (directional derivative) at each phase pair. The desired phase difference cycle lies along the “low energy trough.” (c) The reference field  $R(\phi_1, \phi_2) = k_1 - k_2 \nabla V(\phi_1, \phi_2)$  derived from  $V$ . The legs are driven (at a rate proportional by constant  $k_2$  to the gradient magnitude) out of phase down onto the “trough” along which the phase velocities are pulled at the constant rate  $k_1$ .

straightforward manner to handle the exigencies of the dynamical regime represented by our more complex BSLIP model.

The method, summarized in figure 18.11, is based on defining an *artificial energy function*,  $V$ , on the difference in the phases of the two legs. Specifically, if we let  $\phi_1$  and  $\phi_2$  be the phases of two legs, one suitable definition for  $V$  is  $\cos(\phi_1 - \phi_2)$ . Minimizing the energy between the two phases via changing *PEP* in a suitable fashion requires that  $\phi_1$  and  $\phi_2$  differ by  $\pi$ —that is, that the two legs be exactly out of phase with each other. The resulting dynamics bear a strong resemblance to the provably stable system described in Klavins, Ghrist, and Koditschek (2000). We note that in our simulations we do not have direct access to

the actual phase of a leg. Rather, we define an approximation to phase based on the length of the leg rather than on time. Unless the body moves backward, this quantity increases monotonically with time and is an effective stand-in for phase in the present application.


In simulation, the system regulates the forward velocity of the body to a steady state and regulates the legs to an alternating gait *exactly* 180° out of phase. The alternating gait appears to be stable.

The energy function may be defined by  $\max\{\alpha, \cos(\varphi_1 - \varphi_2)\}$  for some number  $\alpha$  between  $-1$  and  $0$ , to simulate the third mechanism described by Cruse (1990), and results in neutrally stable gaits centered around 180° out of phase. In figure 18.9 we compare the return maps, obtained numerically, of the model with different energy functions and with Cruse's third mechanism in place. We conclude that the phase-regulation method can mimic Cruse-style coordination and provides an obvious way to specify gaits exactly by providing the appropriate energy function.

**Dynamical Reflexes Support “Running” in the BSLIP Model** An alternative means of regulating the legs in this model is to use the fact that phase regulation is based on a *continuous* reference field, which is obtained from the energy function  $V$  as in figure 18.11. Thus, the legs may be synchronized at times other than liftoff events. Specifically, instead of changing  $PEP$ , which, as we have shown, substantially affects the motion of the body, we change the swing-phase dynamics by modifying the stiffness of the hip springs,  $k_i$ , to be  $k_{nom} - k_{sync} \sin(\theta_i - \theta_j)$ —that is, a nominal stiffness modified by a gain times the negative gradient of the “pseudo-energy” function (see the Appendix for an explanation of the gain  $k_i$ ). Thus, we suppose that the leg angles correspond roughly to phase and use the fact that  $k_i$  controls the velocity of leg  $i$  during swing. Because we only change the speed of swing, we do not affect the underlying control, which affects the touchdown angle and the leg-spring stiffness during decompression. Data from a typical run are shown in figure 18.12. It is not obvious how such a coordination controller could have been constructed using the Cruse-style reflex framework directly and in a natural way, although careful consideration of the effects of changing the liftoff length on the horizontal velocity may prove fruitful. We believe the proper route is to use the mechanisms described by Cruse as inspiration, but to work in the more informative language of coupled oscillators when designing controllers.

### Exploring the Feedback/Feedforward Axis of Dynamical Coordination Control

**CPG-Based Coordination Control** We aim to develop a means of interpolation between pure feedforward and pure feedback coordination



**Figure 18.12** A simulation of the coupled oscillator style reflex coordination mechanism applied to the BSLIP model using the “pseudo-energy” function described in figure 8.11. After an initial transient period, the system settles down into an alternating gait and the characteristic body motion of a spring loaded inverted pendulum. See figure 18.10a for an explanation of the axes.

mechanisms. We have, just above, established a pure feedback mechanism that implements (Cruse-inspired) leg angle corrections by means of phase-velocity corrections (as described in figure 18.11). To establish the desired correspondence with pure feedforward coordination, it is helpful to take the obverse step of reworking an absolute reference-velocity-based controller (inspired by the RHex coordination mechanism of figure 18.2) into an absolute reference-position-based controller as follows. One can effectively realize the two-speed RHex velocity profile of figure 18.6 by running a constant speed clock and issuing a “leg liftoff” command as its internal reference phase crosses a marker at  $\phi = \phi_{pep}$ . To coordinate two legs of an alternating gait, one places two markers,  $\phi_{pep1}$  and  $\phi_{pep2}$ , on the circle that are  $180^\circ$  out of phase.

As a feed-forward mechanism, such a position-based clock has no better “idea” of the leg’s true state than its velocity-reference-generating equivalent of figure 18.6. But the physical fact of whether or not the leg is in contact with the ground determines the physical consequences of the command. For example, the leg liftoff command has a physical effect if it arrives while the leg is in retraction. That command can have no physical effect if the leg is already in protraction when it is received. This observation should now make it clear that notwithstanding the absence of an explicitly implemented sensory pathway from the leg’s mechanical oscillator back to the CPG, the following two relationships between actuation effect and phase difference provide a “hidden” mechanical feedback channel—the “preflexive” mechanism. First, the leg liftoff command may cause the leg to begin protracting at any point in retraction and thus alters the effective *PEP* length and changes the leg cycle period. Second, the change in phase difference is a monotonic function of the difference between the periods of the leg and CPG. This


mutual dependency causes the phase difference, sampled at *AEP* for example, to converge to a steady-state constant value. We observe that the losses in the mechanical system and the above phase-actuation relationships are sufficient for asymptotic convergence to a steady-state operating point.

Applied to a single-legged version of the BBR model, this feedforward controller would yield a system very similar to one analyzed formally in Komsuoglu and Koditschek (2000), wherein sufficient conditions are given for the stability of a feedforward “clock”-controlled hopping gait. In that work, stabilization can be formally demonstrated to obtain via the implicit interaction of the mechanical oscillator with the clock—the “prelexive” mechanism intuitively described above. Applied to the complete BBR model of interest here, the feedforward-controlled two-legged systems are best thought of as a pair of one-way coupled oscillators—one internal (the CPG) and one mechanical (the leg)—as depicted in figure 18.2, corresponding to the setup in figure 18.7, where the connection between the body dynamics block and the coordination controller (that block now occupied by the coordinated pair of reference-marked clocks from figure 18.6) has been severed.


Our simulation study of the CPG mechanism applied to the BBR model suggests that feedforward coordination of a pair of prelex-stabilized legs functions as effectively as a prelex-stabilized single leg.<sup>4</sup> Each leg converges to its respective “synchronized” operating point, the transient effects of the other leg are rejected as essentially external disturbances, and the body converges to a steady velocity. At the steady-state operation, the legs are 180° out of phase in accordance with the markers on the CPG (figure 18.13a).

**Comparing CPG and Reflex Coordination of BBR** Having related reflex and clock-style coordination through the language of coupled oscillators, we are now in a position to compare their performance in various versions of our BBR and BSLIP models. Although a detailed comparison across many parameters of interest lies beyond the scope of this chapter, it is worth presenting one specific example of the questions that might now be addressed, both analytically as well as computationally.

As might be expected, the CPG suffers serious limitations on the achievable gaits in an uncertain world. For a single-legged version of the BBR model, one can compute quite straightforwardly a CPG period that has a fixed operation point yielding a desired body velocity: Choose a *PEP* leg length that does not yield net body speed change and integrate the differential equation over a complete leg cycle. This computation assumes there is at most one leg applying force to the body at any given time—hence, absent a “double stance” juncture, it applies as well to the full (two-legged) model, provided there is no double stance




(a)



(b)

**Figure 18.13** Simulations of the BBR model using a CPG for coordination as described in the text. The horizontal axis is time, the vertical axis contains the body’s forward velocity and the positions of the two legs. (a) The CPG period is tuned to the natural frequency of the system as defined by the stiffness of the leg springs. Notice that the gait and the forward velocity stabilize. (b) The CPG is too fast: the legs cannot realize the signals it sends, resulting in an undesired, “limping” gait.

at the desired steady state. As this tuning procedure suggests, the CPG period cannot be chosen arbitrarily, but rather needs to be in proper accord with such task parameters as the natural stiffness of the legs and rail width. In support of this, our numerical studies, sampled in figure 18.13a, b, indicate that the resulting fixed operation may be stable or unstable depending on the period of the CPG relative to the natural period of the BBR model. In general, the feedforward CPG controller cannot dictate an oscillation that is faster than the legs are physically capable of realizing. The equivalent problem in RHex would occur if its CPG oscillated faster than its motors were capable of spinning: The legs would end up playing a never-ending, possibly unstable, game of catch-up with the CPG.




**Figure 18.14** The relationship between the steady state behavior of the CPG coordinated BBR model and the environment into which it operates. The abscissa represents the steady state velocity of the body,  $\dot{x}^*$ , sampled at the AEP event of a specified leg. The ordinate represents the distance,  $w$ , from the central rail to the outer guide rails. The black regions indicate that there is no fixed point for the  $(w, x^*)$  pair. Gray regions indicate that there is a fixed point that is unstable. And, finally, the white region represents the stable fixed points.

The previous intuitive arguments may be confirmed by direct simulation. Figure 18.14 depicts the relation of the steady-state response (forward velocity  $x^*$ ) of the BBR system with CPG control to the distance  $w$  from the rail to the outer guide rails. This figure exhibits the fixed-point dependence to the environmental settings in a feedforward control system. In contrast, figure 18.15, exhibiting the same relationship for the reflex-coordinated BBR, suggests the value of feedback as a hedge against environmental uncertainty. In contrast to the previous case, we see that the model can operate successfully at every rail width, with no a priori information about what that parameter might be.

#### CONCLUSION: STABILIZING FEEDBACK PATHWAYS FOR MECHANICAL LOCOMOTION

RHex (figure 18.1) functions quite ably in unusually demanding terrain on the strength of its “preflexes” alone. At present, RHex’s CPG (figure 18.2) receives no information regarding the state of the actuators, the



**Figure 18.15** The function relating steady state behavior of the Reflex coordinated BBR model to the environment into which it operates. The abscissa and ordinate are the reverse from figure 18.14. As expected, feedback confers significant robustness against environmental variation: there is a unique asymptotically stable body velocity for each environmental condition.

legs they move, the body that translates in consequence, nor the terrain that body visits. It seems clear that significant advantage can derive from feedback, but it is not yet clear how best to insert sensory information in the purely feed-forward control scheme. In this chapter, we have reported on work in progress addressing this question—the beginnings of a simulation study that promises to shed light on how to do so.

Specifically, we have introduced two simple computer models of bipedal locomotion: One, the BBR model, where the exchange of potential and kinetic energy of the body does not significantly impact the locomotion task; and the other, the BSLIP model, where it matters a great deal. We have explored within these two contrasting examples the applicability and efficacy of two general paradigms of biological motor control: CPGs and reflexes. To carry out this exploration we have found it useful to modify the quasi-static representation of reflex coordination originally proposed by Cruse (1990), bringing it into correspondence with the traditional literature on coupled oscillators. Whereas the past representation succeeds in coordinating the BBR model, it leads to tripping and crashing in the BSLIP. The revised representation we introduce here affords a version of the reflex mechanism capable of coordinating the total energy as well as kinematic variables relevant to running.

It would be most desirable to select from a spectrum of possible coordination models that introduce more or less sensory information at will. However, engineers will confidently deploy an appropriate combination of these general coordination mechanisms only when similarly general principles addressing the efficacy of such combinations have

been developed. The daunting modeling and analysis task requisite to such general understanding is arguably best approached by initial appeal to “cartoon” instances developed that capture the essence of the problems, yet which are more amenable to formal inquiry. Such a formal inquiry is the target of our longer-term research effort in this area.

## ACKNOWLEDGMENTS

The Computational Neuromechanics Project is supported by DARPA/ONR under grant N00014-98-1-0747. The CNM Hexapod project is supported by DARPA/SPAWAR under contract N66001-00-C-8026. We thank Martin Buehler, Uluc Saranli, Phil Holmes, Devin Jindrich, and the other members of these projects for their shared insights and help with the specific questions addressed in this paper. We thank Art Kuo as well for a number of useful discussions bearing on the topics addressed and Holk Cruse for reading a draft of this paper and providing several useful comments and suggestions.

## APPENDIX: DETAILS OF THE SIMULATION MODELS

### The Bipedal Bead: A Simple Model

This model consists of a point mass, constrained to lie on a rail, with two massless legs connecting the body to one of two parallel rails, as in figure 18.3. The legs act as springs and pull the body along the rail. A leg can be either protracting (swinging forward, not grasping an outer rail) or retracting (swinging backward, grasping the rail and, therefore, pulling the body). In protraction, the length of a leg is less than  $AEP$  and increasing. In retraction, the length of a leg is greater than  $PEP$  and decreasing. (Note that in this model  $AEP$  and  $PEP$  are not angles but instead are considered, equivalently, to be leg lengths.)

A leg protracts until its length is greater than  $AEP$ , at which point it attaches to the rail and begins to retract. It retracts until its length is less than  $PEP$ , at which point it releases the rail and begins to protract again. The length of time between the beginning of the  $k$ th retraction and the  $(k + 1)$ st retraction is called the *period* of the  $k$ th cycle. The length of the  $k$ th period is a function of the  $PEP$  of the leg and the velocity of the body at the beginning of the  $k$ th retraction. Thus,  $PEP$  is a natural candidate for a control input as it affects the period of leg, which is essential for gait regulation.

We denote by  $s_i[\{p, r\}]$  the discrete state of leg  $i$  and by  $x$  the position of the body. The symbols  $l_1$  and  $l_2$  denote the lengths of the legs, and  $x_{toe,1}$  and  $x_{toe,2}$  denote the positions of the toes of the legs. Finally,  $\theta_1$  and  $\theta_2$  denote the angles from the horizontal to leg one and leg two, respectively. We assume that the body has a mass of one and that it is



constrained in the plane to move along the line  $y = 0$  and the toes along the lines  $y = -w$  and  $y = w$ , respectively, where  $w$  is the distance from the center rail to an outside rail.

Let  $l_0$ ,  $PEP$ ,  $AEP$ , and  $l_{ex}$  be constant parameters such that

$$l_0 < w < PEP < AEP < l_{ex}$$

The constant  $l_0$  is the rest length of a leg in retraction. The constant  $l_{ex}$  is the rest length of a leg in protraction. The constant  $PEP$  is the nominal posterior extreme position of a leg. The actual posterior extreme position of leg  $i$  is  $PEP_i$ —a value that can be used as a control input and which is constrained to lie between  $w$  and  $AEP$ . The constant  $AEP$  is the anterior extreme position. If leg  $i$  is protracting ( $s_i = p$ ), then  $l_i$  increases until it exceeds  $AEP$ , at which time it begins to retract,  $s_i = r$ , until its length is less than  $PEP$ , at which point it begins to protract again.

When a leg is in protraction, its movement does not affect the movement of the body. The leg simply extends and the toe glides along the rail away from the body. Because the legs are massless, we model a protracting leg as a first-order spring:

$$\text{if } s_i = p \quad \text{then } \dot{l}_i = -k_{pro}(l_i - l_{ex})$$

where  $k_{pro}$  is a constant spring stiffness. When a leg is touching the ground, its movement and that of the body are coupled. We have

$$\ddot{x} = \sum_{s_i=r} [k_{ret}(l_i - l_0) \cos \theta_i - b\dot{l}],$$

and the length of leg  $i$  is just a function of the position of the body:

$$l_i = \sqrt{w^2 + (x_{toe,i} - x)^2}.$$

Here,  $k_{ret}$  is the spring stiffness and  $b$  is a damping constant.

### The Bipedal SLIP: A More Complex Model

The bipedal SLIP model of figure 18.4 may be familiar in its resemblance to Raibert's two-legged hopping robot (Raibert, 1986). However, with a view toward relevance to RHex, this model differs from Raibert's hoppers in that the legs must circulate backward and return above the body toward their next touchdown.

Let  $(x, y)$  denote the position of the center of mass of the BSLIP robot, which we assume to have a mass of one. Let  $\theta_1$  and  $\theta_2$  denote the angles of the two legs, which we assume to be essentially massless. The legs may not exceed a length of  $L$ . We denote by  $x_{toe}$  the horizontal position of the toe of the robot when it is in stance and by  $r$ , the distance  $\|(x, y) - (x_{toe}, 0)\|$ . A leg can either be in stance or in swing mode. In stance, it is not controlled directly but instead a slave to the body dynamics. Thus, if leg  $i$  is in stance, then  $\theta_i = \sin^{-1}[(x_{toe} - x)/r]$ .

In flight, the dynamics of the center of mass are

$$\begin{pmatrix} \ddot{x} \\ \ddot{y} \end{pmatrix} = \begin{pmatrix} 0 \\ -g \end{pmatrix} \quad (18.A1)$$

and while in stance they are

$$\begin{pmatrix} \ddot{x} \\ \ddot{y} \end{pmatrix} = -\frac{DU(r) + b\dot{r}}{r} \begin{pmatrix} x_{toe} - x \\ y \end{pmatrix} + \begin{pmatrix} 0 \\ -g \end{pmatrix}. \quad (18.A2)$$

Here  $U(r)$  is the spring potential law. For our simulations we chose  $U(r) = k/r^3$  where  $k$  is the stiffness of the spring, tunable online. In particular, we may change  $k$  during stance to adjust the amount of thrust supplied by the leg. The constant  $b$  determines the amount of damping in the spring and  $g$  is the gravitational constant. We assume that the masses of the legs are much less than the mass of the robot, so that their velocities in swing are directly controllable:  $\dot{\theta}_i = u_i$ .

The discrete dynamics are as follows. The body's dynamics are ballistic (A1) and both legs swing according to their control inputs until one of the legs touches down, which happens when  $y < L \cos(\theta_i)$  for some  $i$ . Then the dynamics of the body are given by equation 18.A2 and the leg that touched down becomes a slave to the body dynamics. The leg that is still swinging is still controllable. When the leg reaches maximal compression, when  $\dot{r} = 0$ , then the spring stiffness  $k$  may be adjusted. When the value of  $r$  exceeds  $L$ , the leg lifts off again and the cycle repeats. In the version of the synchronization controller where  $PEP$  can be altered, liftoff occurs when  $r$  exceeds  $PEP(t)$ , a control input.

Discussion of the details of the body-control algorithm can be found in Saranli, Schwind, and Koditschek (1998). Essentially, we choose the next touchdown angle to be some nominal value plus  $k_v(v^* - v_{prev})$  where  $k_v$  is a gain,  $v^*$  is the desired horizontal velocity during flight corresponding to the nominal touchdown angle and  $v_{prev}$  is the horizontal velocity during the last flight mode. Similarly, we choose the spring stiffness  $k$  to be some nominal value plus  $k_h(h^* - h_{prev})$  where  $h^*$  is the vertical hopping height corresponding to the nominal value and  $h_{prev}$  is the hopping height of the last flight mode.

The leg control during swing is via an artificial spring. That is,

$$\dot{\theta}_i = k_i(\theta_{td} - \theta_i)/(\|\theta_{td} - \theta_i\| + \epsilon)$$

where  $\theta_{td}$  is the desired touchdown angle of the leg and  $k_i$  is a tunable gain. In the continuous control version of the synchronization problem, we have  $k_i(t)$  as a control input, essentially affecting the velocity of the swing.

## NOTES

1. This models exactly the present-day RHex architecture. From the view of biology, the approach can be justified as reflecting the (plausible) presumption that the settling time of

the electrical-chemical energy components inherent in the more complex (second- and higher-order) models of the biological pattern generators should be much faster than those of the body mechanics. This first-cut simplification significantly eases the translation of the Cruse reflex model into the language of coupled oscillators that we introduce in the section entitled Dynamical Reflexes Use the Language of Coupled Oscillators. That translation in place, a more physiologically sophisticated view of the biological pattern generator can be substituted in future studies as needed.

2. To be precise, we simulate the legs using first-order velocity dynamics when they are not in contact, modeling the common situation that they have negligible mass and encounter significant viscous damping. Adding second-leg flight dynamics would have a negligible effect on the qualitative outcome but add significantly to the complexity of the simulation.

3. Even though the Raibert-style architecture decouples the two controller degrees of freedom, the body degrees of freedom are intrinsically coupled through the dynamics of a spring-loaded inverted pendulum. If space and the reader's patience allowed, we would substitute for each double circle in figure 18.7 a copy of the block diagram itself to suggest the manner in which the Raibert controller cleverly coordinates the phases of the vertical and horizontal single-leg degrees of freedom.

4. We have conducted extensive simulations of the single-legged bead on a rail, but have not thought the results worth reporting in the present chapter because they are so reminiscent of the single vertical-hopping leg studied in Komsuoglu and Koditschek (2000). We believe, but have not yet demonstrated, that the stability proof for that case can be readily applied to the single-legged bead on a rail as well. In the sense we have tried to depict in figure 18.7, the single-legged SLIP represents a parallel composition of this vertical hopping leg together with the single-legged bead on a rail. The formal question of whether a CPG controller can stabilize this two-degrees-of-freedom composed mechanism (i.e., the single-legged SLIP) is a matter of ongoing mathematical study.

## REFERENCES

Berkemeier, M. D., and Desai, K. V. (1999). Control of hopping height in legged robots using a neural-mechanical approach. Paper presented at the International Conference on Robotics and Automation, Detroit, Mich., May.

Buehler, M., Koditschek, D. E., and Kindlmann, P. J. (1994). Planning and control of robotic juggling and catching tasks. *Int. J. Robotics Res.* 13(2): 101–118.

Buehler, M., Saranli, U., and Koditschek, D. E. (2000). Design, modeling, and preliminary control of a compliant hexapod robot. Paper presented at the International Conference on Robotics and Automation, San Francisco, April.

Calvitti, A., and Beer, R. D. (2000). Analysis of a distributed model of leg coordination: I. Individual coordination mechanisms. *Biol. Cybern.* 82(3): 197–206.

Chiel, H. J., Beer, R. D., Quinn, R. D., and Espenschied, K. S. (1992). Robustness of a distributed neural network controller for locomotion in a hexapod robot. *IEEE Trans. Robotics Automation* 8(3): 293–303.

Cohen, A. H., and Boothe, D. L. (1995). Sensorimotor interactions: Principles derived from central pattern generators. In M. Arbib (ed.), *Handbook of Brain Theory and Neural Networks*. Cambridge: MIT Press.

Cohen, A. H., Ermentrout, G. B., Kiemel, T., Kopell, N., Sigvardt, K., and Williams, T. (1992). Modelling of intersegmental coordination in the lamprey central pattern generator for locomotion. *Trends Neurosci.* 15: 434–438.

- Cohen, A. H., Holmes, P. J., and Rand, R. H. (1981/82). The nature of the coupling between segmental oscillators of the lamprey spinal generator for locomotion: A mathematical model. *J. Math. Biol.* 13(3): 345–369.
- Cohen, A. H., Rossignol, S., and Grillner, S. (1988). *Neural Control of Rhythmic Movements in Vertebrates*. New York: Wiley.
- Cruse, H. (1990). What mechanisms coordinate leg movement in walking arthropods? *Trends Neurosci.* 13(1): 15–21.
- Cruse, H., Kindermann, T., Schumm, M., Dean, J., and Schmitz, J. (1998). Walknet—A biologically inspired network to control six-legged walking. *Neural Networks* 11: 1435–1447.
- Delcomyn, F. (1985). Walking and running. In *Comprehensive Insect Physiology, Biochemistry, and Pharmacology. Nervous System: Structure and Motor Function*. New York: Pergamon Press, pp. 439–466.
- Full, R. J., and Farley, C. T. (1999). Musculoskeletal dynamics in rhythmic systems—A comparative approach to legged locomotion. In J. M. Winters and P. E. Crago (eds.), *Bio-mechanics and Neural Control of Movement*. New York: Springer-Verlag.
- Full, R. J., and Koditschek, D. E. (1999). Templates and anchors: Neuromechanical hypotheses of legged locomotion on land. *J. Exp. Biol.* 202(23): 3325–3332.
- Full, R. J., Saranlı, U., Buehler, M., Brown, B., Moore, N., Koditschek, D., and Komsuoglu, H. (2000). Evidence for spring loaded inverted pendulum running in a hexapod robot. Paper presented at the International Symposium on Experimental Robotics, Honolulu, Hawaii.
- Full, R. J., and Tu, M. S. (1990). Mechanics of six-legged runners. *J. Exp. Biol.* 148: 129–146.
- Grillner, S. (1985). Neurobiological bases of rhythmic motor acts in vertebrates. *Science* (12 April): 143–149.
- Holmes, P., and Taylor, D. A. (1997). Simple models for excitable and bursting neural networks. *J. Math. Biol.* 12: 37–44.
- Klavins, E., Ghrist, R., and Koditschek, D. E. (2000). Toward the regulation and composition of cyclic behaviors. Paper presented at the Proceedings of the Fourth International Workshop on Algorithmic Foundations of Robotics, Dartmouth, N.H.
- Koditschek, D. (1999). Computational neuromechanics: Programming work in biological and machine systems. University of Michigan. On-line: <<http://ai.eecs.umich.edu/CNM/>>. August 1999.
- Komsuoglu, H., and Koditschek, D. E. (2000). Preliminary analysis of a biologically inspired 1-DOF “clock” stabilized hopper. Paper presented at the World Multiconference on Systemics, Cybernetics, and Informatics, Orlando.
- Kopell, N. (1995). Chains of coupled oscillators. In M. Arbib (ed.), *Handbook of Brain Theory and Neural Networks*. Cambridge: MIT Press, pp. 178–183.
- Kopell, N. (2000). We got rhythm: Dynamical systems of the nervous system. *Notices AMS* (January): 6–16.
- Kubow, T. M., and Full, R. J. (1999). The role of the mechanical system in control: A hypothesis of self-stabilization in hexapedal runners. *Philos. Trans. R. Soc. London B Biol. Sci.* 354: 849–862.
- Kubow, T., Schwind, W., Full, R., Garcia, M., and Koditschek, D. E. (In preparation). Posture selection in a running cockroach.

- Kuo, A. D. (In preparation). The relative roles of feedforward and feedback in control of rhythmic movements.
- Leob, J., and Brown, T. G. (2000). A reductionist approach to creating and using neuromusculoskeletal models. In J. M. Winters and P. E. Crago (eds.), *Biomechanics and Neural Control of Posture and Movement*. New York: Springer-Verlang.
- Lewis, M. A., Etienne-Cummings, R., Cohen, A. H., and Hartmann, M. (2000). Toward biomorphic control using custom VLSI chips. Paper presented at the International Conference on Robotics and Automation, San Francisco, April.
- Nakanishi, J., Fukuda, T., and Koditschek, D. E. (2000). A brachiating robot controller. *IEEE Trans. Robotics Automation* 16(2): 109–123.
- Orlovsky, G. N., Deliagina, T. G., and Grillner, S. (1999). *Neural Control of Locomotion*. New York: Oxford University Press.
- Pearson, K. (1976). The control of walking. *Sci. Am.* 72–86.
- Pratt, J., Dilworth, P., and Pratt, G. (1997). Virtual model control of a bipedal walking robot. Paper presented at the International Conference on Robotics and Automation (1997): 193–198.
- Raibert, M. H. (1986). *Legged Robots That Balance*. Cambridge: MIT Press.
- Raibert, M. H., Brown, H. B., and Chepponis, M. (1984). Experiments in balance with a 3D one-legged hopping machine. *Int. J. Robotics Res.* 3: 75–92.
- Ringrose, R. (1997). Self-stabilizing running. Paper presented at the International Conference on Robotics and Automation, Albuquerque, N. Mex.
- Ritzmann, R. E., Quinn, R. D., Watson, J. T., and Zil, S. N. (2000). Insect walking and biorobotics: A relationship with mutual benefits. *BioScience* 50(1): 23–33.
- Rizzi, A. A., and Koditschek, D. E. (1996). An active visual estimator for dexterous manipulation. *IEEE Trans. Robotics Automation* 12(5): 697–713.
- Rizzi, A. A., Whitcomb, L. L., and Koditschek, D. E. (1992). Distributed real-time control of a spatial robot juggler. *Computer* 12–24.
- Saranli, U., Schwind, W. J., and Koditschek, D. E. (1998). Toward the control of a multi-jointed, monopod runner. Paper presented at the International Conference on Robotics and Automation, Leuven, Belgium.
- Swanson, P. J., Burrige, R. R., and Koditschek, D. E. (1995). Global asymptotic stability of a passive juggling strategy: A possible parts feeding method. *Math. Probl. Eng.* 3(1).
- Vogel, S. (1998). *Cats' Paws and Catapults: Mechanical Worlds of Nature and People*. New York: Norton.
- Wilson, D. M. (1996). Insect Walking. *Annu. Rev. Entomol.* 11: 103–122.

# 19

## A Decentralized, Biologically Based Network for Autonomous Control of (Hexapod) Walking

Holk Cruse, J. Dean, V. Dürri, Th. Kindermann, J. Schmitz, and M. Schumm

All animals, even so-called simple animals, are autonomous systems. They find their food without external help. They can cope with unpredictable environments. They surmount various types of obstacles and they make appropriate decisions under constrained situations. At any given time, animals may have to select between different goals, such as feeding or mating. In sharp contrast to all existing animal-like robots, so-called animats, real animals survive without any external controller. Therefore, a promising strategy for constructing autonomous robots is to study how animals control their behavior. Insects and crustaceans in particular provide suitable models because their behavior is moderately complex and their segmented body plan lends itself to mechanical analysis.

Autonomy is the ability of an agent to decide on its own what to do next—that is, to select on its own among different behavioral actions. In other words, an autonomous system can decide the intensity and the temporal order for activating the modules that are responsible for the control of different behaviors or behavioral elements. This decision process cannot simply be attributed to a central executive, a “homunculus,” because this would immediately raise the question of how this homunculus would make its decision. Therefore, a crucial problem in controlling an autonomous system is the question of how different modules can be coordinated in such a way that decisions arise by some self-organizing process. This coordination may be organized via simple chain reflexes—that is, via loops through the physical environment. In this case, the behavior is strictly determined in a reactive way by the direct, physical action of the world on the body of the animal and by its indirect influences via sensory inputs. The alternative at the other extreme is to control the coordination of these modules by a hardwired central system in the sense of fixed action patterns and central pattern generators organized in a hierarchical structure. An intermediate solution is to apply architectures that use some kind of “soft” coupling between modules. In this case, coupling strength may depend on sensory inputs and it may also depend on specific internal states—often


described as motivational states or, in higher animals, emotions. In such systems, the response to a specific external stimulus may differ depending on the current internal state.

Control of multilegged walking is a good model system for studying biological solutions to these coordination problems. A walking system consists of a number of different and to some extent independent modules that nevertheless have to be coupled to cooperate in a sensible way. Functional modules performing distinct tasks can be identified on different levels of integration. Ordered from more global to more local, levels to be considered here include: movement of the whole body (e.g., control of speed and of walking direction); coordination among the legs (gait control); control of different actions by a single leg (e.g., stance phase, swing phase, avoidance reflexes, or searching for ground contact); the control of the movement of the joints of a single leg performing a given action; and control of synergistic and antagonistic muscles affecting movement of a single joint. (The list could be extended to the coordination of motor units making up an individual muscle and further to the activation of individual motor units, but these issues will not be addressed in this chapter.)

On all these levels, biological research provides examples of mutual competition or cooperation among modules that is based on coupling in the form of information transmitted in both directions between modules. Another control solution is hierarchical, allowing information from other higher-level modules in a unidirectional influence to affect lower-level modules—for example, by influencing the gain of specific connections within them. The activity of these hierarchically superior modules may be considered to represent internal states. These internal states also depend on external influences, but these changes generally occur on a slower time scale than the control of movement. Coupling influences occur via neuronal connections within the central nervous system (CNS) or via the physical environment.

Therefore, control of walking is an interesting paradigm for the investigation of how to control behavior in general. Moreover, the specific ability to walk in unpredictable terrain is interesting in its own right, as it could be utilized in many applied tasks. Examples might be moving in cluttered and unpredictable terrain such as that found in underwater situations, in forestry, or in environments that have been destroyed by catastrophes.

In the following, we will concentrate on hexapod—that is, insectlike—walking. The pattern of leg movement in insect walking is usually described as tripod or tetrapod gait (figure 19.1). These terms may suggest a rigid central control structure. However, both gaits should rather be considered as extremes of a continuum (e.g., Graham, 1985). Actually, very different step patterns can be observed; for example, after a brief disturbance of the movement of a single leg or when ani-



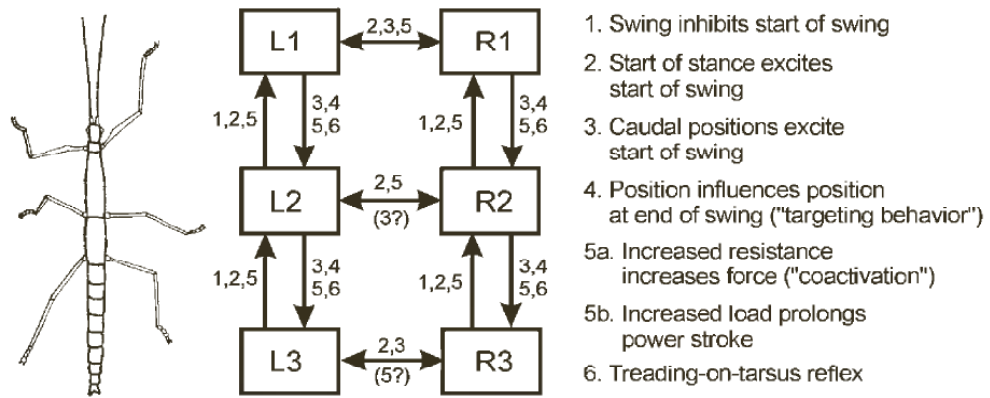
**Figure 19.1** The step patterns of a tripod (above) and a tetrapod (below) gait as produced by a stick insect. The latter is also referred to as a wave gait. The six traces represent the six legs. Black bars correspond to swing movement. Legs are designated as left (L) or right (R) and numbered from front to rear. Left and right legs on each segment (e.g., L1 and R1) always have a phase value of approximately 0.5. The phase value of adjacent ipsilateral legs (e.g., L1 and L2) is 0.5 in the tripod gait but differs in the tetrapod gait. (After Graham 1985.)

mals start walking from different leg configurations (Graham, 1972; Dean and Wendler, 1984). Therefore, insect gaits may be better described by the term *free gait* (McGhee and Iswandhi, 1979). The usually observed tripod or tetrapod patterns represent limit cycle solutions that are only apparent in undisturbed situations (Cymbalyuk et al., 1998). For insects and crustaceans, it has been shown that a small number of local rules acting between neighboring legs suffice for the emergence of different gaits and the recovery from different disturbances. In the following section, these rules will be summarized briefly.


## LEG COORDINATION

In all, six different coupling mechanisms have been found in behavioral experiments with the stick insect. These are summarized in figures 19.2 and 19.3 (three similar, yet different rules have been found for the crayfish; for a review see Cruse, 1990). In stick insects, one mechanism (number 6 in figure 19.2) serves to correct errors in leg placement; another (5) has to do with distributing propulsive force among the legs. These will not be considered here. The other four are used in the present model. The beginning of a swing movement, and therefore the end point of a stance movement (PEP, or posterior extreme position) is modulated by three mechanisms arising from ipsilateral legs (figure





**Figure 19.2** Summary of the coordination mechanisms operating between the legs of a stick insect. The leg controllers are labelled R and L for right and left legs and numbered from 1 to 3 for front, middle, and hind legs. The different mechanisms (1 to 6) are explained in the text.



**Figure 19.3** Illustrations of the mechanisms 1 to 3 (see figure 19.2) as shown from above to below.

19.3): (1) a rostrally directed inhibition during the swing movement of the next caudal leg, (2) a rostrally directed excitation when the next caudal leg begins active retraction, and (3) a caudally directed influence depending upon the position of the next rostral leg. Influences 2 and 3 are also active between contralateral legs. The end of the swing movement (AEP, or anterior extreme position) in the animal is modulated by a single, caudally directed influence (4) depending on the position of the next rostral leg. This mechanism is responsible for the targeting behavior—the placement of the tarsus at the end of a swing close to the tarsus of the adjacent rostral leg.

These interleg influences are mediated in two parallel ways. The first pathway comprises the direct neural connections between the step pattern generators as illustrated in figure 19.2. The second pathway arises from the mechanical coupling among the legs. That is, the activity of one step pattern generator influences the movements and loading of all legs and therefore influences the activity of their step pattern generators via sensory pathways. This combination of mechanisms adds redundancy and robustness to the control system of the stick insect.

Computer simulations (stick insect: Dean 1991a, b, 1992a, b; Cruse et al., 1998; crayfish: Cruse and Müller, 1986) have shown that these rules suffice to describe the spatiotemporal leg pattern observed in different walking situations including reactions to disturbances. For the stick insect, a detailed study has been performed that is based on behavioral investigations and which uses artificial neural networks as a basis. In the following, the results will be described, using this simulation as a basis. As will be illustrated below, the simulated system is able to walk on uneven surfaces, to negotiate curves, to walk after partial amputation of legs, and to recover from a fall.

Investigation of these higher-level coupling rules is only possible after a decision has been made concerning the lower level—that is, the movement of the single leg. How can this be done? A walking leg can be regarded as being in one of two states—namely, performing a swing movement or a stance movement. During stance, the leg is on the ground, supports the body and, in the forward-walking animal, moves backward with respect to the body. During swing, the leg is lifted off the ground and moved in the direction of walking to where it can begin a new stance. The anterior transition point (i.e., the transition from swing to stance in the forward-walking animal) is called the anterior extreme position (AEP) and the posterior transition point is called the posterior extreme position (PEP). Differences in the constraints acting during the two states and in the conditions for their termination suggest that the leg controller consists of three separate control networks. The two *microbehaviors*, swing and stance, are mutually exclusive. A leg cannot be in swing and in stance at the same time, a situation also holding for many *macrobehaviors* such as fight or flight,

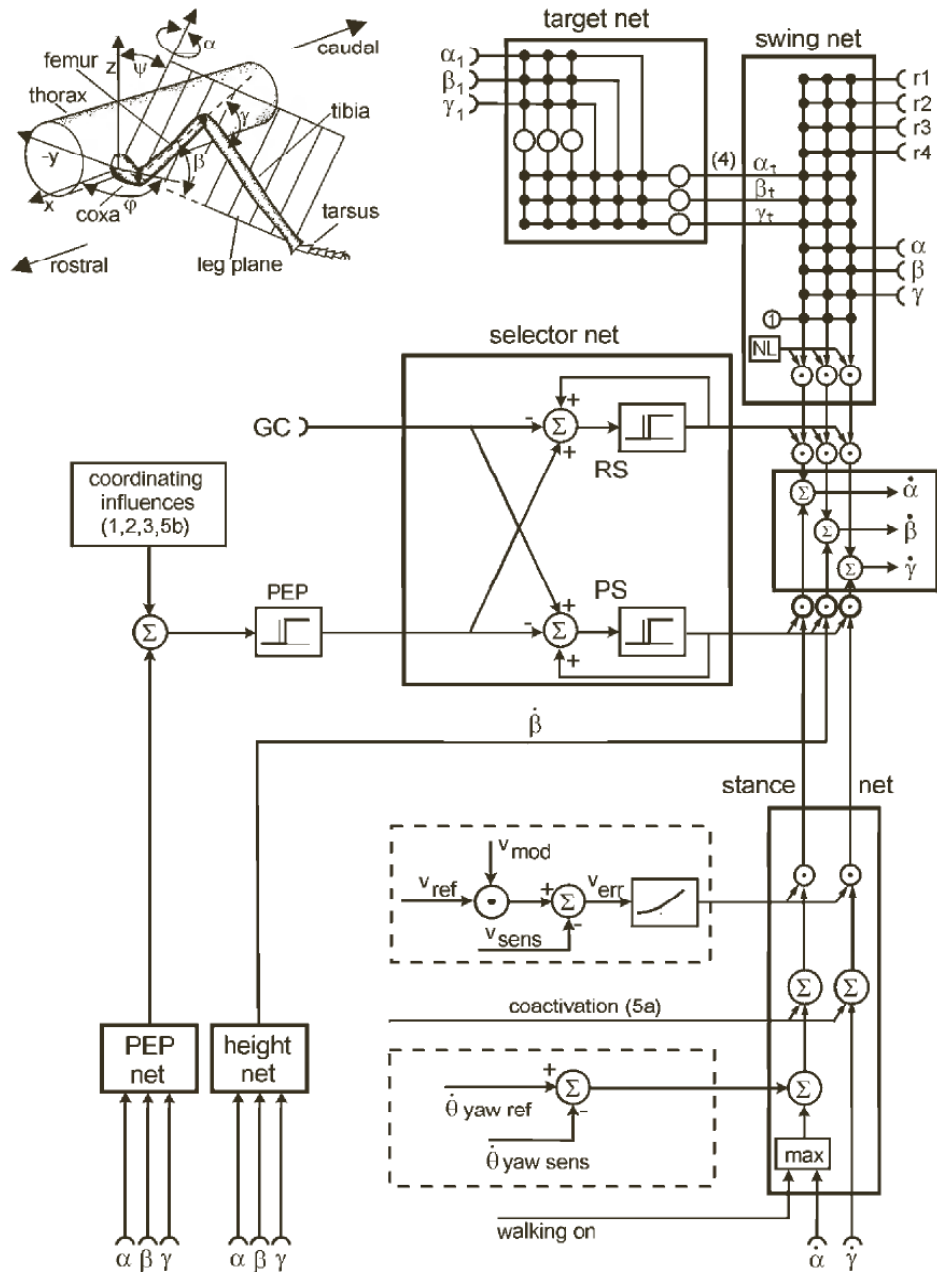
for instance. Therefore, in our simulation network called Walknet, two low-level networks, a swing network and a stance network, control the movement of the leg during swing and stance, respectively (figure 19.4b). The transition between swing and stance is controlled by a selector network (figure 19.4b). (This separation may only be justified on a logical level; it need not correspond to entirely distinct circuits in the nervous system.) The swing network and the stance network are always active, but the selector network determines which of the two networks controls the motor output. To date, there is little evidence for a strong central pattern generator and robust endogenous rhythms for walking in stick insects (Bässler and Wegner, 1983; Büschges, Schmitz, and Bässler, 1995). Therefore, the simulation does not contain a central oscillator. The quasi-rhythmic movement of the leg is based on sensory feedback.

In the selector net, positive feedback from each of the two output units stabilizes the activity of whichever output unit is active. Activity in a sensory input causes a transition by inhibiting the active output unit and exciting the inactive one. The sensory input is either ground contact or the position of the leg tip. The latter depends in a nonlinear way on the position of the leg joints. The results obtained in behavioral experiments have been implemented in a two-layer feedforward net (PEP net), which contains three input units for the three joint angles, two hidden units, and one output unit determining the distance to the PEP threshold value.

Control of the leg movement requires controlling the movement of the different leg joints (figure 19.4a). The three major joints are the thoracic-coxal joint, the coxa-trochanter joint, and the femur-tibia joint; these will be referred to by  $\alpha$ ,  $\beta$ , and  $\gamma$  joint in the following, respectively. As the task of the leg is quite different during stance and during swing, each problem will be treated separately.

## CONTROL OF STANCE MOVEMENT

The task of controlling the stance movements of all the legs on the ground poses several major problems. It is not enough simply to specify a movement for each leg on its own: The mechanical coupling through the substrate means that efficient locomotion requires coordinated movement of all the joints of all the legs in contact with the substrate—that is, a total of 18 joints when all legs of an insect are on the ground. However, the number and combination of mechanically coupled joints varies from one moment to the next, depending on which legs are lifted. The task is quite nonlinear, particularly when the rotational axes of the joints are not orthogonal, as is often the case for insect legs, and for the basal leg joint in particular. A further complication occurs when the animal negotiates a curve, which requires the different legs to move at different speeds.



**Figure 19.4** Summary of leg geometry and control networks for the model. (a) Schematic model of a stick insect leg showing the arrangement of the joints and their axes of rotation. (b) The leg controller consists of three parts: the swing net, the stance net, and the selector net which determines whether the swing or the stance net can control the motor output, i.e., the velocity of the three joints  $\alpha$ ,  $\beta$ , and  $\gamma$ . The selector net contains four units: the PEP unit signaling posterior extreme position, the GC unit signaling ground contact, the RS unit controlling the return stroke (swing movement), and the PS unit controlling the power stroke (stance movement). The target net transforms information on the configuration of the anterior, target leg,  $\alpha_1$ ,  $\beta_1$ , and  $\gamma_1$ , into angular values for the next caudal leg which place the two tarsi close together. These desired final values ( $\alpha_t$ ,  $\beta_t$ ,  $\gamma_t$ ) and the current values ( $\alpha$ ,  $\beta$ , and  $\gamma$ ) of the leg angles are input to the swing net together with a bias input (1) and four sensory inputs ( $r1$ – $r4$ ) which are activated by obstructions blocking the swing and thereby initiate different avoidance movements. A nonlinear influence (NL) modulates the velocity profile. For details see Cruse et al. (1998).

In machines, these problems can be solved using traditional, though computationally costly, methods, which consider the ground reaction forces of all legs in stance and seek to optimize some additional criteria, such as minimizing the tension or compression exerted by the legs on the substrate. Due to the nature of the mechanical interactions inherent in the search for a globally optimal control strategy, such algorithms require a single, central controller; they do not lend themselves to distributed processing. This makes real-time control difficult, even in the still-simple case of walking on a rigid substrate.

Further complexities arise in more complex, natural walking situations, making a solution difficult even with high computational power. These occur, for example, when an animal or a machine walks on a slippery surface or on a compliant substrate, such as the leaves and twigs encountered by stick insects. Any flexibility in the suspension of the joints further increases the degrees of freedom that must be considered and the complexity of the computation. Further problems for an exact, analytical solution occur when the length of leg segments changes during growth or their shape changes through injury. In such cases, knowledge of the geometrical situation is incomplete, making an explicit calculation difficult, if not impossible.

Despite the evident complexity of these tasks, they are mastered even by insects with their “simple” nervous systems. Hence, there has to be a solution that is fast enough that on-line computation is possible even for slow neuronal systems. How can this be done? Several authors (e.g., Brooks, 1991, 1997) have pointed out that some relevant parameters do not need to be explicitly calculated by the nervous system because they are already available in the interaction with the environment. This means that, instead of an abstract calculation, the system can directly exploit the dynamics of the interaction and thereby avoid a slow, computationally exact algorithm. To solve the particular problem at hand, we propose to replace a central controller with distributed control in the form of local positive feedback (Cruse et al., 1996). Compared to earlier versions (Cruse et al., 1995), this change permits the stance net to be radically simplified. The positive feedback occurs at the level of single joints: The position signal of each is fed back to control the motor output of the same joint (figure 19.4b, stance net). How does this system work? Let us assume that any one joint is moved actively. Then, because of the mechanical connections, all other joints begin to move passively, but in exactly the proper way. Thus, the movement direction and speed of each joint does not have to be computed because this information is already provided by the physics. The positive feedback then transforms this passive movement into an active movement.

There are, however, several problems to be solved. The first is that positive feedback using the raw position signal would lead to unpredictable changes in movement speed, not the nearly constant walking

speed that is usually desired. This problem can be solved by introducing a kind of band-pass filter into the feedback loop. The effect is to make the feedback proportional to the angular velocity of joint movement, not the angular position. In the simulation, this is done by feeding back a signal proportional to the angular change over the preceding time interval.

The second problem is that using positive feedback for all three leg joints leads to unpredictable changes in body height, even in a computer simulation neglecting gravity. Body height of the stick insect is controlled by a distributed system in which each leg acts like an independent, proportional controller. However, maintaining a given height via negative feedback appears at odds with the proposed local positive feedback for forward movement. How can both functions be fulfilled at the same time? To solve this problem we assume that during walking, positive feedback is provided for the  $\alpha$  joints and the  $\gamma$  joints (figure 19.4b, stance net), but not for the  $\beta$  joints. The  $\beta$  joint is the major determinant of the separation between leg insertion and substrate, which determines body height. The value for the  $\beta$  joint is given by a three-layered feedforward network (height net) with three input units ( $\alpha$ ,  $\beta$ ,  $\gamma$ ), five hidden units, and one output unit. This net has been trained using the known leg geometry and approximates data from Cruse, Riemenschneider, and Stammer (1989), where force-height characteristics of the standing animal have been measured.

A third problem inherent in using positive feedback is the following. Let us assume that a stationary insect is pulled backward by gravity or by a brief tug from an experimenter. With positive feedback control as described, the insect should then continue to walk backward even after the initial pull ends. This has never been observed. Therefore, we assume that a supervisory system exists that is not only responsible for switching on and off the entire walking system, but also specifies walking direction (normally forward for the insect). This influence is represented by applying a small, positive input value (figure 19.4b, walking on) which replaces the sensory signal if it is larger than the latter (the box "max" in figure 19.4b, stance net).

To permit the system to control straight walking and to negotiate curves, a supervisory system was introduced that, in a simple way, simulates optomotor mechanisms for course stabilization that are well known from insects and have also been applied in robotics. This supervisory system uses information on the rate of yaw ( $yaw_{sens}$ , figure 19.4b, stance net), such as visual movement detectors might provide. It is based on negative feedback of the deviation between the desired turning rate and the actual change in heading over the last time step. The error signal controls additional impulses to the  $\alpha$  joints of the legs, these additional impulses have magnitudes proportional to the deviation and opposite signs for the right and left sides. With this addition

and  $yaw_{ref}$  set to zero, the system moves straight (see below, figure 19.7a) with small, side-to-side oscillations in heading such as can be observed in walking insects. To simulate curve walking, the reference value is given a small positive or negative bias to determine curvature direction and magnitude.


Finally, we have to address the question of how walking speed is determined in such a positive feedback controller. Again, we assume a central value that represents the desired walking speed  $v_{ref}$ . This is compared with the actual speed, which could be measured by visual inputs or by monitoring leg movement. This error signal is subject to a nonlinear transformation and then multiplied with the signals providing the positive feedback for all  $\alpha$  and  $\gamma$  joints of all six legs (figure 19.4b, stance net).

## CONTROL OF THE SWING MOVEMENT

The task of finding a network that produces a swing movement seems to be simpler than finding a network to control the stance movement, because a leg in swing is mechanically uncoupled from the environment and therefore, due to its small mass, essentially uncoupled from the movement of the other legs. Nevertheless, the network for the control of the swing movement, though still simple, is more complicated than that for the stance. The geometry of the leg is shown in figure 19.4a. The coxa-trochanter and femur-tibia joints—the two distal joints—are simple hinge joints with one degree of freedom corresponding to elevation and extension of the tarsus, respectively. The subcoxal joint is more complex, but most of its movement is in a rostrocaudal direction around the nearly vertical axis. The additional degree of freedom allowing changes in the alignment of this axis is little used in normal walking, so the leg can be considered as a manipulator with three degrees of freedom for movement in three dimensions. Thus, the control network must have at least three output channels, one for each leg joint.

A simple, two-layer feedforward net with three output units and six input units can produce movements that closely resemble the swing movements observed in walking stick insects (Cruse and Bartling, 1995). The inputs correspond to three coordinates defining the actual leg configuration and three defining the target—the configuration desired at the end of the swing. In the simulation, the three outputs, interpreted as the angular velocities of the joints,  $d\alpha/dt$ ,  $d\beta/dt$ , and  $d\gamma/dt$ , are used to control the joints. The actual angles are measured and fed back into the net.

Through optimization, the network can be simplified to only 8 (front and middle leg) or 9 (hind leg) nonzero weights (for details see Cruse et al., 1996). We believe this represents the simplest possible network for the task; it can be used as a standard of comparison with physiological



**Figure 19.5** Vector field representing the movement of the tarsus of a left front leg produced by the swing net. Projection of a parasagittal section ( $y = 12$  mm; for coordinates see figure 19.4a). Left is posterior, right is anterior. The average posterior extreme position (start of swing movement) and of the average anterior extreme position (end of swing movement) are shown by an open square and by a closed square, respectively.

results from stick insects. Despite its simplicity, the net not only reproduces the trained trajectories, it is able to generalize over a considerable range of untrained situations, demonstrating a further advantage of the network approach. Moreover, the swing net is remarkably tolerant with respect to external disturbances. The learned trajectories create a kind of attractor to which the disturbed trajectory returns. This compensation for disturbances occurs because the system does not compute explicit trajectories, but simply exploits the physical properties of the world. The properties of this swing net can be described by the 3-D vector field in which the vectors show the movement produced by the swing net at each tarsus position in the work space of the leg. Figure 19.5 shows the planar projections of one parasagittal section through the work space. The complete fields are similar to those shown by Giszter (chapter 11 of this volume) for the frog.

This ability to compensate for external disturbances permits a simple extension of the swing net in order to simulate an avoidance behavior observed in insects. When a leg strikes an obstacle during its swing, it initially attempts to avoid it by retracting and elevating briefly and then renewing its forward swing from this new position. In the augmented swing net, an additional input similar to a tactile or force sensor signals such mechanical disturbances at the front part of the tibia (figure 19.4b, r1) or the femur (figure 19.4b, r2). These units are connected by fixed weights to the three motor units in such a way as to produce the brief retraction and elevation seen in the avoidance reflex. Other reflexes can be observed when the tibia is mechanically stimulated laterally (r3) or when the femur is touched dorsally (r4). These reflexes have been implemented in an analogous manner (figure 19.4b).



In the model, the targeting influence reaches the leg controller as part of the input to the swing net (figure 19.4b). These signals can be generated by a simple feedforward net, with three hidden units and logistic activation functions (figure 19.4b, “target net”), which directly associates desired final joint angles for the swing to current joint angles of a rostral leg such that the tarsus of the posterior leg is moved in the same direction as that of the anterior leg. Compared to a first version (Dean, 1990) the new target net has direct connection between the input and the output layer. There is no explicit calculation of either tarsus position. Physiological recordings from local and intersegmental interneurons (Brunn and Dean, 1994) support the hypothesis that a similar approximate algorithm is implemented in the nervous system of the stick insect.

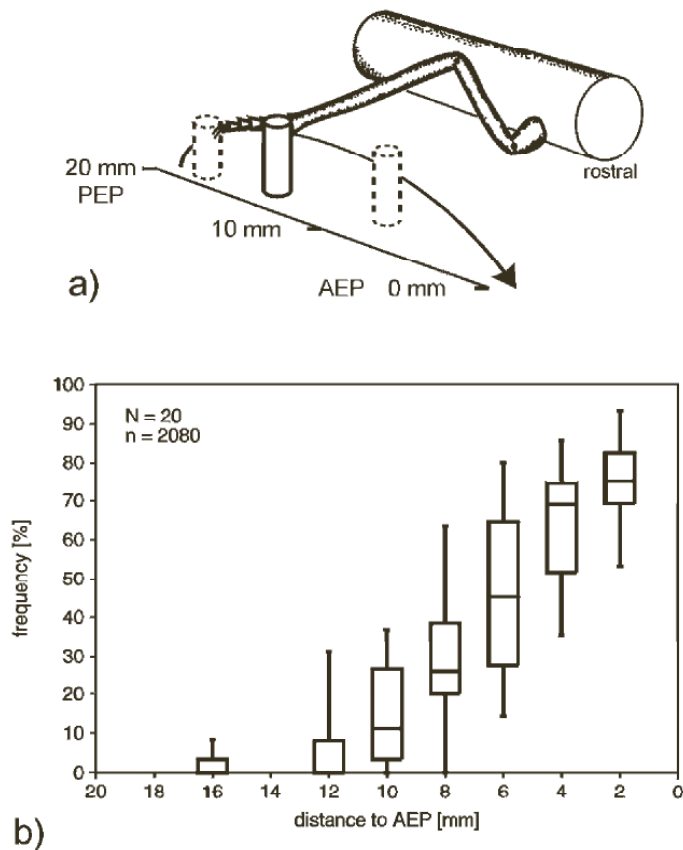
In Walknet, it is assumed that the end of swing is determined by a sensory stimulus affecting the tarsal mechanoreceptors (and possibly load receptors, too). As soon as this stimulus is above a given threshold (figure 19.4b, GC), the selector net is forced to switch from swing to stance. In the animals, however, the situation is not as simple. This should not be surprising, because any arbitrary mechanical contact in the early swing, for example, should be considered an obstacle rather than a stimulus to finish swing. So one might assume that there is some internal state, an *expectation value* or conversely, a value for *swing motivation*, which tells the system whether a given stimulus should be treated as an obstacle or as ground contact.

How do the animals behave? When a possible substrate—for example, a round wooden rod—is held at different positions in the swing trajectory of a leg, the probability it will be grasped by the tarsus to end swing is very low when the rod is in the caudal 40% of the step range, but the probability increases about linearly to 100% as the position approaches that of the normal AEP (figure 19.6). This raises the question of how this decreasing swing motivation is produced. Experimental results showed that it is not simply the position of the leg.

These results also show that another possible parameter, namely, time elapsed since the beginning of the swing, cannot be the only relevant parameter. The results suggest that it is the distance between front and middle leg that determines the “motivation” for the middle leg to accept a stimulus as ground contact, a hypothesis that can quantitatively describe all our data. Thus, this *motivational state* can be traced back to a geometrical parameter.

## RESULTS

As a first step (i.e., before the planned implementation in a real robot), the behavior of Walknet was tested quantitatively in a software simulation. Because the control principle relies on the existence of a body,




**Figure 19.6** The dependency of the frequency of the grasping reaction of the middle leg tarsus on the test substrate's position relative to its normal AEP (b, distance in mm) as illustrated in (a).


the mechanics of the body had to be simulated too. For this purpose, a recurrent network, called MMC, has been developed (see Steinkühler and Cruse, 1998), which will not be explained here.

Like the earlier version (Cruse et al., 1995), Walknet shows proper coordination of the legs for straight and curved walks at different speeds on a horizontal plane (figure 19.7a) (for animation see <http://www.uni-bielefeld.de/biologie/Kybernetik>). Steps of ipsilateral legs are organized in triplets forming *metachronal waves*, which proceed from back to front, whereas steps of the contralateral legs on each segment step approximately in alternation. With increasing walking speed, the typical change in coordination from this tetrapod to a tripodlike gait is found: Front and rear legs on each side step together with the contralateral middle leg. Step amplitude is constant (approximately 20 mm) for all speeds. In relation to the distance between front and middle leg coxa (16 mm) and that between middle and hind leg coxa (11 mm), these step amplitudes are considerably larger than in most robots.

It has been argued that an important advantage of the positive feedback principle is that it can cope with changes in body geometry such




**Figure 19.7** Simulated walk by the basic six-legged system with negative feedback applied to all six  $\beta$  joints and positive feedback to all  $\alpha$  and  $\gamma$  joints as shown in figure 19.4b. Movement direction is from left to right (arrow). Leg positions are illustrated only during stance and only for every second time interval in the simulation. Each leg makes about five steps. Upper part: top view; lower part: side view. (a) Straight walking ( $\theta_{yaw\ ref} = 0$ ). (b) Curved walking ( $\theta_{yaw\ ref} \neq 0$ ).



**Figure 19.8** Simulated walk over an obstacle. Movement direction is from left to right. Leg positions, as viewed from the side, are illustrated only during stance and only for every fifth time interval in the simulation. (Top) The first part of the walk until both front legs reach the top of the obstacle. (Bottom) Descent from the obstacle until both front legs and one middle leg touch the lower ground.

as sudden changes of leg length, for example. This could be shown in simulation by cutting the tibia of the front, middle, or hind leg. For these tests, stability was considered during about 10 steps following a test “amputation” performed during walking. In a similar manner, walking over obstacles of different heights was investigated (figure 19.8).

Finally, curve walking has been investigated and compared to results obtained from stick insects. A kind of minimal model presented in an earlier simulation (Cruse et al., 1996) showed that curve walking was possible when the input from the curve controller was given only to the  $\alpha$  joints of both front legs. Detailed investigation, however, showed that the behavior of the simulation closely resembles that of the stick insect when the  $\alpha$  joints of all three leg pairs are used (figure 19.7b). With respect to the movement of the individual leg, the changes in the AEP



**Figure 19.9** Righting behavior. (a) By clamping the tarsi to the ground (arrowheads), the system is made to fall leading to disordered arrangement of the legs (b). Nevertheless, the system stands up without help and resumes proper walking (c).

and PEP, which may or may not change step amplitude, and the form and speed of leg movement are of interest. Comparison of leg speed and step amplitude between these experimental results and the behavior of the simulation shows good agreement in the range investigated. Outer legs move faster and have a greater step amplitude than inner legs. Despite the larger step amplitudes, the stance duration of the outer legs in the model is shortened compared to straight walking, but this shortening is fairly well compensated by a correspondingly longer swing duration. This means that changes in the temporal parameters induced by curve walking are much smaller than changes in geometric parameters like step amplitude.

Unexpectedly, the following interesting behavior was observed. A massive perturbation, for example by clamping the tarsi of three legs to the ground, can make the system fall. Although this can lead to extremely disordered arrangements of the six legs, the system was able to stand up on its own and resume proper walking (figure 19.9). This means that the simple solution proposed here also eliminates the need for a special supervisory system to rearrange leg position after such an emergency.

## DISCUSSION

To control the individual leg, no explicit central controllers in the sense of central oscillators are used. Rather, the oscillations result from the cooperation of different subsystems and the physical world. Using a central controller might even worsen performance if such an autonomous central controller sometimes favors actions conflicting with those favored by the present control model relying directly on the physical properties of the system. Our results are sensible for slow walkers, where slow neuronal conduction speeds do not play a critical role. For fast walkers, which cannot rely on temporally adequate sensory input, central oscillators might be the better solution. On the other hand,

Jindrich and Full (1999) show that in a fast walker, the mechanical—in particular, the dynamical properties of the system—may well be exploited to substantially decrease the control effort (chapter 18 of this volume). This may also be true in slow walking systems where the dynamics are important, for example, in larger animals or in humans walking in a predictable environment.

One subsystem, the swing net, consists of an extremely simple feed-forward net. The swing net exploits the recurrent information via the sensorimotor system as the leg's behavior is influenced by the environment to generate swing movements in time. The system does not compute explicit trajectories; instead, it computes changes in time based on the actual and the desired final joint angles. More difficult problems have to be solved by the stance system, the second subsystem of the leg controller, particularly if the geometry of the system may change due to growth, injury, or nonrigid suspension of the joints. Solving these complex tasks represents quite a high level of "motor intelligence." However, this does not require a complicated, or a centralized control system. Due to its extremely decentralized organization and the simple type of calculation required, the solution using local positive feedback permits very fast computation. This simplification is possible because the physical properties of the system and its interaction with the world are exploited to replace an abstract, explicit computation. Thus, "the world" is used as "its own best model" (Brooks, 1991). Our solution also provides an example of how information about the environment may not be available to a passive system waiting for sensory input, but it is available to an active system.

Some central commands from a superior level are still required. These are necessary to determine the beginning and end of walking as well as its speed and direction. However, these commands do not have to be precisely adjusted to the particular configurations and states of all the legs; an approximate specification is sufficient for the positive feedback control. For example, turning does not require explicit calculation of leg trajectories using, for example, the desired curve radius and the geometrical parameters of the legs. Simple proportional commands to the legs are sufficient.

One major disadvantage of our simulation is its purely kinematic nature. To test the principle of local positive feedback, at least for straight walking, we have performed a dynamic simulation for the six-legged system under positive feedback control during stance. The basic software was kindly provided by F. Pfeiffer, TU Munich. No problems occurred. Nevertheless, a hardware test of the walking situations is necessary and is planned with M. Frik, University of Duisburg, and his robot TARRY (Frik et al., 1999).

Another problem with the idea of positive feedback is the following. Although the simulation has shown positive feedback to be an elegant

solution for several problems, and although there are biological experiments showing the existence of positive feedback (Bässler, 1976; Schmitz et al., 1995; Bässler and Büschges, 1998), conflicting experimental results point to the realization of negative feedback in some walking situations (e.g., Bartling and Schmitz, 2000). To date, it is not clear how this apparent contradiction may be solved.

Further work is also necessary with respect to the control of the swing movement. Qualitative observations show that at least one property of the swing net appears not to agree with that of real legs performing a swing: In the swing net, the step height decreases about linearly with step length. In the animals, however, step height appears to be more or less independent of step length (Schumm, unpublished observation). Therefore, the swing net has to be adapted to describe these behavioral details.

## REFERENCES

- Bartling, C., and Schmitz, J. (2000). Reaction to disturbances of a walking leg during stance. *J. Exp. Biol.* 203: 1211–1233.
- Bässler, U. (1976). Reversal of a reflex to a single motoneuron in the stick insect *Carausius morosus*. *Biol. Cybern.* 24: 47–49.
- Bässler, U., and Büschges, A. (1998). Pattern generation for stick insect walking movements—Multisensory control of a locomotor program. *Brain Res. Rev.* 27: 65–88.
- Bässler, U., and Wegner, U. (1983). Motor output of the denervated thoracic ventral nerve cord in the stick insect *Carausius morosus*. *J. Exp. Biol.* 105: 127–145.
- Brooks, R. A. (1991). Intelligence without reason. In IJCAI-91, Sydney, Australia. pp. 569–595.
- Brooks, R. A. (1997). From earwigs to humans. *Robotics Autonomous Syst.* 20: 291–304.
- Brunn, D., and Dean, J. (1994). Intersegmental and local interneurons in the metathorax of the stick insect, *Carausius morosus*. *J. Neurophysiol.* 72: 1208–1219.
- Büschges, A., Schmitz, J., and Bässler, U. (1995). Rhythmic patterns in the thoracic nerve cord of the stick insect induced by pilocarpine. *J. Exp. Biol.* 198: 435–456.
- Cruse, H. (1990). What mechanisms coordinate leg movement in walking arthropods? *Trends Neurosci.* 13: 15–21.
- Cruse, H., and Bartling, C. (1995). Movement of joint angles in the legs of a walking insect, *Carausius morosus*. *J. Insect Physiol.* 41: 761–771.
- Cruse, H., Bartling, C., Brunn, D. E., Dean, J., Dreifert, M., Kindermann, T., and Schmitz, J. (1995). Walking: A complex behavior controlled by simple systems. *Adaptive Behav.* 3: 385–418.
- Cruse, H., Bartling, C., Dean, J., Kindermann, T., Schmitz, J., Schumm, M., and Wagner, II. (1996). Coordination in a six-legged walking system: Simple solutions to complex problems by exploitation of physical properties. In P. Maes, M. J. Mataric, J.-A. Meyer, J. Pollack, and S. W. Wilson (eds.), *From Animals to Animats 4*. Cambridge: MIT Press, pp. 84–93.

- Cruse, H., Kindermann, T., Schumm, M., Dean, J., and Schmitz, J. (1998). Walknet—A biologically inspired network to control six-legged walking. *Neural Networks* 11: 1435–1447.
- Cruse, H., and Müller, U. (1986). Two coupling mechanisms which determine the coordination of ipsilateral legs in the walking crayfish. *J. Exp. Biol.* 121: 349–369.
- Cruse, H., Riemenschneider, D., and Stammer, W. (1989). Control of body position of a stick insect standing on uneven surfaces. *Biol. Cybern.* 61: 71–77.
- Cymbalyuk, G. S., Borisyuk, R. M., Müller-Wilm, U., and Cruse, H. (1998). Oscillatory networks controlling six-legged locomotion. Optimization of model's parameters. *Neural Networks* 11: 1449–1460.
- Dean, J. (1990). Coding proprioceptive information to control movement to a target: Simulation with a simple neural network. *Biol. Cybern.* 63: 115–120.
- Dean, J. (1991a). A model of leg coordination in the stick insect, *Carausius morosus*: I. A geometrical consideration of contralateral and ipsilateral coordination mechanisms between two adjacent legs. *Biol. Cybern.* 64: 393–402.
- Dean, J. (1991b). A model of leg coordination in the stick insect, *Carausius morosus*: II. Description of the kinematic model and simulation of normal step patterns. *Biol. Cybern.* 64: 403–411.
- Dean, J. (1992a). A model of leg coordination in the stick insect, *Carausius morosus*: III. Responses to perturbations of normal coordination. *Biol. Cybern.* 66: 335–343.
- Dean, J. (1992b). A model of leg coordination in the stick insect, *Carausius morosus*: IV. Comparisons of different forms of coordinating mechanism. *Biol. Cybern.* 66: 345–355.
- Dean, J., and Wendler, G. (1984). Stick insect locomotion on a wheel: Patterns of stopping and starting. *J. Exp. Biol.* 110: 203–216.
- Frik, M., Guddat, M., Karatas, M., and Losch, C. D. (1999). A novel approach to autonomous control of walking machines. In G. S. Virk, M. Randall, and D. Howard (eds.), *Proceedings of the Second International Conference on Climbing and Walking Robots CLAWAR '99*, 13–15 September, Portsmouth, UK. Bury St. Edmunds: Professional Engineering Publishing, pp. 333–342.
- Graham, D. (1972). A behavioural analysis of the temporal organization of walking movements in the first instar and adult stick insect. *J. Comp. Physiol.* 81: 23–52.
- Graham, D. (1985). Pattern and control of walking in insects. *Adv. Insect Physiol.* 18: 31–140.
- Jindrich, D. L., and Full, R. J. (1999). Many-legged maneuverability: Dynamics of turning in hexapods. *J. Exp. Biol.* 202: 1603–1623
- McGee, R. B., and Iswandhi, G. I. (1979). Adaptive locomotion of a multilegged robot over rough terrain. *IEEE Trans. Syst. Man Cybern.* SMC-9(4): 176–182.
- Schmitz, J., Bartling, C., Brunn, D. E., Cruse, H., Dean, J., Kindermann, T., Schumm, M., and Wagner, H. (1995). Adaptive properties of “hard-wired” neuronal systems. *Verh. Dtsch. Zool. Ges.* 88(2): 165–179.
- Steinkühler, U., and Cruse, H. (1998). A holistic model for an internal representation to control the movement of a manipulator with redundant degrees of freedom. *Biol. Cybern.* 79: 457–466.

## Toward the Development of Agile and Mission-Capable Legged Robots

Roger D. Quinn, Gabriel M. Nelson, and  
Roy E. Ritzmann

### GOALS OF BIOROBOTICS AT CWRU

The Biorobotics group at Case Western Reserve University has two interrelated goals. First, we strive to design and implement agile and useful legged robots. Second, we attempt to further our understanding of legged locomotion in insect species. The two goals are mutually supportive. Information we gain from studying insect locomotion guides our designs of legged robots and efforts to design control systems for our robots inspire useful hypotheses for studying legged locomotion.

In our definition of an agile robot, we mean a legged device that can move rapidly over complex and unpredictable terrain. This, after all, is the promise of legged vehicles. On relatively smooth terrain, wheeled vehicles are faster and cheaper to run and require far fewer degrees of freedom than current legged devices. However, they have limitations in complex terrain. Even treaded wheeled devices have difficulty traversing natural uneven terrains with rocky barriers and varying substrate characteristics. Legged animals have little difficulty moving rapidly through such environments and it is this observation that encourages us to develop legged vehicles. Nevertheless, it is clear that legged devices with designs that inherently limit their agility will never supplant wheeled vehicles.

In our definition of useful vehicles, we mean the robots that we develop should ultimately be capable of carrying out a mission. This means that a robot should be able to lift a considerable payload. Even a small robot that might be used in reconnaissance must be powerful enough to carry a sufficient suite of sensors to perform that function. A useful robot should also be autonomous. In order to carry out a mission, all control and power sources must be on board. However, commands to the robot could be provided by a remote supervisor and carried out by local controls on board.

In the development stage of the robotic system, some of these characteristics can be compromised as long as it is clear that reasonable adjustments can be made later. However, some characteristics are



fundamental to an individual robotic system and alteration would result in a total redesign of the vehicle. For example, it is possible to use tethered control during development so that the control architecture can be optimized before designing onboard circuits. Indeed, given that these robots are pushing the forefront of legged vehicle design, it makes little sense to manufacture onboard control systems until they have been optimized through considerable research. Constraining power supplies to onboard systems early in development is also inappropriate. However, considerations for developing onboard systems should start from the outset, so that solutions are available when control and mechanical problems are solved. On the other hand, choice of actuators must be made carefully at the outset. The type of actuator used on a robot dictates many control issues. Once the robot is fully developed, a change in actuators mandated to generate useful forces would compromise the entire design and require the development team to essentially start from scratch. For example, a switch from electric motors to pneumatic actuators to increase power cannot be accomplished without considerable redevelopment of control systems. Thus, it would seem prudent to start development of the vehicle with the actuators that are appropriate for the robot's ultimate missions.


## STAGES OF PROGRESS

Each of the robots that were developed in our group sought to solve a specific problem leading to the development of useful and agile devices. Our first robotic effort sought to demonstrate that kinematic models of insect locomotion based upon existing literature (Beer et al., 1992) could be used to control a real device. Members of our team (Quinn and Espenschied, 1993) built a simple 12-degrees-of-freedom (DOF) vehicle (Robot I, or R-I) that used DC motors (figure 20.1). Each leg was identical and had two DOF—one revolute and one linear. Its control systems (Beer et al., 1992; Quinn and Espenschied, 1993; Espenschied et al., 1993) were inspired by biological observations on insects, and it was successful in that it could walk in a range of hexapod gaits.

Robot II (R-II) improved on the leg and control designs (figure 20.2). Although each leg was identical, they had three DOF and more animal-like revolute joints, which were driven by DC motors. Its controller was distributed and included modules for control of posture, gait, leg coordination, leg reflexes, and individual joints. Its gait controller consisted of a network of influences and coordination mechanisms described by Cruse (1990) and Dean (1992) that were based upon observations of insect and crustacean movement. R-II was also a successful robot. As with R-I, it could walk with a range of hexapod gaits. It could also traverse a series of slats using a searching reflex. Upon reaching an edge it would determine that no solid surface was present and reverse direc-



**Figure 20.1** Robot I had legs with two degrees of freedom each: one rotational and one linear.



**Figure 20.2** Robot II could walk over slatted surfaces using a searching reflex.

tion. It could also climb over barriers as high as its normal body height using an elevator reflex, and it could adjust to movement of the substrate using its stepping reflex (Espenschied et al., 1996).

Although R-II successfully addressed many issues in legged robotics, it had inherent limitations. The choice of DC motors for R-II simplified the control systems. However, the power-to-weight ratio of its gear motors, the best commercially available at the time, still severely limited the capabilities of the robot. Even with its structural mass minimized,


R-II could not lift a payload consistent with most missions. Although the robot could climb over significant objects, the leg design did not take advantage of specialization found in virtually all legged animals. Without more efficient leg designs and more powerful actuators, we felt that R-II was as good as we could do. Indeed, to date, we know of no other legged vehicle powered by DC motors that significantly surpasses the agility of R-II, although some recent designs have accomplished similar agility with a greater degree of autonomy via more-refined gear motors.

In Robot III (R-III) we sought a design that could ultimately develop enough power to carry a realistic payload. After investigating several possibilities that were currently available, we chose double-acting pneumatic cylinders. These powerful actuators had the additional benefit of allowing the structure to be made out of more-durable materials than was the case in previous robots.

In order to enhance the leg designs and increase agility, we turned to cockroach behavior for inspiration. Insects are among the most efficient legged devices in the world. However, many aspects of biological design are totally inappropriate for robotic design. For example, insect legs have 8 or more degrees of freedom. The efficient power-to-weight properties of muscle and insect cuticle makes this possible, but any attempt to blindly capture anywhere near 48 degrees of freedom in the robot's legs using artificial actuators would probably fail. The challenge to bio-robotics is to capture the useful aspects of biological design, while intelligently removing properties that are not necessary for the vehicle to capture the desired movements.

To create efficient leg design, we used high-speed video systems to measure the actions of leg joints as cockroaches ran on a treadmill or climbed over barriers. These data were ported into a custom-made dynamic simulation that allowed us to limit degrees of freedom and determine the consequences to walking and climbing. Hind legs push the animal forward with powerful extensions, but there is little use of the three degrees of freedom that are available at the body-coxa joint. In contrast, the front legs act like arms to investigate the surroundings and provide braking forces during walking. Here, five degrees of freedom are essential. Middle legs require four degrees of freedom to rotate the leg during climbing movements. Where there was no engineering reason for altering biological parameters (e.g., joint excursions or leg segment length) we accepted those values as default.

The final design of Robot III is 30 inches long, its leg segments are made of aluminum tubes and its joints are fabricated of machined aluminum with roller bearings and hardened steel shafts (Bachmann et al., 1997) (figure 20.3). The fore-aft location of its center of mass is at the body-coxa joints of the rear legs, as it is in the animal. At 30 inches in body length, the robot is 17 times larger than a typical adult *Blaberus*




**Figure 20.3** Robot III has its kinematics scaled up from the cockroach. Its rear legs are much larger than its front legs. It is activated by pneumatic cylinders. Three of the six valve blocks can be seen on the rear of the robot. The white hoses transmit compressed air to the actuators.

cockroach. Its leg segment lengths are also scaled to be 17 times larger than those in the animal. The robot has 24 active DOF and it is controlled with 48 three-way valves, one valve on each side of a double-acting air cylinder (or pair of cylinders). In this configuration, the valve is always either pressurizing one side of a cylinder or exhausting it, and thus no air storage is possible. Unfortunately, the inability to store air eliminates the useful property of passive stiffness discussed later in this paper.

The choice of pneumatic actuators resulted in a strong and powerful robot. However, the benefits of pneumatic actuators came at a cost in controller design. Most pneumatic devices act in a steplike manner, moving rapidly from one extreme to the other. In order for R-III to move efficiently, we had to develop pneumatic control systems that could smoothly control posture and extend leg joints through a wide range of motion.

### PROGRESS IN OVERCOMING CONTROL ISSUES FOR R-III

Legged locomotion control can be broadly divided into three interconnected subsystems: posture control, swing control, and stance control. Posture control seeks to efficiently move the animal's or vehicle's center of mass (CoM) to appropriate positions and to reject external disturbances. Postural control is active both when the animal is walking and when all legs are on the ground. Swing control cycles the legs periodically and is most important during unloaded conditions. Stance control reinforces the posture control at the local level. The alternation of swing and stance makes up the standard leg cycle during walking or running. Although one can focus on these systems as separate entities,



**Figure 20.4** A cockroach rotates its middle legs (arrow) such that they are perpendicular to the substrate as it climbs a block. The middle legs then rotate the body upward for the climb.

they clearly interact. For example, during climbing, the CoM must be moved upward and, in the cockroach, that occurs by rearing the body to allow the hind legs to push the animal over the barrier. The middle legs play a critical role in the rearing movement (Watson et al., 1998). To do so, they swing the tibia to a position that is more perpendicular to the substrate than it is during walking (figure 20.4). This postural adjustment is made by altering the body-coxa joint during swing, presumably to avoid the difficulty of changing leg position under load. So, here the posture control system must interact with the swing control system. During turning, the posture control system may interact with stance control to generate asymmetrical forces in the left and right legs.

A robot such as R-II, activated with DC motors, can stand and walk using only swing control implemented as proportional control with fixed gains, whereas a pneumatically actuated robot such as R-III cannot practically do so. A pneumatic system is much slower than an electrically actuated system. The time to move air from the reservoir to an actuator is many orders of magnitude larger than the time to activate an electric motor from a battery. Neuromuscular systems also have significant delays, especially in the activation of muscle. For this reason, the proportional control gain for a pneumatic control system (or a biological system) must be much lower to avoid instability. Therefore, the feedback torque that can be generated at the joint will be much less. This suggests a more biologically based control system in which a

global posture controller produces most of the force needed for standing and walking and a stance controller that reinforces the posture controller at the local leg level.

### Posture Control—Control CoM

The role of the posture controller is to position the CoM of the body by directing forces with the legs that are in contact with the ground. When the robot is standing, postural adjustments need not be exceedingly rapid. However, during walking or running, it needs to interact with swing or stance controllers and must act reasonably rapidly, although not necessarily on a cycle-by-cycle basis.

Biological data on mammals suggest that robust posture control is essential for locomotion and that the higher centers of the nervous system are important for posture (Horak and Macpherson, 1995). Posture control appears to be the orchestration and tuning of reflexes in lower regions of the central nervous system according to some desired behavior. Robot III's posture controller reflects these ideas (Nelson and Quinn, 1999): It is a centralized controller based on the virtual model approach (Raibert and Hodgins, 1993; Pratt, 1995; Pratt, Dilworth, and Pratt, 1997).

The input to the robot's posture controller is the desired body position and orientation. The desired forces  $F$  (in the  $x$ ,  $y$ , and  $z$  directions) and moments  $M$  (about the  $x$ -,  $y$ -, and  $z$ -axes) on the body required to cause this motion are calculated based on the robot being modeled as a body attached to six springs, one for each of the body's six DOF, for example:

$$F_i = K_i(\text{actual}_i - \text{desired}_i)$$

where  $\text{actual}_i$  is the actual body position in the  $i$  direction,  $\text{desired}_i$  is the desired body position in the  $i$  direction, and the subscript  $i$  denotes any of the three directions  $x$ ,  $y$ , or  $z$ . In Robot III, as in the cockroach, body position is measured indirectly using proprioceptive data: in the case of the robot, from potentiometers on the stance legs.

The posture controller must assign the load-bearing responsibilities to the stance legs such that their summation causes  $F$  and  $M$  to be exerted on the body. As the robot walks, the stance legs typically number from 3 to 6. Therefore, the number of equations needed to solve this load distribution problem varies with the gait and the phase of the gait.

The problem is separated into a solution for the vertical force distribution ( $z$  direction) and a solution for the horizontal force distribution ( $x$  and  $y$  directions) (Nelson and Quinn, 1999).  $F$  and  $M$  are used to determine the desired center of pressure (CoP) for the robot. Equations are developed that constrain the sum of the vertical forces from the stance legs to equal  $F_z$  and the CoP to be positioned as desired.

Typically, depending on the number of legs in the stance phase, there are not enough constraint equations to solve directly for the vertical force required from each stance leg. Therefore, in addition to the constraint equations, an optimization problem is solved that encourages an equal sharing of vertical load among the stance legs. Note that normally the legs will not share an equal vertical load because of the enforcement of the constraint equations.

Independent constraint equations based on satisfying  $F$  and  $M$  on the body are formed to solve for the horizontal forces. Again, the problem is typically underconstrained and an optimization is used in the solution for the horizontal forces that encourages minimization of the joint torques in each leg. This has been shown to cause horizontal ground reaction forces directed in toward the body in a distribution similar to that measured in walking cockroaches (Full, 1993).

With this posture controller, the robot stands and resists large disturbances. When pushed from the side repeatedly while standing, it appears to exhibit a swaying reflex similar to that observed in animals. The posture controller also enables the 30 lb Robot III to perform push-ups while carrying a 30 lb payload, demonstrating its power (Nelson and Quinn, 1999).

### Swing-Cycle Control

During walking, running, climbing, or any other locomotion behavior, positioning of the legs at the onset of stance is critical. In walking or running, the tarsus (foot) is merely placed at the anterior extreme position (AEP) of the leg cycle. However, in climbing, turning, or reaching, that position may change. For example, during a climb, the middle leg is repositioned so that the tibia is more perpendicular to the substrate than during walking. As a result, extension of the middle leg rears the animal upward in preparation for the climb.

In all cases, actuators acting upon individual leg joints ultimately control the movement of the tarsus. Multiple leg joints acting in concert can swing the tarsus to the desired location. In the inverse kinematics problem, the joint angles are determined to place the foot at a desired location. However, if a leg has more than three degrees of freedom, there are typically many different sets of joint angles that can position the foot at a given location. Legs such as these are said to be kinematically redundant, although those DOF are not functionally redundant. For example, in climbing, each middle leg could not be rotated so that its tibia is perpendicular to the substrate without this "redundance."

R-III's front and middle legs are kinematically redundant. We solved the inverse kinematics problem based upon an idea described by Mussa-Ivaldi and Hogan (1991), which has been supported by observations of mammals. The solution is chosen that places the joints as

close as possible to their mid-range locations. This has the advantage of keeping the joints away from their mechanical limits and thus away from those singularities. The optimal solution was solved off-line for points distributed throughout the work space of each leg. Then a neural net was trained using these data. The neural net successfully interpolates between these points to provide a continuous solution for the inverse kinematics problem throughout the work space. The higher-center controller issues commands describing where the tarsus should be placed along with a few intermediate milestones. The local control network then moves the foot through those way points to the desired location. The result is a smooth, efficient movement. Even the front legs, with their five degrees of freedom, move the foot smoothly through the desired range of motion.

The implementation of the controller for R-III reflects its hierarchical organization. The control system is implemented in two off-board computers that communicate via a serial link: a higher-center computer and a local controller computer. The local computer communicates with the robot's sensors and valves via high-speed I/O (input/output), whereas the higher-center computer can not directly communicate with the robot. The swing controller and stance controller reside in the local computer, whereas the posture controller is implemented in the higher-center computer. The local computer monitors the robot's 24 joint angle positions and 18 leg loads and performs the analog-to-digital conversions. It also performs the swing- and stance-control calculations. The local controller determines the duty cycles for the valves based upon commands from the posture controller, the swing controller, and the stance controller. It then performs 80 Hz pulse width modulation (PWM) control of the 48 valves.

This PWM implementation of proportional joint control overcomes one of the major problems that face pneumatic devices. R-III's joint movements are smoother than most pneumatic devices and, in fact, are similar to the movements of R-II's joints, which were powered by DC motors.

Yet another problem to be solved by the swing controllers is coordinating the legs in an efficient tripod gait. A separate swing pattern generator controls each leg and we showed that individually, each leg could move smoothly. We next established that the control system has the bandwidth to move all six legs smoothly at once. Finally, we implemented the stick insect coordination mechanisms as defined by Dean (1992) to coordinate the six legs. As a result, we showed that the swing controller could move R-III's legs in a tripod gait while it is suspended in the air. As in the animal, the rear, middle, and front legs each act differently from one another. Middle and hind legs extend the coxa-femur and femur-tibia joints in synchrony, while front legs make a complex movement where these joints move out of phase. Nevertheless, the



front and rear tarsi on the right side move in concert with the middle tarsus of the left side and this tripod alternates with the tarsi of the remaining legs (figure 20.5).

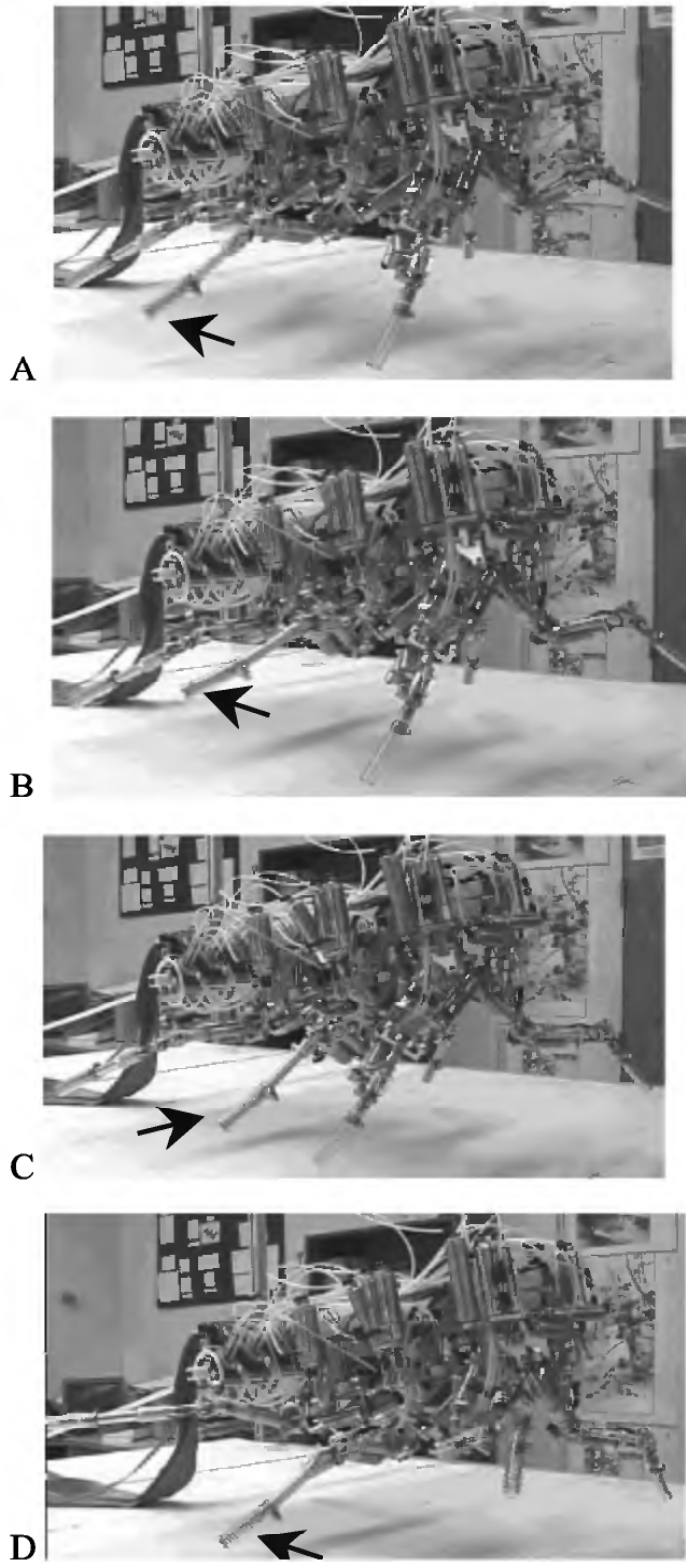
### Stance-Cycle Control

The third control system is the stance controller that develops ground reaction forces to assist the posture controller. Stance control is implemented on the local computer, which is wired directly to the robot via high-speed I/O connections, whereas the posture controller is implemented on a separate computer, which acts through a serial connection to the local computer. As a result, the local stance controller can react much more quickly than the posture controller to reject disturbances and increase stance forces as needed.

In stance, the controller must cause a forceful extension of joints until the tarsus attains a predetermined posterior extreme position (PEP). The stance cycle could be preprogrammed if all possible conditions could be predicted. However, for complex terrain that is not possible. It is more feasible to monitor load sensors to determine the amount of force necessary to continue extension until the appropriate joint movement is attained. When a cockroach climbs over a barrier, the leg joints extend to the same angle that is attained in walking. However, it takes longer to reach that point when the animal is moving against gravity. A reasonable explanation for this observation is that load sensors on the cuticle (campaniform sensilla) detect the strain associated with the leg extension against gravity and provide a positive feedback signal to Ds motor neurons to increase the forceful extension until joint-angle detectors report that the appropriate PEP is achieved.

This form of positive load feedback has been documented in both cockroach and cat locomotion. It is well known that positive feedback can cause a system to become unstable. However, here, several factors prevent the animal's movement from becoming unstable. First, the properties of muscle are self-limiting. The length-tension relationship of striated muscle shows that as the muscle shortens, its capacity to generate tension declines. In addition, proprioceptors that monitor joint angles provide negative feedback to turn off the stance phase when the appropriate extensions occur and switch the leg control from stance to swing.

In R-III, strain gauges were placed on four sides of each tibia near the femur-tibia (FT) joint and two sides of each femur near the coxa-femur (CF) joint to monitor bending of these segments. Signals from these sensors are connected to sensory integration circuits (bridges) for each leg, which are mounted on the robot. The outputs of each of these circuits are input to the local control computer, where the ground reaction force vector at each stance foot is calculated.



**Figure 20.5** Robot III moving its legs in a tripod gait using the Cruse mechanisms and the swing controller. (a) Right middle leg entering swing phase, (b) Middle leg in swing phase, (c) Middle leg entering stance phase, (d) Middle leg in stance phase.

We have demonstrated a slightly different form of positive load feedback on one leg of the robot by modifying the negative feedback proportional controller used in swing. The swing controller is active throughout the walking cycle, including the stance phase. First, the joint-control feedback gains are set to their maximum levels, which avoid instability during the swing phase. The torques that can be produced with these gains are not nearly enough to make the robot stand or walk. When a leg enters the stance phase and begins lifting the body, the gains can be increased because the apparent inertia that the joints are moving includes part of the inertia of the body, which is much larger than that of the legs alone. The load at that joint, calculated from the ground-force measurements, can establish how much each gain can be safely increased. As the load at a joint increases, its gain can be increased even further, thereby increasing its load-bearing capability. We use a positive-load feedback strategy wherein the gain for each negative feedback joint controller increases linearly with sensed load at the joint. Using this strategy, we demonstrated that a front leg could increase its load-carrying capacity by about a factor of two.

## FUTURE IMPROVEMENTS IN CONTROL SYSTEMS

The positive feedback aspects of stance control can be improved with a different type of pneumatic actuator. McKibben artificial muscles, or braided pneumatic actuators, have several advantages over other actuators, including pneumatic cylinders (Chou and Hannaford, 1996). They consist of an inflatable elastomer tube surrounded by a braided sheath that is longitudinally stiff. The ends are plugged and the tubes are crimped on the plugs. One plug has a port through which air enters the elastomer tube, inflating it and causing the actuator to expand circumferentially and shorten longitudinally. Several important properties of these actuators are similar to muscle. They are tension-only devices and their force output decreases as they shorten. They are also flexible, and their strength-to-weight ratio is favorable in comparison to muscle. Moreover, and most important, when acting in antagonistic pairs, their passive stiffness can be tuned independently of their stroke. Tunable passive stiffness is essential for efficient legged locomotion at various speeds. Finally, their maximum strain is about 30%, which compares favorably with muscle. Because of these musclelike properties, we will use McKibben actuators in our next generation vehicle (Robot IV) along with a more cockroachlike exoskeleton structure with scaled inertia properties. These considerations will result in R-IV having dynamics similar to the cockroach, so that it can walk, run, and climb efficiently. This attention to detail is not done for the sake of mimicking the animal, but to give the robot attributes that are important for a mission-capable robot.

New actuator technologies are under development (see chapters 7 and 8, this volume) and offer exciting advances for the future. Those being developed by SRI (chapter 7, this volume) and Honeywell (chapter 8) are activated electrically and they can be constructed to have tunable passive stiffness. When these actuators have been developed to the point where they have sufficient power to weight and size, we will certainly look into incorporating them into our future robots. However, we anticipate that controllers for these new technologies will have to overcome fundamental problems similar to those that we have overcome in controlling our robots.

Fine tuning of control architecture can be improved by returning to the animals for guidance. The neural networks for swing control were inspired by work on control of limb trajectory that has been conducted in amphibian systems (see chapter 11, this volume). Similar motor control systems may exist in insects, and the capacity to study these circuits in animals that allow one to record from identified neurons could provide new insights as to how these circuits respond to higher commands and control leg movement. It is also possible to monitor local control systems during walking and climbing. Sensory inputs and motor outputs can be recorded extracellularly. Intracellular recording techniques can be applied to both motoneurons and interneurons as the animal walks in a tether such as one that was recently developed by Andrew Tryba while working in our group (Tryba and Ritzmann, 2000).

These neurobiological and behavioral studies can guide development of control systems. However, as with the initial design of R-III's legs, one cannot simply attempt to characterize the neural circuitry in the animal and expect to insert it into the control system for the robot. It is unlikely that the entire circuit will be identified any time soon. Even if a complete characterization were attained, mismatches of size, scale, and materials between the animal and the robot would limit the success of such an effort. A more reasonable strategy is to assay the rules by which nervous systems control locomotion, taking advantage of the remarkable sensitivity of extracellular and intracellular neurobiological techniques to formulate specific rules. Once these rules are established, modeling techniques can be used to implement them into improved control systems.

## TOWARD AUTONOMY


Ultimately, a truly capable robot must be fully autonomous. A device that is attached to a tether either for computational purposes or for power cannot carry out most missions. Many of these issues are best addressed after the design of the vehicle has been optimized. If a research robot is to push the boundaries of current design, it is illogical to hamstring the research team by constraining them to limited onboard

devices even before solutions are found. For example, control circuits are easier to manipulate in a fully capable desktop computer than in a preprogrammed chip. Once designs have been debugged and optimized, it is then possible to lay out chips that reduce the size and weight to onboard dimensions. Likewise, once the robot design and actuator and control technologies have been optimized for efficient locomotion, the energy requirements for the robot will have been minimized. This is the time to install the power source onboard the vehicle. However, waiting until then to develop onboard capable control and power systems would cause delays in the development of the robot. For this reason, we and our colleagues are developing compressor, valve, and control systems that would be appropriate for onboard implementation and we have established team relationships with commercial producers of the desired technologies.

We have begun to work on these issues of autonomy in developing another robotic device, which we refer to as the *cricket robot* (Birch et al., 2000). The goals for this vehicle are that it be fully autonomous and approximately two inches long. Power is clearly an issue. Small legged animals must cycle their legs much more rapidly than larger animals to move at the same speed. This limits their speed and is not energy efficient. Some insects overcome this problem by jumping. We chose the cricket as our inspiration for this device, because it is an efficient jumper, yet it also walks very effectively.

The leg segments for the cricket are very different from the cockroach. For example, the coxa, which is a large and powerful leg segment in cockroaches, is very small in crickets. The femurs are the largest and most powerful segments on the rear legs of the cricket because they are used to propel the animal in jumping. As in locusts, the jump is accomplished by cocking the rear leg, flexing the extensor muscle to full tension and then suddenly releasing it. The cricket robot is also being built with large strong rear legs. Furthermore, based upon the cricket's design, we have eliminated the coxa segments from all of the legs to simplify the robot. The front and middle legs will have, at most, three active joints, and the rear legs have two active joints. The rear legs have been constructed and have been demonstrated to be powerful and capable of rapid motion.

The cricket robot uses custom McKibben artificial muscles to extend its leg segments. Flexion is accomplished using springs. A custom-built compressor made from a Smoovy motor and a commercial air cylinder has been designed and built to allow the power to be onboard. Thirty-two on-off valves are needed to control the sixteen actuators. Commercial valves that are small and light enough for this purpose are not available. Some of our colleagues in the cricket project are developing valves for this purpose using MEMS (microelectromechanical systems) fabrication techniques (Geon et al., 2000). MEMS joint-angle position



**Figure 20.6** A wheeled cart was constructed to test the rear legs of the cricket robot. A  $\frac{1}{2}$  inch diameter steel bolt that is  $1\frac{1}{2}$  inches long was placed in the cart as a payload. Small McKibben actuators with white braiding activate the legs.

sensors are also being developed by other colleagues (Fedder, personal communication). Neural network locomotion controllers are being evolved using genetic algorithms (Beer and Gallagher, 1992) and will be implemented in analog VLSI (very large scale integration) circuits (Brown, 2000). A prototype system, consisting of two hind legs attached to a wheeled cart (figure 20.6), has been built to test the hind legs and associated technologies. When the remaining legs and supporting devices are complete, the robot will be fully autonomous.

## CONCLUSION

Development of a mission-capable legged robot is a serious challenge. To accomplish this goal, we are attacking the problem at several levels that center on control, power, actuation, and mechanical design issues. We have outlined some of the problems that we have sought to overcome. We will only make real progress by accepting these challenges and realizing that progress toward a mission-capable robot will take significant effort. Animal systems can be useful resources for solutions. The process by which useful information is extracted and confounding details are left behind is a nontrivial task. However, the benefits can be great.

## ACKNOWLEDGMENTS

The work that is described in this paper was supported by ONR grants N0014-99-1-0378 and N0014-96-0694, DARPA/ETO DAAN02-98-C-

4027, NSF/LIS grant DMI97-20309, and NASA Marshall Space Flight Center grant NGT-52832.

## REFERENCES

- Bachmann, R. J., Nelson, G. M., Flannigan, W. C., Quinn, R. D., Watson, J. T., and Ritzmann, R. E. (1997). Construction of a cockroach-like hexapod robot. In *Proceedings of the Eleventh VPI and SU/AIAA Symposium on Dyn. and Cont. of Large Structures*, Blacksburg, Va.
- Beer, R. D., Chiel, H. J., Quinn, R. D., Espenschied, K., and Larsson, P. (1992). A distributed neural net architecture for hexapod robot locomotion. *Neural Comput.* 4: 356–365.
- Beer, R. D., and Gallagher, J. C. (1992). Evolving dynamical neural networks for adaptive behavior. *Adaptive Behav.* 1(1): 91–122.
- Birch, M. C., Quinn, R. D., Hahm, G., Phillips, S., Drennan, B., Fife, A., Verma, H., and Beer, R. D. (2000). Design of a cricket microrobot. *Proceedings of IEEE Conference on Robotics and Automation*, San Francisco.
- Brown, B. E. (2000). An analog VLSI implementation of a continuous-time recurrent neural network. M.S. thesis, Case Western Reserve University.
- Chou, C. P., and Hannaford, B. (1996). Measurement and modeling of McKibben pneumatic artificial muscles. *IEEE Trans. Robotics Automation* 12(1): 90–102.
- Cruse, H. (1990). What mechanisms coordinate leg movement in walking arthropods. *Trends Neurosci.* 13(1): 15–21.
- Dean, J. (1992). Coding proprioceptive information to control movement to a target: Simulation with a simple neural network. *Biol. Cybern.* 66: 335–343.
- Espenschied, K. S., Quinn, R. D., Chiel, H. J., and Beer, R. D. (1993). Leg coordination mechanisms in stick insect applied to hexapod robot locomotion. *Adaptive Behav.* 1(4): 455–468.
- Espenschied, K. S., Quinn, R. D., Chiel, H. J., and Beer, R. D. (1996). Biologically based distributed control and local reflexes improve rough terrain locomotion in a hexapod robot. *Robotics Autonomous Syst.* 18: 59–64.
- Full, R. J. (1993). Integration of individual leg dynamics with whole body movement in arthropod locomotion. In *Biological Neural Networks in Invertebrate Neuroethology and Robotics*. Boston: Academic Press.
- Gallagher, J. G., Beer, R. D., Espenschied, K. S., and Quinn, R. D. (1996). Application of evolved locomotion controllers to a hexapod robot. *Robotics Autonomous Syst.* 19(1): 95–103.
- Hahm, G., Kahn, H., Phillips, S. M., and Heuer, A. H. (2000). Fully microfabricated, silicon spring biased, shape memory actuated microvalve. In *Proceedings of 2000 Solid-State Sensor and Actuator Workshop*, Hilton Head Island, S.C.
- Horak, F. B., and Macpherson, J. M. (1995). Postural orientation and equilibrium. In *Handbook of Physiology*. New York: Oxford University Press.
- Mussa-Ivaldi, F. A., and Hogan, N. (1991). Integrable solutions of kinematic redundancy via impedance control. *Int. J. Robotics Res.* vol. 10, no. 5 (October): 481–491.
- Nelson, G. M., and Quinn, R. D. (1999). Posture control of a cockroach-like robot. *IEEE Control Syst.* 19(2): 9–14.

- Pratt, J. (1995). Virtual model control of a biped walking robot. Master's thesis, Massachusetts Institute of Technology.
- Pratt, J., Dilworth, P., and Pratt, G. (1997). Virtual model control of a bipedal walking robot. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA '97)*, Albuquerque, N.Mex.
- Quinn, R. D., and Espenschied, K. S. (1993). Control of a hexapod robot using a biologically inspired neural network. In R. D. Beer, R. E. Ritzmann, and T. McKenna (eds.), *Biological Neural Networks in Invertebrate Neuroethology and Robotics*. Boston: Academic Press.
- Raibert, M. H., and Hodgins, J. K. (1993). Legged Robots. In R. D. Beer, R. E. Ritzmann, and T. McKenna (eds.), *Biological Neural Networks in Invertebrate Neuroethology and Robotics*. Boston: Academic Press, pp. 319–354.
- Tryba, A. K., and Ritzmann, R. E. (2000a). Multi-joint coordination during walking and foothold searching in the *Blaberus* cockroach: I. Kinematics and electromyograms. *J. Neurophysiol.* 83: 3323–3336.
- Tryba, A. K., and Ritzmann, R. E. (2000b). Multi-joint coordination during walking and foothold searching in the *Blaberus* cockroach: II. Extensor motor pattern. *J. Neurophysiol.* 83: 3337–3350.
- Watson, J. T., Pollack, A. J., Zill, S. N., and Ritzmann, R. E. (1998). Kinematics and leg muscle activity in cockroaches climbing over obstacles. In *The Fifth International Congress of Neuroethology*, 23–28 August.



**This page intentionally left blank**

## A Biologically Inspired Approach Toward Robust Real-World Locomotion in Legged Robots

Frank Kirchner, Dirk Spenneberg, and  
Ralf Linnemann

Even the most primitive biological systems solve problems that reach far beyond the capabilities of today's technical systems (Ayers et al., 1994; Fearing et al., 2000; Calvitti and Beer, 2000). The biomimetic approach to robotics is the attempt to apply solutions created by evolution to technical systems. This approach is not restricted to mechanical engineering but includes and puts emphasis on the behavior of autonomous systems (i.e., the algorithms that map from sensory stimuli to motor acts). We are convinced that the study of animal behavior leads to models and hypotheses that help create new technical solutions in analogy to their successful counterparts in nature. The benefits of these studies will not only be harvested in the field of mechatronic design but will also help to achieve algorithmic solutions to the questions of generation and organization of behavior.

The true value of studying animal design and behavior is thus to be found in hints and clues to bias our technical systems. If we learn from this research the hows, whens, and wheres to use bias, we will eventually be able to come up with a new class of algorithms that are closer to BI (biological intelligence).

The long-term goal of the biomimetic approach toward autonomous robots is thus to develop systems that interact for an unrestricted time with a real-world environment and enhance their level of performance with experience gained from ongoing interaction with the world (figure 21.1). The following list highlights some design considerations for such machines:

1. Robust, lightweight design, in order to withstand environmental influences on the hardware (e.g., rain, snow, dust, heat, as well as mechanical stress induced by shock, e.g., the system crashing down from a rock while trying to climb it).
2. Fault-tolerant design. The failure of a single component should not be fatal to the robot's ability to complete its mission. This refers to the hardware design (e.g., multiple independent actuators), the electronics design (e.g., multiple independent processing units), and finally to




**Figure 21.1** The first prototype of the Scorpion robot during an autonomous exploration into a rock channel. Note that obstacles have been up to three times as high as the robot itself. The obstacle course was 3 m wide and approx. 9 m long. The task was to complete the course within 20 min time.

the software design (e.g., multiple independent modules in a layered structure) with many-to-many communications and sensor reading redundancy.

3. Energy autonomy. This refers to the fact that the system should be able to carry its own energy supply; moreover, it should be able to recharge this supply by its own means. There are many different approaches to enable robots to do so, most of them consist of enabling the robot to harvest energy from its environment (e.g., solar, wind, and thermal energy are among the promising concepts).

4. Adaptivity and learning capabilities. We can not foresee the course of a robotic system in any given real-world environment. It is, in the worst case, an unpredictable chain of events dependent on and sensitive to changes in the set of start-up assumptions we have made. However, it is not complete magic, nonetheless. We almost always know a set of situation-reaction rules<sup>1</sup> for nearly all environment-system pairs and we can program these for a start. Yet, a system that is able to adapt these rules, and that may learn new ones from the experience it gains interacting with the world, certainly increases its chances for survival.

Reviewing the above list may remind one of the design specifications for complex systems that we know from space shuttles, planes, and submarines, to nuclear power plants. We want them to be robust enough to withstand the stresses, fault tolerant to minimize failure, and



**Figure 21.2** The mechanical design of the Scorpion legs. This front view of the robot shows left and right side legs with the body in the center. Each leg consists of four parts: (1) thoracic joint, (2) basalar joint, (3) middle joint, and (4) distal joint. Except for the distal joint each joint is actuated.


energy efficient to extend mission time. What we do not want such systems to be is autonomous, adaptive, or learning things on their own. But building a biomimetic robot is much of the systems design, as we know it, plus the built-in feature to change this design.

## MECHATRONICS

This section will start with the description of the most demanding parts in a legged machine. The leg modules are constructed from four components:

1. The thoracic pivot, which mediates protraction and retraction.
2. The basalar segment, which contains the pivot mediating elevation and depression.
3. The middle segment, which contains the pivot mediating extension and flexion.
4. The distal segment, which terminates in the limb tip and its sensors.

A leg provides 3 degrees of freedom and is actuated by standard DC geared motors. The rotation of the legs around the horizontal and vertical plan is realized in the thoracic and basalar part (figure 21.2), by two orthogonal connected actuator components, which provide protraction/retraction as well as elevation/depression movements of the limb. The third degree of freedom (figure 21.2) is realized by a geared DC motor, which is integrated into the middle segment of the leg. A major




**Figure 21.3** The distal segment contains a spring damped compliant element with a built in potentiometer to measure contact and load on individual legs. Poti, potentiometer to measure varying load on the spring damped foot. Igus bearing is a special material used to reduce friction.

challenge in the design of the leg modules was the constraint that we wanted to build an outdoor-capable walking machine while achieving a good ratio between the leg weight and its lift capacity. In our design, we have achieved a ratio of 1:8, incorporating shielding from the environment. This was possible by using high-ratio planetary gears in combination with a powerful DC motor. The increased ratio was necessary as we intended to be able to climb obstacles that would exceed the robot's own height. Thus, in certain situations, single legs have to be able to pull/push at least three-quarters of the robot's weight.

Another aspect in leg design is the speed of the leg. A robot walking over uneven outdoor terrain also needs certain reactivity (Cruse et al., 1998; Blickhan and Full, 1998). Simply increasing the gear ratio does help to gain the torques desired but it also decreases the reactive speed of the module. In the Results section, we will describe data on reflex loops that are built into the standard motor control loop. The reflexes are fast reaction schemes to sensorized stimuli and are extremely important to stabilize walking over uneven terrain. These reflexes have to be fast. At least, they have to be fast enough to compensate for the robot's own inertial energy when exposed to gravity. A third aspect in leg design for outdoor walking machines comes in the form of impulse-like stress signals acting on the mechanical design. There is no way to design a leg module that resists these stresses over a reasonable time. As in the biological examples, we have incorporated compliant elements in the leg design. These elements are meant to absorb some of the high peak energy impulses coming from hitting the foot on the ground in the late swing phase (figure 21.3).

Another important part of the leg is the distal segment of the leg. Most approaches use switches to measure ground contact of the foot. There are two drawbacks to this approach: The switch can break, as it is literally hammered on the ground each time the robot makes a step; the switch only closes if the foot hits the ground in a defined position. In this approach, the distal segment terminates in a sensor that is used to measure contact and pressure (load) on the particular leg while also providing solid stand and traction. The distal segment is manufactured



**Figure 21.4** The external sensors of the Scorpion robot comprise an ultrasound transducer in the robot's head. This sensor detects obstacles up to 80 cm in front of the robot. Above the ultrasound is a color CCD camera, which transmits pictures via a wireless video link. A wireless link (DECT is a European standard for wireless communication) is used to send status data and receive commands from a human operator.

as a spring-damped chamber with a built-in potentiometer (see figure 21.3). This approach enables us to measure, contact, and load on individual legs, while the spring mechanism acts as a damping component to reduce the impact of high forces on the leg.

### Sensors and Communications

The front part of the robot's body is equipped with ultrasound/infrared transducers (figure 21.4) and two bumpers that are mounted at the head segment. These sensors are used to measure close contacts in critical maneuvering situations (narrow passages, etc.). The distance sensors (ultrasound) are mounted on small DC motors to allow the robot to actively sample data from locations in the environment. This issue, also referred to as active perception (Whitehead and Ballard, 1990; Kirchner, 1995), is important for autonomous robots for two reasons: (1) It allows the robot to double scan (though from a slightly different angle) certain environment locations—thereby increasing probabilities to correctly map external environment states onto the set of internal states of the robot. A problem that is referred to in the literature as hidden-state or perceptual aliasing. (2) It allows the robot to sample the environment from different angles without moving the whole body.

**Table 21.1** List of proprioceptive and exteroceptive sensors of the scorpion robot

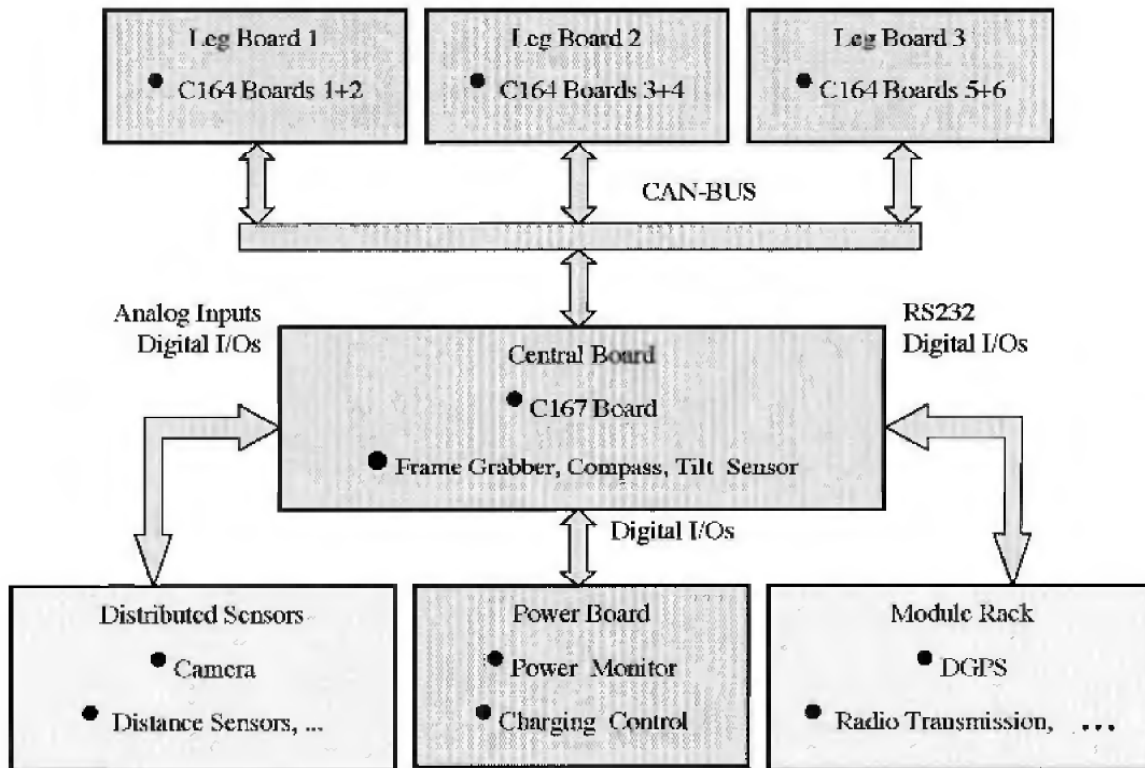
Proprioceptive Sensors	Exteroceptive Sensors and Communications
Relative joint angle measurement devices	Tactile sensors for close contact measurements
Motor current measurement devices	Ultrasound distance measurement devices
Two dimensional Inclinometers	Infrared emitting/receiving diodes
Pressure sensors in foot compartments	CCD-camera, for supervised autonomous behavior
Battery status measurement	Compass for absolute positioning and sense of direction and navigation
	A DECT radio link for data and command transmissions
	A wireless video/audio link for supervised autonomous mode of operation

There are two classes of sensors integrated in the system: Proprioceptive sensors, which are used to measure the internal status of the robot (e.g., used to derive models of spatial displacement of the robot's kinematics in relation to the external environment); and exteroceptive sensors, which will be integrated into the system in order to map states of the external environment onto internal states of the robot (table 21.1).

During autonomous operation of the robot, a high bandwidth radio communications link (see figure 21.4) is used to transmit video images and other status data to a remote operator desk. We are using two wireless systems in parallel: a 28 KBaud bidirectional DECT link on the robot, and a 5 GHz video/audio link for image transmission. The DECT radio link interfaces directly to the main onboard processing unit. Commands received from the operator desk are translated into internal function calls. The video link is only passive at this point, which means that the pure video signal, as delivered by the CCD camera in the tail, is directly transferred to the receiving unit at the operator's desk.

### Computational Hardware

We designed a network of one central and many local computing units (figure 21.5). While the local controllers perform all the low-level tasks to drive legs, read in sensor data, and even react (in a reflexive manner) to smaller complications (an obstacle that blocks a single leg, etc.), the main control unit runs the main program—where these details can be ignored. This reduces processing time and keeps the main processor free from low-level control tasks while it can focus on higher-level, global decision-making aspects of the overall task. The main advantage of this design, besides the aspect of fault tolerance, is that the global control program runs only on the master controller, which interfaces



**Figure 21.5** The electronic design of the Scorpion robot consists of one central controller (C167) and 8 local controllers (C164). Two C164s are integrated into one Leg board, which additionally provides, external power, com ports and a CPLD-chip, which implements a total of 8 quadrature counters and interfaces directly into the C164 memory banks.


with the elementary levels through standard protocols in both hardware and software.

The local and the central computing boards are connected through a serial CAN-bus system, which allows individual addressing and bidirectional data exchange between all components of the network at up to 1 Mbit/sec. This bus system is well suited in distributed embedded control environments and is common in industrial applications.

The leg boards are equipped with special high-amp H-Bridge drives, which allows them to control up to four motors with a maximum of 5 amps each. At the same time, they can monitor the positions of each motor and read the current drawn from each motor via a CPLD<sup>2</sup> chip integrated into the design. In combination with the computing power, they provide the full range of capabilities for real-time feedback control applications.

The central control board implements the master control program. This program is the application-specific implementation of a set of high-level behaviors of the robot. These behaviors typically rely on complex information that is sampled from the environment by a range of different sensors (see table 21.1).





**Figure 21.6** Three different designs for the 8-legged model of the Scorpion robot. The parametric approach allows us to easily change designs in the simulation.


In a second version of this hardware, we are planning to use the Motorola MPC55 Microcontroller version of the PowerPC in combination with an FPGA (free programmable gate array) that will take over the tasks now implemented in the local controllers. As these systems get smaller, yet more and more powerful, it becomes affordable even for a mobile robot to employ a true full-scale processor and also to think about using classical operating systems such as the real-time version of Linux.

This section gave a brief outline of the mechatronic design of the first prototype robot, which we built and evaluated during several in- and outdoor experiments (see Results section for details). However, prior to this work, we implemented a full parametric model of the system in the simulation environment ADAMS. The next section will outline the modeling process and describe a set of initial experiments with that model. These results have been used during the design of the real prototype and the implementation of the behavioral control architecture on the real-world machine.


## A PARAMETRIC MODEL

The simulation model of our walking machine Scorpion describes its basic mechanical features and is based on the following assumptions (see figures 21.2 and 21.6):

1. Eight legs with three active revolute joints each. One joint for horizontal rotation and two joints for vertical rotation. The horizontal and the first vertical joint are the “hip,”<sup>3</sup> the other vertical joint the “knee” of Scorpion. Furthermore, each leg has a passive joint between leg and foot, the “ankle.”
2. A leg consists of:
  - A link—a cylinder that is firmly attached to the body and houses the horizontal joint.
  - A “hip”—another cylinder housing the first vertical joint, which can be rotated by the link.



**Figure 21.7** Implementing a particular design starts with design points: A design point is an anchor for a particular part or joint of the design. The design for the Scorpion robot consists of 43 parts and 42 joints.




**Figure 21.8** Walking gait of a real scorpion (Bowermann, 1975).

- An “upper leg”—a cylinder rotated by the hip.
  - A “lower leg”—a cylinder rotated by the knee.
  - A spherical foot.
3. Lateral symmetry.
  4. Longitudinal symmetry.

The model consists of 43 parts and 42 joints. Most parts consist of only one block, but a few parts comprise up to nine different blocks. All parts and joints (and consequently their interconnections) are fully built by parametric information. Parts are placed in a global coordinate system by means of design points, which serve as anchor points for them. The current model consists of 84 parametric design points with three components each (figure 21.7).

### Posture and Initial Angles

Living scorpions walk almost always with the same gait (figure 21.8): L4-R3-L2-R1 and R4-L3-R2-L1, where L and R stand for left and right side and 1–4 is the enumeration of the legs from front to end. That means when the hind leg starts to move, leg number 3 on the opposite side repeats the motion pattern a short time afterward, then leg number 2 on the same side and leg number one on the opposite side follow.



**Figure 21.9** Initial postures are defined by six variables (see text for details). Here three different postures are presented. (A)  $v\_hip\_angle = 45^\circ$ ,  $knee\_angle = -45^\circ$ ; (B)  $v\_hip\_angle = 45^\circ$ ,  $knee\_angle = 0^\circ$ ; (C)  $v\_hip\_angle = 0^\circ$ ,  $knee\_angle = 0^\circ$ .

**Table 21.2** Design variables in the Scorpion model

Joint	Design variables, group 1 (R4-L3-R2-L1)	Design variables, group 2 (L4-R3-L2-R1)
Horizontal hip	$h\_hip\_angle1$	$h\_hip\_angle2$
Vertical hip	$v\_hip\_angle1$	$v\_hip\_angle2$
Knee	$knee\_angle1$	$knee\_angle2$

Note:  $h\_hip$ , thoracic;  $v\_hip$ , basalar, and  $knee$ , distal joints.

Therefore, the legs of the walking robot are grouped into two design sets according to the explained motion.

The initial posture is defined by six design variables for the joints (figure 21.9, table 21.2). A group of four legs needs three values for the horizontal and vertical hip joints and for the knee joints. The default initial position (angles of zero degree) is rectangular with the upper legs stretched horizontally. Positive values stand for upward/forward, negative values for downward/backward.

The grouping of the eight legs allows for one group of legs not touching the ground (figure 21.10) in the starting position—thereby enabling the simulation to initiate a coordinated movement with one group of legs in stance and the other group in swing mode.

### Ground Contact and Friction

Limitations in computer power make it nearly impossible to establish an overall check for ground contact during simulation. Therefore, we opted for a limited check only. All eight feet are of spherical shape and hence suited for this check. A marker in the origin of the global coordinate system spans an  $x$ - $z$  plane for  $y = 0$ . The repulsive force during ground contact consists of two components—stiffness and damping.



**Figure 21.10** Due to the parametrical design approach it is possible to define two groups of legs in different starting positions each, thereby enabling the simulated Scorpion robot to initiate a walking maneuver.


The equation is described by four design variables. The contact forces for different feet are alike.

Furthermore, static and dynamic friction calculation during ground contact is established. Three design variables are sufficient for both forms of friction and the transition from one to another. In order to integrate mass and inertia, the ADAMS simulation environment offers the definition of mass, density, and mass moments of inertia for each part separately. At the moment, there are four design variables for density of body parts, links and hips, upper legs, and lower legs. Mass moments have to be changed manually and individually.


Some versions of the current model incorporate spring dampers between feet and legs and between body and link. The spring dampers act as shock absorbers, which faintly resemble the biomechanics of animal legs, in a system of otherwise rigid parts. Each set of spring dampers has two design variables for stiffness and damping. In the experiments described here, the stiffness was relatively big to be close to physical behavior and to prevent unnecessary wiggling. In addition to the spring dampers, translational joints create one degree of freedom for the motion between the parts.

### **Canonical Motion with Swing, Stance, and Phase**

One can distinguish two modes of motion for the leg of a scorpion (figure 21.11). During the stance phase, the ground contact of the leg contributes to the forward motion of the entire scorpion. The leg itself moves backward with respect to the body. During the swing phase, the leg is lifted from the ground, it moves forward with respect to the body



**Figure 21.11** Swing and Stance alternations in the legs of a real scorpion. Data according to Bowerman, 1975. The simulation is based on these data sets.




**Figure 21.12** Joint *hhip\_14* (thoracic joint on leg 4) during a swing and stance cycle. Shown are angle (dark curve) and angular velocity (light curve) values.

and overtakes it before it eventually touches the ground again. The sum of time for swing and stance is called *period*. The reverse ( $1/\text{period}$ ) is defined as the step frequency.

The motion functions comprise two design variables for swing and stance. These variables express the duration of each stage and can be changed independently. In a previous section, we introduced the two groups of legs, L4-R3-L2-R1 and R4-L3-R2-L1 (see figure 21.8). Another design variable called *phase* reflects the delay of motion patterns between two contralateral adjacent legs (e.g., L4 and R3). Phase is a fraction of period. A typical value for living scorpions is about 0.1. It is possible to apply individual motion functions for each leg. But a uniform approach is selected at the moment to systematically research the basic properties of the physical behavior. This includes the use of the design variables for swing, stance, and phase. All motions are based on sine functions, because they avoid abrupt transitions and allow for smooth transitions, as outlined in the real-world Results section for the real prototype.

In this example (figure 21.11) the starting position is identical for L4-R3-L2-R1 and R4-L3-R2-L1. There are several methods to realize the




**Figure 21.13** A group of four thoracic joints in the initial phase (upper trace) and late phase (lower trace) of a coordinated forward movement.

transition to the phase shift during motion. Figure 21.11 gives an example (swing = 1.5, stance = 2.5 and phase = 0.1). The first step establishes the phase shift and is a half swing for L4-R3-L2-R1 and a half stance for R4-L3-R2-L1. It takes a quarter of the period (4.0) for the hind legs and subsequently “phase  $\times$  period” ( $0.1 \times 4.0 = 0.4$ ) time units longer for each consecutive pair of legs. The second step establishes the phase shift of 50% between both groups: R4-L3-R2-L1 conducts a swing in half a period, whereas L4-R3-L2-R1 conducts a stance already according to the design parameter. The desired motion pattern is active at the end of this step, and each leg swings and stances rhythmically.

**Basic Experiments and Results** The model was simulated with a set of design parameters to test its general practicability. The body was constructed as a box and all the legs had the same lateral distance from the body. They had an initial M-shaped posture perpendicular to the sides of the body. The length of the upper leg is 136 mm, the length of the lower leg 237 mm.<sup>4</sup> The motion parameters were swing = 1.5 sec, stance = 2.5 sec, phase = 0.1.

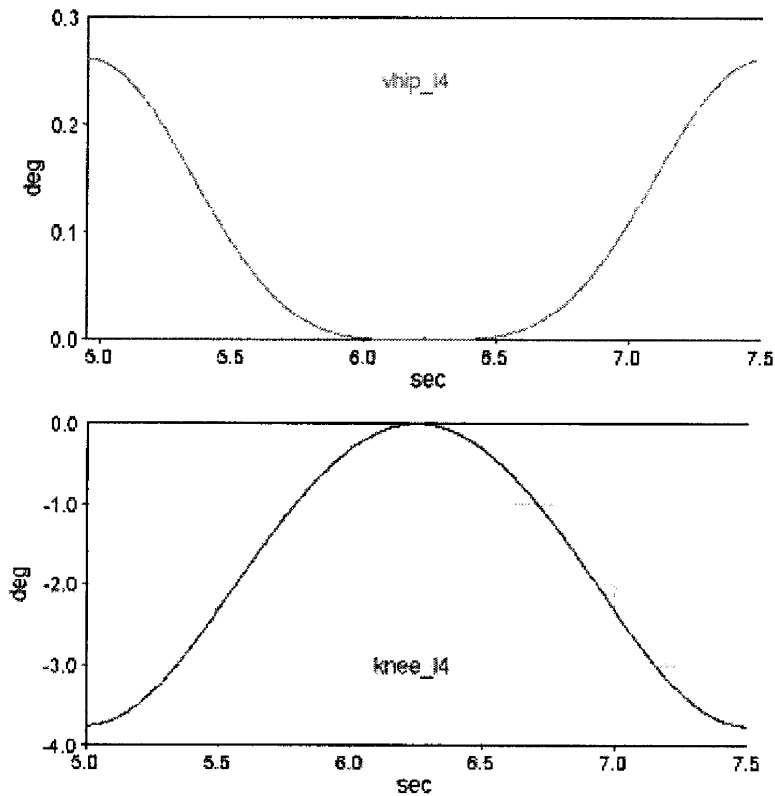
During a stance, a horizontal joint propels the body forward. The associated vertical joints adapt their angles to support the forward



**Figure 21.14** Basalar joint (*vhip\_14*) and distal joint (*knee\_14*) during swing (left) and stance (right) phase.

motion (figure 21.12). As intended for the horizontal joints, a stance or a swing is started and halted smoothly to avoid too-abrupt accelerations and forces during the transitions, which might cause the simulation to stall. A disadvantage might be the changing velocity during a cycle. The smaller the time or the bigger the amplitude for a swing or stance, the faster the rotation of the joint and the bigger the torque that has to be applied.

Figure 21.13 shows the transition from identical start conditions to the phase shift between legs of a group during the initial phase. The sine curve for horizontal motion is stretched according to the leg number. When the phase shift is established, identical patterns are repeated. Swing and stance are clearly distinguished for the vertical joints (figure 21.14). During a stance, the angles are adapted to the motion of the horizontal joint (*hhip\_14*), during a swing, the hip is lifted by  $10^\circ$  and the knee angle is fixed. The adapted vertical angles are relatively small:  $-0.27^\circ$  (*vhip*) and  $3.8^\circ$  (*knee*). They decrease smoothly, drop to zero for  $hhip = 0^\circ$  and increase smoothly to the maximum value again at the end of a stance. Their shape resembles the sine curve of the horizontal angle (figure 21.15).




**Figure 21.15** Stance in the basalar (*vhip\_14*, upper trace) and distal (*knee\_14*, lower trace) joints.

Figure 21.16 shows that during a swing, the leg is lifted by a vertical hip rotation. The knee joint remains in its position. The amplitude depends on the hip lift (*vhip\_up*) and on the maximum (negative) angle during a stance (*vhip\_down*). Again, a sine shape is applied to avoid abrupt accelerations.


The coordination of the joints results in a forward motion (figure 21.17) of Scorpion. At first, the robot accelerates and establishes the motion pattern. The arrow in figure 21.17 marks roughly the point of transition to full speed. Thereafter, the average speed remains rather constant. On the other hand, oscillations can be observed. The oscillation of the motion (figure 21.18) is revealed by observing the total velocity. It alters between 53 mm/sec and 99 mm/sec. The same pattern is repeated every half period. The main cause is the uneven rotation of the horizontal joints. This is superposed by the overall motion pattern of the eight legs (figure 21.19). Minimum speed corresponds to even contributions of both groups to the forward motion (left vertical line) (figures 21.19 and 21.20), whereas maximum speed is reached when one group is in complete stance position.

The sum of the velocities of joints in stance does not vary much compared with the body velocity. It seems that sometimes the joints are






**Figure 21.16** Basalar joint (*vhip\_14*) during a swing movement.



**Figure 21.17** Body displacement relative to an origin over time as a result of a coordinated leg movement.



**Figure 21.18** Undulation of motion caused by friction and the inherent propulsion effects of groups of legs during swing and stance movements.




Figure 21.19 Undulation plotted over the underlying motion pattern.





Figure 21.20 Body velocity (upper trace) vs. sum of joints (in stance phase) velocities (lower trace).



**Figure 21.21** Trace of one foot during a stance movement (left) and the related x and z components with ground contact forces (right).

impeding each other. The minimum body speed meets the minimum sum-of-joints speed quite well, but the maxima lie way apart. Further investigation revealed motion of the feet during stances, although they were supposed to remain fixed to their ground position. In addition, the contact force during ground contact varies remarkably. During a stance phase, the foot is dragged roughly in the direction of motion (figure 21.21).

The motion of Scorpion undulates about the forward direction (figures 21.22 and 21.23). In the acceleration phase (left), asymmetries lead to deviations from the z-axis being the direction of motion. Later on, an equilibrium is established, which keeps the general direction of motion. In the bottom picture of figure 21.22, the lower curve, the mirror ( $-f(y)$ ) of the angular velocity of the center of mass, is a perfect copy of the original function upper curve, shifted half a period. The undulation follows a pattern, which is repeated every half a period with alternating sign.

### Joint Coordination

To enable an efficient stance without up/down or left/right shift, the body should only move in the direction of motion. Therefore, the three joints of a leg have to harmonize (figure 21.24) their rotations at every moment of a stance. The angles of the vertical joints,  $v$  for the vertical hip-joint and  $k$  for the knee joint, shall be expressed as functions of the horizontal joint and its angle  $h$ . The link between the variables can be derived from the starting position with  $h = 0^\circ$ . The entire system is determined by the length of both parts of the leg (*up*, *down*), the initial angles ( $v1$  and  $k1$ ), and the radius ( $r$ ) of the foot. Please note that the initial posture of R4-L3-R2-L1 with *hip\_angle1* and *knee\_angle1* defines  $s$  and  $H$ .





Figure 21.22 Body rotation (y-axis) during forward motion in start (upper trace) and in the later stages of the movement (lower trace).


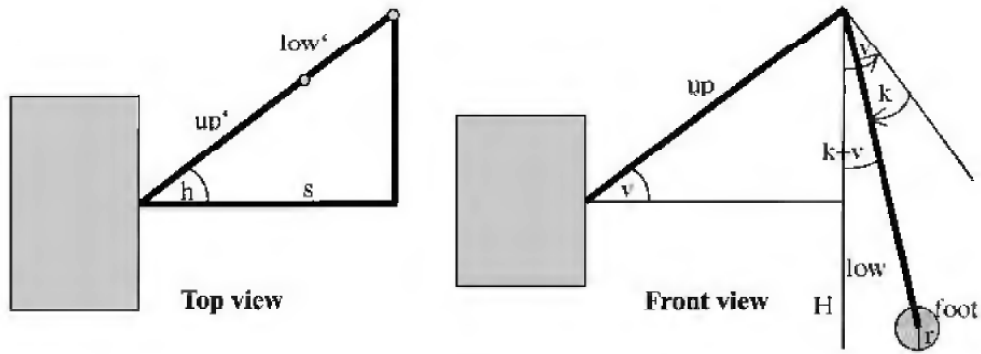


Figure 21.23 The undulation of the center of mass during forward motion.



**Figure 21.24** Harmonizing angle relations among the joints of an individual leg during a stance motion. See also formula below for efficient computation.

The lateral distance between foot and hip is:

$$s = up \cdot \cos(v_i) + low \cdot \sin(k_i + v_i) = \cos(h) \cdot (up \cdot \cos(v) + low \cdot \sin(k + v))$$

$$\sin(k + v) = \frac{s/\cos(h) - up \cdot \cos(v)}{low}$$

where  $v_i$  and  $k_i$  are the vertical angles for  $h = 0$ . The vertical distance between floor and hip is:

$$H = -up \cdot \sin(v) + low \cdot \cos(k + v) + r$$

$$\cos(k + v) = \frac{H + up \cdot \sin(v) - r}{low}$$

Solving the equations for  $v(h)$

$$1 = \sin^2(k + v) + \cos^2(k + v)$$

$$1 = \left( \frac{s/\cos(h) - up \cdot \cos(v)}{low} \right)^2 + \left( \frac{(H - r) + up \cdot \sin(v)}{low} \right)^2$$

$$low^2 = \frac{s^2}{\cos^2(h)} + up^2 \cdot \cos^2(v) - \frac{2 \cdot up \cdot s}{\cos(h)} \cdot \cos(v) + (H - r)^2 + up^2 \cdot \sin^2(v) + 2 \cdot (H - r) \cdot up \cdot \sin(v)$$

$$\frac{2 \cdot up \cdot s}{\cos(h)} \cdot \cos(v) = up^2 - low^2 + (H - r)^2 + \frac{s^2}{\cos^2(h)} + 2 \cdot (H - r) \cdot up \cdot \sin(v)$$

$$\cos(v) = \frac{up^2 - low^2 + (H - r)^2}{2 \cdot up \cdot s} \cdot \cos(h) + \frac{s}{2 \cdot up} \cdot \frac{1}{\cos(h)} + \frac{H - r}{s} \cdot \cos(h) \cdot \sin(v)$$

$$a := \frac{up^2 - low^2 + (H - r)^2}{2 \cdot up \cdot s}$$


For simplification we define:

$$\begin{aligned}
\sqrt{1 - \sin^2(v)} &= a \cdot \cos(h) + \frac{s}{2 \cdot up} \cdot \frac{1}{\cos(h)} + \frac{H-r}{s} \cdot \cos(h) \cdot \sin(v) \\
1 - \sin^2(v) &= a^2 \cdot \cos^2(h) + \frac{s^2}{4 \cdot up^2} \cdot \frac{1}{\cos^2(h)} + \frac{(H-r)^2}{s^2} \cdot \cos^2(h) \cdot \sin^2(v) \\
&\quad + \frac{a \cdot s}{up} + \frac{2 \cdot a \cdot (H-r)}{s} \cdot \cos^2(h) \cdot \sin(v) \\
&\quad + \frac{(H-r)}{up} \cdot \sin(v) \left( \frac{(H-r)^2}{s^2} \cdot \cos^2(h) + 1 \right) \cdot \sin^2(v) \\
&\quad + \left( \frac{2 \cdot a \cdot (H-r)}{s} \cdot \cos^2(h) + \frac{H-r}{up} \right) \cdot \sin(v) + a^2 \cdot \cos^2(h) \\
&\quad + \frac{s^2}{4 \cdot up^2} \cdot \frac{1}{\cos^2(h)} + \frac{a \cdot s}{up} - 1 = 0 \\
\sin^2(v) &+ \frac{(2 \cdot a \cdot (H-r) \cdot up \cdot \cos^2(h) + (H-r) \cdot s) \cdot s}{up \cdot ((H-r)^2 \cdot \cos^2(h) + s^2)} \cdot \sin(v) \\
&+ \frac{s^2(4 \cdot a^2 \cdot up^2 \cdot \cos^2(h) + s^2/\cos^2(h) + 4 \cdot a \cdot s \cdot up - 4 \cdot up^2)}{(4 \cdot up^2) \cdot ((H-r)^2 \cdot \cos^2(h) + s^2)} = 0 \\
\sin_{1,2}(v) &= - \frac{(2 \cdot a \cdot (H-r) \cdot up \cdot \cos^2(h) + (H-r) \cdot s) \cdot s}{2 \cdot up \cdot ((H-r)^2 \cdot \cos^2(h) + s^2)} \\
&\quad \pm \left( (4 \cdot s^2 \cdot up^2 \cdot ((H-r)^2 \cdot \cos^2(h) + s^2) \right. \\
&\quad \quad \left. - 4 \cdot a^2 \cdot s^4 \cdot up^2 \cdot \cos^2(h) - 4 \cdot a \cdot s^5 \cdot up - s^6/\cos^2(h) \right) / \\
&\quad (4 \cdot up^2 \cdot ((H-r)^2 \cdot \cos^2(h) + s^2))^2)^{1/2}
\end{aligned}$$

Because  $up$ ,  $low$ ,  $r$ ,  $H$ ,  $s$ , and their derivative  $a$  are constants,  $v$  depends only on  $h$ . With the knowledge of  $v$ ,  $k$  can be calculated by either of both equations at the beginning of this chapter.

## IMPLEMENTATION

The last section outlined the process of building a fully parameterized model of the scorpion robot. This model was designed so that it matched as closely as possible the real robot we intended to build (e.g., the dimensions of the motor gear units have been determined by simulations carried out with the model). This section is dedicated to the control architecture we designed to implement robust outdoor locomotion in the real robot. The architecture has been evaluated in a set of lab and field experiments that are described in the Results section.



**Figure 21.25** The scorpion robot during an autonomous exploration into a sand bed. The beachlike sand bed was 3 m wide and 9 m long. The robot's feet penetrated the sand for approx. 3–5 cm. A reflex mechanism helped to overcome the obstacle.

### The Control Architecture

Our architecture is based on two approaches to robust and flexible real-world locomotion in biological systems, which seem to be contradictory at first sight. These are the central pattern generator (CPG) model and the pure reflex-driven approach.

A CPG is able to produce a rhythmic motor pattern even in the complete absence of sensory feedback. The general model of a CPG has been identified in nearly every species, even though the concrete instantiations vary among the species to reflect the individual kinematical characteristics in the animals. The idea therefore seems to be very promising as a concept to realize locomotion in kinematically complex robotic systems (figure 21.25), as it resembles the divide-and-conquer strategies that are reflected in nearly all solutions to complex control problems.

Another model for the support of robust locomotion is also provided by evolution in the animal kingdom. This is the concept of reflex-based control (Delcomyn, 1980). A reflex can be viewed as a closed-loop control system with fixed input/output characteristics. In some animals, like the locust, this concept is said to actually perform all of the locomotion control and no further levels of control, like the CPG, are involved (Cruse, 1982). Whether or not complex motion control can be achieved only via reflex systems is subject to further discussion; however, the concept of a set of fixed wired reactions to sensory stimuli is







Figure 21.26 The overall architecture for low level actuation is depicted here.

of high interest to roboticists who aim to gain stability in the system's locomotion.

The design of the control architecture described here was thus driven by these two concepts. The CPG approach appeared to be interesting to generate rhythmic walking patterns, which can be implemented as computationally efficient, while the reflex-driven approach seemed to provide a simple way to stabilize these walking patterns by providing a set of fixed situation-reaction rules to external disturbances and as a way to bias leg coordination among multiple independent legs (Cruse, 1978). Figure 21.26 outlines the general idea.

On the global level (light gray area of figure 21.26), we have implemented locomotion behaviors (LBs—typically forward, backward, and lateral locomotion). These global behaviors are connected to all local leg controllers and activate (with continuous strength) the local (per leg) motion behaviors. At the same time, they implement the interleg phase relation by setting/resetting the local clocks. The local level (dark gray area of figure 21.26) implements rhythmic motion behaviors (RMBs) and postural motion behaviors (PMBs). These behaviors simultaneously influence the amplitude and frequency (figures 21.27 and 21.28) parameters of three oscillating networks ( $OS^T$ ,  $OS^B$  and  $OS^D$ ). The oscillators are connected to a common clock, which is used for local





and global (in relation to other legs) synchronization purposes. The oscillators' output is a rhythmic, alternating flexor and extensor, stimulation signal (see callout box), which is implemented as splined sine waves. This activation signal represents the desired behavioral locomotion, which is translated into PWM signals via the motoric end path. In line with the output of the motoric end path are a set of perturbation-specific reflexes that are implemented as "watchdogs." They override the signals on the end path with precompiled activation signals if the sensor information from the physical joints meets a set of defined criteria.


## RESULTS

This approach was implemented and successfully tested on our Scorpion integration study (see figure 21.27). In figure 21.27 A–C, data of the performance of one leg is shown. The solid line is the real angle of the leg, measured with the motor encoders. The angles for the distal and basilar joints increase during elevation, while the angle for the thoracic joint increases during protraction. The frequency was set to 1.3 Hz (19 time units on the x-axis). The data were taken every 1/25 sec and the curves are directly computed from the raw data. The mean starting position is at 37° for the thoracic joint, 121° for the basalar joint, and about 30° for the distal joint. At first, only a local Forward behavior is stimulated (until  $t = 1375$ ), then the Lateral behavior is activated simultaneously. Because of the equal strength of the activation, the system now tries to walk forward as well as laterally, which results in diagonal walking. At approximately time  $t = 1460$ , the activation of the Forward behavior is set to 0, which leaves only the Lateral behavior to influence the oscillator networks (see figure 21.26). Thus the system walks laterally, which can be observed from the data as when the amplitude of the thoracic joint is 0 while the basalar and especially the distal joints perform large amplitude oscillations. Subsequently, the above described process is reversed.

### Local Reflexes

The described generation of rhythmic motion deals very well with plain surfaces without obstacles. The rhythmic locomotion generated by the oscillators can be seen as a kind of plan to walk stably if there


**Figure 21.27** Traces A through C show the movements of the thoracic (A), the basalar (B) and the distal (C) joints during a transition from forward walking (pure FORWARD activation), to diagonal (equal activation of FORWARD and LATERAL), to lateral walking (pure LATERAL activation) and back to forward walking. See figure 21.26 for the pathways of activation.



**Figure 21.28** Frequency change at the thoracic joint. Frequency increases from left to right, thereby increasing the overall speed of the robot.

are no or little perturbations. In the case of an obstacle blocking the normal swing phase or in the case of uneven ground, poor results would be expected.

Our approach to deal with uneven terrain was to implement a set of reflexes in parallel to the motor end path (see figure 21.26) that override, for a short and predefined period of time, the rhythmic activity of the oscillators (e.g., figure 21.29)—if the current values of the thoracic joint increase steeply and a significant angular displacement error is detected at the same time, it is assumed that the “planned” trajectory is blocked. This triggers a reflex, which moves the leg backward and upward (via joint activity in the thoracic and basalar joints) and then forward at maximum speed. This reflex is illustrated in figure 21.29 A–C (start point at  $t = 46$ , stop at  $t = 58$ ). The reaction time of the reflexes are as fast as  $1/100$  sec because they are directly in line with the motor-control signals. The three pictures (A–C) also illustrate how fast the motor controller returns to the pattern given by the oscillator after the reflex is no longer active. It is important to notice that the action of the reflex does in fact sit right on top of the ongoing rhythmic activity. As can be seen in figure 21.29 A and B, the oscillatory activity is still present in the background (light gray lines). As soon as the reflex action is terminated, the locomotion returns to the oscillation.



**Figure 21.29** A reflex initiated at a leg during a course through a rock bed. The current in the thoracic joint (trace C) increased as a result of the obstacle blocking the way. At the same time the angular displacement error (trace A) in the thoracic joint increased, indicating an exception in the regular swing cycle. As a result of these factors the basalar joint controller initiated the reflex (trace B) that elevated the leg further, thereby overcoming the obstacle.

## OUTLOOK

We intend to further develop the system presented here, with respect to the mechanism, electronics, and also the software architecture. In the field of mechanics, we are currently in the process of redesigning the legs. We intend to further decrease the weight (approx. 150 g/leg) thereby increasing the weight/lift ratio to 1:9 without losing stability or protection against environmental influences. This will be achieved by means of optimizing the design in addition to material changes. Some of the parts have proven to be overdesigned in aluminum,<sup>5</sup> therefore those parts can now be manufactured in POM, which has a 1:1.4 density advantage. In the electronics, we are also in the process of redesign. The new design will incorporate the Motorola MPC555 as the master controller and an FPGA, which will be programmed to perform the tasks of the 8 local (C164) microcontrollers. This design will improve the motor control, as it allows for higher sampling speed and more-sophisticated controller design (e.g., pid instead of only proportional controllers).

On the MPC555 we will be able to explore the use of more complicated algorithms and even learning and optimization procedures due to the integrated FPU and the higher clock speed. The biggest effort however, has to be put on the autonomous operation of the system, which has to go beyond low-level locomotion. Here we are planning for high-level task planning and scheduling algorithms as well as on-line learning and adaptation methods. We intend to incorporate a separate vision-processing module, which will analyze the video taken onboard and evaluate it according to object recognition and classification procedures. These high-level data will then be transferred to the master controller (MPC555), which can incorporate the information in the higher-level task planning and navigation procedures.

## ACKNOWLEDGMENTS

This research is sponsored by DARPA under research grant #N00014-99-0483.

## NOTES

1. These are also called behaviors, which are just useful reactions to given sensorized stimuli.
2. CPLDs are programmable gate arrays—hardware that can be configured to implement finite state machines.
3. Also referred to as *h\_hip\_xn* for the thoracic and *v\_hip\_xn* for the basalar joints (with  $x = [R, L]$  and  $n = [1, \dots, 4]$ , denoting right and left legs with number 1 through 4).

4. Note that these values are closely reflected in the real prototype in order to achieve a good comparability between simulation and real system.
5. This was not initially evident, which is the reason we first built in aluminum.

## REFERENCES

- Ayers, J., Kashin, S., Blidberg, D. R., and Massa, D. (1994). Biologically based underwater robots. *Unmanned Syst.* 12: 30–36.
- Blickhan, R., and Full, R. J. (1998). Similarity in multilegged locomotion, bouncing like a monopode. *J. Comp. Physiol. [A]* 173: 509–517.
- Bowerman, R. F. (1975). The control of walking in the scorpion: I. Leg movement during normal walking. *J. Comp. Physiol. [A]* 100: 183–196.
- Brooks, R. A. (1986). A robust layered control system for a mobile robot. *IEEE J. Robotics Automation* 2: 14–23.
- Calvitti, A., and Beer, R. D. (2000). Analysis of a distributed model of leg coordination: I. Individual coordination mechanisms. *Biol. Cybern.* 82: 197–192.
- Cruse, H. E. S. (1982). Peripheral influences on the movement of the legs in a walking insect *Carausius morosus*. *Exp. Biol.* 101: 161–170.
- Cruse, H., et al. (1998). Walknet—A biologically inspired network to control six-legged walking. *Neural Netw.* 11: 1435–1447.
- Delcomyn, F. (1980). Neural basis of rhythmic behavior in animals. *Science* 210: 492–498.
- Fearing, R. S., Chiang, K. H., Dickinson, M., Pick, D. L., Sitti, M., and Yan, J. (2000). Wing transmission for a micromechanical flying insect. In *IEEE International Conference on Robotics and Automation*.
- Hebb, D. (1949). *The Organization of Behavior (A Neuropsychological Theory)*. New York: Wiley.
- Kirchner, F. (1995). Automatic decomposition of reinforcement learning tasks. In D. Cohn and D. Lewis (eds.), *Proceedings of the AAAI '95 Fall Symposium Series on Active Learning*, 10–12 November, Massachusetts Institute of Technology, Cambridge. AAAI Press, pp. 56–59.
- Kirchner, F. (1998). Q-Learning of complex behaviors on a six-legged walking machine. *J. Robotics Autonomous Syst.* 25: 256–263.
- Kupferman, I., and Weiss, K. R. (1978). The command neuron concept. *Behav. Brain Sci.* 1: 3–39.
- Spenneberg, D., and Kirchner, F. (2000). Omnidirectional walking in an eight legged robot. In *Proceedings of the International Symposium of Robotics and Automation (ISRA2000)*, Monterrey, Mexico.
- Steels, L. (1994). A case study in the behavior-oriented design of autonomous agents. In D. Cliff, P. Husbands, J.-A. Meyer, and S. W. Wilson (eds.), *From Animals to Animats 3, Proceedings of the Third International Conference on Simulation of Adaptive Behavior*. Cambridge: MIT Press/Bradford Books.
- Stein, P. S. G. (1978). Motor systems with specific reference to the control of locomotion. *Annu. Rev. Neurosci.* 1: 61–81.
- Whitehead, S., and Ballard, D. (1990). Active perception and reinforcement learning. In *Proceedings of the seventh International Conference on Machine Learning*.

**This page intentionally left blank**

# *Flying Robots*



**This page intentionally left blank**

## The Organization of Insect Locomotion Systems and Computer-Based Flight Control in the Tobacco Hawkmoth *Manduca sexta*

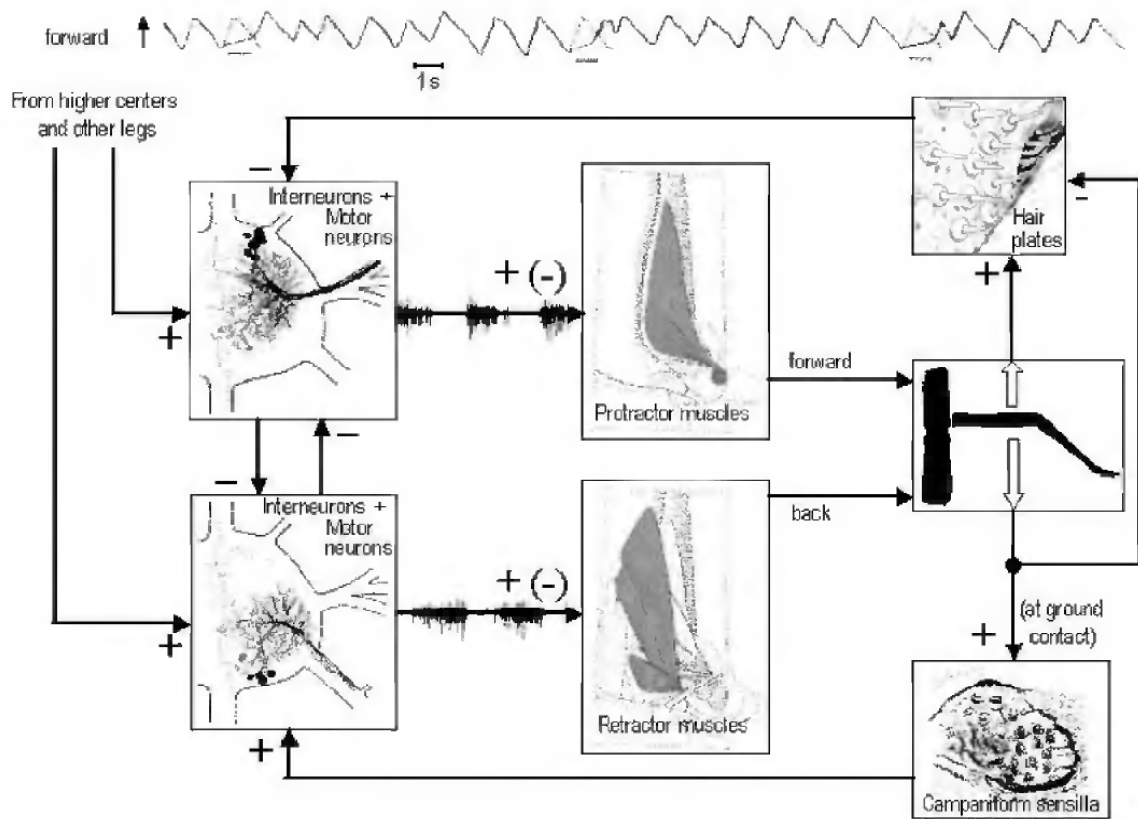
Gernot Wendler

As the design and the construction of nature-inspired locomotory robots were the main topics of the Neurotechnology for Biomimetic Robots symposium, it may be helpful to present a brief outline of the organization in natural locomotory systems—especially in insects. The first part of this chapter is therefore dedicated to a brief and very simplified comparison of the organization, motor control, and signal processing in insect walking and flight systems.

### WALKING

Insect walking systems consist of six coupled relaxation oscillators, one for each leg (see also chapter 19 of this volume). The evidence for the existence of six oscillators results from observations in which the legs of an insect eventually step at different frequencies (e.g., von Holst, 1943; Wendler, 1964; Bässler, 1995). In our current understanding, each of the oscillators, the step generator, is a distributed system. It consists of intraganglionic interneurons, motoneurons, muscles, tendons, the leg with its mechanical properties, and a variety of leg mechanoreceptors.

Figure 22.1 illustrates this organization in a simplified manner for the proximal joint in a single leg of the stick insect *Carausius morosus*. Extracellular recordings from the protractor and retractor nerves in a walking animal reveal a rhythmic spike activity. These nerves contain several motoneurons that produce different spike amplitudes. Each leg muscle is innervated by more than one motoneuron, and, depending on the muscle, their number is between 3 and 14. Each has a distinct and different function, including inhibition of contraction and long-term facilitation. Furthermore, individual muscle fibers receive different combinations of motoneurons. The control of muscle contraction is therefore extremely complicated (see Wendler, 1999 for a review). The resulting leg movements are monitored by a number of sense organs. These sense organs play an essential role in that a full step with coordinated movements of all joints requires their feedback. They influence a network of interneurons and motoneurons. In this network,



**Figure 22.1** Elements of a step generator in the stick insect. The figure is confined to some essential connections for pro- and retraction in the proximal joint. The main retractor muscles and protractor muscles are both situated anterior to the leg. Gray: nonspiking interneurons (Büschges and Wolf 1995). Running through the connectives: a descending visual interneuron (courtesy of J. P. Jander). The negative signs in brackets refer to the action of inhibitory neurons. Experiments with pilocarpine treated and de-afferented stick insects indicate that the step generator can be subdivided into separate oscillators for each leg joint (Büschges, Schmitz, and Bässler, 1995; Bässler and Büschges 1998).

nonspiking interneurons play a major role (Büschges and Wolf, 1995; Wolf and Büschges, 1995). These neurons operate in a pure continuous, or in other words, an analog fashion. Signal transfer between the other elements of the walking system is also mainly analog—in particular, by frequency coding—because graded muscle contraction is largely dependent upon the momentary spike rate of motoneurons.

When designing a walking robot, it is not necessary or even feasible to adopt all properties of insect walking systems, such as action potentials and muscle control by their frequency or the extremely complex muscle innervation. It is, however, very useful to remember that muscles normally consist of a number of muscle fibers, each of them being very weak. Depending on the specific muscle task, these fibers are arranged in specific ways (e.g., in a pinnate arrangement to produce a greater tension on the tendon). Muscle control must be able to produce

an extremely wide range of muscle tensions and contraction speeds. Different tensions are important to overcome obstacles or to walk uphill. As long as an insect does not jump, it can only increase its speed by increasing the step frequency (the observed changes in gait have other consequences), because the leg is mechanically coupled to the ground during the power stroke, and its range of movement is limited. A range of contraction speeds is therefore absolutely necessary in order to be able to walk at different speeds (note the completely different situation in flight).


It is also necessary to study the role of all leg mechanosensors carefully. Hair plates and chordotonal organs monitor joint position and speed, and campaniform sensilla measure direction and amount of strains in the cuticle. Their signals are required to control mechanical stability and the appropriate motor output on uneven grounds. Equivalent signals are expected to be indispensable in walking robots.

## FLIGHT

Most (see also legend to figure 22.2) insect flight systems operate in a different way, as compared to walking systems. The main power muscles do not move the wings directly. Instead, they move the wings by rhythmic deformations of the thorax (figure 22.2A). When the large dorsolongitudinal power muscles contract, the segment is shortened, and the tergum bends upward. This causes a downward rotation of the wing about the wing hinge. A contraction of the dorsoventral muscles move the wings upward. Only some of the steering muscles are directly attached to the wings. Most power muscles or their separate portions are innervated only by a single motoneuron. These general features are common to the two main functional groups of insects that are able to fly: the *neurogenic* and the *myogenic* fliers. Wing-beat frequencies range from about 15–35 Hz in neurogenic fliers to about 30–1000 Hz in myogenic fliers.

### Neurogenic Flight

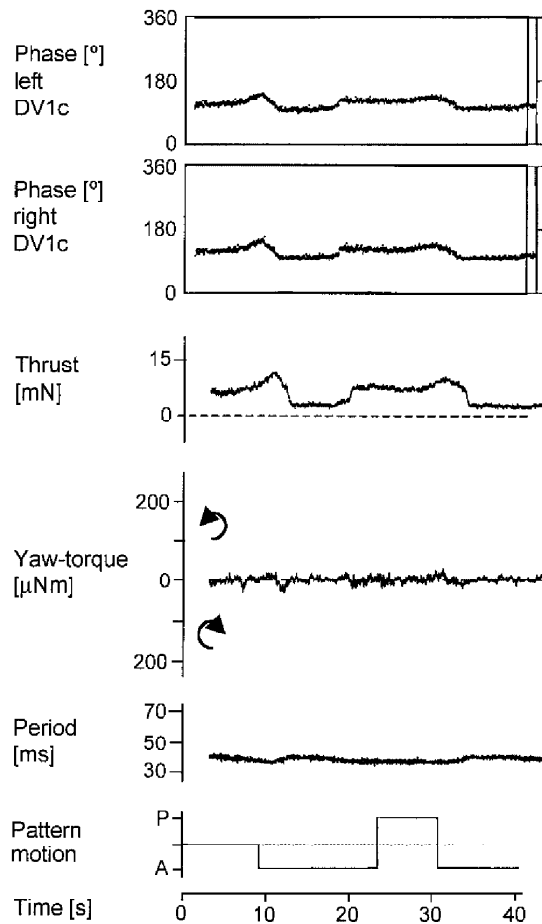
In neurogenic fliers, the wing beat is under full neuronal control (figure 22.2B). A spike in a motoneuron of a main power muscle always elicits a full contraction and produces no graded contraction. Because muscle contraction is always maximal, the neurogenic fliers need other control mechanisms than walking insects. It is mainly the *phase of activation* between muscles, which is defined as the position of this action potential within the period of another (figure 22.2B). Figure 22.3 shows an example. A tethered flying tobacco hawkmoth, *Manduca sexta*, is made to vary its thrust by optomotor stimulation. Thrust is reduced when the action potentials of a main elevator muscle (DV1c) of right and left side



**Figure 22.2** Indirect effects of flight power muscles on wing movements and control of muscle contraction in neurogenic and myogenic fliers. (A) Lateral view of a wing bearing segment. Contracting muscles are dark. The wing hinge has a very complicated anatomy. (B, C) Simultaneous measurements of muscle potentials and wing flapping, both on the same time axis. The Odonata are neurogenic fliers. However, their muscular morphology is different from the drawing in (A). After Pfau 1991 and Wendler 1999.

occur early within the period of the main depressor muscle. The later they occur, the higher the resulting thrust. There is a striking correspondence between the phase change and the change of thrust. Similar results from other muscles also indicate that the prevailing control mechanism for flight maneuvers is the variation of phase between spikes.


The overall organization of a neurogenic flight system is shown in figure 22.4. A number of interneurons produce a basic rhythmic motor



**Figure 22.3** Variations of muscle coordination at flight maneuvers in a neurogenic flier (*Manduca sexta*). The moth is ventrally tethered to a two-component flight balance (Dombrowski 1989). Optomotor stimulation by moving gratings in anterior (A) or posterior (P) direction elicits changes in thrust, but not in yawing. Definition of phase is shown in figure 22.2B. Each point in the phase diagrams is the phase of a muscle potential of a dorsoventral muscle within the period of the left dorsolongitudinal muscle.

pattern. They are all spiking. No nonspiking interneurons have been found. They activate motoneurons, which in turn activate muscles and produce the wing movements. Their effect is measured by sense organs at several levels. Some receptors monitor the movements of each of the four wings. One example: The stretch receptor is activated during the upstroke with one to several spikes. Any disturbance results in a changed timing of all wing receptors. The same is true for the wind hairs, which monitor the rhythmic air stream that is produced by the wings. All these receptors are part of the flight motor, and within this flight motor signal, the transfer is mainly based on the relative timing of events—in other words, phase changes on the basis of a carrier frequency, the carrier frequency being the basic wing-beat frequency.

Other feedback systems are not part of the motor itself, and here signal transfer is analog (graded versus phasic) as in the walking



**Figure 22.4** Organization of the flight system in a neurogenic flier, the locust. White: sketch of the morphology of an element. Gray: input signals. Light gray: output signals. In the thoracic ganglia the motoneuron 113 is shown, which innervates a hindwing elevator muscle, and the interneuron 301, which takes part in the production of the flight rhythm and connects both wing bearing segments (see Burrows 1996 for a general review). (After Wendler 1999.)

system. However, in order to understand how the animal controls the output, one should concentrate on the control of the timing of spikes. The functional separation into power muscles and steering muscles is incomplete in neurogenic fliers. Power muscles such as the dorso-longitudinal and the dorsoventral muscles provide most of the power. However, they also change their relative timing in flight maneuvers. Basalar and subalar muscles have both functions, while the PD2 muscles have mainly steering functions. In the locust, a corresponding gradation can be measured in threshold and latency of muscle reaction to stimulation of campaniform sensilla on the wings (Wendler, 1982).

### **Myogenic Flight**

Myogenic fliers use stretch activation, a property of their power muscles. They respond with an active contraction to a rapid mechanical stretch. When the dorsolongitudinal muscles (DLMs) contract and bend the tergum upward, the dorsoventral muscles (DVMs) will be stretched by the distortion of the thorax. This stretch causes a single contraction of the DVMs, which will stretch the DLMs, whereupon they will contract, and so on. This system operates in *resonance*, and the motor neurons fire at rates much less than the wing-beat frequency (figure 22.2C). Their activity controls the  $\text{Ca}^{++}$  ion concentration in the muscle, which allows the stretch activation to occur.

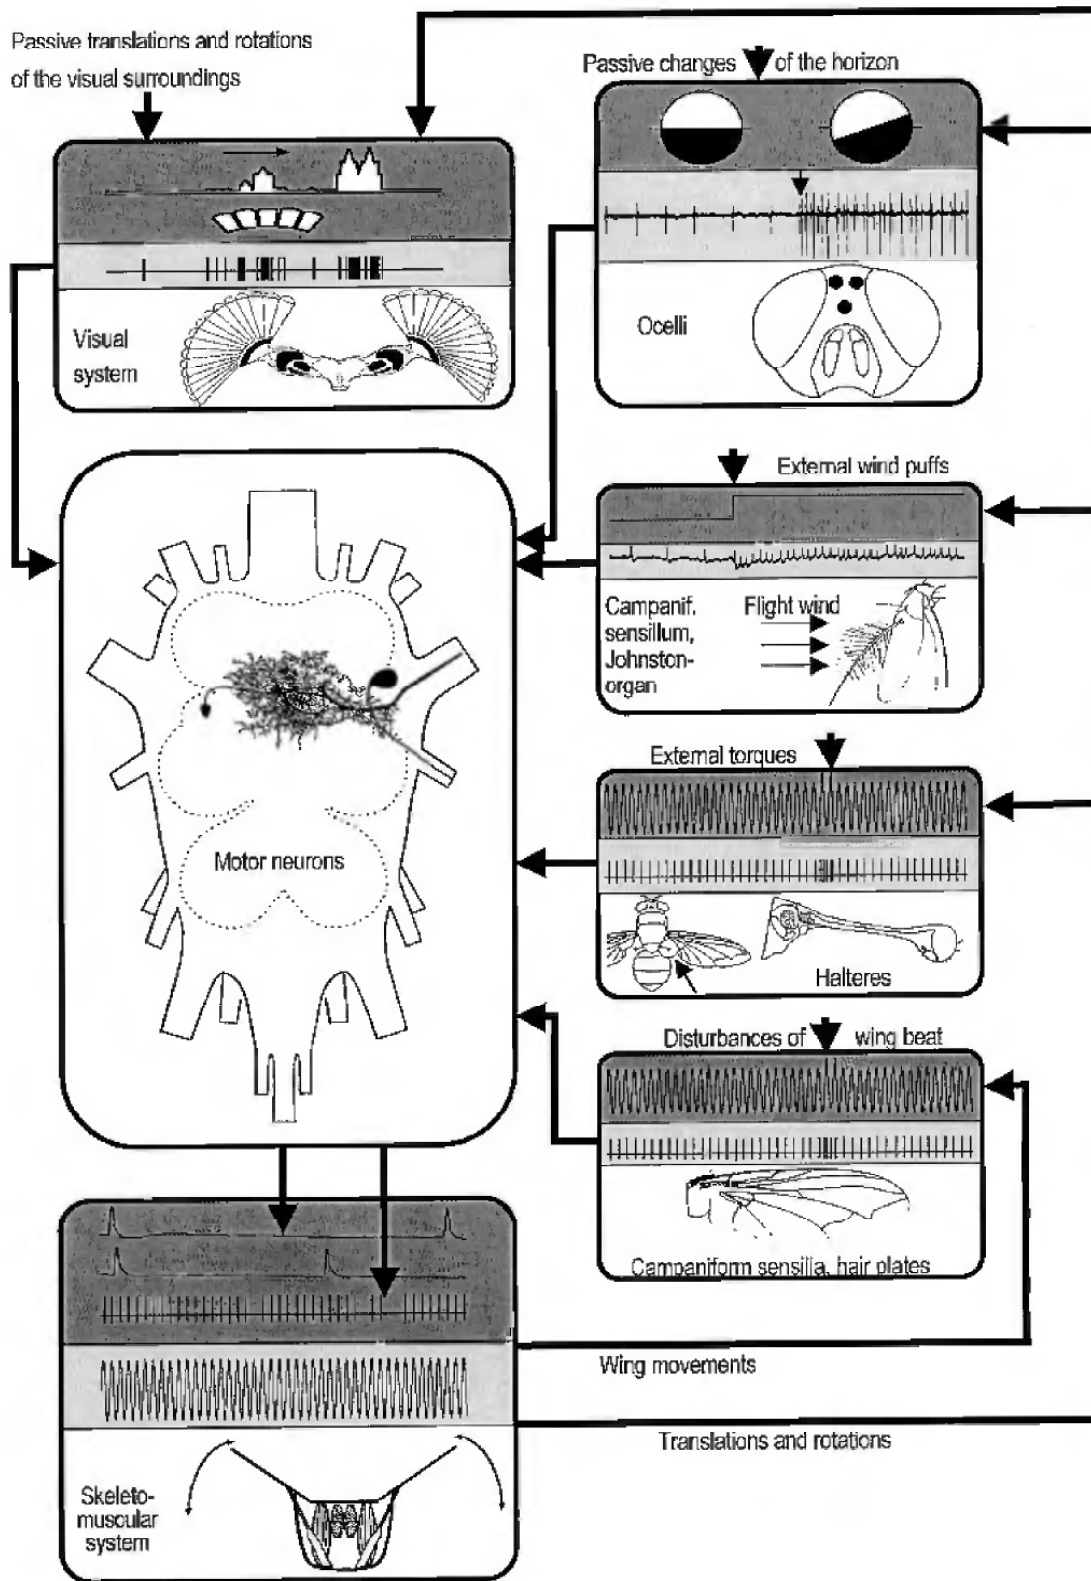
Myogenic fliers such as flies have developed a complete separation between power muscles and steering muscles. Flight maneuvers are controlled by separate steering muscles, which are activated by motoneurons continuously or only during the flight maneuver in wing beat frequency. In the basalar muscle b1, it has been shown that the phase of neuronal activation is changed at flight turns, and this phase controls the stiffness of the muscle (Tu and Dickinson, 1994).

The overall organization of a myogenic flight system is shown in figure 22.5. Despite the differences in the generation of the motor output, there are important similarities in the organization of neurogenic and myogenic fliers in so far as feedback control systems are involved at different levels of the system. As an example, the halteres of the dipterans act as gyroscopes. The sensory input from their campaniform sensilla guarantees flight stability and furthermore is the source of the wing-beat-coupled activity of motoneurons that control the steering muscles.

### **Flight Control in the Hawkmoth *Manduca sexta***

The hawkmoth *Manduca sexta* is able to perform magnificent flight maneuvers such as hovering in front of a blossom or elaborate odor-directed orientation flights. It is a neurogenic flier like the locust; but





**Figure 22.5** Organization of the flight system in a myogenic flier, *Calliphora erythrocephala*. Symbols as in figure 22.4. Two motoneurons are shown in the thoracic ganglia: the DLM5, which innervates a unit of the dorsolongitudinal muscle, and the DVM3, which innervates a dorsoventral muscle (Schlurmann 1997).

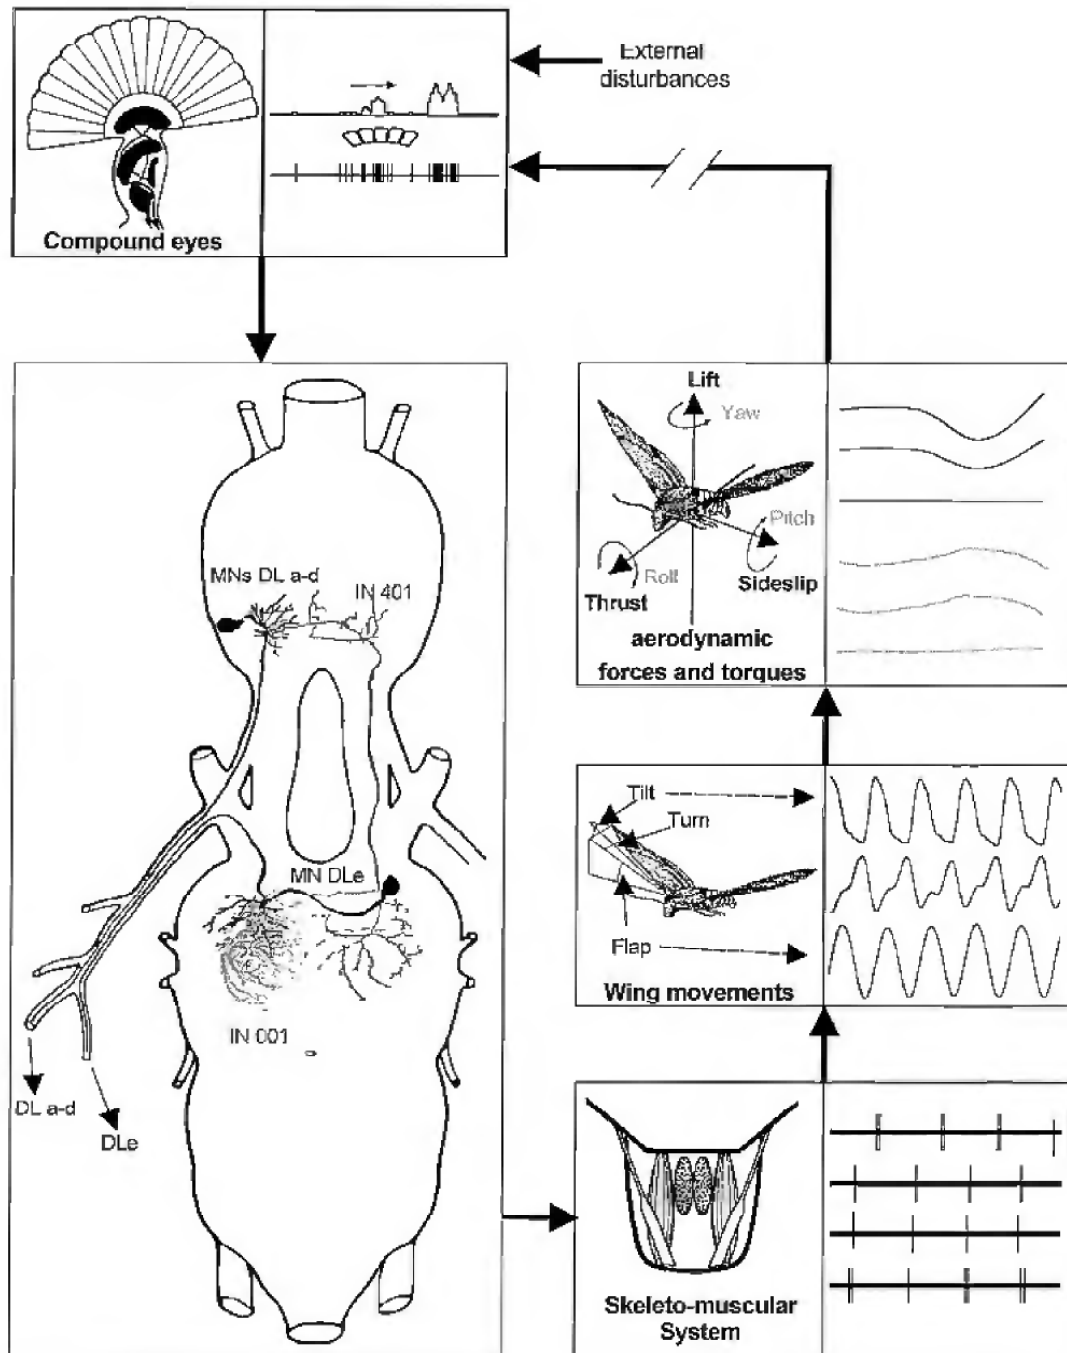
unlike locusts, which actively control the four wings for flight, the hind wings in *Manduca* are mechanically coupled to the front wings. The hindwing muscles are very small and weak. *Manduca* is therefore a functionally two-winged insect. This greatly simplifies the analysis of the neuronal control of flight maneuvers.

**Analysis of Motor Output Patterns** The strategy of the analysis is indicated in figure 22.6. The animal is tethered. Both eyes receive optomotor input from moving black and white gratings, which simulates passive roll or yawing, or changes of lift or height (Scharstein, 1989). These stimuli are presented in order to elicit reproducible flight maneuvers. Because the resulting flight maneuvers have no influence upon the visual input, it is an open-loop situation. The flight maneuvers are simultaneously monitored at three levels: (1) the motor output, (2) the resulting forewing movements, and (3) the resulting aerodynamic forces and moments.

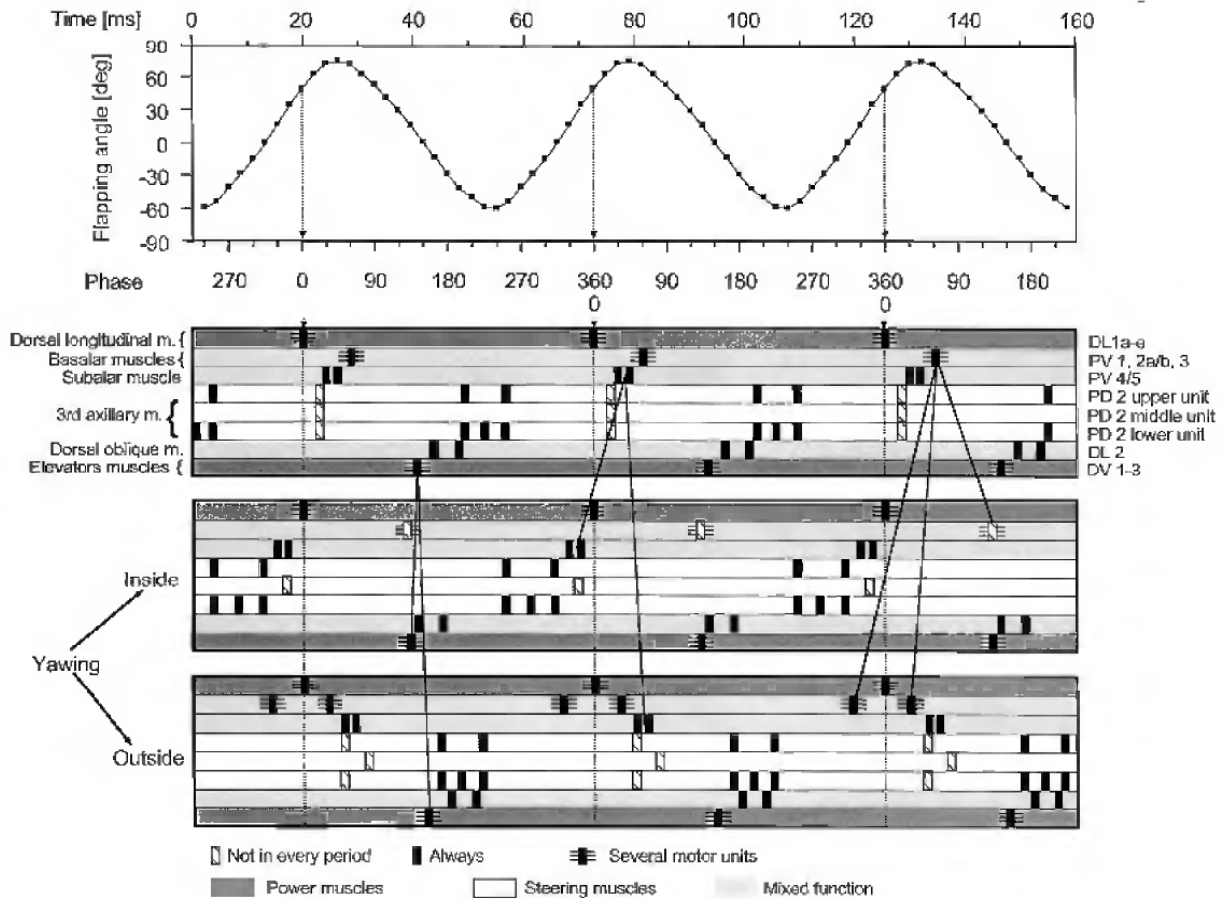
1. The motor output. The muscle activity pattern, the “score,” is derived for each flight maneuver from myograms from about 12 flight muscles. Figure 22.7 shows examples for straight flight and for a moderate turn. In straight flight, the elevator muscles are activated approximately in antiphase to the depressor power muscles. During a turn, the elevators of the wing on the inside are activated slightly earlier, indicating a reduction of thrust on this side (see figure 22.3). The PV4/5 shifts its phase forward on the inner side and backward on the outside, whereas both basilar muscles (PV1, 2a/b, 3) shift in the opposite direction. The PV muscles have a mixed function as wing depressors and steering muscles. The steering muscles (PD2), which correspond to the third axillary muscles in the locust, undergo dramatic changes that we do not yet understand.

2. Simultaneously, the three resulting parameters of the wing movements are measured (figure 22.6): flapping, turning, and tilting. We use a system of miniature angle detectors, mounted on the wing bases of the forewings (Dierkes, 1995). It consists of very small induction coils that move in two oscillating magnetic fields of 150 and 180 kHz, produced by two pairs of Helmholtz coils of 30 cm diameter.

3. The wing oscillations produce aerodynamic forces, which result in movements of the animal. Because a flying object has 6 degrees of freedom, 3 of rotation and 3 of translation, we measure the forces and moments with a 6-component flight balance (von Falkenhausen and Wendler, 1995). The animal is tethered in the middle of a floating system that allows only very small movements (maximal 0.5 mm and 0.2°). The restoring forces for lift, pitch, and roll are provided by dashpots, and for thrust, sideslip, and yaw by a set of springs. Underneath the animal and attached to it, a system of small mirrors and prisms is



**Figure 22.6** Overview of the parts of the *Manduca* flight system that are used for the analysis. The isolated nervous system can produce a long lasting flight rhythm (see Wendler 2000 for a review).

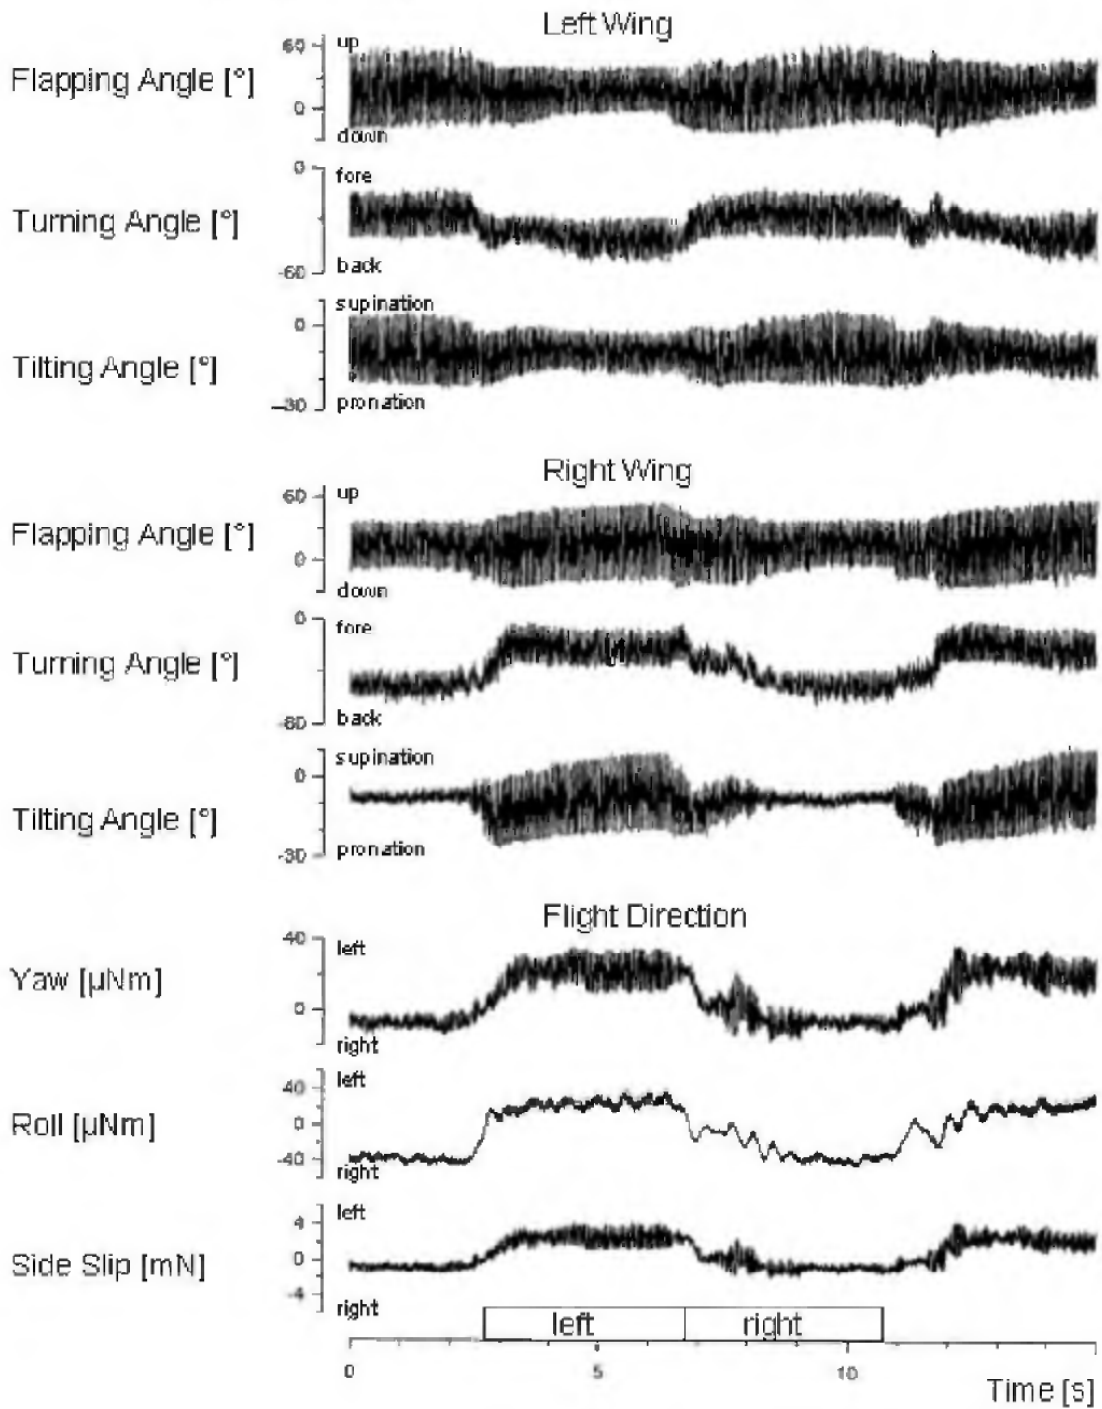


**Figure 22.7** Motor output patterns in the tethered *Manduca*. (Top) Flapping angle of a front wing from a high speed film by U. Dombrowski. One period is repeated three times. (Middle) Relative positions of muscle potentials at straight flight. Only one side is represented. Gray background: power muscles; white: steering muscles; light gray: muscles with both functions. The phases of muscle potentials are mean values from many flights. (Right) Nomenclature for *Manduca*. (Left) Corresponding nomenclature for the locust. (Bottom) Output patterns at moderate yawing.

mounted, which deflects the beams of 3 laser sources. The deflections, caused by the flying *Manduca*, are measured by 6 linear photodiodes.

Figure 22.8 shows an example of an experiment with yaw stimuli with measurements of the resulting wing movements and three of the six parameters that consistently change in yawing experiments. When the optomotor input moves to the left, the animal produces yaw-torque and roll to the left, known as bank turn. Furthermore, a significant sideslip occurs, while thrust remains constant (not shown). This maneuver can be often seen in free flight. It is produced by the following changes in wing movements: At turning to the left, the left wing reduces flapping and tilting amplitudes. Then, the animal moves the left wing backward. At the same time, the wing on the right side performs the opposite changes. This result is not trivial, because such asymmetrical maneuvers

## Wing Movements During Curve Flight



**Figure 22.8** Wing movements and aerodynamic output at yawing stimuli. The high frequency oscillations result from the wing strokes.


could, in principle, be performed by only one wing. When the animal turns to the right, the wings exchange their roles.

When observing the animals during yawing, the most striking effect is the asymmetrical positioning of the wings. The wings on the inner side are shifted far backward, while the other wings are drawn forward. Sometimes, forewing and hindwing on one side are mechanically decoupled. In such cases, it becomes obvious that the shifts are mainly due to shifts of the forewing. The result is a drastic reduction of total wing area on the inner side and an increase on the outside. This effect has already been predicted on the basis of measurements in the non-flying *Manduca* (Wendler, Müller, and Dombrowski, 1993). It raises the questions: To what extent does the normal mechanical coupling between fore- and hindwing allow such relative shifts and to what extent do the hitherto unknown mechanisms contribute to flight maneuvers? Further observations reveal drastic postural changes of the abdomen. The role of such abdominal movements has been studied extensively in locusts. Head movements are also monitored during our experiments. The head reliably turns and rolls in the direction of the movements of the whole animal.

Because all measurements are made simultaneously, it is possible to correlate the motor pattern with the pattern of wing movements and with the aerodynamic output for each flight maneuver. However, such correlations can only be made between complex patterns as a whole. Because most other muscles also change their phase at the same time, the contribution of a specific, individual muscle to a flight maneuver remains unknown. It is also not possible to draw any conclusions from the amount of phase shifts in a single muscle.

The classical way to solve this problem would be to stimulate an individual flight muscle electrically in a resting animal and to observe the effect on the wing movements. During flight, however, the effect of a muscle contraction will depend upon the momentary state of the whole system, which oscillates. In other words, the effect is context dependent. Therefore, one would like to gain full control over this context.


**Computer-Controlled Motor Output** In order to gain control over the motor output, the whole ventral nerve cord is removed, while thorax and wings are left intact. The main *power muscles* are then stimulated by a pattern of computer-controlled electrical pulses of variable amplitude and duration. In figure 22.9 (top), the dorsolongitudinal muscles and the elevator muscles on both sides are stimulated at such a low frequency that the contraction effects of both muscle groups are separated in time. About 8 msec after depressor stimulation, the wing moves downward. The active depression lasts about 23 msec. Afterward, the wing rebounds passively to its resting position due to the elasticity of the thorax. Elevator stimulation has similar effects. At the stimulus



**Figure 22.9** Elastic rebound at electrical stimulation of main power muscles. The stimulation is controlled by a LabVIEW 4.0-based computer program (Weber and Wendler 1996). Steel wire electrodes of 50  $\mu\text{m}$  were chronically implanted into the muscles. Frequency, number of pulses per burst, delay or phase (precision of 100  $\mu\text{s}$ ) and pulse width (1–10 ms) are adjustable independently for each of the eight channels.

frequency of 15 Hz, the active movements of depressor and elevator muscles come closer together and add to the elastic restoring forces. At higher stimulus frequencies, the active movements overlap: The antagonistic muscles act against each other, which results in a decrease of the flapping amplitude.

The frequency response of wing movements in the de-efferented moth is shown in figure 22.10. The variation of the stimulus frequencies reveals maximum amplitudes at 15 Hz for flapping, 17 Hz for turning, and 20 Hz for tilting. There is also a sharp resonance of the thorax at 15 Hz, which can be determined by calculating the frequency spectrum of passive wing oscillations after being released from a deflected position. The experiments have been carried out at room temperature (23°C). We expect that higher temperatures will shift the resonance frequencies to match the wing beat frequency of the intact animal (about 25 Hz) because of its increased body temperature of about 36°C during flight (Heinrich, 1971), which shortens the contraction times. *Manduca* there-




**Figure 22.10** Frequency response in the de-efferented hawkmoth. The experiments were performed with increasing and immediately following decreasing stimulus frequencies. Phase of elevator stimulus in depressor stimulus  $144^\circ$ .

fore probably flies at or near the resonance frequency of the thorax including muscles, wings, and the surrounding air. Flying at resonance allows the most efficient use of energy (e.g., Scharstein, 1998). The flight systems of *Manduca* and perhaps the majority of neurogenic fliers are therefore not as different from the myogenic systems as was believed for a long time.

The results raise a number of new questions. The wing beat is fully controlled by the nervous system in neurogenic fliers. How is the frequency match between the neuronal and mechanical parts of the system achieved? Here, new aspects arise related to the role of sense organs and of small muscles that might modify the mechanical resonance.

Although only the main power muscles are stimulated, the resulting wing movements are similar to normal wing beats including turning and tilting movements of the wings. The basic wing path is therefore produced by the morphological arrangement of the main power muscles and the geometry of the wing articulation and does not require additional activation of steering muscles.





**Figure 22.11** The role of the PV4/5 muscle in the de-efferented hawkmoth. Wing paths elicited by additional (to situation in figure 22.9) stimulation of the subalar muscle on one side. The movement of the corresponding wing is projected to the surface of a sphere with the wing hinge in its centre. The muscle is activated 18 degrees of phase before or 25.2 degrees after stimulation of the dorsolongitudinal muscles. At +18 degrees, the path is 8-shaped, as has been measured in the intact animal and in the score on the wing on the outside during a turn.

Stimulation of *steering muscles* reveals their function. Figure 22.11 shows an example, in which the subalar (PV4/5) muscle has been stimulated in addition to the main power muscles. In straight flight, this muscle is activated slightly after the dorsolongitudinal muscle (figure 22.7). This activity is shifted forward on the inner side during yawing, and backward on the outer side. Computer-controlled phase variations of the muscle on one side in the same range results in drastic changes of the wing path as we find during turns in the intact animal.

In the next step of the analysis, the de-efferented preparation will be mounted onto the six-component flight balance. Here, we can learn to manipulate the motor system in order to obtain the desired flight maneuvers. We know that *Manduca* is able to perform all flight maneuvers that a flying robot should be able to perform. One of the advantages of this preparation is that we already know its capabilities, whereas the design of a robot does not necessarily guarantee the desired performance. One should also keep in mind that all flying insects make extensive use of sensory feedback, not only for the spatial orientation, but in order to stabilize the flight. Flight stability of a flying robot will also depend upon such feedback.

## ACKNOWLEDGMENTS

I am indebted to R. DiCaprio and to J. Schmidt for constructive discussions. R. DiCaprio corrected the English text.

## REFERENCES

- Bässler, U. (1995). Konzepte der Ethologie im Licht neurophysiologischer Forschung an Stabheuschrecken. *Biologie in unserer Zeit* 25: 170–177.
- Bässler, U., and Büschges, A. (1998). Pattern generation for stick insect walking movements—Multisensory control of a locomotor program. *Brain Res. Rev.* 27: 65–88.
- Burrows, M. (1996). *The Neurobiology of an Insect Brain*. Oxford: Oxford University Press.
- Büschges, A., Schmitz, J., and Bässler, U. (1995). Rhythmic patterns in the thoracic nerve cord of the stick insect induced by Pilocarpine. *J. Exp. Biol.* 198: 435–456.
- Büschges, A., and Wolf, H. (1995). Nonspiking local interneurons in insect leg motor control: I. Common layout and species-specific response properties of femur-tibia joint control pathways in stick insect and locust. *J. Neurophysiol.* 73: 1843–1860.
- Dierkes, S. (1995). Three-dimensional measurements of the wing movements and six-component measurement of the aerodynamic resultants during flight of the tobacco hawkmoth *Manduca sexta* (L.). In *Göttingen Neurobiology Report*. Stuttgart: Georg Thieme Verlag, p. 211.
- Dombrowski, U. (1998). Analysis of visually guided behavior in *Manduca sexta*. In *Neural Mechanisms of Behavior*. Stuttgart: Georg Thieme Verlag, p. 168.
- Dombrowski, U., Milde, J. J., and Wendler, G. (1990). Visual control of compensatory head movements in the sphinx moth. In F. G. Gribakin, K. Wiese, and A. V. Popov (eds.), *Sensory Systems and Communication in Arthropods*. Basel: Birkhäuser Verlag, pp. 127–133.
- Falkenhausen, C., von (1995). Eine 6-Komponenten-Flugwaage zur Bestimmung der kinetischen Parameter des Insektenfluges. Dissertation Köln.
- Falkenhausen, C., von, and Wendler, G. (1995). A new six-component flight balance: Correlation between muscle activity, aerodynamic output, and optomotor stimulation of the tobacco hawkmoth, *Manduca sexta* (L.). In N. Elsner, and R. Menzel (eds.), *Proceedings of the twenty-third Göttingen Neurobiology Conference*. p. 211.
- Heinrich, B. (1971). Temperature regulation in the sphinx moth, *Manduca sexta*: I. Flight energetics and body temperature during free and tethered flight. *J. Exp. Biol.* 54: 141–152.
- Holst, E., von (1943). Über relative Koordination bei Arthropoden. *Pflügers Arch.* 246: 847–865.
- Pfau, H. K. (1991). Contributions of functional morphology to the phylogenetic systematics of Odonata. *Adv. Odonatology* 5: 109–141.
- Scharstein, H. (1989). A universal projector for optomotor stimulation. In N. Elsner, and W. Singer (eds.), *Dynamics and Plasticity in Neural Systems*. Stuttgart: Thieme Verlag, p. 116.
- Scharstein, H. (1998). Forces and power necessary for artificial motion of single insect wings. In W. Nachtigall, and A. Wisser (eds.), *Biona Report 12*. Stuttgart: Gustav Fischer Verlag, pp. 257–270.

- Schlurmann, M. (1997). Anatomische und elektrophysiologische Untersuchungen zur indirekten Flugmotorik von Fliegen (*Calliphora*) und ihrer aminergen Modulation. Dissertation Köln.
- Suder, F., Dierkes, S., and Wendler, G. (2000). Moth flight by wire. In preparation.
- Tu, M. S., and Dickinson, M. H. (1994). Modulation of negative work output from a steering muscle of the blowfly *Calliphora vicina*. *J. Exp. Biol.* 192: 207–224.
- Weber, A., and Wendler, G. (1996). Artifizielle Kontrrolle des Insektenfluges. In R. Jamal (ed.), *Virtuelle Instrumente in der Praxis*. München: National Instruments, pp. 195–201.
- Wendler, G. (1964). Über die Fortbewegung der Larven von *Cantharis fusca*. *Z. vergl. Physiol.* 48: 283–294.
- Wendler, G. (1982). The locust flight system: Functional aspects of sensory input and methods of investigation. In W. Nachtigall (ed.), *Biona Report 2*. Stuttgart: G. Fischer, pp. 113–125.
- Wendler, G. (1999). Fortbewegung und sensomotorische Integration. In K. Dettner, and W. Peters (eds.), *Lehrbuch der Entomologie*. Stuttgart: Gustav Fischer Verlag, pp. 197–244.
- Wendler, G. (2000). Steuerungsmechanismen bei Lauf und Flug von Insekten: ein Vergleich. Verh. Westd. Entom. Tag, Löbbecke-Mus., Düsseldorf 1999. In press.
- Wendler, G., Müller, M., and Dombrowski, U. (1993). The activity of pleurodorsal muscles during flight and at rest in the moth *Manduca sexta* (L.). *J. Comp. Physiol. [A]* 65–75.
- Wolf, H., and Büschges, A. (1995). Nonspiking local interneurons in insect leg motor control: II. Role of nonspiking local interneurons in the control of leg swing during walking. *J. Neurophysiol.* 73: 1861–1875.

Ronald S. Fearing, Srinath Avadhanula,  
Domenico Campolo, Metin Sitti, Joseph Yan, and  
Robert Wood


Flapping flight for microrobots is not only an intriguing mode of locomotion, but provides maneuverability not obtainable with fixed- or rotary-wing aircraft. Insects can fly with a payload equal to their body mass, and have peak accelerations approaching  $10 \text{ m/sec}^2$  (May, 1991). Although they require relatively still air, flying microrobots can fly over terrain that would be impassable for a legged microrobot. Following the initial vision of Flynn (1987), pioneering work in microbotic flight was started by Shimoyama (Shimoyama et al., 1993; Kubo et al., 1994; Miki and Shimoyama, 1998) and more recently, millirobotic flapping flight by Cox, Garcia, and Goldfarb (1998).

This chapter considers the kinematic and power requirements for a microrobotic flying device (RoboFly) using beating wings, and presents an initial design of a thorax for the device. A conceptual drawing of the micromechanical flying insect (MFI) is shown in figure 23.1. The planned device has a 25 mm wingspan, driven by piezoelectric actuators. We will be using flexural four-bar elements to provide sufficient wing stroke. A summary of the MFI component dimensions is given in table 23.1. Further work is required with models and prototypes before solving the critical issues of control, sensing, or power supply.

As a design target for the MFI, we are using the blowfly *Calliphora*, which has a mass of 100 mg, a wing length of 11 mm, a wing beat frequency of 150 Hz, and actuator power of about 8 mW. At this size scale, the current best understanding of nonsteady-state aerodynamics comes from experimental observations of real insects and kinematically similar mock-ups (Ellington et al., 1996; Dickinson and Goetz, 1996).

#### INSECT UNSTEADY AERODYNAMICS

The RoboFly apparatus (Dickinson, Lehmann, and Sane, 1999) consists of a two-winged system driven by three stepping motors that can closely mimic the stroke kinematics of a fruit fly (or other arbitrary kinematics). Strain gauges are used to measure instantaneous wing forces, and the integral of forces around a closed wing beat cycle can be



**Figure 23.1** Overview of proposed micromechanical flying insect components.


**Table 23.1** Components of final size MFI

MFI Component	Size	Total Mass
4-bar frames (2 per wing)	Links 5, 5, 4, 0.7 mm	20 mg
Base frame triangle beam	5 mm × 10 mm	20 mg
Piezo actuators (2 per wing)	0.12 × 0.7 × 8 mm	20 mg
Wings (polyester)	5 × 10 × .005 mm	2 mg
Total structure		62 mg

*Note:* Structure will be made from 12 micron thick stainless steel.

measured to determine net flight forces. RoboFly running with a wing beat of 1/6 Hz in oil has the same Reynold's number as a fruit fly with a wing beat of 220 Hz in air. Flow was visualized using air bubbles in the oil tank and particle image velocimetry (PIV).

Results with the RoboFly apparatus have identified the three key aerodynamic mechanisms used by insects: delayed stall, wake capture, and rotational circulation (Dickinson, Lehmann, and Sane, 1999). Dickinson and colleagues have found wing trajectories that generate peak lift forces of four times the equivalent insect weight. The timing of an equivalent of a backspin motion at the bottom of the wing stroke can change the net lift from positive to negative. The second key finding was the significant forces generated by wake capture at the top and bottom of the stroke. These results are directly relevant to the wing kinematics of the MFI, as it appears that a rapid wing rotation of 90° needs to occur before the end of the down stroke to create adequate lift. This wing rotation and the resulting flow patterns can be seen in a sequence of images taken of *Drosophila* (figure 23.2).



**Figure 23.2** Flow for tethered *Drosophila*, with images (left to right, top to bottom) starting 1.4 ms from top of wing stroke.

## THORAX AND WING DESIGN


We know that insect flight at the centimeter scale requires both large stroke amplitude and wing rotation (Dickinson, Lehmann, and Sane, 1999). *Drosophila* uses a wing stroke of  $160^\circ$  combined with wing rotation of more than  $90^\circ$ . Wing rotation is the challenging part of the design. The insect thorax uses a complicated arrangement of linkages and cams (Nachtigall, Wisser, and Eisinger, 1998), which is not yet fully understood and is likely too difficult to fabricate.

Our thorax design uses separate four-bar frames to control the leading and trailing edges of a wing. The motion magnification of the four bar will allow a piezoelectric unimorph with small angular deflection to be used as the actuator.

### Experimental Flow and Force Results with 5X Mock-Up

To more conveniently study fabrication and sensing issues, a 5X scale thorax was constructed, as shown in figure 23.3, using folded stainless steel sheets. The wing for the 5X mock-up is made from folded 25 micron-thick polyester bonded to plastic wing spars. The 5X wing mock-up can be driven statically with a stroke of  $\pm 60^\circ$  or a rotation of  $\pm 45^\circ$ . Figure 23.3 shows PIV images of this wing being driven with voice coil actuators at  $f = 17$  Hz in a wind tunnel operating with a flow of 0.14 m/sec (to the right). Ellington (1984) approximates the mean Reynolds number during a wing stroke as

$$Re = Sf\phi/\nu$$



**Figure 23.3** Flow measured using PIV for 5X mockup beating at 17 Hz.

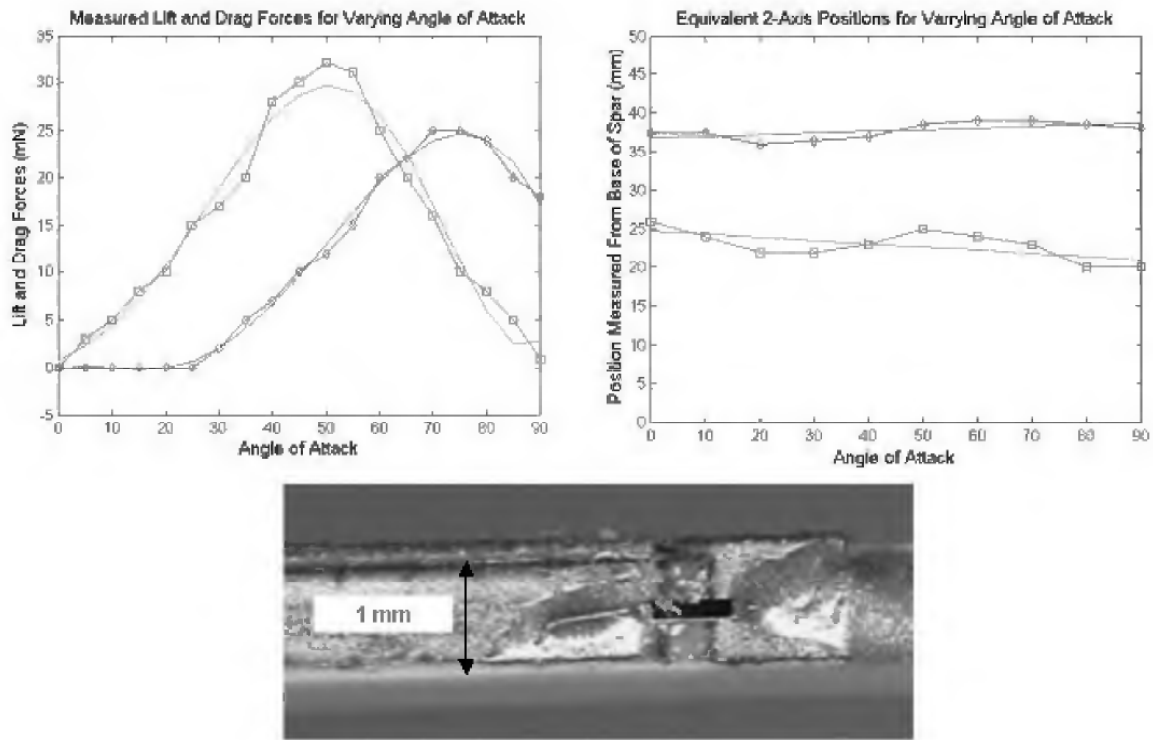
where  $\nu = 1.6 \times 10^{-4} \text{ m}^2/\text{sec}$  is the kinematic viscosity of air,  $S = 32 \text{ cm}^2$  is the wing surface area for both wings, and  $\phi = 1 \text{ rad}$  is the stroke amplitude. In this case,  $Re = 340$ , which is comparable to the blowfly. Processing of the PIV images shows that induced velocities as high as 1.65 m/sec and 0.79 m/sec were recorded during the upstroke and downstroke, respectively, with peak wing-tip speeds of 2 m/sec.

Visually, there are two interesting results that can be seen from the PIV data. First, there is a significant momentum change in the particles directly under the wing from the end of the upstroke to the beginning of the downstroke and this would indicate a period of large lift forces. Second, vortex shedding at the trailing edge of the wing is readily apparent as the wing finishes the downstroke and begins the upstroke.

Unfortunately, our PIV data is only suitable for qualitative estimation of fluid flow. For actual wing-force measurements, wing spar force sensors were developed using conventional semiconductor strain gauges as shown in figure 23.4. A pair of gauges is mounted on the top and side surface of the 1 mm wing spar to measure both net lift and drag force as well as center of force. Preliminary measurements for a wing driven with 1 DOF by a single four bar are shown in figure 23.4.

Each gauge requires two wires, thus eight wires must be brought from the base of the spar, through the four-bar mechanism, to the gauge amplifier. Currently, at 5X scale, we use two four-conductor ribbon cables and string them through the four bar. This presents a problem with the top-link flexures. Ideally, the flexures would be sufficiently compliant; however, the addition of a ribbon cable stiffens the flexures considerably. The voice-coil actuator, which drives the wing at 5X scale, can provide enough force to run this system; however, lighter, more-flexible cable has to be found for the final scale.

Using 12 micron stainless steel, it is possible to construct a spar that is 0.5 mm in diameter and capable of a sensitivity in the order of 10  $\mu\text{N}$ , which is sufficiently close to the calculated flight forces the wing struc-



**Figure 23.4** Wing spar force sensor and measured wing forces for a fixed angle of attack using 5X mockup.


ture will experience. The mass, and thus the inertia of the spar, can be reduced by tapering the beam. Because the gauges are mounted as close to the base as possible, tapering the beam has little effect on the construction of the force sensors.

### Thorax Construction

Thorax construction materials should be chosen for high stress limits, high endurance, and low loss. The insect thorax is constructed from cuticle with resilin connections (Jensen and Weis-Fogh, 1962), which compares very well in stress limits with stainless steel. Polysilicon is a good structural material, but is harder to handle and process than stainless. Note that peak stress must be limited to maintain life. This can be done with thin materials, and by keeping the overall system  $Q$  low because high  $Q$  increases dynamic stresses and reduces controllability. Although the overall  $Q$  of the actuator-wing system is predicted to be only about 2 due to wing and actuator dissipation, for reasonable efficiency the internal dissipation of the structural material (Hosaka, Ito, and Kuroda, 1994) should be kept low.

The four-bar frames need to have a very high strength-to-weight ratio. As pointed out by Yeh and colleagues (1994), hollow beams can be several orders of magnitude stiffer than a solid beam of the same





**Figure 23.5** Laser machined pattern cut from 12 micron stainless for folded 4 bar structure, and (on right) a thorax structure with actuator at base.

mass, and we use the same principle. Solid rectangular and triangular beams of width  $w$  and height  $h$  have respective cross-sectional moments of inertia (Marks' Standard Handbook):

$$I_{rect} = wh^3/12 \quad \text{and} \quad I_{tri} = wh^3/36$$

The stiffness of a cantilever beam of length  $l$  and modulus  $E$  in simple bending due to a concentrated force at the end is given by:

$$k = 3EI/l^3$$

A hollow beam has cross-sectional moment of inertia given by the difference in moment of inertia between a solid beam and a beam smaller by the wall thickness. Consider the links of a four-bar structure with a mass of 0.24 mg/mm length. For the same mass, we could use either a 30 micron-thick cantilever, 1 mm wide, with cross-sectional moment of inertia  $2.3 \times 10^{-18} \text{ m}^4$ , or a 1 mm triangular beam constructed by folding a 12 micron-thick sheet, with cross-sectional moment of inertia  $700 \times 10^{-18} \text{ m}^4$ . Hence, the hollow triangular beam is approximately 300 times stiffer than the solid beam for the same mass. An 8 mm-long box beam made from 302 stainless would have a stiffness of 7000 N/m. A peak load on the structure from the actuators of 100 mN would cause a deflection of less than 130 microns—thus this beam is stiff enough.

Because pin joints are difficult to place in a 1 mm structure without friction, wear, reinforcement, and alignment problems, we use flexural joints. As shown in figure 23.5, the four-bar mechanism, including box beam links and flexural joints, can be fabricated by folding a sheet of photo-etched or laser-cut stainless steel. The side and top links are triangular beams and the base is a box beam for easier attachment to the MFI frame. We have used 12 micron-thick polyester at the joints to in-



**Figure 23.6** Flexural differential mechanism to provide both wing stroke and rotation by coupling two 4 bar structures.

crease lifetime to better than 1 million cycles with large bend angles (Shimada et al., 2000). Currently, the structures are folded manually using fixtures and bonded using cyanoacrylate adhesive. Automatic folding of microstructures is quite feasible, particularly using simple fixtures and a motion-planning approach such as described by Lu and Akella (1999).

Automatic folding is particularly needed for structures, as seen in figure 23.6, which can take a full day to manually fold. In this structure, a differential mechanism between the 2 four bars converts out-of-phase motion into wing rotation.

#### **ACTUATION, TRANSMISSION, AND POWER**

The previous sections have discussed the design and fabrication of a prototype thorax for the MFI. The next critical issue is the delivery of sufficient mechanical power to the wings. Insect flight muscle delivers on the order of 100 W/kg to the wings (Ellington, 1985). To avoid losses associated with accelerating inertial loads, real insects, such as blowflies and fruit flies, drive their wings at mechanical resonance (Sotavalta, 1952). We use the same principle in our thorax design.

For a piezoelectric actuator, the DC mechanical energy density is independent of actuator form factor. For low internal damping piezo actuators operating below their natural resonant frequency, and driving high damping loads such as a wing, the available AC mechanical power is also independent of actuator form factor. Hence, for a given load power requirement, the actuator dimensions can be chosen for ease of mechanical drive, fabrication, and drive voltage requirements.




Figure 23.7 Actuator and wing rotary model.

Consider a micromechanical insect modeled after a blowfly, with a mass of 0.1 g, a wing beat of 150 Hz, and a wing-stroke angle of  $70^\circ$ . The net wing lift force must match the insect weight of 1 mN. Although in the quasi-steady-state, lift and drag forces are generally proportional to the square of velocity, we choose a linear damper with a force at peak wing velocity of twice the weight of the micromechanical flying insect (MFI) as an upper bound. (Note that the linear damper overestimates the damping force for all wing velocities less than the peak velocity.) Hence, the wing damping  $B_w$  (at the wing hinge) can be estimated from:

$$B_w = 2m_{MFI}gl_w/(\omega\theta_w)$$

where  $m_{MFI}$  is the mass,  $l_w$  is the length of the wing,  $\omega$  is the wing beat frequency, and  $\theta_w$  is the wing stroke amplitude (1/2 the total wing stroke). With  $l_w = 10$  mm,  $B_w = 13$  nN-s-m. A simplified linear mechanical model is shown in figure 23.7.

The  $Q$  of a resonant system is defined as the ratio of stored energy to energy dissipated per radian. With proper actuator and transmission design, energy dissipation for the MFI is work done on moving air—that is, useful work. A high  $Q$  hence implies large internal stored energy, and poor controllability of wing amplitude and phase due to actuator saturation. As shown in the data by Sotavalta (1952), blowflies have a relatively low  $Q$ , estimated on the order of 1–3. For the MFI, we choose  $Q = 2.5$ , as a higher  $Q$  system requires a lower transmission ratio and less actuator motion at DC. To have a low- $Q$  maneuverable

wing, the wing inertia is:

$$J_w = QB_w/\omega = 3.3 \times 10^{-11} \text{ kg-m}^2$$

where  $J_w$  is the wing inertia. The actuator stiffness, as seen at the wing hinge, must resonate at  $\omega$ , hence:

$$K_w = J_w\omega^2 = 30 \text{ }\mu\text{Nm}$$

The four-bar transmission converts the small rotation of the actuator  $\theta_a$  to the wing rotation  $\theta_w$  by a transmission ratio  $T$ . The actuator stiffness for resonance with the wing inertia is  $K_a = T^2K_w$ . At DC, the displacement of the wing is just

$$\theta_{wDC} = T\tau_a/K_a$$

At resonance, the reactance of  $J_w$  and  $K_a/T^2$  cancel, so the magnitude of the wing displacement is

$$\theta_{wAC} = \tau_a/(TB_w\omega)$$

The AC motion amplification at resonance is

$$\theta_{wAC}/\theta_{wDC} = Q$$

where we used  $Q = J_w\omega/B_w = K_w/\omega/B_w$ . Hence, the required DC actuator displacement is just

$$\theta_a = \theta_w/(QT)$$

where we assume that the actuator with no load can displace  $\pm\theta_a$ . For a given transmission ratio  $T$ , the required actuator displacement  $\theta_a$  and torque  $\tau_a = K_a\theta_a$  can now be used to specify actuator dimensions. For a unimorph with elastic layer stiffness equal to the piezo stiffness, the required length and width can be found in terms of the the actuator thickness  $h$  (Smits, Dahlke, and Cooney, 1991), which can be chosen for ease of fabrication and available voltage supply. The length  $l$  and width  $w$  are:

$$l = 4h\theta_a/(3d_{31}E_3),$$

and

$$w = \{16\tau_a/(3d_{31}E_3Y_{11}h^2)$$

where  $E_3$  is the electric field, and  $Y_{11}$  is the elastic modulus.

## THORAX POWER DISSIPATION

A critical question for the range of the MFI is the efficiency of power conversion from electrical to wing power. An overall power flow model is shown in figure 23.8. We plan to use high-voltage solar cells (at about 100 V) so that we can drive the piezoelectric actuators directly without step-up converters. We plan to drive the piezoelectric actuators in a

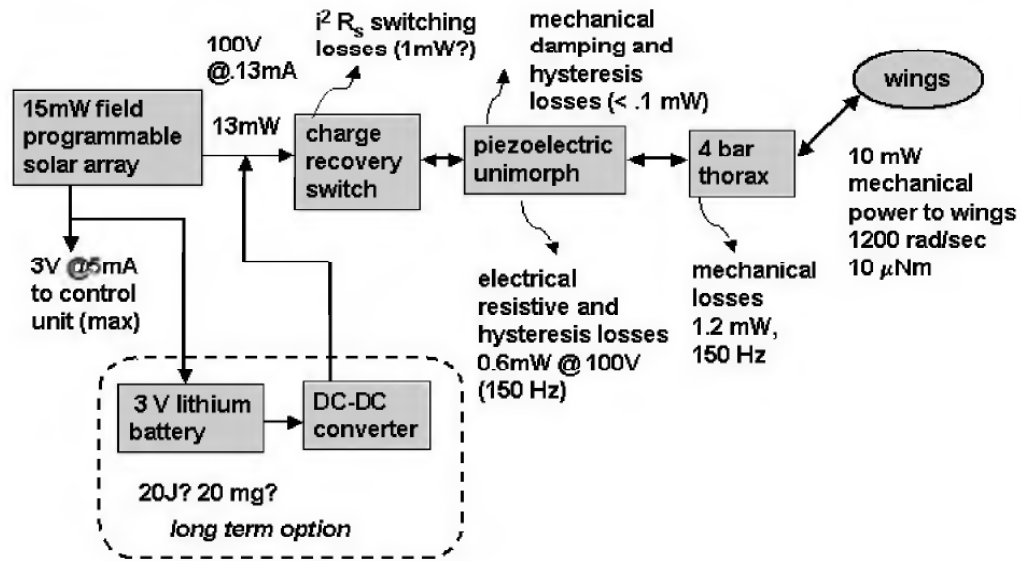


Figure 23.8 Power flow model for piezo actuator and wing transmission structure.

binary fashion to reduce amplifier losses. Due to the large capacitance of piezoelectric actuators, charge recovery is essential to obtain high driving efficiency. We anticipate that efficient transistors will reduce switching losses to less than 1 mW. Piezoelectric actuators have intrinsic electrical and mechanical losses. Finally, the four-bar transmission structure, which amplifies the motion of the piezo actuator, has internal losses due to dissipation at joint flexures and air damping. We measured thorax dissipation of only 1.2 mW at 150 Hz and 150° stroke using a resonance test. We project that total losses, with our current unimorph actuators from electrical power to mechanical wing power are less than 2 mW. This compares favorably to gear losses for conventional motors. (One cautionary note is that we have not yet measured wing or actuator output power. We have just measured the losses. It is possible that under full load conditions, the actuator and thorax might dissipate more power.)

## CONCLUSION

At this stage in the MFI design, we have a good understanding of the necessary wing kinematics, forces, velocities, and power from measurements on RoboFly and real insects. We have designed kinematic structures using folded stainless steel that give adequate wing motion when driven by external (high power) actuators. Using appropriate fixturing, folding stainless steel is feasible at the final size. Current flexure lifetime for the 1.3X structure exceeds 1 million cycles using polyester flexures. Our plan is to develop closed-loop wing controllers that can react to wing forces and modify wing stroke patterns as needed to achieve stable flight.

## ACKNOWLEDGMENTS

The authors thank M. Dickinson and D. Pick for helpful discussions and insights, and Jeff Thompson for laser cutting tools. This work was supported by ONR MURI N00014-98-1-0671, DARPA, and NSF KDI.

## REFERENCES

- Cox, A., Garcia, E., and Goldfarb, M. (1998). Actuator development for a flapping microbotic microaerial vehicle. SPIE Microrobotics Symposium, Boston, November.
- Dickinson, M. H., and Goetz, K. G. (1996). The wake dynamics and flight forces of the fruit fly, *Drosophila melanogaster*. *J. Exp. Biol.* 199: 2085–2104.
- Dickinson, M. H., Lehmann, F.-O., and Sane, S. P. (1999). Wing rotation and the aerodynamic basis of insect flight. *Science* 284 (18 June): 1954–1960.
- Ellington, C. P. (1984). The aerodynamics of hovering insect flight: VI. Lift and power requirements. *Philos. Trans. R. Soc. B Biol. Sci.* vol. 305, no. 1122: 145–181.
- Ellington, C. P. (1985). Power and efficiency of insect flight muscle. *J. Exp. Biol.* 115: 293–304.
- Ellington, C. P., van den Berg, C., Willmot, A. P., and Thomas, A. L. R. (1996). Leading edge vortices in insect flight. *Nature* 384 (19/26 December): 626–630.
- Flynn, A. M. (1987). Gnat robots (and how they will change robotics). IEEE MicroRobots and Teleoperators Workshop, Hyannis, Mass., 9–11 November.
- Hosaka, H., Ito, K., and Kuroda, S. (1994). Evaluation of energy dissipation mechanisms in vibrational microactuators. *IEEE Micro Electro Mechanical Systems*, Oiso, Japan, 25–28 January. pp. 193–198.
- Jensen, M., and Weis-Fogh, T. (1962). Biology and physics of locust flight: V. Strength and elasticity of locust cuticle. *Philos. Trans. R. Soc. B Biol. Sci.* vol. 245, no. 721 (October): 137–169.
- Keller, C. (1998). *Microfabricated High Aspect Ratio Silicon Flexures*. El Cerrito, Calif.: MEMS Precision Instruments. (Also see <[www.memspi.com](http://www.memspi.com)>).
- Kubo, Y., Shimoyama, I., Kaneda, T., and Miura, H. (1994). Study on wings of flying microrobots. In *IEEE International Conference on Robotics and Automation*, San Diego, 8–13 May. pp. 834–839.
- Lu, L., and Akella, S. (1999). Folding cartons with fixtures: A motion planning approach. In *IEEE International Conference on Robotics and Automation*, Detroit, Mich., May. pp. 1570–1576.
- Marks' Standard Handbook for Mechanical Engineers*, New York: McGraw-Hill.
- May, M. L. (1991). Dragonfly flight: Power requirements at high speed and acceleration. *J. Exp. Biol.* 158: 325–342.
- Miki, N., and Shimoyama, I. (1998). Analysis of the flight performance of small magnetic rotating wings for use in microrobots. *IEEE International Conference on Robotics and Automation*, Leuven, Belgium, May. pp. 3065–3070.
- Nachtigall, W., Wisser, A., and Eisinger, D. (1998). Flight of the honey bee: VIII. Functional elements and mechanics of the “flight motor” and the wing joint—One of the most complicated gear-mechanisms in the animal kingdom. *J. Comp. Physiol. [B]* 168: 323–344.

- Oberg, E., Jones, F. D., and Horton, H. L. (1988). *Machinery's Handbook*. New York: Industrial Press.
- Shimada, E., Thompson, J. A., Yan, J., Wood, R., and Fearing, R. S. (2000). Prototyping millirobots using dextrous microassembly and folding. In Proc. ASME IMECE/DSCD Orlando, Fla, 5–10 November.
- Shimoyama, I., Miura, H., Suzuki, K., and Ezura, Y. (1993). Insect-like microrobots with external skeletons. *IEEE Control Systems Magazine* 37–41.
- Smits, J. G., Dahlke, S. I., and Cooney, T. K. (1991). The constituent equations of piezoelectric bimorphs. *Sensors Actuators [A]* 28: 41–61.
- Sotavalta, O. (1952). The wing stroke frequency of insects in wing mutilation and loading experiments at subatmospheric pressure. *Ann. Zool. Soc.* vol. 15, no. 2: 1–67.
- Yeh, R., Kruglick, E. J., Klitzke, M., and Pister, K. S. J. (1994). Towards an articulated silicon microrobot. Winter Annual Meeting, ASME, December.

# 24 The Entomopter

Robert C. Michelson

The Entomopter is a multimode autonomous robot that is capable of flight, ambulatory locomotion, and swimming behaviors in a single vehicle. Autonomous navigation is based on a combination of attraction and avoidance behaviors deriving input from both an integrated optic olfactory sensor for detection of chemical species (or alternately a sensor for a specific type of radiation), and an ultrasonic swept-beam ranging device.

For indoor flight operations, the flapping wing is considered the best approach. Other modes of locomotion (crawling or swimming) are based on the same actuation system as that used for flapping-wing flight, but do not involve the complexity, precision, or energy expenditure associated with flight. For this reason, current development has concentrated on the problems associated with flight using flapping wings.

Entomopter feasibility was established under contract to the DARPA/DSO Mesomachines program where it is referred to as a Mesoscaled Aerial Robot (MAR). The Air Force Research Laboratory (AFRL) has issued a contract under its Revolutionary Technology Program to extend the flight muscle into the fourth generation of size reduction and performance enhancement. Of particular interest to the Air Force and others is the potential for swarms of Entomopter vehicles to rapidly penetrate denied areas such as deeply buried underground facilities. This is possible because of the Entomopter size, multimode locomotion, and its anaerobic propulsion system, which allows covert ingress through sewer pipes, ducts, or electrical conduits. The NASA Institute for Advanced Concepts (NIAC) has funded a study to look at the Entomopter for use as a Mars flier for exploration. The Entomopter's ability to fly in low Reynolds number conditions without need for air-breathing propulsion makes it a natural candidate for a Mars flier, which will operate with a larger vehicle size in the rarefied Martian atmosphere (at Reynolds numbers similar to those encountered by its miniature earth-bound predecessor). Unlike fixed-wing fliers, an Entomopter-based Mars flier would be able to cover a wide area while still being able to



fly slowly, land, crawl, obtain surface samples for analysis, and return to a base station for refueling. As of this writing, one patent has been issued for the Entomopter concept and a second is pending for the reciprocating chemical muscle design.

Biomimetics was the starting point for the Entomopter design, but rather than attempting to replicate biological kinematics and aerodynamics, improved systems have been devised to leverage what is observed in biological systems to produce a machine that is manufacturable, controllable, and able to generate the power necessary to fly from onboard energy sources. Just as wheels are superior locomotors under certain circumstances, but are not a common form of locomotion for biological systems, so the Entomopter has extended its design beyond the biological baseline in some areas.

The hawkmoth (*Manduca sexta*) was chosen as a baseline model for the wing aerodynamics. The flapping mechanism for the Entomopter has been extended beyond that of the hawkmoth to provide a resonant single-piece construction that takes advantage of torsional resonance in the Entomopter fuselage to recover flapping energy as is common to flying insects, which temporarily store potential energy in either the muscles or exoskeletal parts (resilin).

The same structure that provides wing flapping also scans a frequency-modulated continuous wave (FMCW) ultrasonic beam to provide front, side, and down-looking range measurements for obstacle avoidance and altimetry. Also, it has the potential to track and follow free-moving agents in a fashion similar to that employed by bats.

Stability and control in flight as well as navigation is achieved by actively modifying the lift of each wing on a beat-to-beat basis using pneumatic control of the air circulating over the beating wing. Positive lift can be generated on not only the downbeat but the upbeat as well. Coefficients of lift exceeding the theoretical maximum by 500% for the given wing shape can be achieved, without the complexity of active angle-of-attack modulating mechanisms.

A chemically fueled reciprocating chemical muscle has been designed and is in its fourth generation of development at the time of this writing. This actuator system has demonstrated 70 Hz reciprocation rates with throws and evolved power levels necessary to support flight, crawling, or swimming of a self-contained fully autonomous Entomopter system.

## TERRESTRIAL ENTOMOPTER MISSIONS

Assets exist for outdoor aerial reconnaissance. Experiments have shown on numerous occasions, including operations such as Desert Storm, that *conventionally* sized UAVs (unmanned aerial vehicles) often go unnoticed when flying at several thousand feet over their targets. The state of the art in optics is such that high resolution video and infrared

images can be gathered unobtrusively from platforms that exist today (e.g., Predator, Global Hawk). “Tiny” has additional stealth advantage, but carries an unwarranted cost when considered as a replacement for existing assets performing outdoor missions. This cost is not only in dollars, but in the ability to perform missions under even moderate weather aberrations. The real issue to be addressed is how to get the intelligence to be gathered under the control of, and into the hands of the individual war fighter. “Tiny” is only one of several possible solutions to addressing that goal. For example, multiplexing the Predator sensors to take snapshots of specific regions of the battlefield and to deliver them in near real time to the individual users who request them, is probably an easier and better integrated approach to C<sup>3</sup>I than putting tiny personal air vehicles in the hands of the war fighter.

The real mission niche for tiny intelligent unmanned aerial vehicles (UAVs), or mesoscaled aerial robots (MARs) as we shall call them, will be indoors, where the environment is controlled, and there are *no* existing reconnaissance craft that can rapidly negotiate hallways, crawl under doors, or navigate ventilation systems in an attempt to quickly penetrate deep into a building to complete a reconnaissance mission. It is the urban *indoor* mission that will ultimately justify the development expense. The very nature of an indoor mission will necessitate multimode vehicles (flying/crawling/rolling/swimming) and autonomous navigation.

It is desirable that these miniature automatons be not only small but possess multimode locomotion, like insects, so they can move more freely and rapidly through a maze of obstacles, under doors, or inside otherwise inaccessible conduits. Size is important to a multimode robotic vehicle when negotiating indoor terrain. But it is also important when considering that operations in the indoor environment place the mesoscaled aerial robot in much closer proximity to its target than would ever be expected during outdoor operations. Therefore, the stealthiness of small size becomes a significant feature for indoor operations.

A multimode MAR ought to be able to transition between modes. A flying vehicle may need to land to crawl under a door, but then be able to take off again. This requirement leads to propulsion techniques that can generate lift sufficient for flight in confined spaces. Unlike fixed-wing MAR implementations, rotary wing and flapping wing solutions are not only able to land and take off vertically, but they can also fly slowly enough to navigate successfully indoors. Of conceivable rotary-wing and flapping-wing solutions, the flapping wing has the potential to be the least complex and most energy efficient if approached correctly.

An autonomous, homing, multimode insectlike MAR addresses a number of deficiencies that exist for indoor reconnaissance missions. A reciprocating chemical muscle is key to the development of a biomimetic

MAR. Indoor reconnaissance, and a number of other missions of military and intelligence importance that are not currently possible, could be considered were a MAR of even limited endurance in existence.

## ENTOMOPTER SYSTEM DESCRIPTION


An autonomous insectlike robotic vehicle capable of rapidly and covertly placing sensors in remote, inaccessible locations within buildings is under development. Known as an "Entomopter" (Gk: *en*, in + *temnein*, to cut (in reference to an insect's segmented body) + *pteron*, wing "insect wing"), this mesoscaled aerial robotic vehicle will be multimode in that it will not only crawl, but fly. Other embodiments could be made to swim if a mission required it to move through sewers, for example. However, this discussion will concentrate on a biomimetically inspired MAR capable of only flight and ground locomotion.

The heart of the proposed Entomopter is a reciprocating chemical muscle (RCM). Internally funded efforts at the Georgia Tech Research Institute (GTRI) were conducted to develop a proof-of-principal RCM. The RCM is an anaerobic, ignitionless, catalytic device that can operate from a number of chemical fuel sources. While under funding from DARPA/DSO to show the feasibility of an Entomopter-based MAR, the RCM was refined, reduced in size, and demonstrated to develop the power and speed necessary for flight.

Using acoustic ranging methods similar to bat navigation, the Entomopter is able to recycle waste products from locomotive respiration to create a frequency-modulated continuous wave (FMCW) ultrasonic emission for obstacle avoidance and altimetry. A miniature gas-operated ultrasonic transmitter for the Entomopter has been demonstrated as part of the DARPA/DSO Mesomachines for Military Applications program.

A novel approach to flapping-wing design has been developed that avoids the mechanically complex kinematics of an insect wing. Instead, a simple resonant autonomic wing beat has been shown to produce yaw, pitch, roll, and speed changes through the use of pneumatically generated lift modification of the dual Entomopter wings using only waste gas produced by the RCM. These wings have been modeled after those of the hawkmoth *Manduca sexta* but have since been simplified in morphology to accommodate active flow control and to be made easier to manufacture. Wind tunnel tests show that the wing aerodynamics are not degraded by these simplifications.

During indoor missions, stability and navigation is achieved through instantaneous lift modification of the wings on a beat-to-beat basis. A combination of attractive (homing) and reactive (avoidance) behavior is used to navigate. Solid-state interferometric sensors developed at GTRI are capable of detecting minute concentrations of molecular species to



**Figure 24.1** ABS plastic wing ribs from Fused Deposition Modeling machine.

direct the Entomopter to a source of outgassing. The acoustic obstacle avoidance system overrides the olfactory motivation when temporary course changes are necessary in order to prevent collisions.

The goal of the Entomopter program is to develop, build, and demonstrate a flying, crawling, insectlike MAR that can rapidly penetrate a structure to covertly deliver or retrieve a host of miniature payloads to/from a specific location. First-order computer simulations of a 50 gram MAR have already shown that flights of up to a kilometer are achievable in a labyrinth of hallways of a multistory building.

The reciprocating chemical muscle (RCM) is a regenerative device that converts chemical energy into motion through a direct noncombustive chemical reaction. Hence, the concept of a “muscle” as opposed to an engine. There is no combustion taking place nor is there an ignition system required. The RCM is not only capable of producing autonomic wing flapping as well as small amounts of electricity for control of MEMS devices and the “nervous system” of the Entomopter, but it creates enough gas to energize circulation-controlled airfoils. This means that simple autonomic (involuntary, uncontrolled) wing flapping of constant frequency and equal amplitude can result in directional control of the Entomopter by varying the coefficient of lift ( $C_L$ ) on each of the wings, thereby inducing a roll moment about the body of the Entomopter while in flight. Figure 24.1 shows “milliscaled” wing structures grown in Georgia Institute’s stereolithography and fused deposition modeling machines. Wings like these will not only act as smart structures to create proper angles of attack under opposite aerodynamic loads during the upbeat and downbeat, but the hollow microchannels in the ribs provide circulation control gas from the RCM to




**Figure 24.2** Milli-scaled second generation Reciprocating Chemical Muscle.

“blow” the wings for directional control in flight as well as lift when the wing is at negative angles of attack during the upbeat.

The implementation of a reciprocating chemical muscle is motivated chiefly by the basic necessity for very high rate of energy release from compact energy sources. Electrically-driven systems suffer from the poor energy density of batteries, while electrical actuators are typically dense (heavy), or suffer from insufficient force and motion as in the case of electrostatic or piezoelectric propulsors. To increase motion, piezoelectric ceramics can be stacked, but this leads to greater weight, stiffness, and often higher required voltages. Rheological fluids can be slow to respond and will therefore be difficult to use with flapping-wing implementations requiring beat frequencies of 20 to 50 Hz. Faster-acting polymeric muscles have been demonstrated, but require high actuation voltages, dictating the need for power conversion circuits, which add weight and loss to the already heavy onboard battery pack. Actuators of nitinol wire are totally out of the question due to the significant current requirements and variable performance under environmental extremes.

On the other hand, a milliscaled RCM (figure 24.2) has demonstrated pounds of force at frequencies up to 20 Hz on the workbench [1], and the 2X version (figure 24.3) was able to reciprocate at 70 Hz with throws of 0.5 inches. This supports the idea that a reasonable range of power/motion and frequency can be achieved to flap the wings of a 15–20 cm Entomopter. The goal of creating an Entomopter with a wingspan of less than 20 cm places constraints on the size of the muscle



**Figure 24.3** 2X scaled third generation Reciprocating Chemical Muscle.

mechanism but also creates an environment in which high speed muscle operation is less affected by inertia.

The 2X RCM shown in figure 24.3 is fabricated from components using conventional precision machining techniques. Manufacturing techniques used to produce the present milliscaled RCM remain applicable for an implementation at half the current size. In fact, the 2X scaled RCM of figure 24.3 has internal porting that is actual size (1X) for a 50 gram Entomopter. Even greater performance is expected as the muscle mass and volume are reduced.

As conceived, the RCM-driven Entomopter employs a novel X-wing flapping design that differs from the conventional wing flapping observed in birds and insects. It provides additional longitudinal stability (pitch axis) while simplifying manufacture at the mesoscale, because the flapping wings and body (fuselage) can be produced as a single torsional part.

The Entomopter-based MAR is an advance in the state of the art of robotic reconnaissance vehicles in that it uses flapping-wing technology driven by a novel propulsion system to allow not only ground locomotion, but also flight. Its resonant X-wing flapping is a nonobvious solution that simplifies the manufacture of the MAR at the micro air vehicle scale (under 15 cm in any dimension as defined by DARPA), and is ideally suited for actuation by the RCM.

Active flow control of the wings for stability, control, and flight path adjustment is also a nonobvious solution to the complexity of multiple-degrees-of-freedom flapping-wing dynamics. Through the use of this technique, complete control of the vehicle flight dynamics can be had from a simple single-degree-of-freedom autonomic up-down flapping

motion. Because the wing kinematics are simpler, the ability to produce an affordable Entomopter at these scales is enhanced.

Stored energy becomes a significant impediment as MAR mission duration increases. The present state of the art in battery technology does not allow for long-endurance MAR missions, though it is hoped that someday improved electrical storage media (miniature fuel cells, etc.) will result in the energy densities required for useful long endurance (>1 hour) missions in small insect-sized vehicles. Near-term solutions to onboard energy storage will come from chemical or fossil fuels because of their superior energy density. As a point of contrast, consider the amount of releasable energy stored in a drop of gasoline compared to that stored in a battery the size of a drop of gasoline.

Analysis [2] shows that a successful flapping-wing MAR will have to be less than 50 grams in weight and will expend 4 watts continuously to maintain flight. Merely doubling the vehicle weight to 100 grams would cause the energy required for flight to soar by 800%! As fuel is consumed, the Entomopter will become lighter and easier to fly, unlike battery-driven vehicles, which are just as heavy toward the end of the mission as the beginning. The RCM-driven Entomopter will actually improve in performance as the mission progresses. Current battery technology and electric propulsion is simply too heavy.

## FLIGHT LOCOMOTION

There is, of course, a natural prototype for flying and crawling mesoscale machines—the insects. They are among the most accomplished fliers, and with an alternating tripod gait that is stable at all phases of the cycle, they excel at ground locomotion as well. Their crawling abilities have already caught the interest of planetary rover designers, and we now understand the aerodynamics of their flight sufficiently for exploitation in a mesoscale aerial robot. That insects have solved the problem of flight for the mesoscale size range is readily apparent. Their ability to hover and maneuver slowly, combined with fast flight capabilities, make them unsurpassed as an observation platform. Until recently, however, the secret of their aerodynamic success was unknown.

Aerodynamic analyses of their flight consistently revealed that their wings must produce 2–3 times more lift than conventional wings, and in some cases up to 6–7 times. The extra load-lifting capacity this would offer to Entomopters is highly significant, and indicates that a novel design based on flapping insect flight would outperform a more-traditional aerodynamic approach.

The insect solution is also better than that of larger birds and bats, whose wing performance is quite ordinary by conventional standards. The reason for the high-lift production by insect wings has recently been discovered [6–10]. Insect wings are thin with a “sharp” leading




Figure 24.4 Lift vectors on the upthrust and down thrust wing halves.

edge in an aerodynamic sense, and as air flows around the leading edge it separates and rolls up into a leading-edge vortex (LEV—see figure 24.4). This LEV augments the circulation and hence the lift of the wing, according to normal unsteady aerodynamic principles. However, the LEV should grow in size until it breaks away from the wing, causing a deep stall. This is prevented by span-wise pressure gradients that cause the LEV to spiral out to the wing tip, convecting the vorticity into the tip vortex instead of it breaking away from the wing. Thus the LEV is stabilized by this newly discovered spiral form, giving high lift from flapping wings without attendant stall. This LEV phenomenon can be enhanced through circulation control of the wing aerodynamics, which can increase the LEV augmentation by keeping it attached to the wing for longer periods than otherwise possible.

LEVs are a common unsteady phenomenon in the size and Reynolds number range of insect wings. They result from the well-behaved nature of laminar flow around sharp corners, and small trapped vortices can even be seen between the veins of corrugated insect wings. LEVs are also observed on thin man-made wings, such as the Jedelsky profile, in the mesoscale. They are probably present as well on the wings of very small birds and bats, which are not much larger than some insects.

At even larger sizes, however, the laminar flow becomes turbulent and the LEV breaks away. The outstanding lift performance from a laminar LEV is therefore primarily restricted to the size of insects, assuming that it can be stabilized to prevent deep stall. The spanwise pressure gradient needed to create the spiral, stable LEV is absent from fixed-wing designs. Rotary wings (propellers and rotors) have the necessary pressure gradient, but a span-wise flow has not been reported for them. Our working hypothesis is that the LEV is only marginally stable, even with span-wise flow, so stall must eventually develop.

Flapping flight would therefore be essential to exploit the spiral LEV—the wing motion is periodically reversed just as the LEV starts to become unstable. All of the current evidence points to flapping flight



with a spiral, laminar LEV as a high-lift design that is unique to the mesoscale range. Studies of insect flight also define an optimal size for an Entomopter. Insects are geometrically similar on average and they follow a basic design. The smallest flying insects that have been studied have a wingspan of about 1 mm and the largest about 10 cm. In all cases, the airflow is in the laminar regime, but viscous effects become progressively more important as size decreases. Aerodynamic lift is not greatly affected over this range, but viscous drag forces reduce the lift-to-drag ratio for small insects: It is about 3–4 for medium-size insects but close to unity for the smallest ones.

Relative to their size, smaller insects must therefore pay a higher price for the *profile* power needed to overcome friction drag on their wings. The adverse scaling of profile power for smaller insects is countered to some extent by a beneficial scaling of *induced* power—the rate at which kinetic energy is imparted to the downwash of air in the wake, corresponding to the momentum flux necessary to balance the weight. The wings of larger insects must impart a greater downwash to the air, and this entails a higher induced power requirement. The net result is that for a given mass to support, it can be accomplished with less power using a longer wing length and lower frequency (i.e., a larger design). The power saving can be very substantial, approaching an order of magnitude for designs of different sizes. Therefore, larger machines would be the best starting point for MARs, provided that the airflow remains laminar. Recent discoveries about the aerodynamics of insect flight have come from studies on the hawkmoth *Manduca sexta*—one of the largest insects, with a wingspan of 10 cm. We have characterized the flight of *Manduca* very thoroughly, from high-speed films of free flight to flow visualization around the wings, and it is most fortunate that the first successful Entomopter will be built in a similar size range using the hawkmoth wing as an initial aerodynamic design.

The Entomopter will rely on flapping-wing flight to allow it to negotiate the interior of buildings at flight speeds that are compatible with an enclosed environment. At the mesoscale, the Entomopter will benefit directly from the hawkmoth wing-vortex aerodynamic analyses. Unlike the hawkmoth, however, the Entomopter will employ several quite innovative propulsion and control techniques that are not known to occur in nature, but which simplify the design and facilitate manufacturability.

First, the Entomopter wing will be designed to produce lift on both the downstroke *and* the upstroke. Instead of relying on wing twist under muscular control (a complex action requiring an extra degree of freedom in the wing hinge), the wings will be stiffened with materials that react differently to opposite loads. The flexure of the wing ribs (similar to those shown in figure 24.1) will cause the wing to deform relative to the leading edge spar (which drives the wing up and down)




Figure 24.5 X-wing motion of fore and aft wings as viewed from front of Entomopter.

such that it maintains an angle of attack and camber that provides positive lift on the downstroke. Upon the upstroke, these same ribs will deform under an opposite load to create an angle of attack and camber relative to the leading edge spar that also has an upward-lift vector on the inboard section of the wing for at least a portion of the upbeat. The interstitial material between the wing ribs serves as the aerodynamic lifting surface, but relies on the wing ribs to give it form. This is depicted in figure 24.4. In the circulation-controlled airfoil development work conducted for NASA, generation of positive lift was measured at very large *negative* angles of attack (approaching  $-70^\circ$ ), and was produced by very high supercirculation caused by the trailing-edge-circulation-controlled blowing [11–12]. Coupling the deformation of the wing on the upstroke with intelligent application of circulation control will allow lift to be generated not only on the entire downbeat, but on the upbeat as well. This will result in an efficiency greater than that of a conventional insect wing. Beyond the upbeat lift that can be created, the overall coefficient of lift ( $C_L$ ) of the wings can be augmented by pneumatic blowing to achieve values that are 5–8 times higher than the theoretical maximum achievable by a typical wing planform and camber (which has a  $C_L$  of 1 or less).

Experiments under Georgia Tech’s Entomopter IRAD program have shown that it may be possible at the mesoscale to directly use the ABS plastic output from an FDM machine to form the hollow-rib structure and then to overlay and bond a nonstretching polyester film as an interstitial aerodynamic surface.

Another innovation is the use of an X-wing design instead of the classical opposed “clapping” wing flap mechanism. The X-wing relies on twin wings, one fore and one aft, that pivot like seesaws across the central fuselage. If operated  $180^\circ$  out of phase with one another, there will always be two rising wings and two descending wings during the flapping cycle, as depicted in figure 24.5.

The forces from each wing pair are resolved across the fuselage. If each wing is producing an upward lift vector on both the down- and upstrokes, then a net lift for the entire vehicle will be realized. By separating the pair of wings at either end of a fuselage, longitudinal (pitch) stability is increased over a single flapping wing pair.

If the fuselage separating the fore and aft wing spars is made of a material with a torsional memory, then the wing spars can become integral to the fuselage and can be torqued out of phase such that the fuselage acts as a torsional spring with a resonant frequency depending on the inertia of the wing spars, the flapping drag of the wings (damping), and the spring constant ( $K$ ) of the torsional fuselage.

This implementation achieves two important things. First, it allows the entire system of fuselage, wing spars, and wing ribs to be constructed as a single piece with no hinged parts. This greatly simplifies the construction of the Entomopter and facilitates mass production. It further guarantees a higher level of consistency of performance.

More important, this construction forms a totally resonant structure. The torsional resonance is essential to make a flapping-wing vehicle viable. All insects have resonant structures that store potential energy in either the muscles themselves, or in their exoskeletal parts. Many insects use a substance called resilin to store mechanical potential energy. In order to flap wings, energy must be added incrementally to make flight viable. Any attempt to flap wings using brute force without a tuned resonant structure will require too much energy. This is particularly obvious when one considers that actual insect muscles (e.g., wasp) have been shown to expand and contract only 2% of the overall muscle length during each full beat of the wing [3], while the muscle efficiency and level of resonant energy storage has been estimated for the fruit fly *Drosophila hydei* to be only 10%, with the energy stored elastically for resonant release estimated to be somewhere between 35 and 85% [4].

In another innovation that simplifies the Entomopter wing structure, waste gas from the RCM is valved and can be vented through the hollow wing ribs to slots along the trailing edge and tip of the wing in order to achieve a level of circulation control and to keep the lift-augmenting outward-spiraling leading-edge vortex (LEV—see figure 24.4) attached to the surface of the wing longer than normally possible. Circulation-controlled airfoils have been shown to exceed the theoretical coefficient of lift by more than five orders of magnitude by keeping the air passing over the airfoil from separating into shed vortices.

In either case, intelligently changing the lift of one wing and not the other, will result in an aileron motion that will cause the Entomopter to turn in flight (combined roll and yaw). This technique for flapping-wing flight control eliminates the need to beat wings harder on one side or to modify the angle of attack of only one wing in order to change direction. Employing a simple pair of binary valves to reuse gaseous waste products from the resonant autonomic wing-beating process, results in a significant weight and energy efficiency compared to that of less-innovative wing-manipulative approaches.

Because the fuselage is torqued alternately in opposite directions by as much as  $\pm 45^\circ$  to  $\pm 60^\circ$ , the fundamental resonant mode will have a node exactly half way between the fore and aft wing spars. No twisting of the fuselage will occur at this node. Therefore, the propulsion system, and any payload, will be attached to the fuselage at this point to prevent rotation of these components, and to prevent interference with the resonance of the wing system.

The innovative flapping-wing design of the Entomopter therefore eliminates the need for additional degrees of dynamic freedom found in most insect wings. Twist and differential beating motions are supplanted by a simple constant and equal autonomic wing beat that combines circulation control techniques with the compliance of smart materials to achieve the same effect. This in turn allows for fewer separable moving parts, fewer manufacturing steps, and added robustness of the flight mechanism.

To enhance the effectiveness of the Entomopter as an indoor sensor placement platform, it must be multimodal. In addition, it should be able to crawl, roll, and/or swim. For the purpose of this discussion, alternate modes of locomotion will be limited to crawling.

## SURFACE LOCOMOTION

Crawling will allow the Entomopter to precisely position a sensor once flight has brought the sensor into the general vicinity of where it is to be used. Crawling also allows the Entomopter to pass through small openings (perhaps under doors) where flight is not possible. Such a capability could make the difference in the continued prosecution of a mission when it might otherwise be terminated by an obstacle. Finally, crawling can be a mode of continued locomotion when either too little energy remains to maintain flight, or when the flight system has become damaged and is no longer functional.

When crawling, the same RCM used for flight can actuate sets of legs. As shown in figure 24.6 (left), a set of four legs can be integrated onto the RCM such that they reciprocate in a plane  $90^\circ$  to that of the wing motion. The number of legs is not really an issue, but the mechanical complexity is. Figure 24.6 (left) shows very simple unjointed legs that basically “paddle” like oars. The shape of each leg is bent to catch on rough features when pulling the Entomopter forward, but slides over those features when repositioning for the next step. Certainly, other leg configurations are possible, but this simple concept can be used to drive the Entomopter forward with the fewest moving parts per leg.

During crawling, the rate of reciprocation of the RCM is lowered from tens of hertz, down to 1 Hz or less by appropriately metering the




Figure 24.6 Possible ground locomotor schemes.

fuel. Depending on leg height and wing-flapping angle, the wings could be allowed to remain engaged and continue to flap at this low frequency with little effect. Shedding of wings upon landing, like alates (e.g., winged termites), is another option. The legs, on the other hand, must be disengaged during flight because they would adversely affect the power required for flight as well as the resonance of the system were they allowed to flail at 20–50 Hz.

Another ground locomotion concept is depicted in figure 24.6 (right). A “velcro-foot/friction-pad” pair would allow inchwormlike ground locomotion as the RCM oscillates. This configuration would be useful on relatively smooth surfaces that have enough roughness to allow the velcro hooks to engage the surface in one direction and slide in the other. Applications for this type of ground locomotion might include travel through an electrical conduit inside a wall, or down heating ducts where space is limited. Such missions would favor the ability to go long distances through a guiding pipe but would not require flight (no wings).

The inset in figure 24.6 (left) shows one mechanical method of automatically disengaging the legs from the guide bushing of the RCM during flight. The weight of each leg causes it to pivot away from its driving source to hang motionless during flight, but when the Entomopter lands and its weight is placed on the legs, they reengage automatically.

Rather than making the legs move intelligently with alternating motions to allow steering, they are set only to provide forward propulsion. A single steering device (wheel or skid) forward of the legs, directs

the forward thrusts of the legs by providing a side force to the front of the Entomopter. For simplicity, this guiding device can be triadic in nature such that it exists in one of three positions: left, center, or right. The position is determined by onboard sensors.

For a homing response, the Entomopter could contain a pair of side-mounted monolithic molecular concentration sensors. If the concentration of a target chemical was greater to the right, the guiding device would turn to the right. If the concentration on both sensors was equal, the guiding device would be aligned along the axis of the Entomopter fuselage. In this way, the Entomopter would exhibit a homing behavior as its legs propelled it forward. In flight, the same molecular concentration sensors would modulate the circulation control valves to induce changes in heading until the concentration of the molecular target species was equal on both sensors and increasing with time.

Both the legs and the steering mechanism provide a simple, low-parts-count, low-energy solution to ground locomotion by employing the same RCM used for flight as a primary ground locomotive power source. An analogous example could be conceived for paddles and a rudder when implementing an amphibious multimode robot variant.

## AUTONOMOUS NAVIGATION

An Entomopter operating indoors must be autonomous. Due to the size of the Entomopter, it will be difficult to support the long antennas necessary to communicate at lower frequencies. The higher frequencies, which can be easily supported, are more susceptible to attenuation through walls and become very directional (requiring line-of-sight between transmitter and receiver) as the frequency of operation increases. Teleoperation once within a building is precluded not only by the physics of signal propagation, but also by the inability of a remote operator to see the Entomopter or know its location. The use of the global positioning satellite (GPS) system is not an option for navigation because the required antenna is too large to be carried by a meso-scaled vehicle and the signal strength is too weak to be reliable inside a building.

Through the use of homing behavior, however, the Entomopter can find its way around within buildings as it seeks particular chemical scents, acoustic sources, or strong RF emitters. It could also use dead reckoning to move down corridors and into rooms if there is an a priori knowledge of the building.

Critical to autonomous navigation will be obstacle avoidance. A sense of where the floor and walls are is essential. Knowledge of ceiling height is of lesser importance as the Entomopter could be programmed to fly at a predetermined altitude using the floor as a reference.




Figure 24.7 RCM-driven FMCW ultrasonic acoustic ranging source techniques.

Only short-range obstacle avoidance is really necessary. When flying down a corridor, the fact that a wall is 3 meters from the Entomopter is of little significance. Only when the Entomopter approaches within some minimum distance dictated by the Entomopter's ability to maneuver should a correction be needed. This implies that any ranging system need only operate over a short distance.

Inherent to the reciprocating chemical muscle is the ability to generate a frequency modulated continuous wave (FMCW) acoustic emission that can be used for short-range unambiguous altimetry and obstacle avoidance. In addition, the energy cost to the RCM-equipped Entomopter is negligible because the FMCW transmission is essentially free through the use of RCM waste gas, and energy only need be expended to receive and process the returned acoustic signal.

Figure 24.7 shows how a tuned resonant cavity can be acoustically excited by a waste-gas jet from the RCM to produce an ultrasonic emission (other techniques physically similar to dog whistles can also achieve an ultrasonic acoustic emission at scales compatible with the Entomopter and its RCM). By varying the length of the cavity with a mechanical linkage to the RCM or Entomopter wing system, the frequency of the ultrasonic emission can be swept as a series of eigenvalue frequencies approximating a linear triangle or sinusoidal wave (depending on the linkage used) as shown in figure 24.8. As shown in figure 24.9, such an FMCW waveform can be used to generate a ranging solution that resolves Doppler ambiguities.




Figure 24.8 Empirically derived frequency excursion vs. cavity length.

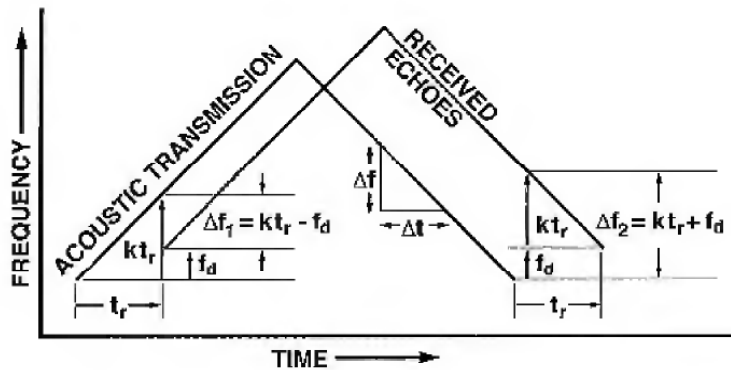


Figure 24.9 FMCW ranging waveform.

A single ultrasonic acoustic source can be vectored by the wing-flapping mechanism of the Entomopter to illuminate not only the floor for altimetry, but also to each side to detect impending collisions with walls. The same beam can be multiplexed as shown in figure 24.10 to create a forward-looking beam. Response to the received ranging signal would control the lateral flight path of the Entomopter as well as its altitude. Active flow control of the wings would effect lateral maneuvers and can be used to temporarily adjust altitude. An increase or decrease in wing-flapping frequency around the resonant point (as controlled by fuel metering/RCM pressure) would effect longer-term vertical maneuvers.

Processing of these return signals could be discrete in that any range to a side obstacle can be ignored so long as it does not intrude within the Entomopter's safety zone. When an obstacle moves too close (either due to the movement of the obstacle or due to the Entomopter closing on the obstacle) the behavior could be a simple avoidance response. Such a simple obstacle-avoidance algorithm can be implemented with a minimum of onboard processing and can in fact be almost a reflex action, but the resulting flight path could be zigzagging, as the vehicle






Figure 24.10 Multiplexing of waste-gas-driven FMCW ultrasonic acoustic ranging source.




Figure 24.11 Directivity of swept ultrasonic source.

alternately tends toward and avoids opposite walls. On the other hand, altitude might be monitored more closely in terms of actual distance to facilitate more-coordinated landings.

Georgia Tech Research Institute (GTRI) has designed and built a mesoscaled ultrasonic ranging transmitter for use on the Entomopter and has performed sound pressure level measurements at RCM waste gas pressures in the range of 40 psi. These tests have shown that sensible ranges of 3 meters are attainable, with greater ranges possible (see sound pressure levels of figure 24.8). Emission patterns have been plotted to show that a directional beam can be created (figure 24.11), and an acoustic mirror scheme has been designed into the Entomopter, which is capable of scanning the output from a single ultrasonic source

ranging from right-horizontal, to head-on, to down-looking (altimetry), to head-on, to left-horizontal on each wing beat (20–30 Hz scan rate).

### **Energy Efficiency**

When all is said and done, a highly innovative reuse of waste products to increase system efficiency is employed, wherein the gas created in the reciprocating chemical muscle is reused (for different functions) six more times after wing-flap excitation prior to expulsion. Decomposition of the chemical fuel source creates gas and heat. A gas generator is situated within the fuel tank to allow it to preheat the fuel for better performance, and to allow the high thermal conductivity of the fuel to provide a heat sink for conformal thermoelectric materials on the outer surface of the gas generator. The first use of the gas generated within the RCM is to drive the wing actuation mechanism. During this process, mainly heat energy is extracted. At the same time, the thermoelectric generator scavenges waste heat from the exothermic decomposition of the fuel to create a small amount of electricity. Waste gas is then used to resonate a cavity to create ultrasonic FMCW emissions for altimetry and obstacle avoidance. The act of wing flapping scans this ultrasonic beam spatially and also modulates it with the FMCW encoding. The gas then passes through a system of ejector nozzles to entrain external air in order to increase the mass flow while dropping the pressure and temperature of the waste gas. The waste gas and entrained air then drives gas bearings at the junction between the fixed-frame fuselage components and the oscillating wing structure. After passing through the gas bearings, the gas is diverted into each wing by means of fast-acting binary valves to effect beat-to-beat active flow control for lift modulation of the wings. Any excess gas is finally dumped as jet thrust. This recycling process is depicted in figure 24.12.

### **EXTRATERRESTRIAL ENTOMOPTER MISSION: MARS EXPLORATION EXAMPLE**

Although the Entomopter was originally sized and designed as a mesoscaled aerial robot for rapid penetration of denied structures, it has found potential utility in an unlikely application: Mars exploration.

#### **Flight on Mars?**

The notion of flight on Mars has been a subject of contemplation since Werner von Braun conceived a rocket plane as a means of Martian exploration in 1953. In the 1950s, Mars flight was purely fancy, but in the 1970s it was revisited more seriously, being spurred on by the successes of the Viking program. The NASA Dryden Research Center,

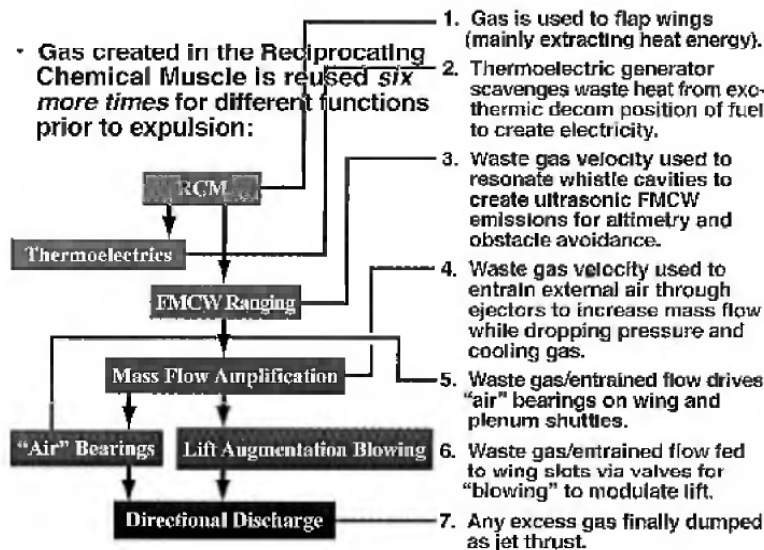


Figure 24.12 Innovative recycling of waste products.

Developmental Sciences, Incorporated, and the Jet Propulsion Laboratory (JPL) proposed unmanned aircraft designs for Mars exploration in 1977 and 1978. Their concept was a propeller-driven fixed-wing aircraft fueled by hydrazine. No follow-up was forthcoming and the ideas died.

A decade later, JPL sponsored a Mars airplane study in which Arora Flight Sciences proposed the electrically propelled "Jason" aircraft. At about the same time, Ames Research Center and Sandia National Labs conceived a high-speed aerospace plane named Aerolus. Unlike the earlier attempts to make a slow-speed aircraft that would be deployed from an aeroshell after touchdown on the Martian surface, Aerolus would make a direct atmospheric entry and then fly through the Martian atmosphere at hypersonic speeds. To date, neither the Jason nor the Aerolus projects have been embraced by NASA's Mars exploration program.

Successes with the Mars Pathfinder and Global Surveyor programs renewed interest in Mars fliers for exploration. In 1995, NASA Dryden and Ames Research Centers once again considered unmanned aerial vehicles to extend the reconnaissance range of Mars landers. The new concept was to launch a small unmanned aerial vehicle (UAV) from the lander after it had stabilized on the surface. The UAV would provide video of the immediate vicinity of the lander (within several thousand meters) to provide feedback as to the most interesting areas for investigation by ground-based rovers. The expendable, one-flight UAV would be electrically powered with rocket-assisted takeoff.

The following year (1996), the Ames Research Center proposed an unmanned Mars aircraft in response to a NASA Announcement of Opportunity for Discovery Exploration Missions. Ames's approach was to

use a propeller-driven, sailplane configuration they called “Airplane for Mars Exploration” (AME); however, it was not selected for the Discovery mission.

On the following NASA Announcement of Opportunity for Discovery Exploration Missions in 1998, JPL submitted a proposal for a multiple glider system (dubbed “Kitty Hawk”), wherein several areas could be investigated during a single mission. Being gliders, the vehicles were obviously limited in endurance, but benefited from the lack of weight and complexity associated with a propulsion system in return for redundancy of numbers. NASA Ames also submitted a proposal to the 1998 Announcement for a motorized UAV named “MAGE.” Both concepts deployed from an aeroshell once it had become subsonic approximately 12,000 meters above the Martian surface. Again, neither concept was selected for the Discovery mission [13].

On 1 February 1999, NASA Director Daniel Goldin announced the Mars Airplane Micromission, which will be the first NASA program to launch on an Ariane 5 rocket. Mars’s next close encounter with Earth would have the first Mars airplane arriving on the red planet around December 2003, coincidentally close to 17 December 2003: the hundredth anniversary of the Wright brothers’ first flight. However, for various reasons including risks associated with the deployment of vehicle types proposed thus far, the Wright brothers’ hundredth anniversary will pass without a Mars flier being deployed.

### **The Mars Environment**

The Martian environment provides a number of significant challenges to atmospheric flight, not the least of which is the lack of oxygen to support combustion for propulsion, a rarefied atmosphere, and extremely cold temperatures. Specifically, the Martian atmosphere is more than 95% carbon dioxide and is less than 0.5% as dense as the Earth’s. The average surface pressure is only 0.007 of Earth’s atmospheres, which is roughly equivalent to Earth’s atmospheric pressure at an altitude of 105,000 ft. The average temperature near the surface of Mars is  $-63^{\circ}\text{C}$ , with diurnal highs and lows ranging from  $+20^{\circ}\text{C}$  down to  $-140^{\circ}\text{C}$ . On the other hand, Mars has only 37% of Earth’s gravity, making flight easier in spite of the thin Martian atmosphere.

A Mars flier faces a number of engineering challenges beyond those of the environment in which it must operate. In particular, one must consider:

- How to deploy the Mars flier.
- The aerodynamics of low Reynolds number flight and the flight dynamics of a vehicle of low weight and potentially large momenta.

- Non-air-breathing propulsion, and operation of moving parts at extremely low temperatures.
- Fully autonomous flight and the ability to land, maneuver on the surface, and take off.
- Navigation while in flight and while on the surface.
- Electrical energy generation/storage/harvesting.
- How to communicate findings.

### **Entomopters as Mars Fliers**

As it turns out, the Reynolds number regime in which the Entomopter operates is nearly identical to that of Mars's atmosphere, were the Entomopter to be upscaled from its present size. In addition, the propulsion technique being used in the earthbound design happens to be anaerobic, and can use multiple fuels including hydrazine.

Unlike prior attempts at Mars airplanes, the Entomopter is a flapping-wing vehicle that can fly slowly and, theoretically, hover. The fact that the wings move rapidly in the supporting atmosphere means that its fuselage can move at slower speeds, thereby facilitating aerial reconnaissance. In contrast, fixed-wing Mars airplane solutions must fly at high speeds (over 100 m/sec) simply to stay aloft.

Images returned from the Viking and Pathfinder landers show a rocky surface that precludes landing an airplane moving 100 m/sec. The slow flight of the Entomopter affords the possibility of landing on the Martian surface to inspect objects, take samples, or to rest during periods of communication blackout, adverse weather, or to harvest energy from the environment.

Because the Martian atmosphere is thin, fixed-wing Mars fliers will have to make tremendous sweeping turns in order to change course. Orbiting a point close to the surface will be difficult. The flapping-wing approach used by the Entomopter allows much tighter maneuvers because directional control and stability are achieved through active lift modification of the wings on a beat-to-beat basis. On Earth, the coefficient of lift can be changed from a maximum of around 1 (based on wing shape) to coefficients of between 7 and 15, and can be done rapidly on a beat-to-beat basis for wings flapping in excess of 25 Hz. This control is derived from the waste gas resulting from the propulsion process and therefore comes almost free of charge.

The Entomopter uses autonomic wing beating as part of a resonant structure, thereby avoiding the complications of traditional biomimetic flapping-wing machines (constraints placed upon an Entomopter Mars flier by the Ariane 5 payload area will complicate the simple design by requiring a folding or inflatable wing structure, however). The simplicity of the wing design will be an advantage in the cold Martian envi-

ronment where moving parts must maintain bearing tolerances and require lubrication. The gas bearings currently designed into the earthbound Entomopter avoid wet lubricants.

Locomotion, either flying or crawling, is energy intensive. The earthbound Entomopter will require between 4 and 8 watts to propel 50 grams in flight. Because of its mission to operate within confined spaces, the Entomopter is necessarily small. A Mars flier based on the Entomopter design benefits from a larger form factor, meaning that there will be a greater physical volume available to carry fuel. Proportionally, the fuel will also weigh 37% less than on Earth.

The exothermic metabolism of the Entomopter affords a means of heating electronics and moving structures. Further, there is the potential for thermoelectric power generation by scavenging waste heat liberated by the propulsion system.

Because the Entomopter can potentially land and take off from the Martian surface, communications may be facilitated. Stored data can be uploaded to orbiting repeaters over periods of hours or days without expending locomotive energy if the Mars flier can land on the surface. Traditional fixed-wing Mars fliers will have to upload all of their information while in flight, or have a survivable electronics package that can continue to operate after the vehicle has crashed into the surface of the planet at over 100 m/sec. The ability to land and take off allows the Entomopter to continue its reconnaissance mission for as long as it has fuel available, but with long intermissions. If energy harvesting is possible by converting surface materials, collecting solar radiation, or harnessing the Martian winds in an onboard generator, then the Entomopter could potentially operate for years.

Autonomous operation of all Entomopter modes has already been anticipated in the present earthbound design, which includes an obstacle avoidance system as well as an altimeter. A Mars flier would likely use an RF interrogator for obstacle avoidance and altimetry while navigating, based on cooperative cues from preplaced radio beacons around the landing area. The magnetic field on Mars is too weak and unreliable for navigation. In a fashion similar to the honeybee, the Entomopter could use solar cues coupled with Martian time and sidereal information to perform gross navigation tasks.

Multiple Entomopters could be deployed from a single lander or "mother ship" (figure 24.13). During daylight, the Entomopters would explore the Mars landscape while solar navigation is possible. Unlike fixed-wing solutions, an Entomopter-based Mars flier would be capable of slow flight in the Martian atmosphere because the rapidly flapping wings can generate sufficient lift even if the fuselage is moving at a relatively slow air speed. The ability to fly slowly also makes landing and takeoff from unprepared surfaces possible. Landing to acquire samples or to wait out the dark phase is possible with the Entomopter



**Figure 24.13** Deployment of Entomopter-based Mars Flier from Mars Lander.

configuration. Periodically, each Entomopter would return to the stationary lander for refueling and to deposit samples.

Because of the latency in transmissions between Mars and Earth, teleoperation of a Mars flier is impractical. Even supervised autonomy is of limited value. A Mars flier will have to be able to carry out its science mission without human intervention while being ever cognizant of its environment to assure that it avoids obstacles, hazards, and situations that would result in starvation.

Motivation for navigation would be based on various remote sensors that will be dictated by the type of science experiments to be performed by each Entomopter. For example, the search for life might entail sensors that can detect traces of water or fossil-bearing rock. Other Mars fliers might be measuring atmospheric species or performing reconnaissance for later close inspection by ground-based rovers. In each case, the Mars fliers would use preprogrammed search patterns initially. When measuring a volume, as in the case of atmospheric sampling, the entire flight might be preprogrammed. When searching for life, a preprogrammed search pattern would be abandoned in favor of following gradients based on the frequency of occurrence of evidence (motivational behavior). During the landing process, obstacles on the surface must be negotiated (avoidance behavior), and the Entomopter must select a spot from which it can launch itself back into the air as it transitions from ground locomotion to flight.




Figure 24.14 Measurements taken during a dust devil (pathfinder data: Schofield [14]).

Due to the occurrence of Martian storms, the Entomopter-based Mars fliers might have to seek shelter on the surface by landing in a self-preservation behavior. Dust storms tend to occur when Mars is near perihelion in its orbit, when the solar intensity is the greatest. It is believed that the greater intensity of solar radiation coupled with variations in the topology of Mars triggers the dust storms. The storms can last up to several months and the opacity of the storms can be quite high, thereby precluding flight and some types of navigation, such as solar tracking. Even local events, such as dust devils could be a threat to a small Mars flier. Dust devils are short-term variations in measured surface pressure, wind velocity, and air temperature over periods of tens of seconds to minutes. This is shown in figure 24.14. Dust devils about 2 km wide and a few kilometers high have been observed in the tropics by the Viking orbiters.

In all cases, however, the Mars fliers would have to be able to find their way back to the lander in order to replenish depleted fuel supplies, as they exhibit a feeding behavior also driven by a self-preservation motivation.

A definition of the requirements for the sensors and the type of science that can be performed on Mars by an Entomopter-based Mars flier has been defined as part of the Mars micromission aircraft program [15–16]. The following is a list of some of the science data that can be collected by the Entomopter.

- High resolution surface imaging
- Surface mineralogy and sampling
- Atmospheric sampling, dust properties, and weather data



- Payload delivery
- Magnetic field mapping
- Water content and distribution
- Infrared mapping
- Radar mapping

The use of an Entomopter-based Mars flier for the exploration of the Martian environment can be summarized as three main mission types: within the range of a central vehicle, independent exploration, and in a tandem system with a rover.

**Exploration within the Range of a Central Vehicle** For this scenario, a lander containing a number of Entomopter vehicles would land on the surface. The Entomopters would then depart this central vehicle and explore the surrounding terrain. All data and samples collected would be brought back to the central vehicle, which would also act as a refueling station for the Entomopters as well as a communications link between the mission and Earth.

The main advantage of this type of mission structure is that the lander provides a number of applications that would otherwise have to be performed by the Mars fliers. The communications requirements for the Entomopters are relaxed because they need only communicate over short (typically line-of-sight) distances to the lander. Also, the lander can provide a source of fuel for the Entomopters, greatly extending their mission lifetime. This fuel can either be carried by the lander from earth or possibly made on site, depending on the type of fuel used.

The main drawback to this type of mission is that the exploration area is limited to the round trip range of the Entomopter about the lander.

**Independent Exploration Using an Entomopter** In this scenario, a lander containing one or more Entomopters lands on the surface. The lander serves solely as a transport vehicle. The Entomopters leave the lander and begin to explore the surrounding territory. Each Entomopter is an independent vehicle and relays its data directly to an orbiting communications platform. Unless fuel can be harvested from the environment and processed by the Entomopter, this second mission scenario is one-way, with the Entomopter being considered expendable.

The main advantage of this type of mission structure is that the Entomopter can double its range of exploration; however, the ability to return to an area for further investigation is not an option once the onboard fuel supply is expended.

**Tandem System, the Entomopter Works in Conjunction with a Rover** In this scenario, a lander contains one or more Entomopter-based Mars fliers as well as a surface roving vehicle. The rover and Entomopters

leave the lander and begin to explore. The lander serves as transport vehicle, communications relay, and refueling station. The Entomopters communicate with the rover, which in turn relays the data to an orbiting communication system. The Entomopters can assist the rover in terrain navigation as the group slowly moves across the surface. The mission capabilities would be similar to the first scenario described—however, in this case, instead of a fixed lander, the mobile rover serves as the home base. The Entomopters would be able to dock with the rover for refueling.

The main advantage of this third deployment scenario is that new territory can be explored each day by the Entomopters, because their home base (the rover) slowly progresses along the surface. This scenario only makes sense if stored fuel is available for the Entomopters on the rover, and the rover itself is able to recharge (solar electric). Alternately, the rover could be a mobile factory to process Entomopter fuel from solar electricity or from synthesis of constituents found in the environment. With rechargeable vehicles, the only limit on mission duration is mechanical failure or physical barriers.

The science to be performed by each Entomopter-based Mars flier could be different. In all three of the scenarios, different Entomopters could return to previously visited locations to perform different investigations. In this way, no single Entomopter has to bear the weight of multiple sensors, yet over a period of time, multiple experiments can be performed.

## CONCLUSION

The Entomopter is a unique aerial robot with biomimetic origins, but which goes beyond biomimicry to incorporate novel control techniques that facilitate both performance and manufacturability. A fueled reciprocating chemical muscle allows the Entomopter to fly, crawl, and swim in environments devoid of oxygen. In its smallest embodiment, the Entomopter reuses the waste products of its own propulsive respiration to control itself in flight as well as to navigate around obstacles. Because flapping-wing propulsion is used, the Entomopter is conducive to indoor terrestrial flight operations and even shows promise as an airborne extraterrestrial exploration vehicle over the landscape of Mars.

The Entomopter was conceived as an autonomous vehicle from its inception. Olfactory inputs are the primary terrestrial motivators for navigation, although any radiator can serve as an attractant. Obstacle avoidance is achieved by means of an active FMCW ultrasonic emission that sweeps a 180° field of regard in front of the Entomopter while in flight. The mechanism that provides propulsive motion also scans the acoustic ranging beam to detect nearby obstacles as well as to provide altimetry information.

Subsystems associated with the reciprocating chemical muscle, the wings, and the ultrasonic ranging system have been demonstrated in hardware at or near the scale of the terrestrial flight vehicle. A chemically fueled reciprocating chemical muscle has been designed and is in its fourth generation of development at the time of this writing. This actuator system has demonstrated 70 Hz reciprocation rates, with throws and evolved power levels necessary to support the flight, crawling, or swimming of a self-contained fully autonomous Entomopter system weighing 50 grams.

## ACKNOWLEDGMENTS

The author wishes to acknowledge contributions by Dr. Charles Ellington of Cambridge University, Dr. Krish Ahuja of the Georgia Tech Research Institute, and Anthony Colozza of NSI.

## REFERENCES

- [1] Michelson, R., and Reece, S. (1998). Update on flapping wing micro air vehicle research—Ongoing work to develop a flapping wing, crawling Entomopter. In *thirteenth Bristol International RPV Conference Proceedings*, Bristol, England, 30 March–1 April.
- [2] Michelson, R., Helmick, D., Reece, S., and Amarena, C. (1997). A reciprocating chemical muscle (RCM) for micro air vehicle “Entomopter” flight. In *1997 Proceedings of the Association for Unmanned Vehicle Systems, International*, June. pp. 429–435.
- [3] Gilmour, K. M., and Ellington, C. P. (1993). *J. Exp. Biol.* 183: 101.
- [4] Dickson, M. H., and Lighton, J. R. B. (1995). Muscle efficiency and elastic storage in the flight motor of *Drosophila*. *Science* 268 (7 April): 87–90.
- [5] Azuma, A. (1992). *The Biokinetics of Flying and Swimming*, Tokyo: Springer-Verlag, pp. 77–154.
- [6] Ellington, C. P., van den Berg, C., Willmott, A. P., and Thomas, A. L. R. (1996). Leading-edge vortices in insect flight. *Nature* 38(D): 626–630.
- [7] van den Berg, C., and Ellington, C. P. (1997). The vortex wake of a “hovering” model Hawkmoth. *Philos. Trans. R. Soc. Lond. B Biol. Sci.* 352: 317–328. Also, ref. 8, The three-dimensional leading-edge vortex of a “hovering” model hawkmoth, pp. 329–340.
- [8] Willmott, A. P., Ellington, C. P., and Thomas, A. L. R. (1997). Flow visualization and unsteady aerodynamics in the flight of the hawkmoth *Manduca sexta*. *Philos. Trans. R. Soc. Lond. B Biol. Sci.* 352: 303–316.
- [9] Willmott, A. P., and Ellington, C. P. (1997). The mechanics of flight in the hawkmoth *Manduca sexta*: I. Kinematics of hovering and forward flight. *J. Exp. Biol.* 200: 2705–2722. Also, ref. 11, The mechanics of flight in the hawkmoth *Manduca sexta*: II. Aerodynamic consequences of kinematic and morphological variation, pp. 2723–2745.
- [10] Liu, H., Ellington, C. P., Kawachi, K., van den Berg, C., and Willmott, A. P. (1998). A computational fluid dynamic study of hawkmoth hovering. *J. Exp. Biol.* 201: 461–477.
- [11] Englar, R. J., Smith, M. J., Kelley, S. M., and Rover, R. C. III. (1994). Development of circulation control technology for application to advanced subsonic transport aircraft: I.

Airfoil development. AIAA paper no. 93-0644, log no. C-8057. *AIAA J. Aircraft* vol. 31, no. 5 (September–October): 1160–1168.

[12] Englar, R. J., Smith, M. J., Kelley, S. M., and Rover, R. C. III. (1994). Development of circulation control technology for application to advanced subsonic transport aircraft: II. Transport application. AIAA paper no. 93-0644, log no. C-8058. *AIAA J. Aircraft* vol. 31, no. 5 (September–October): 1169–1177.

[13] Foch, R. (1999). UAV flight on Mars. *Proceedings of the 1999 Association for Unmanned Vehicle Systems International Symposium*, Baltimore, Md., July. (CD ROM).

[14] Schofield, J. T., Barnes, J. R., Crisp, D., Haberle, R. M., Larsen, S., Magalhaes, J. A., Murphy, J. R., Seiff, A., and Wilson, G. (1997). The Mars Pathfinder atmospheric structure investigation/meteorology (ASI/MET) experiment. *Science* 278(5344): 1752–1758.

[15] McKay, C. P., chair, Mars '03 Airplane Micromission Science Definition Team. (1999). Mars '03 airplane micromission science definition team final report. Document 0907-01-01A, September.

[16] Cutts, J. A., Bauer, J., Blaney, D. L., Lemke, L. G., Smith, S., Stetson, D. S., and Kerzhanovich, V. V. (1998). Prepared in support of the Mars architecture study. On-line: <http://telerobotics.jpl.nasa.gov/aerobot/reports/reports.html>, August 1998.

**This page intentionally left blank**

# **III** Autonomous Behavior

**This page intentionally left blank**

# Autonomous Behavior: Overview

Wendell G. Sykes

There has always been strong interest in the perceptive and cognitive functions that are used by intelligent life to survive in the real world. Given the difficulty of developing an understanding of the detailed working of such functions at the neural level in humans, considerable scientific effort has been devoted to studying these functions in simpler species, as defined by the number of available neurons. There is a logical progression in such studies, starting with the determination of the perceptive functions of a particular species. For example, it is possible to determine with reasonable certainty that a particular species can detect the plane of polarization of incident light [1]. Such determination is, of course, dependent on being sure that the subject of the experiment is actually responding to what the experimenter believes is being tested and not to some unique noise or other effect in the mechanism that creates the sample condition. The story of “Clever Hans,” a horse that was able to give the answer to simple mathematical problems, is worth remembering. The most believable explanation for this capability was that Clever Hans was able to recognize the reaction of members of the audience that knew the answer to the posed question [2].

Given that the species has been trained to emit an observable response when a particular experimental condition is present or absent, the perceptive capability and the limits of this capability can often be determined. Studies of the anatomy of such species can also assist in understanding some aspects of their perceptive systems. Even these approaches may not be effective in some cases. An example is the possibility that some species of birds and insects can sense the structure of the Earth’s magnetic field to provide a directional input. Magnetic materials have been found in conjunction with neurons in these species, though unambiguous results demonstrating sensitivity to magnetic field direction have been difficult to obtain. It is not clear whether magnetic sensitivity is specific to the horizontal direction of the field, like a compass, or its dip angle, which would be related to latitude—or both for that matter. Measurement of directional behavior in the wild may be difficult or impossible. For example, attaching magnets to the heads



of homing pigeons would certainly overwhelm their magnetic-field-sensing capabilities, but other modalities, like vision-based navigation would not be affected [1].

Developing an understanding of how the perceptive capabilities of a species are utilized for those activities associated with its survival is an even more challenging problem and is the subject of the following chapters. There are several aspects to be studied. Given that one or more perceptive capabilities are available, understanding how they are used individually or in concert is a more difficult problem. Beyond this, the structure used to represent and utilize the perceptive results is also difficult to determine. For example, does a species develop a map of its world or does it recognize landmarks in that world?

Chapter 29, by Mandyam V. Srinivasan, is an elegant evaluation of the optical-flow-based navigational capabilities of the honeybee. The experimental conditions chosen allow the experimenter to control an isolated and easily observable behavior that is directly related to the ability of the honeybee to navigate in confined spaces. Changes in experimental conditions provide strong confirmation of one approach to the use of visual information from a multifaceted eye. Extensions of this idea to other guidance applications seem attractive.

Chapter 28 examines the search behavior of the fruit fly, an even simpler organism than the honeybee from the neurological point of view. These insects are quite efficient in localizing food using several perceptive mechanisms. Dickinson, Tammero, and Tarstino have developed an elegant and automated system for tracking the three-dimensional movement of this species in an enclosure that is large compared with their dimensions and local search patterns. The first problem is to gain an understanding of the effects of the experimental system on the behavior of the insect. As the amount of structure or complexity of the walls of the enclosure is increased, the search pattern changes. If landmarks are added to the interior of the enclosure, the observed search behavior changes again, becoming focused on these obstacles. Finally, adding sources of odor produce another change in search behavior. The insect also uses mechanical methods to stabilize its flight. Dickinson, Tammeron, and Tarstino have succeeded in isolating the elements of a complex fused process used by an apparently simple organism.

Another approach to understanding the utilization of perceptive capabilities in a particular species is to develop a robotic system with equivalent capabilities and show that this robotic system can exhibit similar behavior to the modeled species. The remaining three chapters in this section use this approach. Each chapter models a different perceptive mechanism; acoustic, visual, and sensing odor plumes in a fluid medium. Chapter 27, by Grasso, includes an excellent in-depth discussion of the strengths and problems associated with this approach. Other species, such as insects, also use the techniques described by Grasso [3]. Greater understanding of the robust search techniques used

by various odor-tracking species has only been possible as our understanding of the microstructure of odor plumes (as they are shaped by hydrodynamic forces) has improved by sophisticated visualization techniques of such plumes. Grasso shows that the plume-tracking capability of the American lobster can best be accomplished by simultaneous evaluation of odor concentration and hydrodynamic forces. It is probable that species that track odors attached to a substrate, such as dogs and ants, do not require this hydrodynamic understanding.

Webb and Harrison (chapter 26) investigate the acoustic mechanism used by female crickets to home in on the signals emitted by male crickets. This requires that the female select one signal in a background of similar signals from a number of male crickets. Webb and Harrison postulate that female crickets can track using only acoustic signals, but performance is improved when visual information is available to assist the cricket in linear movement. The male cricket is able to generate unique signals that are identifiable by frequency and pulse repetition rate, so the sensing system of the female must be able to utilize these clues. The detection system of the female cricket is much simpler than that of the human ear, using phase difference sensed by a membrane connected to apertures in the legs of the cricket. Webb and Harrison show that an analog of this simple mechanism can detect fine phase differences in the arrival time of a signal, which can be converted into directional information.

There is a continuing debate as to the representational mechanism used by various species when storing visual information that can be used to return to a previously visited location. Does the species develop a template or map of an area, or does it recognize landmarks in the vicinity of the desired location? Möller (chapter 25) has developed a model of the landmark approach based on studies of homing in desert ants. It is believed that these ants must rely on visual clues, because desert conditions will evaporate the odor trails that are used by other types of ants in more-temperate conditions. A number of experiments have shown that landmark-based navigation is the probable mechanism used by this species. Möller assumes that these insects can measure and store the absolute vector angles to two or more landmarks based on the availability of a compasslike directional measurement. Given this information, Möller demonstrates a model algorithm that will generate a vector path to return to a previously recorded location from any point where the landmarks are visible.

## REFERENCES

- [1] Baker, R. R. (1984). *Bird Navigation*. New York: Holmes and Meier Publishers.
- [2] Hauser, M. D. (2000). *Wild Minds*. New York: Henry Holt.
- [3] <http://flightpath.neurobio.arizona.edu/newindex.html>

**This page intentionally left blank**

Ralf Möller

## INTRODUCTION

**Image Matching versus Parameterization**

Insects of many species are repeatedly returning to significant places in their environment, particularly to nest sites and food sources (Collett and Zeil, 1998). The impressive abilities of these animals to precisely locate a goal by using the surrounding landmarks have received interest since the classic experiments by Tinbergen (1932) on digger wasps. A number of studies with hover flies (Collett and Land, 1975), honeybees (Anderson, 1977; Cartwright and Collett, 1983; Gould, 1987), and desert ants (Wehner and R aber, 1979; Wehner, Michel, and Antonsen, 1996) aimed at unraveling which aspects of a landmark configuration are actually memorized and restored during return.

Visual pattern memory, which is underlying both landmark navigation and flower recognition, can be explained on the basis of two different hypotheses: the *template hypothesis* and the *parameter hypothesis* (for reviews, see Heisenberg, 1995; Giurfa and Menzel, 1997). According to the template hypothesis, a photograph-like, retinotopically fixed image is stored, whereas the parameter hypothesis asserts that only specific features of the image, like spatial frequency or orientation, are extracted. From experiments investigating the pattern-recognition abilities of insects, evidence for both template and parameter hypotheses could be gained (Lehrer, 1997; Ronacher, 1998). In landmark navigation, earlier work tended to the parameter hypothesis (Anderson, 1977), but later experiments provided strong support for the template hypothesis (Wehner and R aber, 1979; Cartwright and Collett, 1983; Gould, 1987; Wehner, Michel, and Antonsen, 1996). A link between both hypotheses could be established with the recently suggested *average landmark vector* (ALV) model, which demonstrated that landmark configurations can be compacted into a vector representation without losing information (Lambrinos, 1999; Lambrinos et al., 2000; M oller, 2000).

## Synthetic Modeling Using Analog Hardware

The work presented in this chapter is an example of the *biorobotics* approach to modeling (for reviews, see Beer et al., 1998; Franz and Mallot, 2000; Webb, 2000b), and specifically explores the properties of the ALV model by means of a *neuromorphic* implementation in analog electronic hardware. Neuromorphic implementations (Mead, 1989; Douglas, Mahowald, and Mead, 1995; Indiveri and Douglas, 2000) can be understood as an extremely strict form of modeling, where neural models are implemented in a medium with properties resembling those of biological nervous systems. Specifically, computation in both media is analog, asynchronous, inherently parallel, and a direct consequence of physical operations like summation of voltages or currents (Webb, 2000a). Architectural constraints shared by analog circuits and neural networks will guide the development of models, making a model less arbitrary than a computer simulation. Shifting of image data, for example, appears to be a simple operation in a computer simulation, where it can be accomplished by a single pointer operation, but it proves to be a considerable effort in a neuromorphic implementation.

### Contributions of This Chapter

The analog implementation of the ALV model directly provides insights into the architecture of a neural circuit implementing a visual homing mechanism that is based on the parameter hypothesis. Demonstrating its successful operation on a mobile robot is a rigorous test of the model under real-world conditions. Resemblances between the behavior of the robot and the behavior observed in insect experiments adds validity to the ALV model as a model of insect navigation.

The next section will introduce the snapshot model as an instance of the template hypothesis, and the ALV model as an instance of the parameter hypothesis, and discuss the relation between these models. The analog implementation of the ALV model is presented in the section entitled Analog Implementation on a Mobile Robot; the robot experiments are described in the Experiments section. The Discussion section compares equivalent neural models for the snapshot model and the ALV model, analyzes the differences between both models with respect to the external compass reference, and suggests another homing method, which is, as the ALV model, based on a gradient descent strategy in parameter space.

## MODELS OF INSECT NAVIGATION


### Image-Matching Models

Models of landmark navigation, which are based on the template hypothesis, need to define how a direction of movement can be derived from the discrepancies between a stored image (the template, or “snapshot”) and the current view. An algorithmic model of this matching process, the so-called snapshot model, was suggested by Cartwright and Collett (1983). For simplicity, this model assumes that only the one-dimensional, horizontal portion of each image is considered. Moreover, it is supposed that landmarks can be separated from the background by some means, which leads to a representation consisting of two types of sectors. For each sector in the snapshot, the closest sector in the current view is determined, and each such pair of sectors contributes one or more vectors to a home vector. The vector contributions in the original model were chosen to optimally match the experimental behavior of bees, but mathematically, the model is only one example of a large class of algorithms. Note that all models of this class as well as the ALV model require that the views are aligned to the same coordinate system.

Figure 25.1 (left) shows a version of the snapshot model, called the *difference vector* (DV) model, which has a close relation to the ALV model presented below. In the DV model, the contribution vectors are difference vectors between unit vectors pointing toward the paired landmark sectors in current view and snapshot. If the sector corresponding to the image of landmark  $i$  in the snapshot is in the direction of the vector  $\mathbf{e}_i^0$ , and the closest sector in the current view is the image of the landmark  $M(i)$  in direction  $\mathbf{e}_{M(i)}$ , the sector pair contributes the difference  $\mathbf{e}_{M(i)} - \mathbf{e}_i^0$  to the home vector  $\mathbf{h} = \sum_i (\mathbf{e}_{M(i)} - \mathbf{e}_i^0)$ . Each contribution vector points in a direction of movement that would reduce the difference in bearing between the paired sectors; the home vector is a compromise between all those directions.

### Average Landmark Vector Model

To derive the ALV model from the DV model, it can be assumed that a correct pairing between landmark sectors is established, so that paired landmark sectors are the images of the same landmark. If this condition, expressed as  $M(i) = i$ , is introduced into the home vector equation, the sum can be split into  $\mathbf{h} = \sum_i \mathbf{e}_i - \sum_i \mathbf{e}_i^0$ , which gives the equation for the ALV model (Möller, 2000). The mechanism is graphically explained in figure 25.1 (right): In the target location, the average landmark (AL) vector  $\sum_i \mathbf{e}_i^0$  is determined from the sum of all landmark vectors. This vector is stored instead of the snapshot image. To compute the



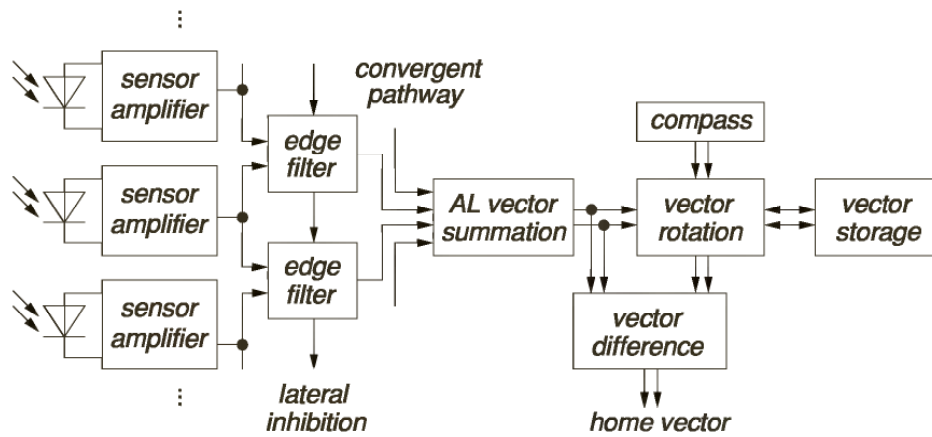
**Figure 25.1** Home vector computation in the difference vector model and in the average landmark vector model, shown for a situation with two landmarks (black circles). (Left) The target location is marked by a cross. The inner ring visualizes the snapshot taken at the target location, the second ring shows the current view as it is visible from the position in the center of the rings. Vectors attached to the outer ring are difference vectors obtained from all landmark pairs. Their sum is the home vector (originating from the center). (Right) The inner ring depicts the current view. At the target location (cross), landmark vectors (attached to the outer ring) are averaged, the resulting average landmark vector (thin, wide head) is stored. The average landmark vector of another location (thin vector, small head) is computed in the same way. The difference between current and stored average landmark vector is the home vector (thick vector).

home vector, the AL vector of the current location  $\sum_i \mathbf{e}_i$  is obtained in the same way, and the difference between the latter and the first yields the home vector.


On the one hand, the ALV model is an instance of the parameter hypothesis, because the two components of the AL vector can be considered as parameterization of the original image information. On the other hand, the ALV model and the DV model yield identical home vectors, at least in the vicinity of the target location where  $M(i) = i$  holds, and no information is lost through this transformation. As can be shown mathematically (Möller, 2000), the AL vector is sufficient to unambiguously characterize each location in the plane. An interpretation of the homing process as gradient descent in a potential function is presented in the Discussion section.

### ANALOG IMPLEMENTATION ON A MOBILE ROBOT

The implementation of the ALV model in analog hardware is based on off-the-shelf electronic components, mainly operational amplifiers. Figure 25.2 gives an overview of the circuit. The details of the circuits are described elsewhere (Möller, 1999, 2000); the Discussion section presents the equivalent neural circuit. Visual input comes from 32 photodiodes, which are mounted in a horizontal ring (figure 25.3). From




**Figure 25.2** Overview of the analog implementation of the average landmark vector model.



**Figure 25.3** Mobile robot with analog implementation of the average landmark vector model.

the amplified photo voltages, edges of one polarity are detected by combining the signals of two neighboring elements; these edges serve as landmark cues. Lateral inhibition between neighboring edge filters ensures that exactly one edge pixel per contrast edge becomes active. This is necessary, because additional or missing edge pixels directly influence the AL vector and therefore entail an erroneous home vector. Up to this point, the circuit is organized in a strictly retinotopic way.






**Figure 25.4** (Left) AL vectors measured while the robot was placed at 64 locations in the arena. Bars on the outside depict landmarks, arrows mark their detectable edges. The vector scaling ( $V$ ) is shown in the upper right corner. (Center) Home vectors measured after storing the AL vector of the target location (cross-circle). (Right) Homing trajectories of the robot towards the target location marked with the cross-circle. Small points are added at the end of the trajectories. At the V-shaped parts of the trajectories, the robot changes the direction.

From the binary outputs of the edge filters, signals converge onto two amplifiers that compute the average landmark vector. This processing step is a weighted summation as it is assumed to be a basic function of biological neurons. The two weights assigned to each edge-filter output encode the landmark vectors in robot-centered coordinates. The subsequent part of the circuit accomplishes the alignment of the AL vectors with an external reference system, which is based on a magnetic compass. The AL vector of the target location is rotated to world coordinates and stored in two capacitors when a button is pressed. After releasing the button, the stored vector is continuously transformed back to robot coordinates and subtracted from the current AL vector, giving the home vector, which directly affects the speed of the two motors.

## EXPERIMENTS


All robot experiments were done in an 1 m by 1 m arena with white walls (30 cm high); black pieces of paper (21 cm by 29 cm) attached upright to the walls served as landmarks (figures 25.4 and 25.5). Each position in the arena is characterized by the corresponding AL vector; figure 25.4 (left) shows the AL vector voltages measured at 64 locations. At the selected target location, the AL vector is stored and subtracted from the AL vector field; the resulting home vector field for that target location is visualized in figure 25.4 (center). Following this home-vector field, the robot will return to the vicinity of the target location from all starting positions in the arena; example trajectories of the robot are



**Figure 25.5** Homing experiment. (Top) The robot was placed at the target position in the center of the arena, and the corresponding AL vector was stored. The other images show the return of the robot from four starting points in the corners of the arena. The white arrow marks the position of the fourth landmark on the inside of the front wall.

plotted in figure 25.4 (right). Figure 25.5 shows the traces of the robot for a target position in the center region of the arena. Systematic experiments with different landmark configurations and target points were conducted to evaluate the precision of homing. For the setup with three landmarks, as in figure 25.4, the mean deviation from the target point was 68 mm; with an additional fourth landmark, as in figure 25.5, the mean deviation was reduced to 48 mm (Möller, 2000). These deviations can be fully explained as an effect of the coarse visual resolution ( $11.25^\circ$  intersensor angle) of the robot (Möller, 2000).

In the experiment shown in figure 25.6, the landmark configuration was changed after the target AL vector was registered. This type of



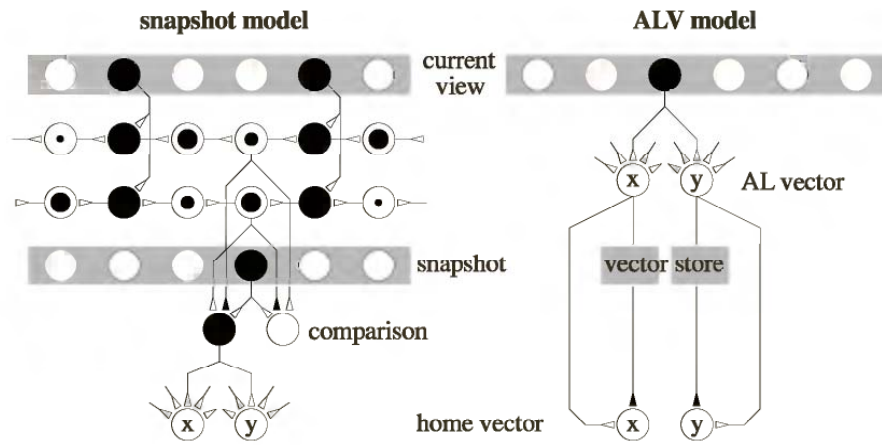
**Figure 25.6** Effect of changes in the landmark configuration. While the robot is in the target location (left), the landmark on the rear wall is removed. The robot moves inside the remaining three landmarks (right). The white arrow marks the position of the fourth landmark on the inside of the front wall.

experiment corresponds to animal experiments where the landmark configuration is modified between training and test. As soon as one of the four landmarks is removed, the robot moves inside the array of the remaining landmarks. This behavior was also observed in bee experiments by Anderson (1977): Bees trained to locate a feeder in the center of a circle of eight landmarks concentrated their search inside the remaining landmarks after three of the landmarks had been removed in the test. Anderson concluded that bees do not store an image but remember the overall gestalt of the landmark array—specifically, the “surroundedness” by landmarks. His model can be interpreted as an instance of the parameter hypothesis; in the ALV model, the parameter surroundedness would be expressed by the AL vector (Möller, 2000).

## DISCUSSION

### Neural Models

A question raised in the discussion of the differences between the template and parameter hypotheses is: Which model requires more effort with respect to the number of neurons and synaptic connections involved? The ALV model as an instance of the parameter hypothesis and the mathematically related snapshot (DV) model as an example for template matching provide a good opportunity to compare the equivalent neural models on equal grounds, because both yield comparable homing behavior. While template matching is sometimes regarded to be a “minimal model for insect visual pattern recognition” (Ronacher, 1998), at least for the task of landmark navigation it seems that template matching requires considerably more-complex neural circuits. Figure 25.7 visualizes the neural circuits suggested for a simple version of the snapshot model (Möller, Maris, and Lambrinos, 1999) and the equiva-




**Figure 25.7** Comparison of the equivalent neural models for a version of the snapshot model and the average landmark vector model. Circles depict neurons, the diameter of the black inner circle shows their activity. Layers with gray background are input layers. Excitatory and inhibitory synapses are shown as white and black triangles, respectively. Gray synapse symbols have weights depending on the spatial position of the presynaptic neuron.

lent neural circuit derived from the analog implementation of the ALV model described above (Möller, 1999). In both cases, circuits required for the detection of landmark features and for rotation are not shown.

In the neural architecture of the snapshot model (figure 25.7, left), correspondences between landmark features in the current view (top gray layer) and in the snapshot (bottom gray layer) have to be established. This can be achieved by injecting activity from the current view into two ring structures. In these rings, activity decays while spreading in one direction. By comparing the activity arriving at a location in the two rings where the snapshot has an active neuron, the direction toward the closest feature in the current view can be determined. In the simplest case, a tangential unit vector can be assigned to one of the two comparator neurons by encoding the components of the unit vector in the two-output weight of the neuron. Two neurons integrate the weighted output of the comparator neurons; their activity corresponds to the  $x$ - and  $y$ -component of the home vector. Details and computer simulations of this model are presented by Möller, Maris, and Lambrinos (1999).

In the equivalent neural architecture for the ALV model (figure 25.7, right), the first processing stage is already a convergent pathway that integrates the weighted outputs of the current-view neurons. In this case, their output weights encode radial unit vectors (landmark vectors), and the summation of the weighted signals yields the two components of the AL vector. The difference between the current and the stored AL vectors gives the home vector. There is experimental evidence that insects are actually equipped with a neural apparatus that can

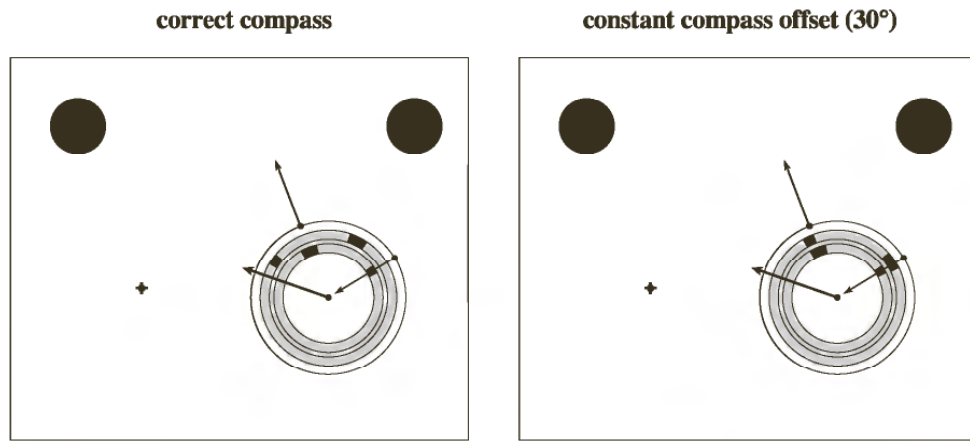


**Figure 25.8** Effect of compass errors on the home position in the ALV model. (Left) Trajectories of the robot, when the arena is rotated clock-wise. Cross-circles mark the starting positions where the target AL vector was stored. (Right) Theoretical iso-angle and iso-length curves for the same landmark configuration. Dots mark the detectable edges of the landmarks. Crosses correspond to the starting points in the robot experiment.

extract parameters like the AL vector: Ernst and Heisenberg (1999) have shown that *Drosophila* flies use a closely related measure—the vertical position of the center of gravity—in pattern-recognition tasks. The simplicity of the ALV model is directly visible in the working real-world implementation: Only 91 operational amplifiers—of which 64 are used to extract landmark features from the signals of the 32 photodiodes—are necessary; in addition, only a few analog switches and multipliers are required for vector rotation and storage.

### External Reference

Extraction of parameters from an image certainly has advantages over a rigid template with respect to reduced storage space and invariance (Heisenberg, 1995), but abandoning image information will always entail some loss of information. This is also the case for the ALV model. The mathematical equivalence between DV and ALV models has been shown under the condition of a perfect alignment between snapshot and current view, or between stored and current AL vector in an external reference system. What is lost in the ALV model is the tolerance that some versions of the snapshot model exhibit against changes in alignment. Alignment errors can be experimentally introduced by rotating the test arena and thus changing the compass signals measured by the robot. Figure 25.8 (left) shows the traces of the robot when the arena is slowly rotated after the AL vector was registered in



**Figure 25.9** Tolerance of a version of the snapshot model against compass errors. Each pair of a landmark image in the snapshot (inner ring) and in the current view (second ring) contributes a radial vector (attached to the outer ring) to the home vector (originating in the center). The direction of the contribution vectors depends on the difference in apparent size between the paired sectors. In the right diagram, a compass offset of  $30^\circ$  leads to a rotation of the current view, but the resulting home vector is identical to the case with correct alignment of the views (left).

the starting point. In the two experiments, the robot roughly moves on different “iso-length” lines of the AL vector: In the transformation from the stored AL vector back to robot coordinates, the angle of the AL vector is changed, but its length remains unaffected. These curves are in accordance with the theoretical iso-curves of the AL vector that are visualized for the same landmark situation in figure 25.8 (right).

That a certain tolerance against compass errors can be observed in template-based models is shown in figure 25.9, for a version of the snapshot model. In this model, only differences in apparent size between the paired landmark sectors are considered; depending on their sign, these differences yield centrifugally or centripetally directed vectors. A rotation of the current view by  $30^\circ$  does not affect the home vector in this case, as long as the pairing between the sectors is not changed.

### Gradient Descent in Parameter Space

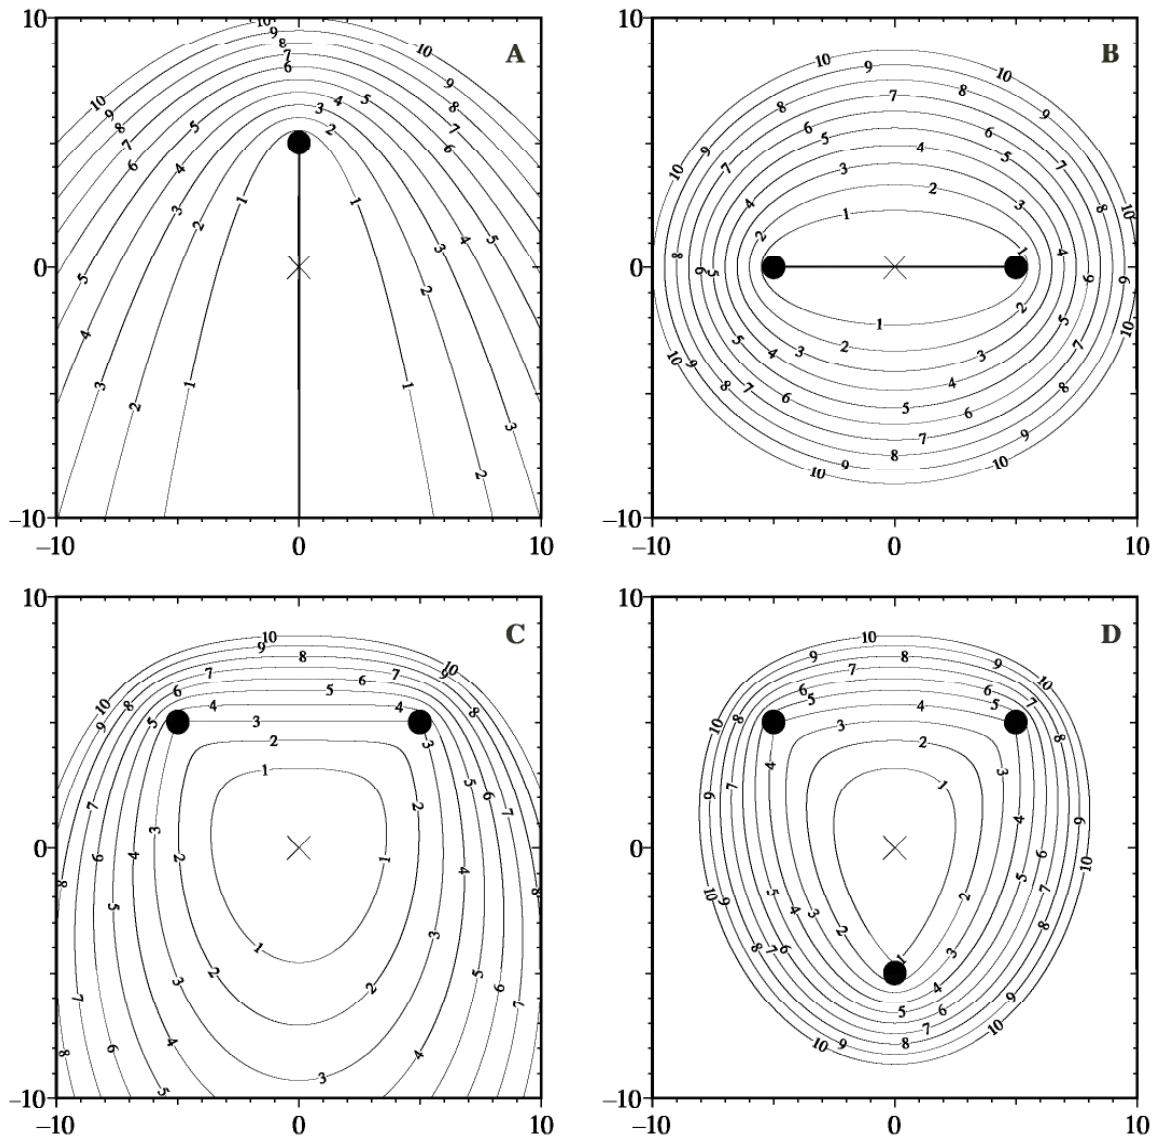
For the landmark navigation task, the matching process between two images results in a vector of movement that will reduce the discrepancies between the images. Similarly, home vectors in a parameter method will point into a direction where some distance measure between stored and current parameters will be decreased. Note that parameterization of images has different goals in pattern recognition: While landmark navigation relies on the spatial variation of the parameters, pattern recognition requires spatial invariance.



The distance between a stored and a current parameter vector can be reduced by following the gradient of the distance function. The ALV model is a somewhat special case of this gradient descent. The home vector  $\mathbf{h}$  is the gradient of a potential function  $U = \sum_i d_i(1 - \mathbf{e}_i^0 \mathbf{e}_i)$ , where  $d_i$  is the spatial distance between the current position and landmark  $i$ . The scalar product relates to the cosine of the angle between the landmark vectors; the potential will only become zero if all landmark vectors are identical with the landmark vectors in the target location. In figure 25.10, this potential function is plotted for different situations. It can be shown that  $U$  is a convex function that, except in pathological cases, has an isolated local minimum in the target location (Möller, 2000), so that the descent along the home vector will lead back to the target location. It is a special property of the ALV model that the gradient can be directly computed from the image information, but it is not possible to compute the potential itself, because it requires knowledge about the distance  $d_i$  to all landmarks and, to compute the scalar product  $\mathbf{e}_i^0 \mathbf{e}_i$ , about their identity.

In contrast to this, a parameter method for landmark navigation will usually compute some distance measure between a stored and a current parameter vector, and use some gradient estimation method to descent in this function to the minimum. Such a parameter model could actually provide a better explanation for some aspects of the behavior of desert ants. Wehner and Räber (1979) trained ants to locate a goal midway between two landmarks. When the distance between goal and landmarks was doubled in the test, the ants exhibited completely different behavior depending on whether landmark size was also doubled or not (figure 25.11, left column). With doubled size, a sharp search-density peak in the correct location is observed similar to the test in the training situation. If, however, the landmark size remains unchanged, ants concentrate their search to a roughly circular area around each of the two landmarks.

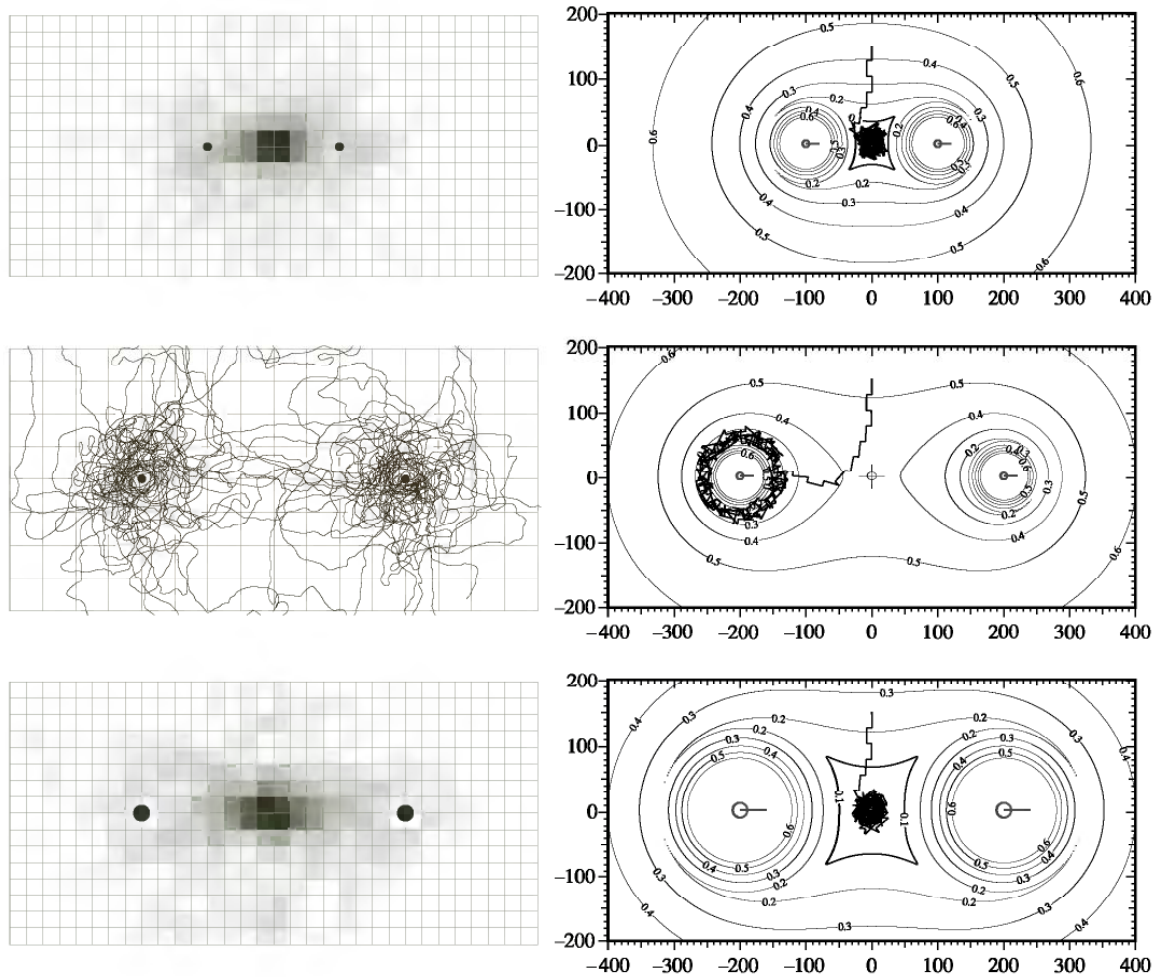
This behavior has recently attracted new interest, because it cannot be reproduced by the template hypothesis. Nicholson and colleagues (1999) and Möller and colleagues (1999) suggested partial image matching as a possible explanation, but according to their models, the ants in the critical experiment should search in two spots halfway between each landmark and the center of the array, which is definitely not visible in the data. The potential function of a parameter model that provides a better explanation is visualized in the right column of figure 25.11. Two parameters are computed from an image of the scene. The first parameter is the length of the landmark contours; the second parameter is the length of the center-of-gravity vector of the landmark contours. The potential function of this model is a weighted sum (ratio 1:4) of the absolute values of the differences between stored and current contour length, and between stored and current length of the center-of-



**Figure 25.10** Potential function of the average landmark vector model for different landmark configurations. Landmark points are marked with black dots, the target location is depicted by a cross. (A and B) Pathological cases with one landmark, and with two landmarks which are in a row with the target location. The potential is minimal (zero) on the thick lines through the target point. (C and D) Typical cases where there is only one isolated point with zero potential in the target location.

gravity vector. The traces shown in figure 25.11 (right) are obtained by a simple descent method, which, after a certain number of steps along the last descent direction, inserts a perpendicular search step that is used to update the estimate of the gradient, and afterward, follows an average of the last and the new direction. This model, which qualitatively reproduces the searching behavior in the critical case, is closely related to the ALV model, but does not require the detection of a fixed number of landmark features. The reproduction of a single experiment





**Figure 25.11** A parameter model as possible explanation for the behavior of desert ants. (Left column) Search density histograms or search paths of three tests with desert ants (original data of experiment by Wehner and Rähler, 1979, used by permission). (Top) Test in the training situation (small, cylindric landmarks close to the goal, indicated by black circles). The goal was located in the center between the two landmarks. (Center) Test with landmark distance from the goal doubled. (Bottom) Test with landmark distance and landmark size doubled. (Right column) Potential function of the parameter model and trajectories (single run, starting at 0, 150) for the setups on the left. Axis labels in cm, scaling identical for left column. Landmark height is depicted by a line.

with this model does not, of course, allow for far-reaching conclusions, but the results indicate that the parameter hypothesis should not be completely neglected as a possible explanation for landmark navigation abilities.

## CONCLUSION

While models of insect landmark navigation currently widely accept the template-matching hypothesis as their basis, this chapter attempts to provide some arguments for the alternative parameter hypothesis.

The arguments are based on the ALV model, which is an extremely parsimonious yet powerful parameter model of navigation. A strictly neuromorphic implementation of the ALV model on a mobile robot demonstrates that the model works under real-world conditions and reproduces some aspects of the behavior of insects. The analog implementation provides direct insight into the architecture and complexity of an equivalent neural circuit. The close relation between the parameter-based ALV model and the template-based snapshot model enables a fair comparison between the two hypotheses. It could be shown that, despite the claim that template matching is the simpler mechanism, the neural circuit of the ALV model is considerably simpler than the circuit of a version of the snapshot model. This became obvious by comparing the analog implementation of the ALV model with a computer simulation of the equivalent neural circuit for the snapshot model. A disadvantage of the ALV model is the lack of tolerance against compass errors, which apparently requires that an image be stored in uncompressed form. Computer simulations show that some observations of the behavior of desert ants, which could not be satisfactorily explained with template-based models, can be reproduced with another parameter model that is closely related to the ALV model.

## ACKNOWLEDGMENTS

Financial support by a personal grant from the Kommission zur Förderung des akademischen Nachwuchses of the University of Zurich is kindly acknowledged. Many thanks to Sepp Ruchti for his valuable contributions and to Verena Hafner for critical comments on the manuscript.

## REFERENCES

- Anderson, A. M. (1977). A model for landmark learning in the honey-bee. *J. Comp. Physiol. [A]* 114: 335–355.
- Beer, R. D., Chiel, H. J., Quinn, R. D., and Ritzmann, R. E. (1998). Biorobotic approaches to the study of motor systems. *Curr. Opin. Neurobiol.* 8: 777–782.
- Cartwright, B. A., and Collett, T. S. (1983). Landmark learning in bees. *J. Comp. Physiol. [A]* 151: 521–543.
- Collett, T. S., and Land, M. F. (1975). Visual spatial memory in a hoverfly. *J. Comp. Physiol.* 100: 59–84.
- Collett, T. S., and Zeil, J. (1998). Places and landmarks: An arthropod perspective. In S. Healy (ed.), *Spatial Representation in Animals*. Oxford: Oxford University Press, pp. 18–53.
- Douglas, R., Mahowald, M., and Mead, C. (1995). Neuromorphic analogue VLSI. *Annu. Rev. Neurosci.* 18: 255–281.
- Ernst, R., and Heisenberg, M. (1999). The memory template in *Drosophila* pattern vision at the flight simulator. *Vis. Res.* 39: 3920–3933.

- Franz, M. O., and Mallot, H. A. (2000). Biomimetic robot navigation. *Robotics Autonomous Syst.*, special issue: *Biomimetic Robots* 30: 133–153.
- Giurfa, M., and Menzel, R. (1997). Insect visual perception: Complex abilities of simple nervous systems. *Curr. Opin. Neurobiol.* 7: 505–513.
- Gould, J. L. (1987). Landmark learning by honey bees. *Anim. Behav.* 35: 26–34.
- Heisenberg, M. (1995). Pattern recognition in insects. *Curr. Opin. Neurobiol.* 5: 475–481.
- Indiveri, G., and Douglas, R. (2000). Neuromorphic vision sensors. *Science* 288(5469): 1189–1190.
- Lambrinos, D. (1999). Navigation in biorobotic agents. Ph.D. thesis, Department of Computer Science, University of Zurich, Switzerland.
- Lambrinos, D., Möller, R., Labhart, T., Pfeifer, R., and Wehner, R. (2000). A mobile robot employing insect strategies for navigation. *Robotics Autonomous Syst.* special issue: *Biomimetic Robots* 30: 39–64.
- Lehrer, M. (1997). Honeybees' visual spatial orientation at the feeding site. In M. Lehrer (ed.), *Orientation and Communication in Arthropods*. Basel: Birkhäuser, pp. 115–144.
- Mead, C. (1989). *Analog VLSI and Neural Systems*. Reading, Mass.: Addison-Wesley.
- Möller, R. (1999). Visual homing in analog hardware. *Int. J. Neural Syst.* 9(5): 383–389.
- Möller, R. (2000). Insect visual homing strategies in a robot with analog processing. *Biol. Cybern.* (to appear).
- Möller, R., Lambrinos, D., Pfeifer, R., and Wehner, R. (1999). Do desert ants use partial image matching for landmark navigation? In N. Elsner and U. Eysel (ed.), *Proceedings of the Neurobiology Conference, Göttingen*. Stuttgart: Georg Thieme Verlag, p. 430.
- Möller, R., Maris, M., and Lambrinos, D. (1999). A neural model of landmark navigation in insects. *Neurocomputing* 26–27: 801–808.
- Nicholson, D. J., Judd, S. P. D., Cartwright, B. A., and Collett, T. S. (1999). Learning walks and landmark guidance in wood ants (*Formica rufa*). *J. Exp. Biol.* 202: 1831–1838.
- Ronacher, B. (1998). How do bees learn and recognize visual patterns. *Biol. Cybern.* 79: 477–485.
- Tinbergen, N. (1932). Über die Orientierung des Bienenwolfes (*Philantus triangulum* Fabr.). *Zeitschrift für Vergleichende Physiologie* 16: 305–335.
- Webb, B. (2000a). Are "biorobots" good models of biological behavior? *Behav. Brain Sci.* (submitted).
- Webb, B. (2000b). What does robotics offer animal behavior? *Anim. Behav.* (to appear).
- Wehner, R., Michel, B., and Antonsen, P. (1996). Visual navigation in insects: Coupling of egocentric and geocentric information. *J. Exp. Biol.* 199: 129–140.
- Wehner, R., and Räber, F. (1979). Visual spatial memory in desert ants, *Cataglyphis bicolor* (Hymenoptera: Formicidae). *Experientia* 35: 1569–1571.

Barbara Webb and Reid R. Harrison

Over the past decade, we have built and tested several robot models to investigate a particular biological behavior, the sound localizing (phonotaxis) ability of the cricket. This work has had several purposes. One is to develop robotic technology, such as novel sensors and control systems, by copying biology. However, the primary motivation is the “reverse”—to use the technology to develop understanding of biological systems—in particular, how neural circuits control sensorimotor behavior. This is effectively a new methodology for biological modeling, discussed in Webb (to appear). In this chapter, the aim is to provide a summary of the problems addressed and the key results to date. (More-detailed presentations of the biological background and the implementations can be found in Lund, Webb, and Hallam, 1997, 1998; Webb and Scutt, 2000; and Webb and Harrison, 2000.)

#### BACKGROUND

Crickets are probably best known for their communication systems, which form an important part of their reproductive behavior and are thus critical to their species survival. We hear male crickets singing; what we usually do not see is the female approaching—able (if necessary) to locate a potential mate by this signal alone. She needs to identify the sound as one produced by a conspecific and not by another species or some other environmental sound source. She needs to determine the direction of the sound despite the noisy environmental conditions and the spatial and temporal limits of resolution of her sensory system. She needs to travel—by flying or walking—over a substantial distance with possibly interfering obstacles and predation risks, without losing track of the target signal. Typically, she is also faced with a choice—several potential mates calling within earshot—that must not result in irresolution but in direct approach to one or another.

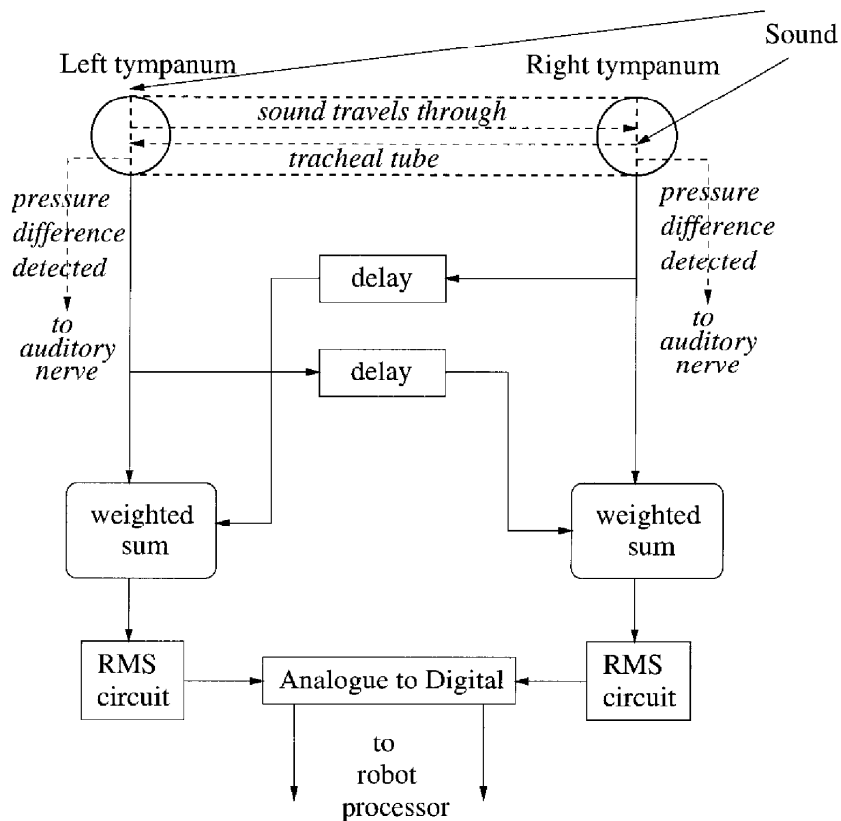
Much behavioral, anatomical, and neurophysiological investigation of this phonotaxis (sound approaching) behavior has been done (e.g., reviewed in Huber and Thorson, 1985; Schildberger, 1988; Weber and

Thorson, 1989; Huber, 1992; Pollack, 1998). It has been established that some of the critical “recognition” cues for the female are the carrier frequency (typically around 4–5 kHz) and temporal pattern (regularly repeated bursts) of the song. The sound is produced by the male moving one wing against the other, which rubs a “file” across a “comb,” the sound amplified by a resonant area of the wing. In the model described below, the main focus is to explain the female cricket’s apparent preference for a particular repetition rate in the signal (typically 20–40 Hz), as established in a number of studies (e.g., Popov and Shuvalov, 1977; Thorson, Weber, and Huber, 1982; Doherty, 1985b; Wendler, 1990). As will be shown, the obvious task division—into first recognizing and then locating the sound—is not the only efficient and effective explanation of the female cricket’s behavior.

In addition, recent work with the robot has studied how the basic sound-localizing behavior might be combined with other sensorimotor tasks. Although the cricket can locate sound in complete darkness, it can also use visual information if it is available. Studies of female cricket tracking have shown apparent improvement in tracking speed and angular accuracy under lit conditions (Thorson, Weber, and Huber, 1982; Atkins et al., 1987; Weber et al., 1987). One explanation might be that in these conditions, the cricket is able to use its optomotor response to maintain a more stable path direction. Evidence for the interaction of optomotor and phonotaxis responses has been provided by Bohm, Schildberger, and Huber (1991), who suggest that these two sensorimotor responses are simply additively combined in the animal’s walking direction. We report some initial experiments on the robot model to examine this idea.

### **Description of the Model**



The “robot cricket” is based on the miniature Khepera robot base.<sup>1</sup> It has an electronic sound-processing circuit, which has been designed to mimic the unique auditory system of the cricket. Crickets have tympanal organs on their legs, connected by tracheal tubes to each other and to other openings on their body (spiracles). Consequently, vibration at the cardrum represents a “pressure difference” between the sound waves incident on the external and internal sides. Because sounds from different directions will travel different path lengths to reach the inside and the outside, the phase cancellation of the signals and resulting amplitude of tympanal vibration will represent the sound direction (Michelsen, Popov, and Lewis, 1994).<sup>2</sup> Thus the cricket has strongly directional hearing (albeit confounded by signal wavelength), despite the relatively small separation of its auditory receptors. The same mechanism is implemented in the robot using programmable electronic delays and subtractions of the signal detected by two microphones placed a



**Figure 26.1** Cricket/robot ears: sound at each eardrum/microphone is combined with delayed sound from the other side. In the cricket (italics and dashed lines) this is via a tracheal tube connecting the tympani; in the robot (solid lines) an electronic circuit serves the same function. The amplitude thus depends on relative phase, which depends on direction and frequency of the sound source. For the robot, the sound was 4.7 kHz, the ear separation 18 mm and the delay 53 msec (i.e., corresponding to  $1/4$  the wavelength).

small distance apart (figure 26.1). One consequence of tuning the separation between the ears and the internal delay to match the carrier frequency of the cricket's calling song is that localization ability is worse for sounds of the wrong frequency (Lund, Webb, and Hallam, 1997). This side-effect is in fact a bonus, as it may help the cricket's selectivity for carrier frequency (e.g., Popov and Shuvalov, 1977; Stout, De Haan, and McGhee, 1983), which is somewhat more sharply tuned than can be explained purely by the frequency tuning in receptors.

The sound signal is processed on the robot using a model neural network that is derived in part from identified neurons in the cricket. A pair of identified ascending interneurons (AN1) in the prothoracic ganglion appear to be critical in controlling phonotaxis (Schildberger and Horner, 1988). They receive direct input from the auditory receptors. Their firing response copies the pattern in normal song, and their firing rate and latency both encode the amplitude. These characteristics are replicated in the robot using a simple state-based "integrate-and-fire"



**Figure 26.2** The four-neuron network used to control behavior. The AN receive sound and the MN control the motor response. Responses of the ears and neurons to a typical input pattern are shown.




Figure 26.2 (continued)

model of membrane potential in single-compartment neurons (Webb and Scutt, 2000). In the model, the two auditory input neurons (AN) are connected to two output or motor neurons (MN) as illustrated in figure 26.2. Each AN makes an excitatory connection to the MN on the same side, and an inhibitory cross connection to the opposite AN-MN synapse. The synapse connections are represented by a single “weight” value, but the weight changes dynamically with the activity of the neuron. Thus the synapses show suppression (the weight is decreased) if a rapid series of spikes arrive at it and require a gap in input to recover. Synapse-on-synapse inhibition similarly decreases the weight at the target synapse.

In this circuit, the spikes in the MNs are used to generate turns by the robot. A sound signal on the left will excite the left AN, which excites the left MN, and also suppresses the excitation of the right MN by the right AN. Effectively, this means that whichever AN fires first will control the response. Thus it is the latency coding of the amplitude that is used, rather than the firing rate. Also, because of the synaptic suppression, only the onset of AN activity contributes significant excitation to MN. Thus the MN response usually requires several onsets, with gaps in between, to reach threshold and fire, signaling a turn. This use of temporal properties of the neural signal to determine turning to the sound inherently makes the neural circuit selective to the pattern in the signal. A continuous or rapidly repeated sound will cause continuous



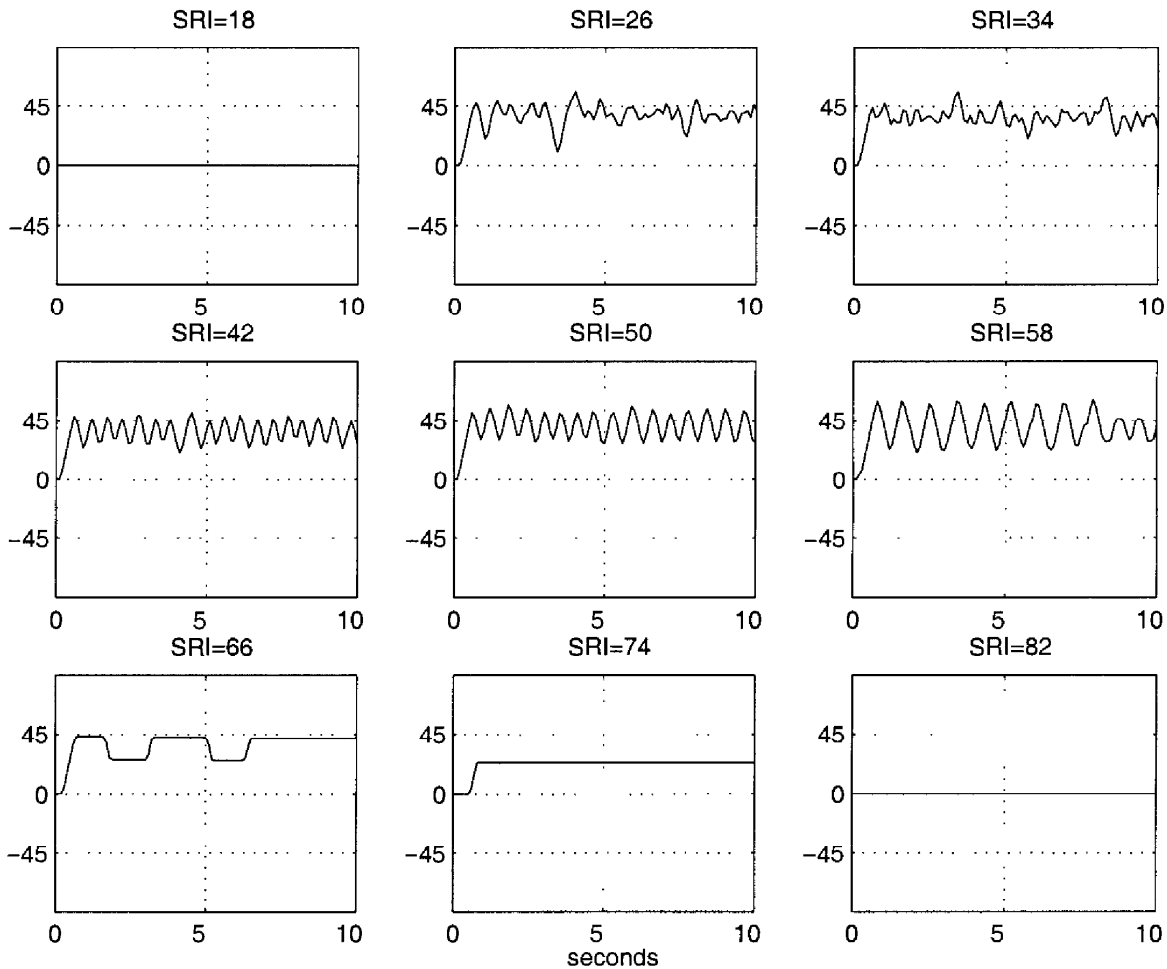
firing in AN, which, due to exhaustion of the synaptic connection, will cease to contribute to MN. A slowly repeating sound will not allow successive activations of MN to summate above threshold and cause a turn. The ideal signal for generating turns will fall within a specific band of repetition rates. Thus in the process of locating the sound, the signal is also “recognized,” in that incorrect signals will fail to generate appropriate movements toward the sound.

## Results

Although this four-neuron network appears very simple, it is in fact able to account for a wide range of observed behaviors in the cricket, including behaviors usually taken as evidence for separate recognition and choice circuits (Webb and Scutt, 2000). We were able to test the robot using stimulus and behavior paradigms comparable to the cricket. For example, constraining the robot to turn only on the spot and not move closer to the sound mimics the behavior of a cricket on a Kramer treadmill (Weber, Thorson, and Huber, 1981). Playing recorded cricket songs from a speaker under these conditions leads the robot to oscillate about the sound direction, reliably following a switch in speaker direction (cf. the criteria for phonotaxis suggested by Thorson, Weber, and Huber, 1982). This behavior can then be used as an assay for the preference shown for simulated signals with varying characteristics.

First the robot was tested with simulated songs of different syllable repetition intervals (SRI). As noted above, the SRI appears to be a critical recognition cue for the cricket. The robot was tested with a 4.7 kHz tone repeated at intervals of 18, 26, 34, 42, 50, 58, 66, 74, and 82 milliseconds, with the duty cycle held constant at 50%. As illustrated in figure 26.3, it showed no tracking to the fastest rate or slowest rates and did track between 26 and 58 msec SRI, with the best (closest) tracking between 34 and 50 msec SRI. Thus this network successfully tracks only the correct syllable repetition rates and effectively ignores other signals. The preference closely resembles that shown for crickets (Thorson, Weber, and Huber, 1982).





It would be fair to argue that this behavior is not “true” recognition: But what is the evidence that the cricket actively recognizes correct signals rather than, like the robot, merely fails to track incorrect ones? One paradigm that has been taken to indicate an independent “recognizer” uses a song played from a speaker directly above the animal, which provides no consistent directional cue (Weber, Thorson, and Huber, 1981; Schmitz, Scharstein, and Wendler, 1982). Crickets nevertheless show “phonotaxis-like” behavior—that is, they modify their walking speeds and make turns in a manner that is not normally seen when not trying to track sound. Moreover, when a continuous signal is then




**Figure 26.3** The robot tracking songs with differing syllable rates. The speaker is at  $45^\circ$ , and the robot's heading over time is plotted; it only tracks syllable repetition intervals between 26 and 58 msec. (Reprinted from Webb and Scutt 2000.)

added from one side of the animal (Stabel, Wendler, and Scharstein, 1989), it attempts to turn *away* from this sound—that is, toward the side that has a clearer song signal rather than the side that has a higher amplitude response. This has led to the suggestion that “the signals ... are filtered individually on each side with respect to the syllable and/or chirp frequency before the right/left comparison occurs” (Stabel, Wendler, and Scharstein, 1989, p. 175).

We tried testing the robot under the same conditions and the results are shown in figure 26.4. Sound from above, relative to no sound, clearly produces a change in the robot's behavior—in which it appears to be tracking, but not in any consistent direction. This response occurs because despite the equal amplitude of sound reaching the microphones, slight differences in the processing circuitry are enough to produce slight differences in the neural response and, consequently, turning










**Figure 26.4** The robot “tracking” sound from directly above: (a) when a recorded song is switched on it starts to meander; (b–c) the tracking direction changes over time; (d–e) the robot tracks 180° away from a continuous sound source added at 45° (cf. Stabel, Wendler, and Scharstein, 1989; reprinted from Webb & Scutt 2000).

behavior. There seems no reason to suppose that the cricket’s auditory and neural systems are more reliable than the robot’s in this situation—that is, that a sound directly above the animal results in precisely equal activation of each side of the sound processing system. Moreover, adding the continuous signal from one side produces a response in the robot directly analogous to the cricket—it turns and tracks 180° away from the continuous sound source direction. This behavior can be explained in just the same way as for the cricket—the song pattern is better represented on the side away from the continuous signal, and as the pattern is necessary to produce a response in MN, the robot turns that way. But neither explicit filtering nor explicit comparison takes place in the robot.

A third set of experiments examined the ability of this simple circuit to re-create the choice capabilities exhibited by female crickets. The fact that crickets can track one from several simultaneous songs and apparently choose between attractive signals has also been cited as evidence for more-complex recognition processes than the simple circuit we have proposed (Doherty, 1985a; Weber and Thorson, 1988). Although it had not been designed with this issue in mind, it was simple to test what the robot would do when played similar or differing signals simultaneously. The results of several experiments are shown in figure 26.5. Faced with identical signals (recorded cricket song) from two speakers, the robot turns and reliably tracks one of them, rather than getting confused. The one chosen can be controlled by varying the amplitude; the louder sound is preferred. When two songs differing slightly in syllable-repetition interval (40 vs. 60 msec) are played, the robot tracks the faster song. The same behavior is observed in the cricket.





**Figure 26.5** The robot “choosing” the louder of two identical sound sources, or the faster of two acceptable syllable rates. (Reprinted from Webb and Scutt, 2000.)

### ADDING AN OPTOMOTOR RESPONSE

So far, we have shown that a simple four-neuron network appears to be sufficient to explain much of the cricket’s phonotaxis behavior. The cricket actually has around 100,000 times more neurons than this—so it might be asked: Why? One obvious difference is that the motor control system of a two-wheeled robot is simpler than the cricket’s six multi-jointed legs. Another is that the robot only does phonotaxis, while the cricket has a wide range of other behaviors and sensorimotor systems. Moreover, these must in some way be coordinated. As a starting point for investigating some of these issues, we have looked at integrating a visual response into the phonotaxis model. If a simple visual reflex such as the optomotor response is added to the robot, how can the behavioral controller best utilize both signals? Can the observed improvement in the cricket’s tracking behavior under lit conditions (Weber, Thorson, and Huber, 1981; Weber et al., 1987) be replicated?

#### Description of the Model

The implementation of the optomotor response was achieved in collaboration with Reid Harrison, who has developed a neuromorphic analog VLSI (very-large scale integration) chip that processes light signals in a manner akin to the known neural circuitry in insect brains (described in chapter 2 of this volume). Briefly, the chip computes a local measurement of motion between adjacent pairs of photoreceptors using delay and correlation, implemented by inherent lag in low-pass filters and multiplier circuitry. The output is a summation across the circuit, which is further low-pass filtered (time-constant = 100 msec) to provide an analog signal that indicates the approximate velocity of the visual field. Using this signal to bias the speed of the robot’s motors, with an appropriate gain, creates an optomotor reflex—that is, when a

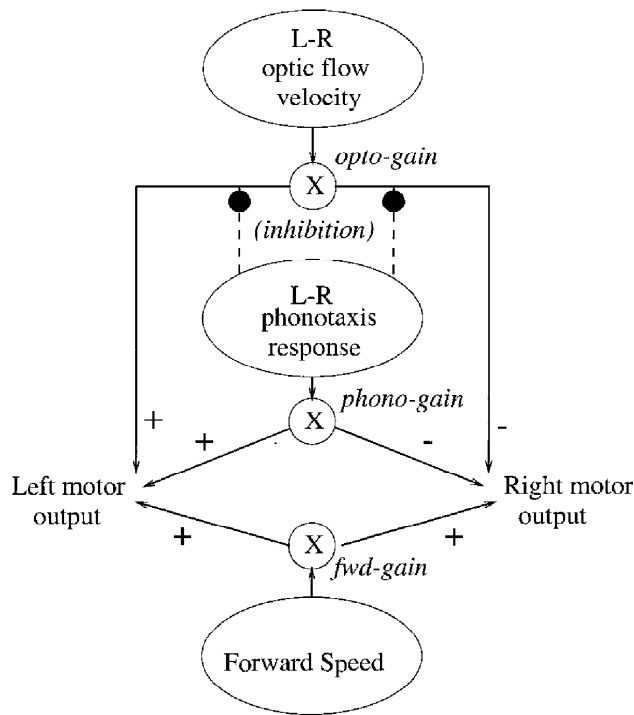
deviation from straight-ahead movement is signalled by a rotating visual field, the robot applies corrective compensation.

So, can this optomotor correction simply be added to the established phonotaxis response and produce improved behavior? To test this, it was necessary to reimplement the phonotaxis system on a larger robot that could also carry the lens and circuitry used with the VLSI chip. We used a Koala robot because it is designed to be closely compatible with the Khepera. The same electronic ear circuitry and spiking-neuron processing program was used for the phonotaxis as described above. Given the output of the phonotaxis controller—a fixed-length turning signal is produced each time an MN neuron spikes—the issue was how to modify this signal so as to integrate the visual response. Initially, this was done by simply adding the optomotor gain directly to the motor commands that had been produced by the sound (figure 26.6).

An immediately apparent problem was that any turn toward the sound generated by the phonotaxis system produced an excellent “visual rotation” signal, which the optomotor system would correct (i.e., turning the robot back away from the sound). This was a convincing empirical demonstration of what Robert and Rowell (1992) describe as an “old conundrum”—how does an animal operating with an optomotor response ever manage to make intentional turns? There are several “old” proposed solutions. Perhaps the best known is the principal of *reafference* (von Holst and Mittelstaedt, 1950)—that when making a turn, a predictive countersignal is sent to cancel out the anticipated optomotor signal. Another suggestion is that the intentional turns might be fast or slow enough to fall outside the range of movement normally detected by the optomotor system. However, having already chosen biologically justified ranges for turning speeds in our system, it seemed inappropriate to change these. A third solution is that the intentional movement simply suppresses or switches off the optomotor response.

Reafference is an attractive principal, but its practical application is more problematic. A perfect reafferent signal would consist of the inverse of the expected optomotor pattern—but how can the reafferent system precisely predict what it will see when the optomotor response depends not only on the animal’s rotational velocity but also on the (unknowable) spatial patterning in the visual field? An approximate signal is easier to achieve, and could be sufficient, but would at minimum require tuning—not only in magnitude but also in timing. The command to turn and the visual feedback are not coincident, so it becomes important to supply the right “cancellation” signal at the right time.

A switching or suppression scheme has the advantage of simplicity over reafference. There is also biological evidence from several different systems that optomotor reaction might be simply switched off during intended turns; for example, in response to “escape” signals in the lo-




**Figure 26.6** Controllers that combine phonotaxis with an optomotor response using simple addition of the two turning tendencies, with the phonotaxis response inhibiting the optomotor response.

cust (Robert and Rowell, 1992), and during pursuit turns in the housefly (Srinivasan and Bernard, 1977). Heisenberg and Wolf (1988) suggest, for *Drosophila*, what could be called a “hybrid” reafferent-suppression scheme, in which visual motion in the expected direction is suppressed but in the opposite direction is still used. It should be noted that the data taken to support the “summation” of the phonotactic and optomotor responses in the cricket (Bohm, Schildberger, and Huber, 1991) is in fact consistent with switching because it is based on average direction of tracking. If the animal alternately responds to the visual and the auditory signal, on average it will steer a course between them.

Thus we decided to test the robot with a simple switching mechanism based on inhibition (figure 26.6). The motor signal is normally modified by the optomotor gain signal. When the phonotaxis gain signals a turn toward the sound, it also inhibits the optomotor response. The robot was first tested without the optomotor response to create a baseline for comparison and ensure that the mechanism that had worked on the smaller robot also worked on this one. We used an “arena” paradigm (cf. Atkins et al., 1987 for crickets) in which the robot was started from several points in the room and had to approach the sound source. One difficulty was that the larger robot was moving a larger distance (several meters) and consequently needed to be able to






**Figure 26.7** Tracks of the Koala robot toward the sound with no optomotor response. The robot is fairly successful in tracking the sound.

deal with a greater range of sound amplitudes while still responding appropriately. We solved this problem by imposing an approximate logarithmic compression to the input amplitude before feeding it to the neural processor.

## Results

With just the phonotaxis system operating, the robot was quite reliably able to find the sound source when started either directly opposite the speaker (10 trials) or at either side of the room (6 trials each). The tracks shown (figure 26.7) were reconstructed from the shaft encoders on the robot. The two failures to reach the sound seemed principally due to adverse echoes in the unsoundproofed environment.


Next, we added the optomotor response, and ran the same sequence of trials. In this case, the robot never failed to find the speaker, and the tracks, on visual inspection, look a little more direct (figure 26.8). However, we found no statistically significant difference in the mean lengths or heading errors of the paths, though there was a difference in the



**Figure 26.8** Tracks of the Koala robot toward the sound with an optomotor response. There appears to be a slight improvement, but not a substantial one.

variance of these measures—that is, without optomotor control, a wider range of paths (both more and less direct) were seen. The lack of obvious improvement can be largely attributed to the fact that phonotaxis control was already fairly good at getting the robot directly to the sound source, because there was little else to put it off course once it had turned to face the right direction. The main function of an optomotor response is to maintain a straight-line heading when subject to disturbance or drift. The Koala robot, when instructed to move in a straight line on a flat floor, has little difficulty maintaining a straight course and hence little to gain from an optomotor reflex.

Hence we decided to add a disturbance to the robot's normal behavior: a systematic bias in its normal motor speeds, such that the left motor would run 20% faster than the right. Most crickets walking on a treadmill show some systematic directional bias (Schul, 1998), and crickets in their natural environment often have motor asymmetries as well as environmentally caused deviations. With this bias, the robot without phonotactic or optomotor control would turn in a circle. With



**Figure 26.9** (Left) Tracks of the robot with a biased motor system and no optomotor response. The robot is less successful in tracking the sound. (Right) Tracks with a biased motor system and an optomotor reflex. The robot is able to track the sound.

optomotor control (and no sound), it was able to correct for the bias to make an approximately straight path (this replicated the experiments described in Harrison and Koch, 1999).

Figure 26.9 (left) illustrates 10 sound-tracking trials with the bias and without the optomotor reflex. The robot still reaches the sound in half the trials, but the tendency to head to the right of the speaker is clear. In one case, it loses the sound altogether; in another, it makes a complete circle in the bias direction before moving to the sound. When the optomotor response is added (figure 26.9, right) there is a clear improvement in the behavior. In fact, the tracks become comparable to those in the original, no-bias conditions (see figures 26.7 and 26.8).

## CONCLUSIONS

Using robot technology, we are able to rigorously test hypotheses about biological sensorimotor control. We can determine, by testing the robot




Figure 26.9 (continued)

in the same experimental paradigms as the animal, whether our understanding of the system is sufficient to replicate the behavior it displays. A robot has the advantage over simulation such that you do not have to simulate the environment in which it operates, which can not only be difficult but positively misleading if some real-world effects are inadvertently left out. From the robotics perspective, the animal provides a number of novel insights into ways of achieving efficient behavioral control.

In the work described here, we have shown that a relatively simple neural control system can suffice to reproduce a surprisingly wide variety of the female cricket's behavior in approaching sound. A key reason for this success is the tuning of sensory mechanisms and motor actions to the task environment. We used sensors that were closely based on the animal's sensors, and thus well matched to the task. Another is the explicit use of some of the temporal processing properties inherent in real neurons (which are often ignored under the traditional firing-rate interpretation of neural function). Because the circuit uses onset latencies, it is inherently sensitive to temporal patterns, and this can suffice

to copy the cricket's apparent preferences for certain sounds. To say whether this is actually how the cricket does it requires further biological experimentation; we claim only that the evidence available so far does not require a more complex explanation than that embodied in the robot.

The extension to incorporate an optomotor response indicated that while simple summation may not suffice to successfully combine the behaviors, simple suppression is adequate and a reafferent signal is not needed. When the robot's motor task is made more difficult, the optomotor response significantly improves the phonotaxis behavior. We plan further investigations in this critical area of multimodal sensor fusion, using additional sensors such as antennae and more-realistic motor situations, including outdoor terrain and legged locomotion.

## NOTES

1. While a legged robot would more accurately represent the cricket, nearly all available data on cricket behavior is at the level of whole-body trajectories (not individual leg movements), which a simple wheeled robot is able to replicate. A legged implementation is one of the intended outcomes of a current research project in collaboration with Quinn and Ritzmann (see chapter 20 of this volume).
2. In fact, the story is a little more complicated, involving membranes within the tracheal system that actively delay the internal signal.

## REFERENCES

- Atkins, G., Atkins, S., Schoun, D., and Stout, J. (1987). Scototaxis and shape-discrimination in the female cricket *Acheta domestica* in an arena and on a compensatory treadmill. *Physiol. Entomol.* 12: 125–133.
- Bohm, H., Schildberger, K., and Huber, F. (1991). Visual and acoustic course control in the cricket *Gryllus bimaculatus*. *J. Exp. Biol.* 159: 235–248.
- Doherty, J. A. (1985a). Phonotaxis in the cricket, *Gryllus bimaculatus* De Geer: Comparisons of choice and no choice paradigms. *J. Comp. Physiol. [A]* 157: 279–289.
- Doherty, J. A. (1985b). Trade-off phenomena in calling song recognition and phonotaxis in the cricket, *Gryllus bimaculatus* (Orthoptera Gryllidae). *J. Comp. Physiol. [A]* 156: 787–801.
- Harrison, R. R., and Koch, C. (1999). A robust analog VLSI motion sensor based on the visual system of the fly. *Autonomous Robotics* 7: 211–224.
- Heisenberg, M., and Wolf, R. (1988). Reafferent control of optomotor yaw torque in *Drosophila melanogaster*. *J. Comp. Physiol. [A]* 163: 373–388.
- Huber, F., and Thorson, J. (1985). Cricket auditory communication. *Sci. Am.* 253: 47–54.
- Huber, F. (1992). Behavior and neurobiology of acoustically oriented insects. *Naturwissenschaften* 79: 393–406.
- Lund, H. H., Webb, B., and Hallam, J. (1997). A robot attracted to the cricket species *Gryllus bimaculatus*. In *Fourth European Conference on Artificial Life*, Brighton, England. Cambridge: MIT Press.

- Lund, H. H., Webb, B., and Hallam, J. (1998). Physical and temporal scaling considerations in a robot model of cricket calling song preference. *Artif. Life* 4: 95–107.
- Michelsen, A., Popov, A. V., and Lewis, B. (1994). Physics of directional hearing in the cricket *Gryllus bimaculatus*. *J. Comp. Physiol. [A]* 175: 153–164.
- Pollack, G. S. (1998). Neural processing of acoustic signals. In R. R. Hoy, A. N. Popper, and R. R. Fay (eds.), *Comparative Hearing: Insects*. New York: Springer, pp. 139–196.
- Popov, A. V., and Shuvalov, V. F. (1977). Phonotactic behaviour of crickets. *J. Comp. Physiol. [A]* 119: 111–126.
- Robert, D., and Rowell, C. H. F. (1992). Locust flight steering: II. Acoustic avoidance maneuvers and associated head movements, compared with correctional steering. *J. Comp. Physiol. [A]* 171: 53–62.
- Schildberger, K. (1988). Behavioural and neuronal methods of cricket phonotaxis. *Experientia* 44: 408–415.
- Schildberger, K., and Horner, M. (1988). The function of auditory neurons in cricket phonotaxis: I. Influence of hyperpolarization of identified neurons on sound localisation. *J. Comp. Physiol. [A]* 163: 621–631.
- Schmitz, B., Scharstein, H., and Wendler, G. (1982). Phonotaxis in *Gryllus campestris* L.: I. Mechanism of acoustic orientation in intact female cricket. *J. Comp. Physiol. [A]* 148: 431–444.
- Schul, J. (1998). Song recognition by temporal cues in a group of closely related bush-cricket species (genus *Tettigonia*). *J. Comp. Physiol. [A]* 183: 401–410.
- Srinivasan, M. S., and Bernard, G. D. (1977). The pursuit response of the housefly and its interaction with the optomotor response. *J. Comp. Physiol. [A]* 115: 101–117.
- Stabel, J., Wendler, G., and Scharstein, H. (1989). Cricket phonotaxis: Localization depends on recognition of the calling song pattern. *J. Comp. Physiol. [A]* 165: 165–177.
- Stout, J. F., De Haan, C. H., and McGhee, R. (1983). Attractiveness of the male *Acheta domestica* calling song to females: I. Dependence on each of the calling song features. *J. Comp. Physiol. [A]* 153: 509–521.
- Thorson, J., Weber, T., and Huber, F. (1982). Auditory behaviour of the cricket: II. Simplicity of calling-song recognition in *Gryllus* and anomalous phonotaxis at abnormal carrier frequencies. *J. Comp. Physiol. [A]* 146: 361–378.
- von Holst, E., and Mittelstaedt, H. (1950). Das Refferenzprinzip. Wechselwirkungen zwischen zentralnervensystem und peripherie. *Naturwissenschaften* 37: 464–476.
- Webb, B. (to appear). What does robotics offer animal behaviour? *Anim. Behav.*
- Webb, B., and Harrison, R. (2000). Eyes and ears: Combining sensory motor systems modelled on insect physiology. In *IEEE International Conference on Robotics and Automation*, San Francisco, 24–28 April. pp. 3913–3918.
- Webb, B., and Scutt, T. (2000). A simple latency dependent spiking neuron model of cricket phonotaxis. *Biol. Cybern.* 82: 247–269.
- Weber, T., Atkins, G., Stout, J., and Huber, F. (1987). Female *Acheta domesticus* track acoustical and visual targets with different walking modes. *Physiol. Entomol.* 12: 141–147.
- Weber, T., and Thorson, J. (1988). Auditory behaviour in the cricket: II. Interaction of direction of tracking with perceived temporal pattern in split-song paradigms. *J. Comp. Physiol. [A]* 163: 13–22.

Weber, T., and Thorson, J. (1989). Phonotactic behaviour of walking crickets. In F. Huber, T. E. Moore, and W. Loher (eds.), *Cricket Behaviour and Neurobiology*. Ithaca: Comstock. pp. 310–339.

Weber, T., Thorson, J., and Huber, F. (1981). Auditory behaviour of the cricket: I. Dynamics of compensated walking and discrimination paradigms on the Kramer treadmill. *J. Comp. Physiol. [A]* 141: 215–232.

Wendler, G. (1990). Pattern recognition and localization in cricket phonotaxis. In *Sensory Systems and Communication in Arthropods*. Basel: Birkhauser Verlag.

# Flow and Chemo-Sense for Robot and Lobster Guidance in Tracking Chemical Sources in Turbulence

Frank W. Grasso

## TWO PURPOSES OF BIOMIMETIC ROBOTS

Biomimetic robots can be put to two uses. They can be employed as tools to answer questions about biological processes and, as the end products of biological research, they can lead to useful technologies that perform tasks with an efficiency approaching that of animals. Once the level of animal competence has been approached or reached, the processes may be refined to exceed in desirable dimensions the capacities of the animal models. My focus is primarily in the first phase of using biomimetic robots as aides to developing a deeper understanding of animal behavior and its neural basis. In this chapter, I will focus on an approach that aims first to bring robot performance to a state where it reproduces the animal behavior, and places the translation of those principles into usable technologies second.

Through this methodology, we cannot “prove” that the algorithms our robots implement are the biological mechanisms. Those mechanisms can only be inferred from biological experiments. But, in the application of the scientific method, we can exclude untenable biological explanations by the simple expedient that they fail to account for the biological behavior. As this process of excluding the biologically implausible proceeds, something unrelated occurs in parallel. The ability of the robots to track chemical plumes to their source will improve as their performance approaches that of real animals. Thus, although there is no scientific guaranty that the robots will ever match their biological models, there will be a systematic approach to an attainable benchmark. The algorithms gleaned from these biologically inspired studies can find practical application as components in autonomous systems intended for exploration.

My use of the term *biomimetic* is not intended to denote a slavish mimicry of all the biological details. The mimetic aspect is conceptual first and structural only as needed to support the conception. The idea is to reproduce only those points that are essential to test a given hypothesis. Thus, to test the power of odor-gated rheotaxis (OGR) or



tropotaxis to explain plume tracing by lobsters, we constrained the critical features of body size and shape, sensor arrangement in space, speed and pattern of locomotion, and the temporal and spatial resolution of the sensors in the robots' hardware and software to known lobster values. (The details can be found in Grasso et al., 1999.)

In the study of biological systems, there are two broad areas in which biomimetic robots contribute. First, they allow for the direct test of biological hypotheses themselves. If an appropriate implementation of a theory fails to replicate the biological process it is intended to explain, that failure constitutes strong grounds for rejecting that hypothesis. Another application is exploring the constraints of the sensory and motor world of an animal. The behavior produced and data collected from a properly scaled autonomous robot in the natural environment of the animal it emulates are about as close as we can come to experiencing the world of another species. This can be very instructive for researchers working with animals that have sensory modalities of motor capabilities that we humans lack. I shall reflect on both of these applications of biomimetic robots in detail in this chapter.

Biomimetic robots naturally aim to close the sensorimotor loop. However, a complete description of the input-output relations implemented by a simple organism like the nematode *Cnaenorhabditis elegans* is not practically workable today with any biological validity. Instead, one must have a well-defined problem and a workable potential solution to that problem. Certainly, animals can be found with capabilities that are desirable in robots. Researchers whose aim is to implement animal algorithms and approach animal efficiency in these areas face the burden of having the animal solution, or an approximation of it, in hand before they begin construction. The extent of this self-imposed burden depends upon the vagueness with which the builder is "inspired" by the natural system. It is clear that an engineer could design a system without reference to biological inspiration. The success of such a project would depend on the engineer's grasp of the problem and the quality of the solutions that are in hand to be brought to bear on the problem. To take advantage of the "existence" proofs of solutions that animals represent, much more biological inspiration, even constraints, must be included. Future designers of practical biomimetic robots will rely on understanding of the model animals gleaned from biological research. For such an understanding to be adequate, it will most likely have one foot in traditional biological research and one foot in previous generations of biomimetic robots used as tools of biological research.

### THREE CRITICAL QUESTIONS

As a result of my experiences in the study of robots and animals, three issues have risen to prominence in my thinking about projects involving biomimetic robots:

- What constitutes a sufficient understanding of the task, the animal, its behavior, and its information-processing abilities to ensure that a biomimetic robotic implementation will lead to useful results?
- What is the proper level of abstraction to represent a biological function in a robot?
- What constitutes good measures of robot performance?

These are good questions for researchers planning on using robotics to understand biological phenomena and for those interested in the development of practical robots from biological inspiration to ask themselves in the early planning phases of a biomimetic robotic project. The answers, even partial answers, to these questions, obtained from the perspective of a particular task or biological system, are more useful to have in hand at the start of a biomimetic robotics project than they are as discoveries at the end of one. I will review the results of studies exploring the nature of lobster chemo-orientation in turbulence with two generations of autonomous underwater biomimetic robots and use them to illustrate the effect answers to these questions have on the biomimetic robotic approach to understanding biological phenomena and the production of useful robots.

## **THE ROBOLOBSTER AND BICSAAR EXPERIMENTS**

The RoboLobster and the (second generation) BICSAAR robots (biologically inspired chemical-sensing aquatic autonomous robot) experiments and design specifications have been described in previous publications. In this chapter, I summarize the key features of that work to illustrate the general points. The reader interested in the details of the robot and chemo-orientation is directed to those papers (Consi et al., 1994, 1995; Grasso et al., 1996, 1997; Grasso, Basil, and Atema, 1998; Grasso et al., 1999).

## **THREE CRITICAL DIMENSIONS OF BIOMIMETIC ROBOTICS**

Marine species appear to have evolved a variety of strategies to deal with chemo-orientation in the sea (Dodson and Dohse, 1984; Moore, Scholz, and Atema, 1991; Dunsenberry, 1992; Weissburg and Zimmer-Faust, 1994; Zimmer-Faust et al., 1995; Baird, Johari, and Jumper, 1996; Cox et al., 1997; Dale, 1997; Moore and Lepper, 1997; Barbin, 1998). The spatial dynamics of the fluid medium and the animal's speed of locomotion appear to dominate the strategy a species uses. Some, and probably many species not yet studied under more than a few fluid dynamics conditions, appear to have evolved multiple strategies that they switch on and off depending on context (Moore, Scholz, and Atema, 1991; Dunsenberry, 1992; Oliver, Grasso, and Atema, 1996; Beglane et al., 1997; Mjos, Grasso, and Atema, 1999).

The aim of the RoboLobster and BICSAAR projects is to evaluate alternative hypotheses of lobster chemo-orientation in the American lobster, *Homarus americanus*. Chemo-orientation trials with robots were one component of three parallel lines of research aimed at explaining lobster chemotaxis. The other two critical components involved understanding the behavioral and neural information-processing abilities of the lobster and achieving a suitable understanding of the cue environment in which the lobster tracks odors to their source. Viewed as a biomimetic robotics problem, the task and its biological solution are a mixture of environmental and animal processes.

This presents something of a paradox: We want to understand the task and its solution well enough to specify a robot, but we simultaneously want the robot to delineate the valid aspects of the animal and its environment so that we can understand the solution. Our resolution of this paradox was to divine constraints from studying both the animal and the task environment and then to use the robot's ability (or failure) to reproduce lobster behavior to discriminate the relevant and irrelevant aspects of the system. Our studies would have been mired in free parameters without these constraints.

One important consideration in our studies has been that the animal behaviors, the fluid mechanics of the odor dispersal, and the robot performance have all been studied under identical physical conditions. Thus, we know that the lobsters are able to track the same plumes with which we test the robots. And, we can use the lobster behavior for quantitative benchmarks of robot performance. With statistical data about the temporal and spatial dispersal of odor, we can interpret both lobster behavior and robot performance in terms of the environmental conditions.

## **The Environment**

To understand the task an American lobster faces as it tracks an odor plume to its source, we must place ourselves in the lobster's situation.

Chemicals or mixtures of chemicals are referred to as odors when we consider them as percepts organized by the groupings imposed by a nervous system. The identities of these odors are a brain function. Our studies do not concern themselves with odor identities and they assume that one odor is transported by turbulent dispersal in the same way as another; and we use tracers as model odors in our studies.

Diffusion is too slow to move chemicals over distances of more than a few microns in the time that a lobster walks a number of meters in distance. This justifies the use of a monomolecular tracer as a model for an odor and points to turbulent dispersal by the ambient fluid flow as the means by which odor is transported from a source to a lobster's chemoreceptors. When tracking a plume, the lobster places its primary

chemoreceptive organs, its lateral antennules, above its carapace to sense the passing chemicals transported by the passing ambient flow.

At the same time, a lobster obtains a frame of reference (that swimming creatures like fish lack or only get from visual cues) to evaluate the relative motion of the flow-borne patches of odor through its contact with the substrate. Turbulence produced by physical features that effect the motion of the seawater generates spatial patterns of flow and therefore spatially heterogeneous distributions of odors. Thus, what can be thought of as noise to a concentration-tracking strategy becomes a potential source of guidance cues to the lobster tracking the chemical to its source.


Studies of chemical distribution in turbulence in the field and in the laboratory provide insights into the nature of the signal environment that animals actually encounter (Moore and Atema, 1991; Murlis, Elkin-ton, and Carde, 1992; Finelli et al., 1999; Grasso, Basil, and Atema, submitted). The period of time required to obtain a statically stationary estimate of concentration depends on the magnitude of the turbulence. In our seawater conditions, the time scale for such estimates is typically several minutes (Grasso, Basil, and Atema, submitted). These studies demonstrate, on purely physical evidence, that the idea that animals (or at least lobsters and similar-size crustaceans) track chemicals to their source by ascending the concentration gradient in environments that contain even moderate levels of turbulence ( $Re > 1 \times 10^2$ ) is untenable.

Studies of the dynamics of the concentration signal have revealed features (such as characteristic rates or patterns of change) that could provide the lobster (or other animals) with guidance cues for certain plumes (Dittmer, Grasso, and Atema, 1995, 1996; Finelli et al., 1999; Grasso, Basil, and Atema, submitted; but see Hanna, Grasso, and Atema, 1999). Examples of cues extracted from the concentration signal include *peak light* (the maximum value of a pulse of odor), *peak slope* (the rate of change of concentration in time on the rising phase of an odor pulse), *interpeak interval* (the delay between the passing of one peak and the arrival of the next at a sensor). There are many others that have been studied in a number of laboratories. Such cues might provide the information that lobster behavior indicates they are using to decrease their heading angles during chemo-orientation (figure 27.1).


## The Animal

In the case of the American lobster, it is clear that, at least for the plumes so far studied, they use two chemoreceptors to inform their chemo-orientation strategies. They perform more efficiently (i.e., more-direct, less-convoluted paths) when they use both lateral antennules, but they can still track a plume to its source with one (Devine and Atema, 1982; Beglane et al., 1997; Mjos, Grasso, and Atema, 1999). Although it is not

**A. Lobsters decrease heading angle error and increase speed the longer they track a plume.**



**B. Heading Angle is measured between the direction of the animal or robot's shortest route to the source and its actual direction of motion.**



**Figure 27.1** (A) Shows schematically the pattern of lobster behavior that indicated the extraction of guidance information during the tracking. OGR, the prediction of odor-gated rheotaxis; RLD, representative lobster data. (A) summarizes the results obtained by Moore, Scholz, and Atema (1991) (B) Illustrates the method of extracting heading angle ( $\alpha$ ) from a digitized path as viewed from overhead. Heading angle may be computed from any pair of position measurements along this hypothetical path.

possible at this time to exclude the possibility that animals using one antennule are employing a degraded form of the two-antennule strategy, it appears that the lobster is capable of using at least two strategies, and multiple, context-triggered strategies appear likely.

Lobster behavior suggests that they might actually be using the turbulence-induced spatiotemporal structure of the odor plume to track odors to their source (Moore, Scholz, and Atema, 1991; Basil, 1994; Atema, 1996). In contrast to the predictions of other strategies, such as odor-gated rheotaxis (OGR) (Kennedy, 1986), the closeness of the lobster's direction of travel (heading angle) to the direction of the source increases as the lobster progresses through the plume (figure 27.1). This unequivocally demonstrates that the animal is extracting some form of information about the direction of the source that is distinct from the direction of mean flow on which OGR is based. If the lobster were to approach the source along a line parallel to the direction of mean flow (the definition of rheotaxis) its heading angle would increase unless it were walking a parallel directly downstream from the source. The likely origin of this guidance is the pattern of intermittent odor patches

produced by turbulence source (Moore, Scholz, and Atema, 1991; Basil, 1994; Atema, 1996).

That said, it seems odd that an animal about the task of tracking odors to their sources would ignore the information contained in the mean flow. After all, the knowledge that the source lies “upstream” allows an agent to exclude at least half of the possible directions of available movement. Lobsters that had the chemoreceptors, but not the flow (mechano-) receptors, on the antennules inactivated show intermediate performance deficits compared to animals that had had the whole antennule removed. This suggests that lobsters in our experiments were using flow and detailed chemical information for guidance to the source.

If we start from the premise that lobsters are extracting and using some plume information beyond mean flow, the natural question is: What information is the lobster using? Recordings of these intermittent odor patches encountered by plume-tracking lobsters indicate that they make turning decisions on the 2–4 second time scale (Grasso, Basil, and Atema, 1998; Basil, Grasso, and Atema, submitted). This is much too rapid an interval for an accurate estimate of the average odor concentration. Neurophysiological studies of the response dynamics of the chemoreceptor neurons on the lobster’s lateral antennules indicate that these cells are capable of following the temporal fluctuations of turbulence-influenced concentration signals (see next section) (Gomez, Voigt, and Atema, 1992, 1994a, b; Gomez and Atema, 1996a, b). Simulation studies based on this data indicate that the temporal filter properties of these neurons selectively preserve this information (Grasso, Voigt, and Atema, 1998, 2000).

### **The Robots**

The constraints supplied from the physics of odor plumes and the behavior of lobsters provide constraints and a large set of testable hypotheses on how lobsters locate odor sources in turbulence. Using these constraints and hypotheses as guides to design, we undertook the construction of robots that could be used to test those hypotheses.

**Why a Robot?** Hypotheses of chemo-orientation mechanisms are difficult to test in animal behavior experiments because measures of the cue stimuli are very difficult to obtain at the same time as behavior of a plume-tracking animal. The viscosity of water guarantees that the effects of physical perturbations, from obstacles or changes in the shape of the channel, persist far downstream from the agents that produce them. Once a perturbation is introduced, say from a plume-generating device or the animal itself, it evolves in a way that is only predictable through stochastic measures. Large numbers of repetitions are required

before statistical stability is achieved. The control of flow through a stimulus delivery device makes the construction of apparatus for testing lobster reactions to specific features problematical, if not technically impossible. To produce odor pulses of a predetermined duration, slope, or other character is exceedingly difficult, and good methods are not yet available to reproduce arbitrary concentration fluctuations on demand (but see Zettler and Atema, 1999). It would require restraining the animal in a manner that would allow tight control of the stimulus delivery to the chemoreceptors but would remove the possibility to observe plume-tracking behavior in the animal. Sequential effects would also be very difficult and time consuming to extract from such a preparation.

At first glance, simulation studies appear a good alternative. An artificial plume environment constructed from calculations within a theoretical framework would provide the experimenter with perfect control of the environment. Such direct numerical simulation (DNS) methods have been developed in the theoretical fluid dynamics. Their computational load on a computer system of a DNS that approaches the appropriate spatial and temporal scale is so great as to prohibit the large repetitions required for evaluation of orientation algorithms (Belanger and Willis, 1996). Further, such models provide temporal average ensemble descriptions at the appropriate time scales but do not contain simultaneous appropriate spatial correlations. This makes the evaluation of multisensor algorithms that use spatial and temporal correlations untenable.

An alternative to theoretical plume models involves capturing movies of turbulent flows. The laser-induced-fluorescence (LIF) technique captures spatial and temporal dynamics of a tracer in a turbulent flow at spatial and temporal scales appropriate for lobster chemo-orientation. We have implemented a number of algorithms in simulation, with "playback" plumes. These plume simulations are based on the playback of data recorded from actual odor plumes rather than the theoretical approach described in the previous paragraph. These simulation studies have identified guidance features, which are available to the lobster through its peripheral chemoreceptors. The use of these features speeds up the simulation's orientation by an order of magnitude compared to concentration-averaging techniques (Grasso, Gaito, and Atema, 1999). Such studies are useful as a means of screening algorithms and tuning parameters prior to biomimetic robotic studies but they cannot, for reasons discussed above, completely capture a number of important environmental features, such as the effect of the robot's or lobster's hydrodynamic shadow on the guidance cues. This is a problem that must be solved in all simulation studies of behavior: The results of the simulation study are only as good as the simulation of the task environment. In the case of chemo-orientation in turbulence, a robotic



**Figure 27.2** (Left) RoboLobster with a standard dye-tinted plume running a tropotaxis algorithm. The robot is inside a 10m  $\times$  2m flume in seawater 44 cm in depth. The robot is designed to serve as a platform to carry chemical sensors, visible above the front end of the robot, through an odor plume. (Right) the BICSAAR Robot “Wilbur” in chemo-orientation trials in the same flume running one of the OGR algorithms.

implementation is a critical and unequivocal test of the quality of an orientation algorithm.

Constructing an artificial lobster and placing it into the same environment as the biological lobster could test hypotheses of chemo-orientation without concern for the formidable difficulties mentioned above. Tom Consi, Jelle Atema, David Mountain, and I undertook a series of experiments to test simple hypotheses of lobster chemo-orientation (Consi et al., 1995; Grasso et al., 1997; Grasso, Basil, and Atema, 1998; Grasso et al., 1999). In these experiments, we studied the limits of purely chemical guidance with chemical concentration-estimation algorithms. In later studies with our BICSAAR robots, we were able to study algorithms that combined a sense of the mean flow of a chemical with a sense of turbulence-induced chemical concentration dynamics.

**Biomimetic Scaling** RoboLobster was the first biomimetic robot we used to study lobster chemotactic strategies. It was designed and built by Tom Consi and Cliff Goudey at the MIT AUV lab (Consi et al., 1994) to permit us to conduct chemo-orientation experiments. It has wheels instead of legs and was not intended to mimic lobster biomechanics or morphology (figure 27.2). Instead, it incorporated those features essential for the evaluation of the chemotactic hypotheses: RoboLobster was built to the proportions of the American lobster and has two chemoreceptors that can be positioned at the same height and separation as the lobster’s two lateral antennules.

In addition to scaling the hardware, we also wrote the control software to scale the robot’s performance to that of the lobster. The RoboLobster’s chemoreceptors acquired sensor data at rates that were much higher than those that could possibly operate in the American lobster’s nose. We used software to scale the data acquisition rates to match those we knew from lobster-chemoreceptor neurophysiology. Simi-

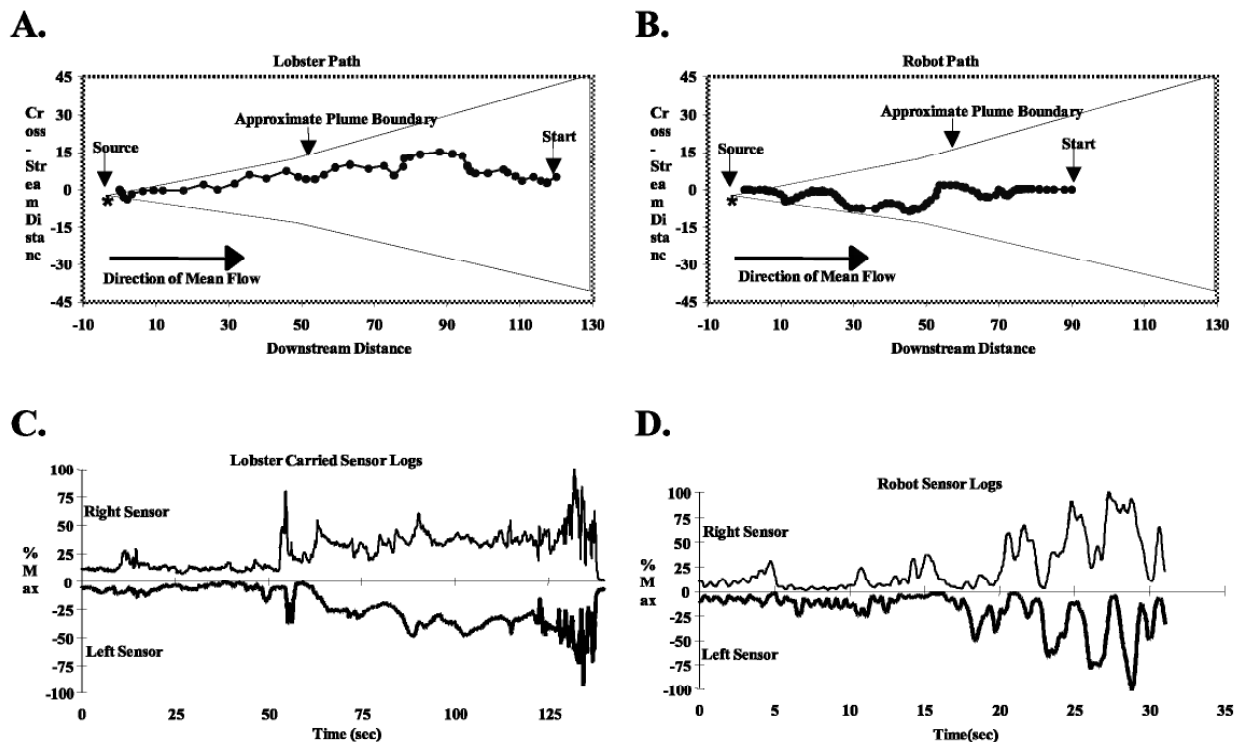


larly, RoboLobster's motor could propel and turn the robot at speeds and angular speeds that exceeded those of the lobster. We programmed lobster-scaled speeds into the algorithms that controlled the behavior during trials.

The BICSAAR robots were designed to include a number of advances over our prototype. They possess more lobsterlike hydrodynamic signature and physical proportions. The sensors use optical methods to detect chemical tracers and are therefore faster and more energy efficient (allowing longer missions). We also incorporated the ability to use flow sensors, which have not yet been installed. In their place, we have used the robots' gyroscopes as a way of providing a frame of reference to the mean flow. Although the BICSAAR robots cannot sense the flow directly, they can operate as if they could, using this reference and an initial cue from the experimenters. With this setup, the BICAAR robots have been recently used to study the effectiveness of odor-gated rheotaxis mechanisms at tracking our plumes.

**What Biomimetic Robots Taught Us about Chemo-orientation in Turbulence** Our ultimate aim is to gain an understanding of lobster sensorimotor integration during plume tracking. First, however, with Occam's razor in mind, we tested simpler algorithms. Our thinking was that, if these simpler algorithms were adequate to provide a satisfactory explanation of lobster plume-tracking behavior, such a conclusion would render more-complex algorithms unnecessary. Because the lobster appears to require both of its antennules for normal plume tracking, we started with the simplest dual-sensor strategy: tropotaxis. Tropotaxis is an orientation strategy that is based on an instantaneous estimate of the local gradient of a cue derived from two sensors (Fraenkel and Gunn, 1961).

The major result of these studies was a solid demonstration that tropotactic concentration-sensing algorithms could not explain the plume-tracking behavior in lobsters—at least not in the type of plume in which we tested lobster and RoboLobster behavior (Moore and Atema, 1991; Dittmer, Grasso, and Atema, 1995). So we are forced to consider other biologically feasible algorithms to find a reasonable explanation. This negative result was significant because the conditions of the series of experiments were carefully constructed to succeed if and only if the plume structure supported the algorithm. The plume structure failed to do so in an interesting way. The behavior of the robot when tracking the plume showed two statistically distinct patterns of locomotion in two regions of the plume using the same algorithm. In the downstream *distal patch field*, the robot did not make significant progress toward the source. In the upstream *proximal cone*, where progress toward the source was possible with a tropotactic algorithm, the paths produced showed



**Figure 27.3** Representative Robolobster and Lobster paths and sensor profiles. These particular paths were chosen because they show the closest similarity between lobster and robot paths obtained. (A) and (B) show the paths, (C) and (D) the corresponding sensor logs. The lobster's path in (A) contrasts that of the robot in (B). Robolobster, running a simple tropotactic algorithm makes a series of sweeping counterturns as it approaches the source (marked by an \*). (C) shows dopamine concentration profiles recorded with a pair of electrochemical electrodes mounted on the back of the lobster and positioned in the olfactory sampling area of the lobster. Dopamine was used as a tracer in the food odor plume the lobster tracked. (D) shows concentration signals taken from the robot's sensor logs. Both (C) and (D) are expressed as percent of maximum signal. RoboLobster's counterturning near the source is evident in the sensor log.

little resemblance to those of lobsters (figure 27.3). Thus RoboLobster revealed to us something about the lobster's world that we had previously only suspected: the need to switch tracking strategies between different regions of the plume.

In one series of experiments, we systematically varied the intersensor separation to ranges that were biologically impossible for the antennules of the American lobster. Because tropotaxis relies on an instantaneous estimate of the gradient, the increased distance between the sensors should have improved performance. Performance improved in the proximal cone but not the distal patch field of the plume. This represents a type of evolutionary experiment in that it tests the sensitivity of an algorithm to morphological variation: the biological road not traveled by the American lobster. It raises the question of why the lobsters possess the intersensor separation that they do (other species

have both greater and less separations) and opens the door to further experiments that can explore the relationship between environment and evolutionarily derived structure.

Recent studies with BICSAAR robots have tested the OGR strategy described earlier. Although OGR does not seem to be a mechanism that lobsters use in its pure form, the fact that mean flow is a potentially powerful guidance cue led us to try it as the simplest algorithm, which combines chemo- and rheo-(flow) sense. As expected from the biological studies, the variants of OGR produced paths that are very inefficient compared to those of lobsters. However, the algorithms are able to track in the distal patch field, or its equivalent, at distances of up to seven meters from the source. This represents a gain in performance over the tropotactic algorithms studied with the RoboLobster prototype.

As a result of the RoboLobster and BICSAAR studies sketched here, we have made progress toward understanding the biology and toward creating practical plume-tracking robots. We have excluded, or at least called into question, two hypotheses of lobster chemo-orientation that a priori appeared to hold great explanatory power. We have also learned when and where these respective algorithms, though not very biological, can be practically applied to general plume-tracking problems.

#### **WHAT WE WOULD LIKE ROBOTS TO BE ABLE TO DO: MISSION DESCRIPTION**

The general aim of our studies has been to understand how biological systems are able to track chemicals to their sources in environments where turbulence dominates the chemical dispersal process. Examples of applications technologies derived from such an understanding abound. Tracking pollutants in estuaries to their sources, locating leaks from ships or aircraft, locating unexploded ordinances or mines—the list is very long. All that is required is a sensor that can measure the chemical of interest at the appropriate time and spatial scale. And further, the principle of dynamic similarity holds that with appropriate scaling, successful algorithms can be translated to operate in any fluid medium, gas or liquid (Vogel, 1994).

Ideally, a derived system should be able to track a diversity of plumes. The robot, armed with algorithms developed in our BICSAAR robots and an appropriate sensor, could be deployed into a field location. It would be able to make an estimate of the fluid dynamic conditions and determine which types of tracking algorithms are appropriate. It would execute a search pattern aimed at making contact with a plume composed of the chemical or chemical mixture its sensors detect. On contact with the plume, the robot would begin a graceful and efficient path to the source, detecting changes in condition and changing tracking strategy as required. At the source, it would mark or signal to a waiting

supervisor the presence of the source and commence a new search for the next source.

Of course, the above paragraph is in the conditional future tense because we are still some distance from knowing how to do this. The essential seeds of this knowledge, however, lie in our understanding of the fluid dynamics of chemical dispersal and the behavior and neural processing of animals that already possess the methods.

## THE THREE CRITICAL POINTS FOR BIOMIMETIC ROBOTICS REVISITED

### Sufficient Description for Biomimetic Robotic Algorithm Implementation

A good practical robot might arise from first principles but a biomimetic robot cannot be created *de novo* from first principles. There must be sufficient biological knowledge about the system to provide appropriate constraints for the design parameters if the robot is to capture animal levels of competence. Without this set of constraints, one who sets about building a loosely constrained biomimetic robot will, unless very lucky, produce neither a practical robot nor good biology. Lacking first principles *and* biological inspiration for guidance, no robot capable of any behavior intended by its builders would probably ever be finished.

These statements revisit the paradox mentioned before. The issue is one of deciding between essential facts and irrelevant details. Platt (1964), in an influential article on the modes of scientific inquiry and discovery, argued that the basis of judging the relevant and irrelevant aspects of a system should be strong theoretical inference. That is, are the included features both necessary and sufficient to explain the process under study?

This is the essence of the biomimetic robotic strategy: Construct a robot that is competent to test a hypothesis or set of hypotheses that have been suggested by the biology and then allow the robot's behavior to inform you of the acceptability of that hypothesis. If you are in possession of *the* biological algorithm you can proceed to the production of practical robots and hypothesis testing is not an issue. If *the* algorithm is not known, one need include only those constraints that are required to test the hypothesis under consideration. I assume here that one undertakes a robotic study to answer a biological question only if it is not possible or too costly to obtain the answer directly from the animal. Constraints or design features that are not required to test the hypothesis (in our case legs, for one example) are likely to cloud the scientific issues and make for greater technical difficulties. In other words, a sufficient database for a biomimetic robotic implementation is

one that leaves at most one or two, preferably no, free parameters in the model suggested by the hypothesis to be evaluated.

### **Appropriate Level of Abstraction and Internal Representation**

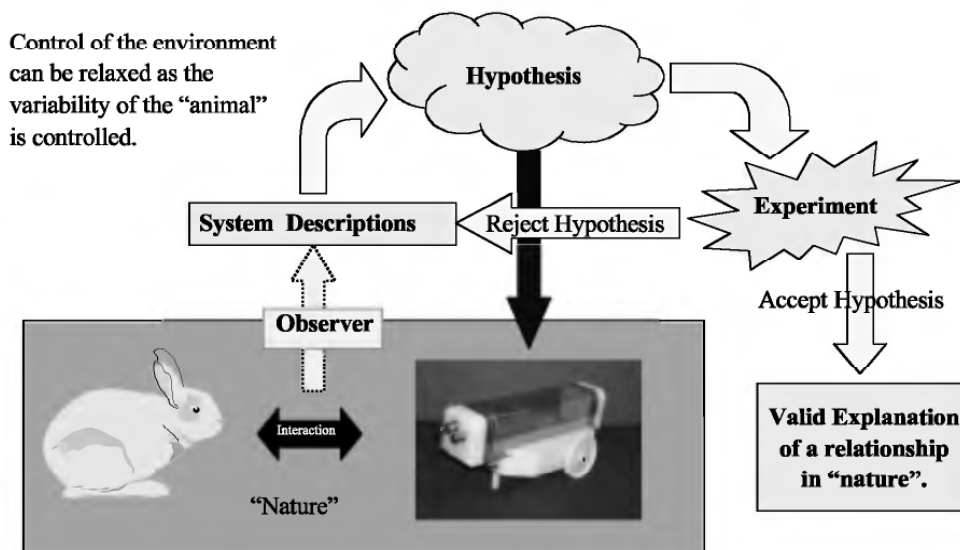
The choice of the sophistication of the algorithms to be tested in an experimental biomimetic robot related to the available degree of knowledge concerning environmental, neural, and behavioral constraints. The biomimetic robot could be designed with elaborate internal representations of the world it will encounter and detailed contingency mechanisms that are activated when certain conditions are encountered. Alternatively, it could be designed so that its behavior is driven by the stimulus world. The balance between these two extremes is an empirical question in the animal but in the biomimetic robot it is a matter of methodology.

If the available information suggests that the animal use some form of internal representation, then biomimetic robotics is an ideal methodology to explore hypotheses of representation. On the other hand, it is possible to generate quite-complex behaviors from a relatively simple algorithm if the environment is complex (Braitenberg, 1984; Kennedy, 1992). The level of abstraction in the representation in a biomimetic robot, therefore, is a literal mapping of what is known about the biological system, and any elaboration beyond what is known becomes a hypothesis to be tested.

This, again, is the principle of Occam's razor and it seems as useful a heuristic in the evaluation of biologically relevant hypotheses with biomimetic robots as it does in the rest of science. This approach also highlights one of the major advantages of biomimetic robots in biology. The biomimetic robot offers a way of systematically evaluating the contribution and influence of environmental structure on behavioral patterns. It can also be used to justify the use of more-complex devices and systems (such as those that include internal representations) in models when simpler explanations are demonstrated to be untenable (figure 27.4). By placing the biomimetic robot in the same environment as the animal, we are in reality placing the hypothesis directly into the environment.

### **Adequate Performance Measures**

If the aim of the biomimetic robot is hypothesis testing, then the performance measures must be sufficiently well understood in the biological system for a meaningful comparison to be made between the robot and the animal. As with constraints on robot design, the hypothesis under test delineates the relevant and irrelevant features of the biolog-



**Figure 27.4** The role biorobots can play in basic biological research. The environment, through evolution or ontogeny, plays a critical role in shaping the behavior of an animal. By placing the robot in the natural environment the contribution of the environment to behavior can be disassociated from that of the animal.

ical behavior in an adequate description of performance. The performance measures should be derived from descriptions of the biological behaviors that are adequate to permit the determination that the algorithm under study is or is not operating. Vague performance measures that do not discriminate between hypotheses are of limited value.

In our case, it was essential to have a good description of the animal behavior. We measured a number of descriptors of the paths RoboLobster took as it traversed the plume. These were measured using exactly the same methods and under exactly the same conditions as those used in lobster studies.

The trials wherein RoboLobster “missed” the source must be due to failures of the implemented algorithms. We cannot be certain the same is true for the lobster. Motivation makes the “hit” rates of the lobsters difficult to contrast with RoboLobster because the robot is programmed to always seek the source while the lobster is not guaranteed to be driven by a single motivation. Thus we used a finer-grained performance measure. We digitized the paths taken by the robot at intervals that are shorter than the variation in lobster locomotor performance. From these we were able to measure, in both the lobster and the robot, distance from the source, time to various points from the source, locomotory speed, and path tortuosity (a measure of the deviation of the path from a straight line). These parameters allowed us to make direct comparisons between lobster and robot performance with direct relevance to the tropotactic hypothesis.

This narrow view of the performance measure has a cost. Given that the performance matches some reasonable criterion of the biological system, there is no guarantee that the function implemented in the robot is *the* strategy used by the animal. It merely constitutes evidence that the two are consistent. The narrowing of the performance measures to just those that are adequate to evaluate the hypothesis may mask inconsistencies in other aspects of performance. In contrast to the robot design, the performance measures should be chosen to *at least* permit the hypothesis testing. A discrepancy in performance in some aspect that is not essential for the hypothesis points to a limitation of the model (maybe the hardware, maybe the software implementation, or maybe the theory itself) that is in need of explanation.

## CONCLUDING REMARKS

Recently, Vogel (1998) has strongly argued against the common prejudice that biological solutions to problems are necessarily optimal. He points out that evolution has shaped biological systems to be adequate to ensure the survival of the individual animal until it can reproduce. This requires biological systems to be merely adequate, not optimal. At the same time, biological solutions exist (to problems such as chemorientation in turbulence) that are superior to our present ones. The animals themselves are existence proofs of these superior solutions. Biomimetic robots are powerful tools that can aid our efforts to uncover natural strategies.

To achieve this, however, we must acknowledge that the science of biology, with its own brand of methods and logic, must be at the center of any such effort. The performance measures, hypotheses, and design parameters must be biologically meaningful and not merely biologically inspired if we are to divine nature's solutions to difficult problems. From this type of solid understanding, we will be able to proceed to "optimal" solutions to problems as we, and not nature, pose them. Today, our jet airplanes exceed the speed and load capacities of birds and flying insects. They have been optimized for speed and load but at the cost of the maneuverability of their biological forebears. In the sense that biological wings inspired the first artificial airfoils, the biomimetic robot is a comparable step on the way to discovering natural principles of neural processing for translation into artificial systems.

## ACKNOWLEDGMENTS

The experimental work summarized here has been previously published. Nevertheless, I must acknowledge the contributions of a number of people. The work was conducted in the laboratory of Professor Jelle Atema and rests largely on the many years of lobster research

conducted in that laboratory in earlier years. Dr. Tom Consi and Professor David Mountain were also central to the RoboLobster project. I must also thank Professor Jennifer Basil for generously supplying the lobster data in figure 27.3. Lisa Borden, Jonathan Dale, Paul DiNunno, Kevin Dittmer, and Rainer Voigt all made significant technical contributions to the project. Steve Gaito, Megan Mahaffy, Christina Manbeck, and Christian Reilly made significant contributions to the development and study of theoretical chemotactic algorithms. Photo credit for figure 27.2 goes to Tom Consi and Marjorie Steele. Portions of the research reported here were supported by two NSF Grants (BES-9315791 and IBN-9631665) to Jelle Atema and by awards to Jelle Atema and Frank Grasso under the DARPA/ONR plume-tracing program administered by Keith Ward and Regina Dugin.

## REFERENCES

- Atema, J. (1996). Eddy chemotaxis and odor landscapes: Exploration of nature with animal sensors. *Biol. Bull.* 191: 129–138.
- Baird, R. C., Johari, H., and Jumper, G. (1996). Numerical simulation of environmental modulation of chemical signal structure and odor dispersal in the open ocean. *Chem. Senses.*
- Barbin, G. (1998). The role of olfaction in homing and estuarine migratory behavior of yellow-phase American eels. *Can. J. Fisheries Aquatic Sci.* 55: 564–575.
- Basil, J. A. (1994). Lobster orientation in turbulent odor plumes: Simultaneous measurement of tracking behavior and temporal odor patterns. *Biol. Bull.* 187: 272–273.
- Basil, J. A., Grasso, F. W., and Atema, J. (Submitted). Odor plume dynamics measured during American lobster orientation behavior. *J. Exp. Biol.*
- Beglane, P. F., Grasso, F. W., Basil, J. A., and Atema, J. (1997). Far field chemo-orientation in the American lobster, *Homarus americanus*: Effects of unilateral ablation and lesioning of the lateral antennule. *Biol. Bull.* 193: 214–215.
- Belanger, J. H., and Willis, M. A. (1996). Adaptive control of odor-guided locomotion: Behavioral flexibility as an antidote to environmental unpredictability. *Adaptive Behav.* 4(3/4): 217–253.
- Braitenberg, V. (1984). *Vehicles: Experiments in Synthetic Psychology*. Cambridge: MIT Press, a Bradford book.
- Consi, T. R., Atema, J., Goudey, C. A., Cho, J., and Chryssostomidis, C. (1994). AUV guidance with chemical signals. In *Proceedings of the IEEE Symposium on Autonomous Underwater Vehicle Technology*, Cambridge, Mass.
- Consi, T. R., Grasso, F., Mountain, D., and Atema, J. (1995). Explorations of turbulent odor plumes with an autonomous underwater robot. *Biol. Bull.* 189: 231–232.
- Cox, C., Hunt, J. A., Lyons, W. G., and Davis, G. E. (1997). Nocturnal foraging of the caribbean spiny lobster (*Panurlis argus*) on offshore reefs of florida, USA. *Mar. Freshwater Resources* 48: 671–679.
- Dale, J. (1997). Chemosensory search behavior in the starfish *Asterias forbesii*. *Biol. Bull.* 193: 210–212.



- Devine, D. V., and Atema, J. (1982). Function of chemoreceptor organs in spatial orientation of the lobster, *Homarus americanus*: Differences and overlap. *Biol. Bull.* 163: 144–153.
- Dittmer, K., Grasso, F. W., and Atema, J. (1995). Effects of varying plume turbulence on temporal concentration signals available to orienting lobsters. *Biol. Bull.* 189: 232–233.
- Dittmer, K., Grasso, F. W., and Atema, J. (1996). Obstacles to flow produce distinctive patterns of odor dispersal on a scale that could be detected by marine animals. *Biol. Bull.* 191: 313–314.
- Dodson, J. J., and Dohse, L. A. (1984). A model of olfactory-mediated conditioning of directional bias in fish migrating in reversing tidal currents based on the homing migration of American shad (*Alosa sapidissima*). In McCleave, Arnold, Dodson, and Neil (eds.), *Mechanisms of Migration in Fishes*. New York: Plenum Press, pp. 263–281.
- Dunsenberry, D. (1992). Guiding. *Sensory Ecol.*: 413–436.
- Finelli, C. M., Pentcheff, N. D., Zimmer-Faust, R. K., and Wethey, D. S. (1999). Odor transport in turbulent flows: Constraints on animal navigation. *Limnology Oceanography* 44(4): 1056–1071.
- Fraenkel, G. S., and Gunn, D. L. (1961). *The Orientation of Animals: Kineses, Taxes, and Compass Reactions*. New York: Dover Publications.
- Gomez, G., and Atema, J. (1996a). Temporal resolution in olfaction: II. Time course of recovery from adaptation in lobster chemoreceptor cells. *J. Neurophysiol.* 76: 1340–1343.
- Gomez, G., and Atema, J. (1996b). Temporal resolution in olfaction: Stimulus integration time of lobster chemoreceptor cells. *J. Exp. Biol.* 199: 1771–1779.
- Gomez, G., Voigt, R., and Atema, J. (1992). High resolution measurement and control of chemical stimuli antennule of the lobster *Homarus americanus*. *Biol. Bull.* 183: 353–354.
- Gomez, G., Voigt, R., and Atema, J. (1994a). Frequency filter properties of lobster chemoreceptor cells determined with high-resolution stimulus measurement. *J. Comp. Physiol.* [A] 1–9.
- Gomez, G., Voigt, R., and Atema, J. (1994b). Tuning properties of chemoreceptor cells of the American lobster: Temporal filters. In *Olfaction and Taste XI*. pp. 788–789.
- Grasso, F. W., Basil, J. A., and Atema, J. (1998). Toward the convergence: Robot and lobster perspectives of tracking odors to their source in the turbulent marine environment. In *Proceedings of the 1998 IEEE*, Gaithersburg, Md.
- Grasso, F. W., Basil, J. A., and Atema, J. (Submitted). Directional information in the dynamics structure of a turbulent odor plume measured with a pair of lobster-scaled “odor” sensors. *J. Exp. Biol.*
- Grasso, F. W., Consi, T. R., Mountain, D. C., and Atema, J. (1999). Biomimetic robot lobster performs chemo-orientation in turbulence using a pair of spatially separated sensors: Progress and challenges. *J. Robotics Autonomous Syst.* 807: 1–17.
- Grasso, F. W., Dale, J. H., Consi, T. R., Mountain, D. C., and Atema, J. (1996). Behavior of purely chemotactic robot lobster reveals different odor dispersal patterns in the jet region and the patch field of a turbulent plume. *Biol. Bull.* 191: 312–313.
- Grasso, F. W., Dale, J. H., Consi, T. R., Mountain, D. C., and Atema, J. (1997). Effectiveness of continuous bilateral sampling for robot chemotaxis in a turbulent odor plume: Implications for lobster chemo-orientation. *Biol. Bull.* 193: 215–216.
- Grasso, F. W., and Di Lorenzo, P. M. (1997). Gussto: A neuronal network model of gustatory processing in the rat NTS and PBN. In J. Bower (ed.), *Computational Neuroscience*. New York: Plenum Press, pp. 321–326.

- Grasso, F. W., Gaito, S. T., and Atema, J. (1999). Behavior, peripheral chemo-reception and information in turbulent odor plume tracking: Results from simulation studies. *Soc. Neurosci. Abstr.* 29(1): #550.3.
- Grasso, F. W., Voigt, R., and Atema, J. (1998). Lobster chemoreceptor dynamics in turbulent odor plumes: Transmission of environmental cues from the periphery to the CNS. *Soc. Neurosci. Abstr.* 28: #711.2.
- Grasso, F. W., Voigt, R., and Atema, J. (2000). Dynamic filter properties of lobster-nose chemo-receptors: Peak performance evaluation. *J. Comput. Neurosci.*
- Hanna, J. P., Grasso, F. W., and Atema, J. (1999). Temporal correlation between sensor pairs in different plume positions: A study of concentration information available to the American lobster, *Homarus americanus*, during chemotaxis. *Biol. Bull.* 197: 26–27.
- Kennedy, J. S. (1986). Some current issues in orientation to odour sources. In T. L. Payne, M. C. Birch, and C. E. J. Kennedy (eds.), *Mechanisms in Insect Olfaction*. Oxford: Clarendon Press, pp. 11–25.
- Kennedy, J. S. (1992). *The New Anthropomorphism*. New York: Cambridge University Press.
- Mjos, K., Grasso, F. W., and Atema, J. (1999). Antennule use by the American lobster, *Homarus americanus*, during chemo-orientation in three turbulent odor plumes. *Biol. Bull.* 197: 25–26.
- Moore, P. A., and Atema, J. (1991). Spatial information in the three-dimensional fine structure of an aquatic plume. *Biol. Bull.* 181: 408–418.
- Moore, P. A., and Lepper, D. M. E. (1997). Role of chemical signals in the orientation behavior of the sea star *Asterias forbsii*. *Biol. Bull.* 192: 410–417.
- Moore, P. A., Scholz, N., and Atema, J. (1991). Chemical orientation of lobsters, *Homarus americanus*, in turbulent odor plumes. *J. Chem. Ecol.* 17: 1293–1307.
- Murlis, J., Elkinton, J. S., and Carde, R. T. (1992). Odor plumes and how insects use them. *Annu. Rev. Entomol.* 37: 505–532.
- Oliver, S. J., Grasso, F. W., and Atema, J. (1996). Filament tracking and casting in American eelers (*Anguilla rostrata*). *Biol. Bull.* 191: 314–315.
- Platt, J. R. (1964). Strong Inference. *Science* 146: 347–352.
- Vogel, S. (1994). *Life in Moving Fluids: The Physical Biology of Flow*. Princeton: Princeton University Press.
- Vogel, S. (1998). *Cats' Paws and Catapults*. New York: W.W. Norton.
- Weissburg, M. J., and Zimmer-Faust, R. K. (1994). Odor plumes and how blue crabs use them in finding prey. *J. Exp. Biol.* 197: 349–375.
- Zettler, E., and Atema, J. (1999). Chemoreceptor cells as concentration slope detectors: Preliminary evidence from the lobster nose. *Biol. Bull.* 197: 252–253.
- Zimmer-Faust, R. K., Finelli, C. M., Pentcheff, N. D., and Wetthey, D. S. (1995). Odor plumes and animal navigation in turbulent water flow: A field study. *Biol. Bull.* 188: 111–116.

**This page intentionally left blank**

Michael Dickinson, Lance Tammero, and Michael Tarstino

## FRUIT FLIES AS MODEL SYSTEM FOR SENSORY FUSION

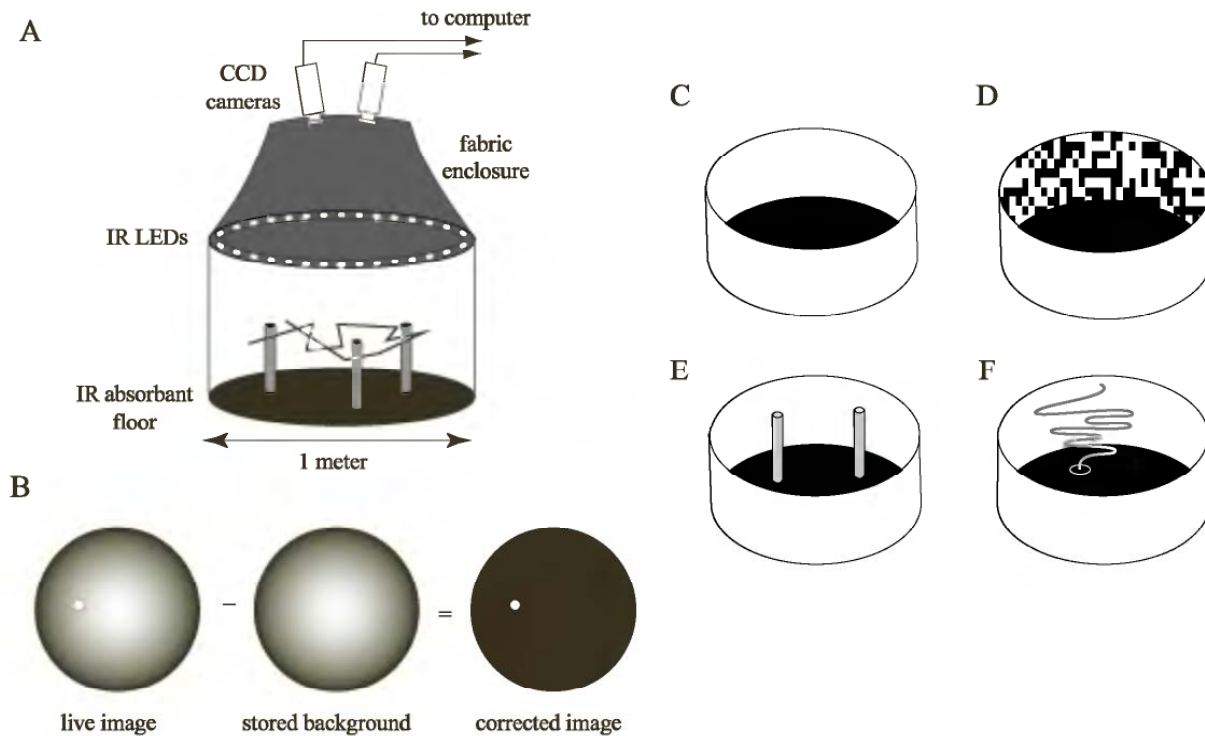
Whether the goal is food, mates, or a place to hide, search behavior is an essential component in the life history of many insects. The successful search and recognition of any given target requires a structured integration or *fusion* of many different sensory cues, some emanating from the target, some embedded elsewhere within the environment, and some intrinsic to the animal itself. We have been using the fruit fly, *Drosophila melanogaster*, as a model system for studying sensory fusion in search behavior. The extraordinary search capacity of these tiny animals is manifest to anyone who has filled a fruit bowl on a warm day. Within minutes, the ripening fruit attracts a swarm of flies that appear to materialize from thin air. Many of the flies that reach the fruit bowl have traveled a great distance, drawn by the characteristic aroma of ripened fruit. The impressive search capabilities of these tiny flies have been honed by selective pressures to find and colonize new food resources. A piece of fruit in a shady spot represents nearly everything to a female fruit fly—food, mates, and a place to lay eggs. Finding and choosing the right piece of fruit is perhaps the most important thing a female fruit fly ever does. The larvae that emerge from her eggs are either doomed or blessed by her choice.

Three sensory modalities are most critical in a fly's ability to localize a food source: olfaction, vision, and mechanoreception. Fly eyes are well known for their elevated flicker fusion rate (Autrum, 1958) and motion-sensitive responses (Borst and Egelhaaf, 1989). Such specializations are undoubtedly related to the role of the visual system in flight behaviors. However, while the temporal resolution of the eyes may be extraordinary, their spatial resolution is quite limited. Each compound eye of *Drosophila* contains approximately 700 ommatidia (Heisenberg and Wolf, 1984), the basic sampling units of the arthropod eye, which is roughly equivalent to a 26 by 26 pixel array. The six photoreceptors that converge onto second-order cells in the first optic ganglion sample light from about a 5° sector of space. Even with hyperacuity, an object

the size of an apple will not be visible over distances greater than a couple meters. Thus, the spatial range over which the fly's visual system could detect a target is quite small, and it is up to the olfactory system to bring the eyes within striking distance of a potential food source. Attractive odorants are detected by chemosensory cells on the antennae and processed within the antennal glomeruli (Laissue et al., 1999). Proper orientation toward an odor source may also require knowledge of wind direction, encoded by mechanoreceptors on the antennae and optic flow sensors within the visual system (Vickers, 2000). In addition, throughout an entire search, the visual system and various mechanosensory modalities provide crucial information for maintaining a stable flight posture and course direction in the face of external perturbations or a changing landscape. Chief among the equilibrium detectors are the halteres, minute hind wings that function as gyroscopes to detect the angular rotation of the body during flight (Nalbach, 1993; Nalbach and Hengstenberg, 1994; Dickinson, 1999). Input from the halteres is essential for stable flight. Flies whose halteres have been removed rapidly corkscrew toward the ground if they attempt to fly (Derham, 1714; Fraenkel, 1939). Efficient search in flies requires the seamless integration of information from the halteres, eyes, and olfactory system in a way that allows the animal to actively maneuver toward a source while remaining stable in the face of external perturbations. One long-term goal of this research is to understand how the fly's nervous system accomplishes this sensory fusion, and to employ the same strategy in the design of control systems for autonomous robots.

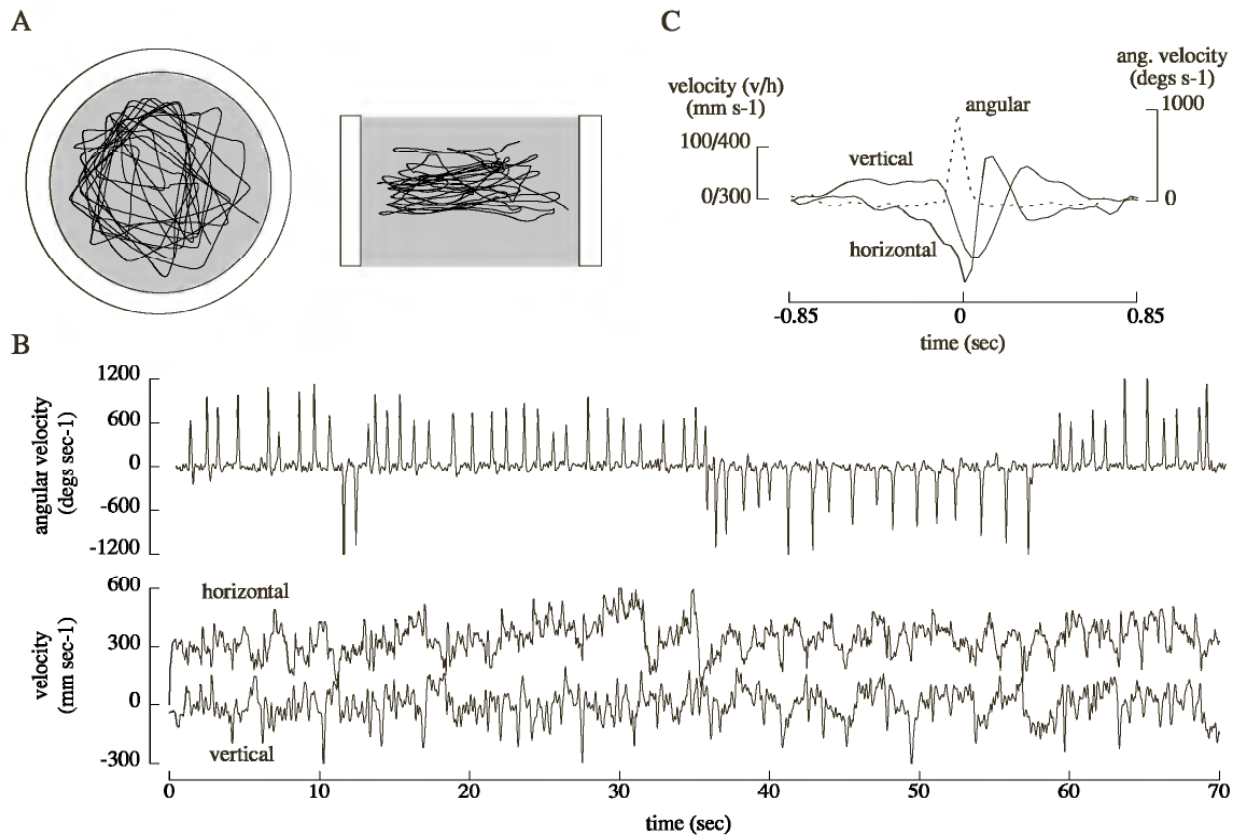
#### **EXPERIMENTAL PARADIGM: TRACKING SEARCH PATHS IN FREE FLIGHT**

To study the interactions among different sensory modalities in the search behavior of fruit flies, we have constructed a large free-flight arena equipped with a 3-D tracking system (figure 28.1). The cylindrical arena is 1 meter in diameter and 0.4 meters high. While this enclosure is still small enough to fit within a laboratory, it represents roughly 18 million cubic body lengths within which the fruit flies must search. The floor of the arena is covered with IR-absorbent flock paper, and the rim of the arena is lined with 200 IR diodes, angled downward to evenly flood the interior with light. Two video cameras placed above the arena detect the animals as reflected spots of light. We use IR (infrared) light, to which the animals are blind, so that bright lights do not interfere with the visually mediated components of their flight behavior. Digital signal processors subtract, frame by frame, a stored background image from the incoming live video signal (figure 28.1B). This step removes the contrast texture of the arena floor from each image,



**Figure 28.1** Apparatus for recording free flight trajectories of fruit flies. (A) Flight arena, see text for details. (B) Method for online background subtraction. (C–F) sample experimental paradigms. (C) Uniform white background. (D) Random checkerboard background. (E) Uniform white background with cylindrical objects. (F) Uniform white background with odor source.

leaving only the image of the fly, which, after adjustment of contrast and brightness, appears as a bright spot on a black background. By processing the image in this fashion, the task of automatically tracking the position of the fly in downstream processing is greatly simplified. After digitizing the planer coordinates of the fly in the two fields of view, the data are transformed into 3-D coordinates. A circular array of lights placed outside the arena is used to project visible light to the interior through the translucent white acrylic walls. A cylindrical curtain runs from the upper rim of the arena to the ceiling, blocking off the view of the laboratory. The top of the enclosure is covered except for two small holes through which the lenses of the video cameras project. When placed in the arena, the fly's visual world consists of a dark ground, a broad white band extending around the horizon, and a dark ceiling. We can manipulate this simple sensory landscape in one of three ways (figure 28.1C–F). First, we can introduce a more complex visual background to the arena by lining the wall with printed patterns. Second, we can place real 3-D objects, such as dark vertical cylinders or square blocks, within the arena. Third, we can introduce sources of odor within the arena. The odor sources consist of small centrifuge tubes filled with vinegar solution. The tubes are placed either on the



**Figure 28.2** Example flight trajectory in uniform white background. (A) Single flight trajectory viewed from the top and side. (B) Time course of animal's angular velocity in the horizontal plane, horizontal velocity, and vertical velocity. Angular velocity is estimated from the change in the animal's heading over time. (C) Saccade-averaged changes in horizontal, vertical, and angular velocities from the traces shown in (B). The angular velocity average was constructed using the absolute value so that saccades to the left and right have been pooled.

3-D objects or within recessed holes in the floor of the arena. The basic paradigm is to manipulate the placement of objects and odor sources within the arena, as well as the pattern of the visual background, and examine the effect on the search behavior of the flies. The air within the arena is still, so the flies cannot use wind as a cue to the location of the olfactory source. Thus, the algorithms that we are attempting to uncover with these particular experiments would be most appropriate for robots operating in indoor settings.

### SACCADES ARE STEREOTYPED FIXED-ACTION PATTERNS

As indicated in figure 28.2, the flight trajectories of flies moving through a simple visual landscape consist of straight flight sequences interspersed with rapid changes in direction, termed *saccades* (Collett and Land, 1975). In tethered flight, saccades are manifest as torque spikes

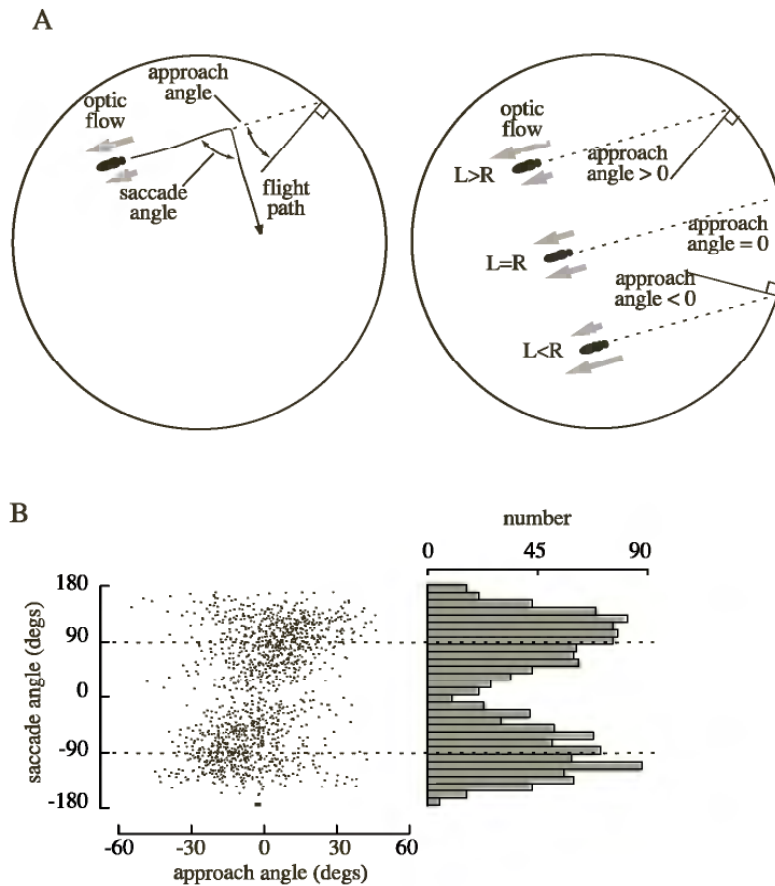
(Heisenberg and Wolf, 1979) or transient changes in wing-beat amplitude (Heide and Götz, 1996). The episodic nature of the saccades is well illustrated by the intermittent spikes in plots of angular velocity (figure 28.2B), which are correlated with the changes in the translational velocities (figure 28.2C). Horizontal velocity decreases gradually before and during the saccade, then increases rapidly as the saccade ends. Vertical velocity often falls below zero during the saccade, but recovers after. These correlated changes in body kinematics are expected, because fruit flies cannot vary lift, thrust, and turning moments independently (Götz and Wandel, 1984; David, 1985). Thus, the change in force production required to generate yaw torque will be manifest as a simultaneous alteration in lift and thrust. The electrophysiological results of Heide and Götz (1996) suggest that a saccade is generated by a brief series of spikes in a subset of steering muscles. This transient burst is probably triggered by a command from an interneuron that descends from the protocerebrum to the thorax.

Because the velocity with which an image moves across the fly's retina depends on the distance to the object, the gross pattern of optic flow preceding each saccade can be conveniently parameterized by calculating the approach angle—the angle at which the fly would intersect a tangent to the arena were it to continue toward the wall (figure 28.3A). A collision angle of  $90^\circ$  indicates that the fly is flying along a trajectory that passes directly through the center of the arena, and the optic flow on its left and right sides should be balanced. Higher and lower collision angles indicate that the fly is near one wall, and thus the flow detected on both sides should be different. Theoretically, the fly might use such asymmetry to modulate the angular magnitude of each saccade. For example, the fly might produce a U-turn when headed directly toward the wall, but generate a glancing turn when moving almost parallel to a wall. However, a plot of saccade angle versus approach angle (figure 28.3B) shows that this is not the case. Saccades are distributed within two clear populations: one centered at  $90^\circ$  to the left, the other at  $90^\circ$  to the right. At an approach angle of  $0^\circ$ , animals do not generate U-turns, but rather saccade either to the left or right with equal probability. Thus, the fly's nervous system dictates the timing and direction of each saccade, but the magnitude is apparently predetermined.

### **What Initiates a Saccade?**

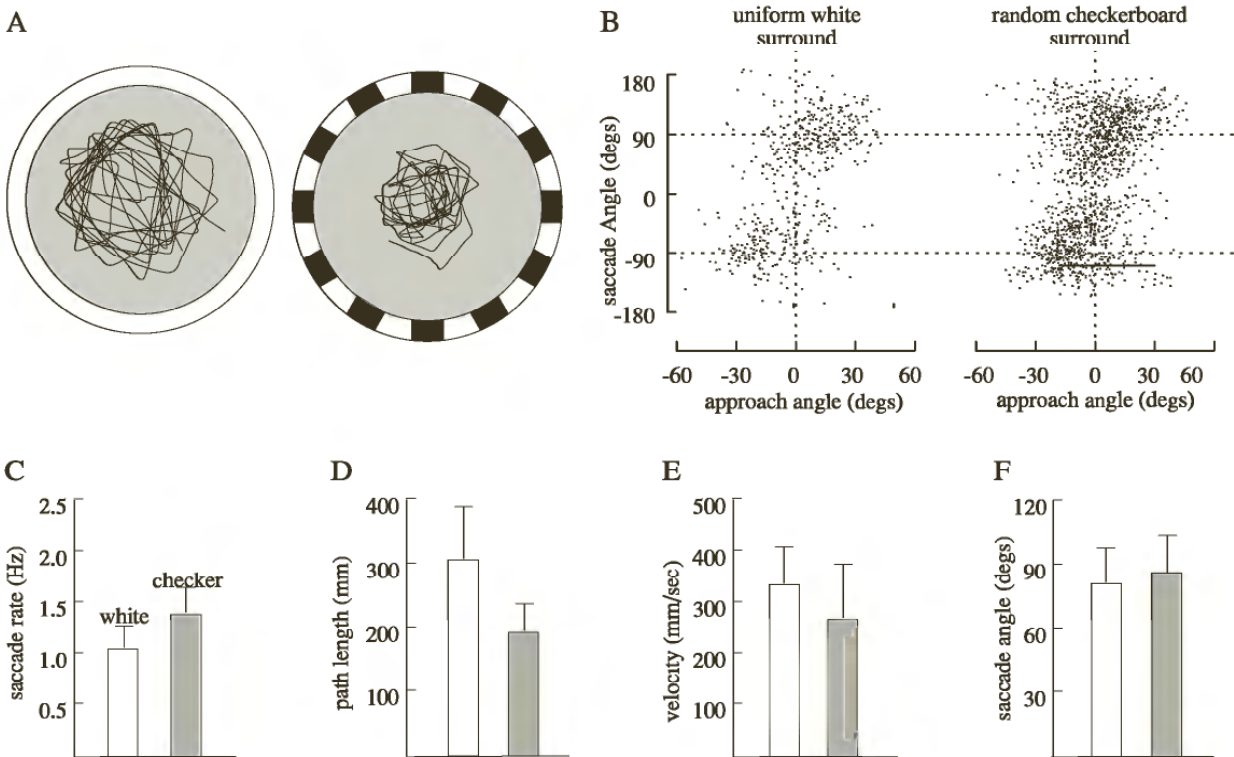
The regularity with which flies turn within a visually sparse arena suggests that the saccades might be initiated by some form of central oscillator or clock. That this is not the case is illustrated by figure 28.4, which compares the flight trajectories within a uniform white surround with those the flies produce when the arena is lined with a random





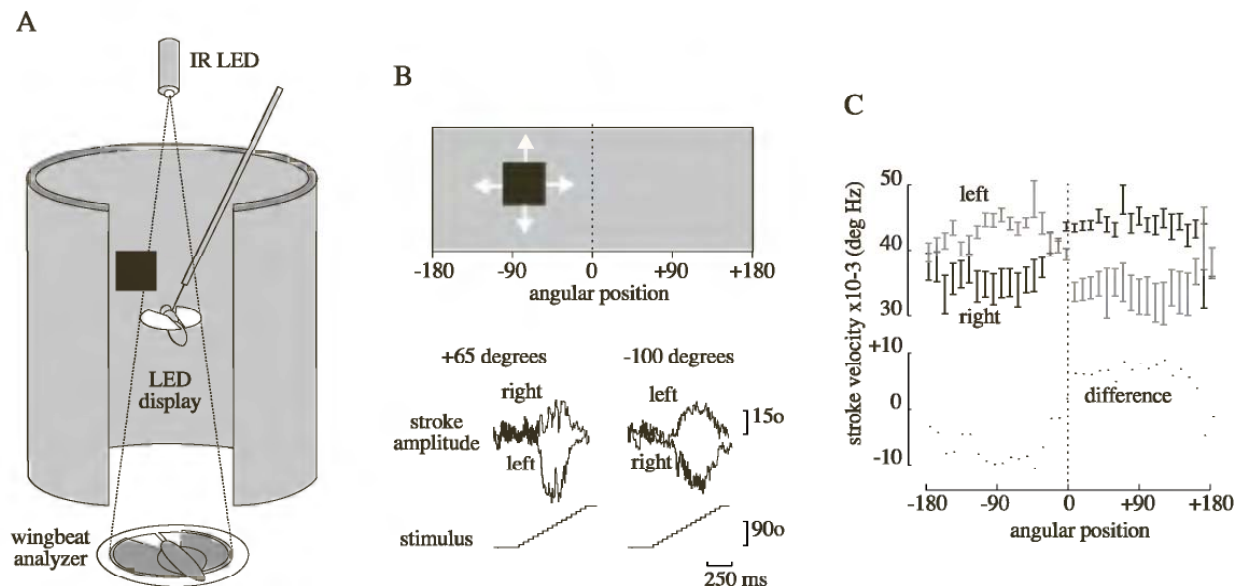
**Figure 28.3** Saccades are binary behaviors. (A) Method for calculating approach angle, which serves as simple indicator of the animal's visual experience prior to each saccade. Approach angles near zero indicate symmetric optic flow on the left and right. Positive and negative collision angles indicate stronger optic flow on the left and right, respectively. (B) Magnitude of saccade plotted as a function of the prior approach angle. The corresponding distribution of saccade angles is shown at the right. Saccades separate into two clusters, corresponding to turn towards the left and right. At an approach angle of zero, the fly is equally likely to turn toward the left or right.

checkerboard pattern. In the presence of a visually more complex background, the flies saccade at a much greater distance from the walls and are thus squeezed into the central portion of the arena (figure 28.4A). In addition, the saccade rate increases (figure 28.4C), indicating that the visual environment strongly influences the timing of saccades. Along with the increase in saccade rate, the length of the straight-flight sequences decreases (figure 28.4D), as does the mean flight velocity (figure 28.4E). One possible explanation for these results is that each saccade is elicited by a visually mediated *collision avoidance* reflex that depends in some way on the structure of the visual environment. The striking regularity of the saccade rate seen in figure 28.2 results not from an internal clock but from a resonance of the animal's flight-control reflexes with its visual environment.



**Figure 28.4** Visual background substantially alters search behavior. (A) Example flight trajectories in uniform white and random checkerboard surrounds. (B) Saccade angle-approach angle distributions for the two visual conditions. The basic distribution of saccade angles is effected by the visual surround. (C–F) Kinematic parameters measured under the two visual conditions. In each panel, left and right histogram shows the results in the uniform white and random checkerboard surrounds, respectively. (C) Saccade rate. (D) Path length, measured as distance between saccades. (E) Path velocity. (F) Absolute saccade angle.

Although a landing response has been extensively documented in flies (Borst, 1990), collision-avoidance reactions of the type suggested above have not. For this reason, we developed an experimental paradigm in a tethered-flight arena to investigate such responses under more controlled stimulus conditions (figure 28.5). Within the tethered flight arena, an optoelectronic device tracks the motion of the fly's wings while the animal responds to a panoramic display of LEDs (Lehmann and Dickinson, 1997). The device may be configured in closed loop, so that the animal can itself control the motion of patterns within the display. To study the collision-avoidance response, the animal is allowed to orient toward a small dark square under closed loop; but at regular intervals, the square expands at a fixed rate. As indicated by figure 28.5B, C, the fly's response to the expansion depends systematically on the position of the square when it starts to expand. If the square expands to the animal's left, the fly generates a transient increase in left-stroke amplitude and a transient decrease in right-stroke



**Figure 28.5** Tethered flight experiments confirm presence of a collision avoidance reflex. (A) Tethered flight arena. An optoelectric wingbeat analyzer tracks IR shadows of the two wings, providing a measure of stroke amplitude and frequency. The arena is operated in closed-loop such that the flies control the angular velocity of a dark square by adjusting the relative stroke amplitude of their two wings. (B) As the animal controls its position, the square intermittently expands, eliciting a transient change in stroke amplitude. The sign of the amplitude change depends upon the position of the square when it begins to expand. (C) Summary of amplitude changes as a function of the square's angular position (mean  $\pm$  SD,  $N = 8$ ).

amplitude—a kinematic alteration that would turn the fly away from the stimulus in free flight. An expansion on the animal's right generates the opposite response (figure 28.5C). The directional tuning is not particularly sharp. Aside from the direction (left vs. right), the magnitude of the changes in stroke amplitude does not vary with stimulus position. This coarse directional tuning is consistent with the properties of free-flight saccades as discussed above and supports the hypothesis that flies possess a robust collision-avoidance reflex.

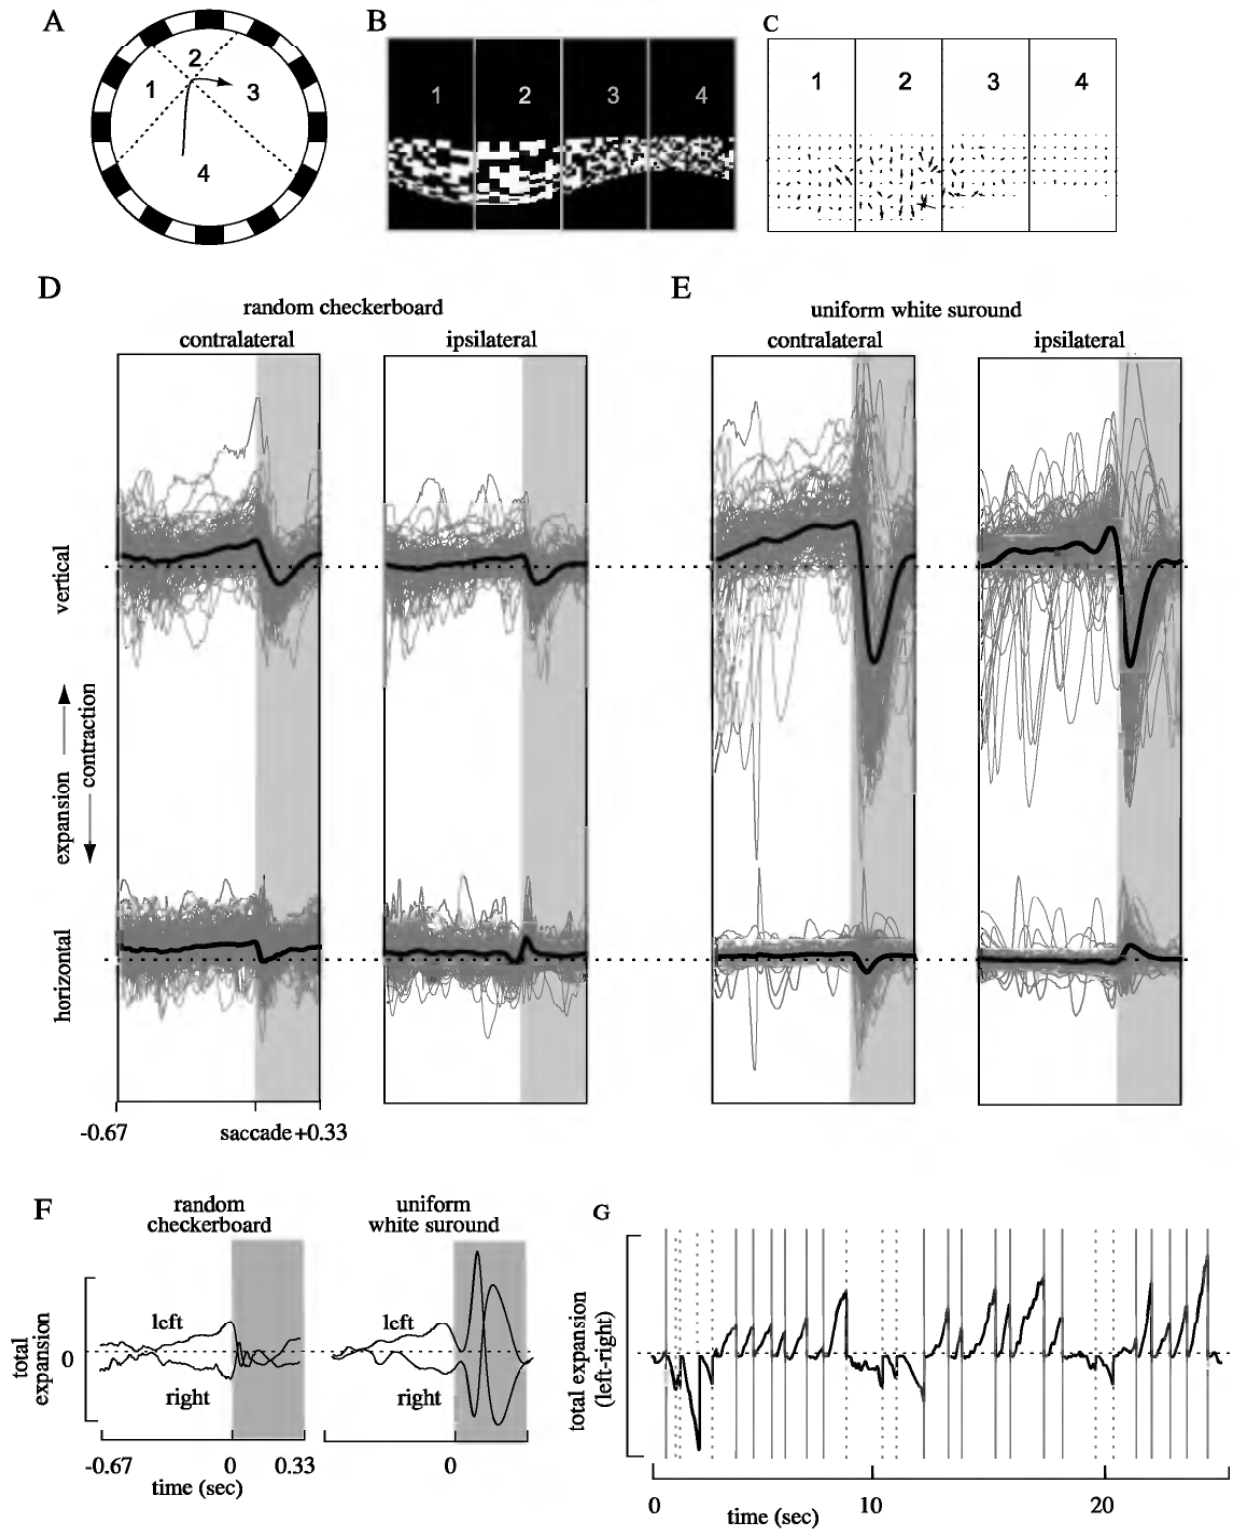
### Reconstructing a Fly's-Eye View

Although the tethered-flight experiments indicate that flies possess a collision-avoidance reflex driven by expansion, it is not clear whether such reactions are actually responsible for triggering free-flight saccades. To gain better insight into what visual stimuli are present immediately preceding each saccade in free flight, we have used the 3-D position of the fly and a mathematical representation of the image background to reconstruct the view of the arena from the fly's perspective (figure 28.6). Because we only have access to the Cartesian coordinates of the fly's body, our reconstruction must make the assumption that

the head (on which the eyes are rigidly fixed) is oriented with the body, and that both head and body are oriented along a tangent to the animal's flight path. The assumptions are, however, consistent with recent free-flight measurements of head position in blowflies (Hateran and Shilstra, 1999) and nevertheless provide a rough estimate of the flow pattern encountered by the fly as it cruises through the arena.

The first step in the analysis is to reconstruct the background image, as seen by the fly, at each time interval in the sequence (figure 28.6A, B). This stream of input images is then processed by a 2-D array of Hassenstein-Reichardt elementary motion detectors (EMDs) to create a signal that is sensitive to the direction and magnitude of optic flow (figure 28.6C). The Hassenstein-Reichardt EMDs use a *delay and correlate* algorithm to compute local flow velocity, and was originally derived to explain the properties of motion detection in insects (Reichardt and Poggio, 1976; Borst and Egelhaaf, 1989). To create a signal sensitive to image expansion, the outputs of the EMDs are summed over eight regions in each visual field, and the difference in summed EMD output across neighboring regions is used as an estimate of vertical and horizontal expansion. The end result of these reconstructions is a series of bilateral signals that are proportional to the visual expansion in the horizontal and vertical directions. To gain insight into what features the fly might use to initiate each saccade, we generated reverse-correlation functions by aligning the visual expansion signals with the occurrence of each saccade (figure 28.6D, E). This technique is analogous to *prespike signal averaging* as a means of estimating the dynamics of sensory cells (Rieke et al., 1997). The reverse correlations are then parsed according to the direction of each saccade (ipsilateral vs. contralateral).

The most obvious features in the reverse correlations are the large transient changes in expansion that occur during the saccade itself (gray areas in figure 28.6D, E). However, it is the time course of expansion *before* the start of the saccade that is most relevant to the properties of the saccade trigger. As expected of a collision-avoidance reaction, the reverse correlations indicate that each saccade is preceded by a stronger expansion on the side contralateral to the direction of the subsequent turn (figure 28.6D, E). The level of expansion rises gently for approximately 500 msec preceding each saccade, as the animal flies across the arena and approaches the wall. The time course and direction of the expansion signals vary with the type of visual background. Vertical expansion is greater in a uniform white surround, whereas horizontal expansion is stronger with a random checkerboard. These differences are expected, because the white background contains two strong horizontal edges, but no vertical edges. However, the time course of total expansion, calculated from the sum of horizontal and vertical components, is quite similar in the two visual backgrounds prior to



**Figure 28.6** Reconstruction of optic expansion preceding each saccade. (A) At each point in time, the fly's visual world is divided into four regions: left rear (1), left front (2), right front (3), and right rear (4). (B) Reconstruction of the arena seen by the fly, based on its instantaneous position. (C) The images (such as that shown in B) are processed by a 2-D array of Reichardt-Hassenstein elementary movement detectors (EMDs). The array computes magnitude and direction of local optic flow based on a delay and correlate algo-


each saccade (figure 28.6F). This suggests that the same neural algorithm is used in both cases, but that the fly must travel faster and fly closer to the wall before its visual system triggers a saccade within a plain white surround.

Based on the results of the reverse correlations of optic expansion, we propose an *integrate-and-fire* model for the saccade generator (figure 28.7). According to this scheme, optic expansion is integrated simultaneously within the left and right visual fields. When the integrator rises above a threshold, it activates a descending command interneuron that elicits a stereotyped change in steering-muscle activity that produces a saccade away from the expansion. Bilateral inhibition of the left and right pathways ensures that the fly cannot attempt to saccade in both ways at the same time. The circuit makes only two decisions: when to turn and whether to do so toward the left or right. It does not compute precisely in which direction the fly should turn. While only a working hypothesis, we are currently attempting to test this model through further manipulation of the visual display within the arena and with electrophysiological recordings.

### WHY DO FLIES FLY STRAIGHT?

So far, our discussion has focused on the sensory stimuli that influence the timing and direction of saccades. However, the straight-flight segments between saccades are important components of the flight-search algorithms used by flies. What sensory cues do flies use to fly straight? Work on two important model systems of insect vision (flies and bees) has led to two rather distinct models of visual flight control.

rithm. The output of the EMDs are spatially summed with the upper and lower halves of regions 1–4 and the resulting eight vectors are used to calculate the horizontal and vertical expansion in the left and right halves of visual space. (D) Vertical and horizontal expansion increases prior to each saccade with a random checkerboard surround. The average vertical and horizontal expansion is calculated on the side away from (contralateral) or towards (ipsilateral) the direction in which the animal turned. The white portion of each plot indicates the time before each saccade, the gray portion indicates the time after each saccade. Both the vertical and horizontal expansion is greater on the contralateral side before each saccade than on the ipsilateral side. (E) Optic expansion preceding saccades in a uniform white surround. Vertical expansion prior to each saccade is much greater than with a random checkerboard due to the presence of the sharp horizon with the uniform white surround, while horizontal expansion is much smaller. As with the random checkerboard, however, the expansion signals are always greater on the contralateral side. (F) Total expansion (sum of horizontal and vertical signals) preceding left and right saccades. The time course of total expansion leading up to the saccades is similar in both visual surrounds. (G) Time course of total expansion within a random checkerboard. Saccades to the right are indicated by solid vertical lines, left saccades are indicated by dotted lines. The direction of optic expansion predicts the direction of each saccade.



**Figure 28.7** “Integrate-and-fire” model for collision avoidance reaction. Schematic representation of model. First, the visual image is processed by an array of EMDs. Next, the output of the EMD layer is spatially summed and processed to construct a total expansion signal on the left and right sides. When an integration of the expansion signal reaches threshold, a saccade is triggered by a descending interneuron and the integrator is reset. Mutual inhibition between the two sides ensures that left and right saccades will not be generated simultaneously.

The prominent model in flies, *optomotor equilibrium*, is based on the optomotor response of tethered animals. In the presence of a rotating visual field, tethered flies modulate yaw torque in an attempt to minimize the velocity of the rotating pattern. Thus, if a striped drum rotates to the animal’s right, the fly attempts to turn toward the right. In free flight, such negative feedback should continuously stabilize flight course. Similar reflexes have been documented for rotations around the pitch and roll axes (Blondeau and Heisenberg, 1982; Götz, 1983). The behavior of tethered flies is in sharp contrast to the centering response of free-flying bees, which has been analyzed by Srinivassen and colleagues (1991). Worker bees turn away, not toward, a moving visual pattern. Such reflexes, acting bilaterally, are sufficient to explain the bee’s robust ability to fly down the geometric center of a narrow channel—the only position where the avoidance reflexes of both side are balanced. The two visual-feedback models (optomotor equilibrium from flies and the centering reflex from bees) predict rather different responses of insects flying within a small uniform circular arena such as the one used in our experiments. In order to stabilize course directions, the optomotor system must be capable of distinguishing rotation from translation. Although this may be possible when the translational signal is equal on the two sides of the body, the task is more difficult



when there is a bilateral asymmetry in optic flow, as would occur when an animal flies close to one wall. According to a simple optomotor-equilibrium model, an animal flying along a path that does not pass through the center of the arena should tend to turn toward the closest wall. In contrast, according to the centering-reflex model, the insect should veer away from the wall. When fruit flies are released into the arena, they neither turn away nor toward the wall, but rather fly along a straight trajectory until the approach of a wall triggers a saccade. The straightness of the flight paths do not appear to depend upon heading or proximity to the wall. Thus, the free-flight behavior of fruit flies is consistent with neither the optomotor-equilibrium nor the centering-reflex models of visual flight control.

If fruit flies do not follow the predictions of the optomotor response, what sensory cues guide them during the straight stretches of their flight path? One possibility is that the optomotor equilibrium circuitry is sophisticated enough to differentiate between rotation and translation, even when the translation signal is greatly skewed by the location of walls (Kern and Egelhaaf, 2000). Another, not mutually exclusive, possibility is that feedback from a different modality overrides the input from the visual system, stabilizing flight in the face of a strong optomotor disequilibrium. The likely source of such input is the haltere system (Dickinson, 1999). Because the detection of rotation is based on mechanosensory neurons, the response dynamics of the halteres are much faster than the visual system and should provide a more finely tuned control of yaw, pitch, and roll. A multimodal hierarchy in which the halteres can override feedback from the visual system could explain some of the differences between tethered- and free-flight experiments. For example, optomotor reflexes may appear strong in tethered flight because the flies are rotationally fixed and the halteres are silent. In free flight, input from the halteres would override the optomotor responses, maintaining a straight flight course in the face of asymmetric optic flow. The input from the haltere to the flight motor is quite strong. Haltere afferents make electrotonic, monosynaptic connections with important steering motoneurons (Fayyazuddin et al., 1994; Fayyazuddin and Dickinson, 1996; Trimarchi and Murphey, 1997). In contrast, there is no direct evidence for a connection between descending visual interneurons and steering motoneurons of the wing, although such connections may be gated off under the conditions in which it is possible to record from the steering motoneurons.

If the flight motor listened only to the haltere input, however, optic expansion could never elicit a saccade. Thus, at certain times the visual system must be capable of overriding the feedback from the halteres. Behavioral, anatomical, and electrophysiological evidence suggest that this could happen in one of two ways. First, the motor pattern elicited by a descending command might temporarily overpower the feedback

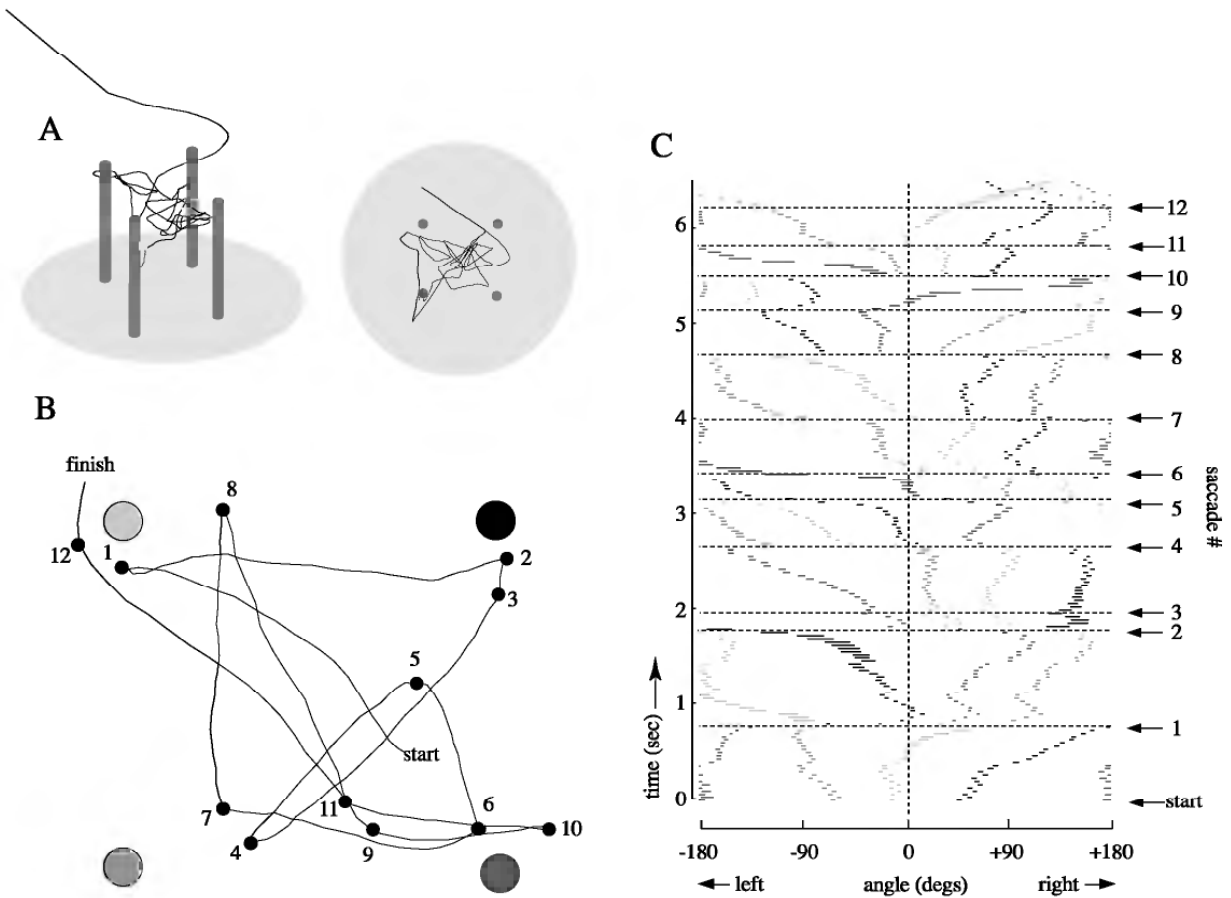


from the haltere. Second, the nervous system might lower the gain of the haltere feedback to allow the visual feedback to take over. Like the wing, the haltere is equipped with a set of tiny control muscles that insert upon sclerites at its base (Bonhag, 1948; Mickoleit, 1962). In blowflies, motoneurons of at least two of these muscles receive input from motion-selective descending visual interneurons (Chan et al., 1998). It is not yet known what effects these steering muscles have on haltere kinematics, but it is possible that they might influence the recruitment or firing phase of haltere sensory neurons. These afferents could then relay these temporal changes to wing muscles via their strong connections with steering motoneurons. Such a pathway provides a means by which the visual system could alter the gain of haltere-mediated reflexes.

## RESPONSES TO REAL VISUAL OBJECTS

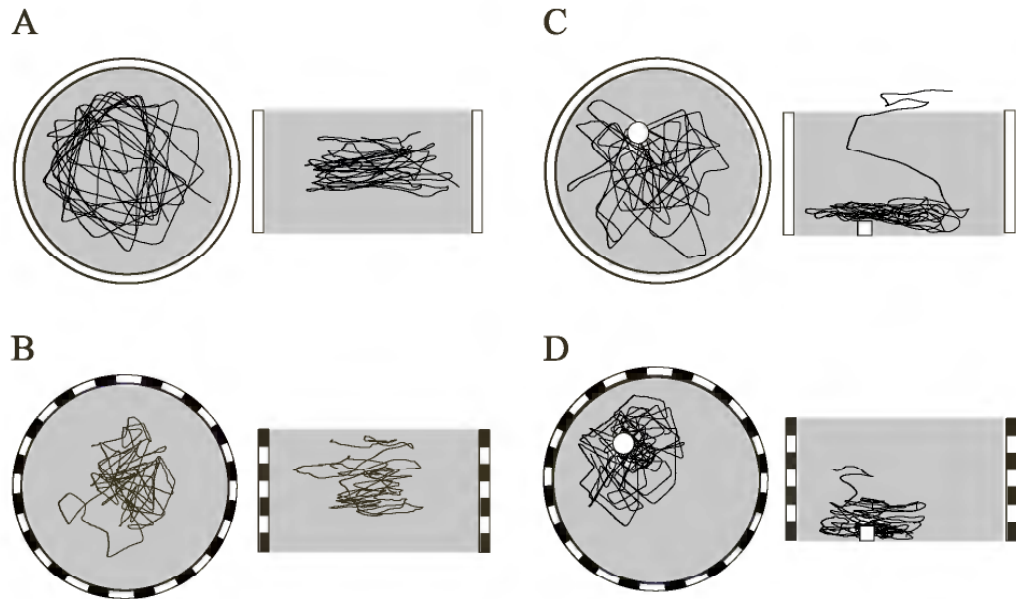
While altering the 2-D background of the arena provides a convenient way for studying visual feedback in flight, in natural environments, flies must navigate among real physical objects. To create a more spatially complex landscape, we placed sets of black cylinders inside the arena. The presence of these objects greatly alters the flight trajectories, as indicated by figure 28.8A. Superficially, the flight trajectories look like those the flies produce in response to a simple background: straight sequences interspersed with rapid saccades. A more detailed analysis, however, indicates that the response to individual objects is more complex (figure 28.8B, C). Immediately after each saccade, the flies appear to orient directly toward an object. So far, we find no evidence that flies direct their saccades toward particular targets. Instead, they orient toward whatever object happens to be in front of them after they saccade. The orientation toward a contrast edge following each saccade is reminiscent of fixation behavior, the robust tendency to track a vertical stripe in tethered flight under closed-loop conditions (Reichardt and Wenking, 1969; Wehner, 1981). If the orientation of flies to the cylinders in free flight is analogous to fixation in tethered flight, then the response begins to weaken as the animal approaches the object. As indicated in figure 28.8C, after initially tracking the edge of an object, the fly begins to deviate such that the object's edge drifts progressively away. This lapse in fixation is followed by a saccade, presumably triggered by the expanding retinal image of the approaching object. In tethered flight, fixation never weakens. Given the opportunity, flies will track a stripe for hours and even days within a tethered-flight arena (Götz, 1987). In free flight, something about the approach of the object must inhibit this strong fixation.

There are several possible explanations for the apparent diminution of object fixation as the animal approaches the target. First, the net motion of all the other objects might elicit an optomotor response that



**Figure 28.8** Free-flying flies temporarily orient towards 3-D objects. (A) Inclined and birds-eye view of flight trajectory among four thin cylinders. Arena background was uniform white surround. (B) Enlarged view of a different flight sequence consisting of 12 saccades. (C) Fixation during free flight sequence. At each time point, the four short horizontal lines indicate the azimuth angles subtended by the four posts throughout the sequence shown in B. Each post is labeled with a different shade of gray. Note that the length of the lines increases in as the animal approaches. The timing of saccades are indicated by the dashed horizontal lines labeled by a number. The vertical dashed line indicates the center point of the fly's visual field. Fixation is indicated by the edge of a post being held close to the center of the visual field. Fixation is strongest immediately following a saccade, but tends to deteriorate as the animal approaches the cylinder.

acts counter to the fixation response. If this were true, one would expect the fly's trajectory to deviate even more strongly as it approaches the object. However, the fly's path is quite straight just prior to each saccade, as if the image of the object exerts very little influence on flight control. An alternate hypothesis is that the fixation response is simply inhibited by the approach of the object, and the fly resorts to flying straight under haltere control. An inhibition of fixation with object proximity it is consistent with the physiology of certain small field (FD) cells identified in *Calliphora*. The FD cells in lobula plate are sensitive to the horizontal motion of small objects, but become less sensitive as the



**Figure 28.9** Response to odor depends on visual background. The fly requires a rich textured background in order to localize the horizontal position of the odor. (A) Flight trajectory in uniform white surround in the absence of odor. (B) Flight trajectory with random checkerboard in the absence of odor. (C) Flight trajectory in uniform white surround with odor. (D) Flight trajectory with random checkerboard with odor.

absolute size of the object grows (Egelhaaf, 1985). If such a cell were to underlie fixation, one would expect the response to attenuate as the object begins to loom.

### A GAIN-CONTROL MODEL FOR VISION-OLFACTORY FUSION

As shown in figure 28.9, the presence of an attractive odor dramatically alters the flight patterns of flies within the arena. While they continue to exhibit the typical pattern of straight flight interspersed with saccades, they center their flight paths directly over the source of the odor. Unless the odor source is associated with an obvious visual object (such as a black cylinder), the flies do not land, but will continue to search for the odor for more than 30 minutes, at which point their behavior gradually transforms to the pattern typical of an arena without odor—an alteration that is most parsimoniously explained by stimulus sensitization typical of many olfactory systems.

There are two general explanations for the transformation of the fly's behavior in the presence of odor. The first possibility is that the detection of odor can independently activate the saccade generator by a convergent or parallel pathway to that of the visual system. The second possibility is that odor alters flight path, not through its direct influence on the flight motor, but through its effect on the visually mediated pathway. To test between these two general alternatives, we presented the

same odor stimulus to flies within the visually sparse uniform white surround (figure 28.9C). In the absence of a complex visual background, flies were not capable of localizing the horizontal position of the odor source. This result is consistent with the hypothesis that odor alters flight path through its influence on the visual pathway. Curiously, although the flies cannot locate the horizontal position of the odor source within the uniform white visual surround, they can localize the vertical position. This observation is significant because it indicates that the failure to localize the odor in the horizontal plane does not result from a general inability to detect the odor. On the other hand, it begs the question as to why the fly can localize the horizontal position of an odor but not the vertical position. One possibility is that the odor can directly influence motor circuits that control altitude. However, even when the arena is lined with a uniform white background, two strong vertical edges are present: one horizon at the border of the dark floor and white wall, and another at the border between the white wall and the black curtain above the arena. Therefore, another possibility is that odor's influence on altitude does operate via its effect on visual control circuits, and that the uniform white background provides visual input for the vertical system but not the horizontal system. The motion of these two horizontal edges may provide adequate input for visual cells that detect vertical motions, whereas the complete absence of vertical edges greatly reduces the input to cells that respond to horizontal motion. At present, there is no way to discriminate among the two hypotheses. The important result, however, is that much of an odor's influence on flight behavior appears to operate through its effect on the gain of visually mediated reflexes.

### **CONCLUSIONS: A MECHANISTIC MODEL OF SEARCH BEHAVIOR**

By studying the free-flight behavior of fruit flies under controlled conditions, we have identified a series of behavioral responses to visual backgrounds, objects, and odors. Can these quite simple behavioral units combine to generate a robust and flexible search behavior? Based on our analysis so far, we propose the following hypothetical model to account for the search behavior of *Drosophila*. The foundation of the model is a set of background feedback circuits that keep the animal stabilized. Chief among these is an equilibrium reflex, mediated by the halteres, which controls rotations around the roll, pitch, and yaw axes. In the absence of all external sensory information, the animal flies straight under the control of the halteres. During straight flight, the visual system constantly integrates the optic flow on the left and right sides of the animal. If the flow exceeds a threshold, indicating the approach of a large obstacle, the animal rapidly turns away from

the expansion using the saccade motor program. This simple system—straight flight and collision avoidance—is sufficient to keep an animal flying throughout complex visual landscape and does not require that the visual system have extreme spatial resolution. The repetitive iteration breaks down if a saccade happens to direct the fly toward an object that stands out from the background. Under these circumstances, the animal will transiently correct its course toward the edge of the object, overriding the yaw control from the halteres. This fixation behavior serves the functional role of attracting the animal toward potentially interesting targets within its visual world. As the fly approaches the object, however, the growth of the object's image inhibits the fixation response, and the animal flies straight once again under haltere control. The expansion of the object will eventually lead to a collision-avoidance saccade that will initiate a new sequence of fixation, fixation inhibition, and collision avoidance. This iterative search pattern is altered by the presence of attractive odors. Odors change the thresholds of both the fixation inhibition and the collision-avoidance saccades. The net effect of these sliding thresholds is that the fly will approach objects and boundaries within its environment much more closely if odors are present. In order to account for the fly's ability to land, we propose that looming objects can either elicit a collision-avoidance reaction, or a landing response, but that the thresholds for the two reflexes vary with odor. In the absence of any odor, looming triggers a collision avoidance response at a lower threshold than it does a landing response. Under these conditions, a hungry fly will turn away from an object and continue its search algorithm. If, however, the fly detects odor, then the threshold for the landing response decreases, while the threshold for collision avoidance increases. In this modified state, the fly will land in response to the same looming stimulus.

The model outlined above represents a working hypothesis in an ongoing investigation of search-control algorithms in fruit flies. We are currently attempting to experimentally verify and refine the model, as well as formalize it mathematically. While still preliminary, this work begins to show how a rather complex and robust process, such as search behavior, might arise from a system of simple reflexes. A common feature of these reflexes is the ability of a stimulus in one sensory modality to alter the gain of another. Such principles should provide useful insight to the goal of creating robotic behaviors with the same degree of robustness and efficiency.

## REFERENCES

- Autrum, H. (1958). Electrophysiological analysis of the visual systems in insects. *Exp. Cell Res.*, supplement 5: 426–439.
- Blondeau, J., and Heisenberg, M. (1982). The three dimensional optomotor torque system of *Drosophila melanogaster*. *J. Comp. Physiol. [A]* 145: 321–329.

- Borst, A. (1990). How do flies land? *BioScience* 40: 292–299.
- Borst, A., and Egelhaaf, M. (1989). Principles of visual motion detection. *Trends Neurosci.* 12: 297–306.
- Collett, T. S., and Land, M. F. (1975). Visual control of flight behavior in the hoverfly, *Syrirtta pipiens* L. *J. Comp. Physiol. [A]* 99: 1–66.
- David, C. T. (1985). Visual control of the partition of flight force between lift and thrust in free-flying *Drosophila*. *Nature* 313: 48–50.
- Derham, W. (1714). *Physico-Theology*. London.
- Dickinson, M. (1999). Haltere-mediated equilibrium reflexes of the fruit fly, *Drosophila melanogaster*. *Philos. Trans. R. Soc. Lond. B Biol. Sci.* 354: 903–916.
- Egelhaaf, M. (1985). On the neuronal basis of figure-ground discrimination by relative motion in the visual system of the fly: II. Figure-detection cells, a new class of visual interneurons. *Biol. Cybern.* 52: 195–209.
- Fayyazuddin, A., Chan, W. P., Jordan, C. E., and Dickinson, M. H. (1994). The role of haltere afferents in the activity of a steering muscle in the blowfly, *Calliphora vicina*. *Soc. Neurosci. Abstr.* 20: 1595.
- Fayyazuddin, A., and Dickinson, M. H. (1996). Haltere afferents provide direct, electrotonic input to a steering motor neuron of the blowfly, *Calliphora*. *J. Neurosci.* 16: 5225–5232.
- Fraenkel, G. (1939). The function of the halteres of flies. *Proc. Zool. Soc. Lond. [A]* 109: 69–78.
- Götz, K. G. (1983). Bewegungssehen and Flugsteuerung bei der Fliege *Drosophila*. In *BIONA-Report 2* W. Nachtigall (ed.), Stuttgart: Fischer, pp. 21–34.
- Götz, K. G. (1987). Course-control, metabolism, and wing interference during ultralong tethered flight in *Drosophila melanogaster*. *J. Exp. Biol.* 128: 35–46.
- Götz, K. G., and Wandel, U. (1984). Optomotor control of the force of flight in *Drosophila* and *Musca*: II. Covariance of lift and thrust in still air. *Biol. Cybern.* 51: 135–139.
- Hateran, J. V., and Shilstra, C. (1999). Blowfly flight and optic flow: II. Head movements during flight. *J. Exp. Biol.* 202: 1491–1500.
- Heide, G., and Götz, K. G. (1996). Optomotor control of course and altitude in *Drosophila* is correlated with distinct activities of at least three pairs of steering muscles. *J. Exp. Biol.* 199: 1711–1726.
- Heisenberg, M., and Wolf, R. (1979). On the fine structure of yaw torque in visual flight orientation of *Drosophila melanogaster*. *J. Comp. Physiol. [A]* 130: 113–130.
- Heisenberg, M., and Wolf, R. (1984). *Vision in Drosophila*. Berlin: Springer-Verlag.
- Kern, R., and Egelhaaf, M. (2000). Optomotor course control in flies with largely asymmetric visual input. *J. Comp. Physiol. [A]* 1186: 45–55.
- Laissue, P., Reiter, C., Hiesinger, P., Halter, S., Fischbach, K., and Stocker, R. (1999). Three-dimensional reconstruction of the antennal lobes in *Drosophila melanogaster*. *J. Comp. Neurol.* 405: 543–552.
- Lehmann, F.-O., and Dickinson, M. H. (1997). The changes in power requirements and muscle efficiency during elevated force production in the fruit fly, *Drosophila melanogaster*. *J. Exp. Biol.* 200: 1133–1143.

- Nalbach, G. (1993). The halteres of the blowfly *Calliphora*: I. kinematics and dynamics. *J. Comp. Physiol. [A]* 173: 293–300.
- Nalbach, G., and Hengstenberg, R. (1994). The halteres of the blowfly *Calliphora*: II. Three-dimensional organization of compensatory reactions to real and simulated rotations. *J. Comp. Physiol. [A]* 174: 695–708.
- Reichardt, W., and Poggio, T. (1976). Visual control of orientation behaviour in the fly: I. A quantitative analysis of neural interactions. *Q. Rev. Biophys.* 9: 311–375.
- Reichardt, W., and Wenking, H. (1969). Optical detection and fixation of objects by fixed flying flies. *Naturwiss.* 56: 424–425.
- Rieke, F., Warland, D., de Ruyter van Steveninck, R., and Bialeck, W. (1997). *Spikes: Exploring the Neural Code*. Cambridge: MIT Press, p. 395.
- Srinivasan, M. V., Lehrer, M., Kirchner, W., and Zhang, S. W. (1991). Range perception through apparent image speed in freely-flying honeybees. *Vis. Neurosci.* 6: 519–535.
- Trimarchi, J. R., and Murphey, R. K. (1997). The shaking-B<sup>2</sup> mutation disrupts electrical synapses in a flight circuit in adult *Drosophila*. *J. Neurosci.* 17: 4700–4710.
- Vickers, N. (2000). Mechanisms of animal navigation in odor plumes. *Biol. Bull.* 198: 203–212.
- Wehner, R. (1981). Spatial vision in arthropods. In H. Autrum (ed.), *Comparative Physiology and Evolution of Vision in Invertebrates. C: Invertebrate Visual Centers and Behavior II*, vol. VII/6C. Berlin: Springer-Verlag, pp. 286–616.

Mandyam V. Srinivasan

Unlike vertebrates, insects have immobile eyes with fixed-focus optics. Therefore, they cannot infer the distances to objects or surfaces from the extent to which the directions of gaze must converge to view the object, or by monitoring the refractive power that is required to bring the image of the object into focus on the retina. Furthermore, compared with human eyes, the eyes of insects are positioned much closer together, and possess inferior spatial acuity (Horridge, 1977). Therefore, the precision with which insects could estimate range through binocular stereopsis would be much poorer and restricted to relatively small distances, even if they possessed the requisite neural apparatus (Srinivasan, 1993). Not surprisingly, then, insects have evolved alternative strategies for dealing with the problems of visually guided flight. Many of these strategies rely on using image motion, generated by the insect's own motion, to infer the distances to obstacles and to control various maneuvers (Horridge, 1987; Srinivasan, 1993, 1998).


Here we describe three ways in which flying honeybees use cues based on image motion to navigate safely to a goal. The specific navigational problems considered are: negotiating narrow gaps and avoiding obstacles, controlling flight speed, and performing smooth landings. We also briefly describe tests of the effectiveness of these strategies by implementation in robots.

#### **NEGOTIATING NARROW GAPS AND AVOIDING OBSTACLES**

When a bee flies through a hole in a window, it tends to fly through its center, without colliding with the rims. How does she gauge and balance the distances to the two rims?

One possibility is that the bee does not measure distances at all, but simply balances the speeds of image motion on the two eyes, as she flies through the opening. To investigate this possibility, Kirchner and Srinivasan (1989) trained bees to enter an apparatus that offered a reward of sugar solution at the end of a tunnel. Each side wall carried a pattern consisting of a vertical black-and-white grating (figure 29.1).





**Figure 29.1** Illustration of an experiment which demonstrates that flying bees infer range from apparent image speed. The shaded areas represent the means and standard deviations of the positions of the flight trajectories, analyzed from video recordings of several hundred flights. (Adapted from Srinivasan et al. 1996.)

The grating on one wall could be moved horizontally at any desired speed, either toward the reward or away from it. After the bees had received several rewards with the gratings stationary, they were filmed from above as they flew along the tunnel. When both gratings were stationary, the bees tended to fly along the midline of the tunnel—that is, equidistant from the two walls (figure 29.1a). But when one of the gratings was moved at a constant speed in the direction of the bees' flight—thereby reducing the speed of retinal image motion on that eye relative to the other eye—the bees' trajectories shifted toward the side of the moving grating (figure 29.1b). When the grating moved in a direction opposite to that of the bees' flight—thereby increasing the speed of retinal image motion on that eye relative to the other—the bees' trajectories shifted away from the side of the moving grating (figure 29.1c). These findings demonstrate that when the walls were stationary, the bees maintained equidistance by balancing the speeds of the retinal images in the two eyes. A lower image speed on one eye evidently caused the bee to move closer to the wall seen by that eye. A higher image speed, on the other hand, had the opposite effect.

Were the bees really measuring and balancing image speeds on the two sides as they flew along the tunnel, or were they simply balancing the contrast frequencies produced by the succession of dark and light bars of the gratings? This question was investigated by analyzing the flight trajectories of bees when the two walls carried gratings of different spatial periods. When the gratings were stationary, the trajectories were always equidistant from the two walls, even when the spatial frequencies of the gratings on the two sides—and therefore the contrast frequencies experienced by the two eyes—differed by a factor of as much as four (figure 29.1d). When one of the gratings was in motion, the trajectories shifted toward or away from the moving grating (as described above) according to whether the grating moved with or

against the direction of the bees' flight (figure 29.1e, f). These results indicate that the bees were indeed balancing the speeds of the retinal images on the two eyes and not the contrast frequencies. The above findings are true irrespective of whether the gratings possess square-wave intensity profiles (with abrupt changes of intensity) or sinusoidal profiles (with gradual intensity changes) and irrespective of whether the contrasts of the gratings on the two sides are equal, or considerably different (Srinivasan et al., 1991). Thus, the honeybee's visual system is capable of measuring the image velocities in the two eyes robustly and independently and using this information to steer a collision-free path through the gap.

The mechanism described above can also be used to avoid obstacles. For example, consider a situation in which an insect flies along a straight line toward a goal. If an object on the left is dangerously close to the intended trajectory, it will generate a high image velocity in the left eye, causing the insect to veer to the right (provided there isn't an equally close object on the other side repelling it back). Thus, the "centering" response described above is also likely to mediate obstacle avoidance. Indeed, it has been shown that flying bees strongly avoid stimuli that present rapid movement (Srinivasan and Lehrer, 1984; Srinivasan et al., 1991; Srinivasan, Zhang, and Chandrashekara, 1993; Srinivasan and Zhang, 1997).

## CONTROLLING FLIGHT SPEED

Do insects control the speed of their flight and, if so, how? Work by David (1982) and by Srinivasan and colleagues (1996) suggests that flight speed is controlled by monitoring the velocity of the image of the environment.

David (1982) observed fruit flies flying upstream in a wind tunnel, attracted by an odor of fermenting banana. The walls of the cylindrical wind tunnel were decorated with a helical black-and-white-striped pattern, so that rotation of the cylinder about its axis produced apparent movement of the pattern toward the front or the back. With this setup, the rotational speed of the cylinder (and hence the speed of the backward motion of the pattern) could be adjusted such that the fly was stationary (i.e., did not move along the axis of the tunnel). The apparent backward speed of the pattern then revealed the ground speed that the fly was "choosing" to maintain, as well as the angular velocity of the image of the pattern on the fly's eyes. In this setup, fruit flies tended to hold the angular velocity of the image constant. Increasing or decreasing the speed of the pattern caused the fly to move backward or forward (respectively) along the tunnel at a rate such that the angular velocity of the image on the eye was always "clamped" at a fixed value. The flies also compensated for headwind in the tunnel, increasing or

decreasing their thrust so as to maintain the same apparent ground speed (as indicated by the angular velocity of image motion on the eye). Experiments in which the angular period of the stripes was varied revealed that the flies were measuring (and holding constant) the angular velocity of the image on the eye, irrespective of the spatial structure of the image.

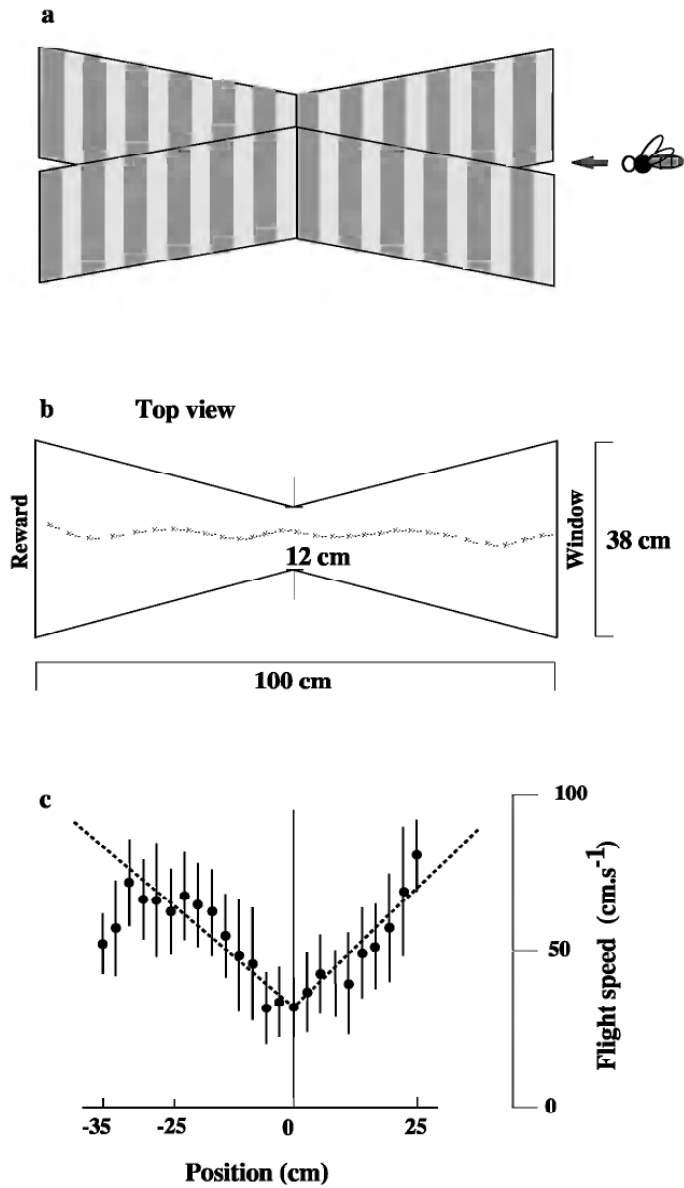
Bees appear to use a similar strategy to regulate flight speed (Srinivasan et al., 1996). When a bee flies through a tapered tunnel, she decreases her flight speed as the tunnel narrows so as to keep the angular velocity of the image of the walls, as seen by the eye, constant at about  $320^\circ/\text{sec}$  (figure 29.2). This suggests that flight speed is controlled by monitoring and regulating the angular velocity of the image of the environment on the eye. (That is, if the width of the tunnel is doubled, the bee flies twice as fast.) On the other hand, a bee flying through a tunnel of uniform width does not change her speed when the spatial period of the stripes lining the walls is abruptly changed (figure 29.3). This indicates that flight speed is regulated by a visual-motion-detecting mechanism that measures the angular velocity of the image largely independently of its spatial structure.

An obvious advantage of controlling flight speed by regulating image speed is that the insect would automatically slow down to a safer speed when negotiating a narrow passage.

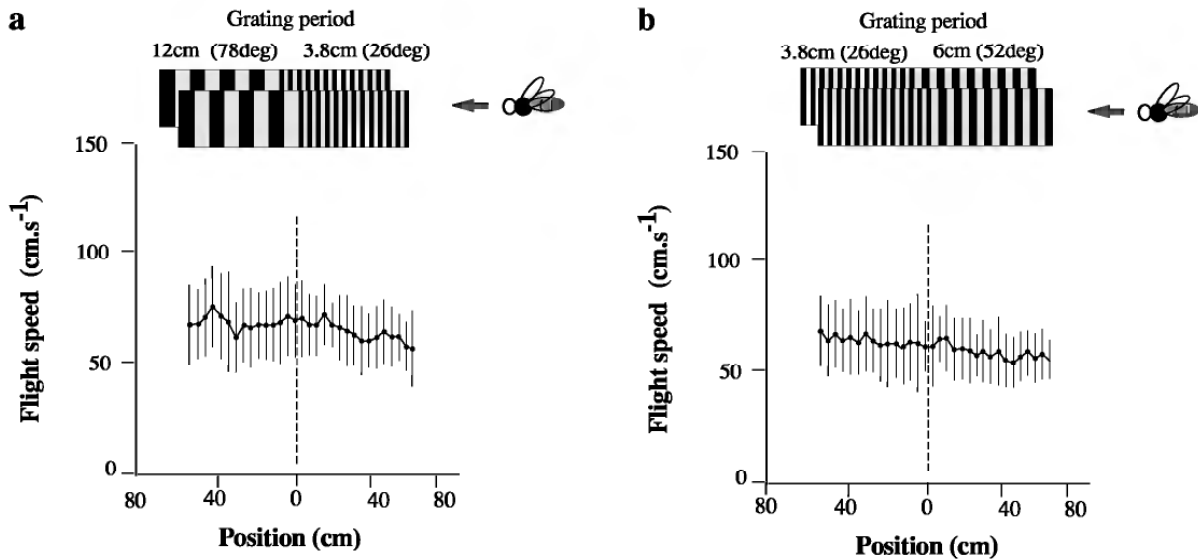
The above findings offer simple strategies for autonomous navigation of robots along corridors. By balancing the speeds of the images of the two side walls, one can ensure that the robot progresses along the middle of the corridor without bumping into the walls. Furthermore, the speed of the robot can be adjusted to a safe value by holding constant the average velocity of the images of the two walls.

Four different laboratories have now built robots that negotiate corridors successfully using the bee-derived principle of balancing lateral image motion (Coombs and Roberts, 1992; Sandini et al., 1993; Duchon and Warren, 1994; Weber, Venkatesh, and Srinivasan, 1997). Computationally, this method is far more amenable to real-time implementation than methods that use stereo vision to calculate the distances to the walls.

The design and performance of one of these robots (Weber, Venkatesh, and Srinivasan, 1997) is shown in figure 29.4. The robot is approximately the size of a small skateboard, with a single video camera mounted facing upward (figure 29.4a). This camera captures views of the side walls (one of each wall) through a mirror assembly positioned above the lens. Video information from the camera is transmitted to a computer, where the image velocities of the two walls, induced by the motion of the robot, are measured. (The algorithms for measuring image motion are described in Srinivasan, 1990, 1994 and Srinivasan et al., 1997.) The computer then issues appropriate steering commands to the



**Figure 29.2** Experiment investigating visual control of flight speed. Bees flying through a tapered tunnel decrease their flight speed when the tunnel narrows, and increase the speed when it widens. The dots and vertical bars represent means and standard deviations of flight speed measured at various locations. The dashed line shows the theoretically expected flight speed profile if the bees were to hold the angular velocity of the images of the walls constant at  $320 \text{ deg}/\text{sec}$  as they fly through the tunnel. The data indicate that bees control flight speed by holding constant the angular velocity of the image of the environment. (Adapted from Srinivasan et al. 1996.)




**Figure 29.3** Two experiments (a, b) examining control of flight speed in tunnels of constant width, each lined with black-and-white gratings whose spatial period changes abruptly in the middle. The dots and vertical bars represent means and standard deviations of flight speed measured at various locations. Bees flying through such tunnels maintain a nearly constant flight speed regardless of whether the period increases (a) or decreases (b). This suggests that the speed of flight is controlled by a movement-detecting system which measures and holds constant the speeds of the images of the walls accurately regardless of their spatial structure. (Adapted from Srinivasan et al. 1996.)


robot to ensure that it stays close to the mid-line of the tunnel. The tunnel-following performance of the robot is illustrated by the examples shown in figure 29.4b–e. In all cases, the robot reliably follows the axis of the corridor. The presence of an obstacle next to one of the walls (figure 29.4e) causes the robot to go through the middle of the gap remaining between the obstacle and the other wall. Additional control algorithms have been developed for controlling the speed of the robot. Speed control is achieved by holding constant the sum of the image speeds from the two walls, thereby ensuring that the robot automatically slows down to a safe speed when the corridor narrows, like the bee in figure 29.2. Thus, the strategies used by flying bees to negotiate narrow gaps and to regulate flight speed do appear to be feasible.

### PERFORMING SMOOTH LANDINGS

The seminal work of Gibson (1950) has highlighted the optic-flow cues that can be brought to bear in controlling the landing of an aircraft. Studies of landing behavior in flies have revealed that, as a surface is approached, the expansion of the image of the surface provides strong cues that are used to control deceleration and trigger extension of the legs in preparation for contact (Goodman, 1960; Eckert and Hamdorf,



a

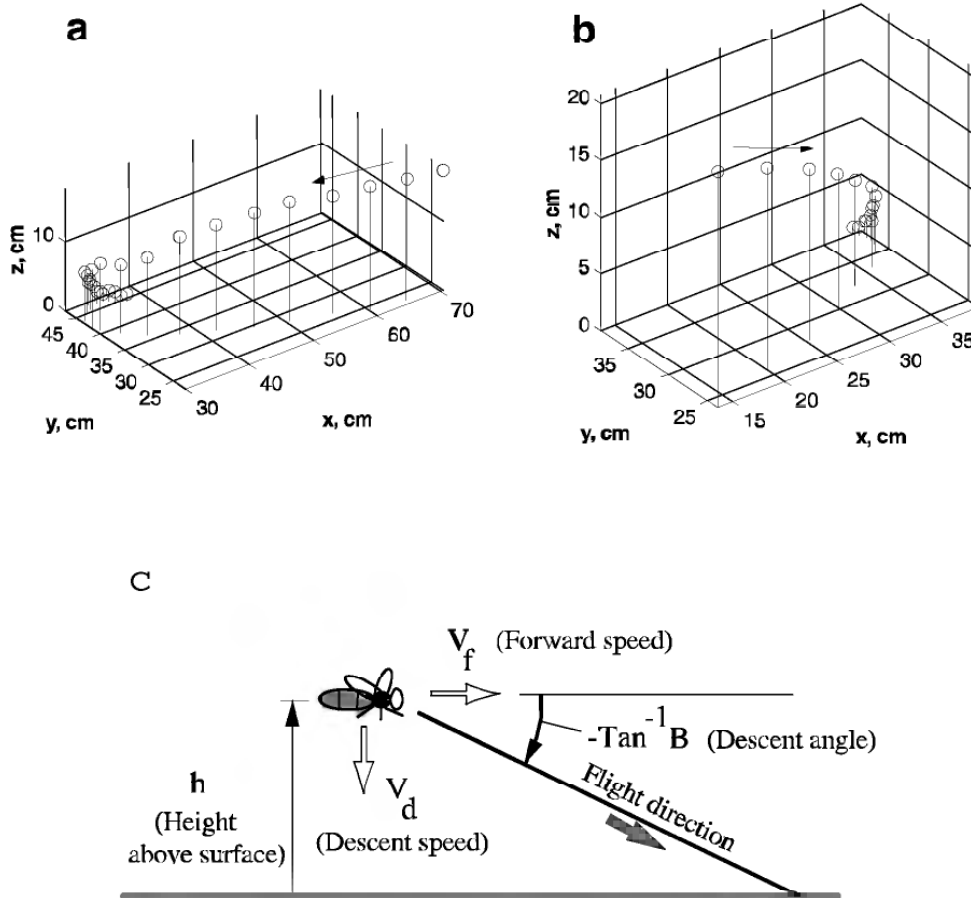


**Figure 29.4** (a) Frontal view of corridor-following robot. (b–e) Performance in variously shaped corridors, of width approximately 80–100 cm. (Adapted from Srinivasan et al. 1997.)

1980; Wagner, 1982; Borst and Bahde, 1988). There is also evidence that the rate of expansion of the image is used to infer the time to contact the surface, even if the insect does not possess explicit information about the speed of its flight or the distance to the surface (Wagner, 1982).

However, when an insect makes a grazing landing on a flat surface, cues derived from image expansion are relatively weak. This is because the dominant pattern of image motion is then a translatory flow in the front-to-back direction. Given that flying insects often make grazing landings on flat surfaces, what are the processes by which such landings are orchestrated? Recently, Srinivasan and colleagues (2000) have investigated this question by videotaping trajectories, in three dimensions, of bees landing on a flat, horizontal surface.

Two examples of landing trajectories, reconstructed from the data, are shown in figure 29.5a, b. A number of such landing trajectories were analysed to examine the variation of the instantaneous height above the surface ( $h$ ), instantaneous horizontal (forward) flight speed ( $V_f$ ), instantaneous descent speed ( $V_d$ ) and descent angle ( $\alpha$ ). These

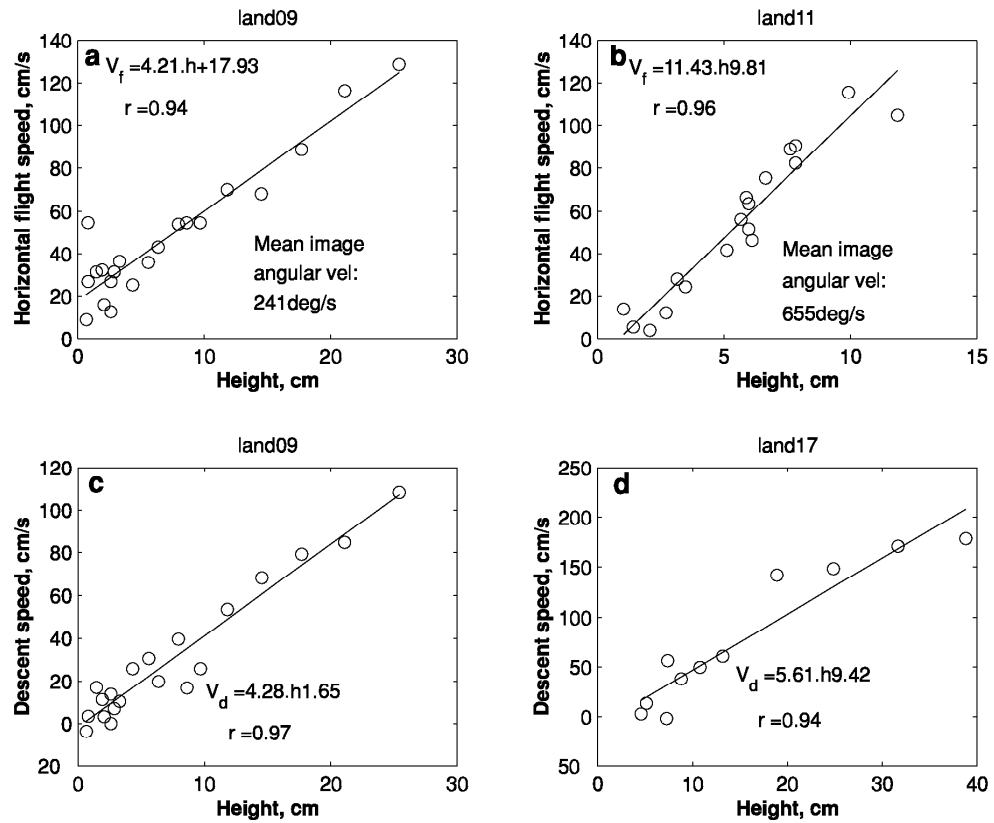


**Figure 29.5** (a, b) 3-D reconstruction of two typical landing trajectories, from video films. Vertical lines depict height above surface. (c) Illustration of some of the variables analysed to investigate the control of landing.  $h$  (cm): height above surface;  $V_f$  (cm/sec): horizontal (forward) flight speed;  $V_d$  (cm/sec): vertical (descent) speed;  $\tan^{-1} B$  (deg or rad): descent angle, where  $B$  is the ratio of instantaneous descent speed to instantaneous forward speed.

variables are illustrated in figure 29.5c. Analysis of the landing trajectories revealed that the descent angles were indeed quite shallow. The average value measured in 26 trajectories was approximately  $28^\circ$  (Srinivasan et al., 2000).

Figure 29.6a, b shows the variation of flight speed with height above the surface, analyzed for two landing trajectories. This data reveals one of the most striking and consistent observations of this study: Horizontal speed is roughly proportional to height, as indicated by the linear regression on the data. When a bee flies at a horizontal speed of  $V_f$  cm/sec at a height of  $h$  cm, the angular velocity  $\omega$  of the image of the surface directly beneath the eye is given by

$$\omega = \frac{V_f}{h} \text{ rad/sec}$$



**Figure 29.6** (a, b) Variation of horizontal flight speed ( $V_f$ ) with height ( $h$ ) above the surface, for two different landing trajectories. (c, d) Variation of descent speed ( $V_d$ ) with height ( $h$ ) above the surface, for four different landing trajectories. The straight lines are linear regressions through the data, as represented by the equations.  $r$  denotes the regression coefficient.

From this relationship, it is clear that if the bee's horizontal flight speed is proportional to her height above the surface (as shown by the data), then the angular velocity of the image of the surface, as seen by the eye, must be constant as the bee approaches it. This angular velocity is given by the slope of the regression line. The angular velocity of the image varies from one trajectory to another, but is maintained at an approximately constant value in any given landing. An analysis of 26 landing trajectories revealed a mean image angular velocity of  $\sim 500^\circ/\text{sec}$  (Srinivasan et al., 2000).

These results reveal two important characteristics. First, bees landing on a horizontal surface tend to approach the surface at a relatively shallow descent angle. Second, landing bees tend to hold the angular velocity of the image of the ground constant as they approach it.

What is the significance of holding the angular velocity of the image of the ground constant during landing? One important consequence is that the horizontal speed of flight is then automatically reduced as the height decreases. In fact, by holding the image velocity constant, the



horizontal speed is regulated to be proportional to the height above the ground, so that when the bee finally touches down (at zero height), her horizontal speed is zero, thus ensuring a smooth landing. The attractive feature of this simple strategy is that it does not require explicit measurement or knowledge of the speed of flight, or the height above the ground. Thus, stereoscopic methods of measuring the distance of the surface (which many insects probably do not possess) are not required. What is required, however, is that the insect be constantly in motion, because the image motion resulting from the insect's own motion is crucial in controlling the landing.

The above strategy ensures that the bee's horizontal speed is zero at touchdown, but does not regulate the descent speed. How is the descent speed controlled? Plots of descent speed versus height reveal a linear relationship between these two variables as well. Two examples are shown in figure 29.6c, d. This finding implies that landing bees control their forward flight speed to hold the image velocity of the ground constant and couple the descent speed to the forward speed, so that the descent speed decreases with the forward speed and also becomes zero at touchdown. These two rules, operating together, ensure a smooth landing.

Given these two rules for landing, can we incorporate them into a model to make some predictions about the landing characteristics? We have seen from figure 29.6a, b that the speed of forward flight is controlled by holding the angular velocity of the image of the ground constant. As discussed above, this ensures that the instantaneous forward flight speed,  $V_f(t)$ , is proportional to the instantaneous height  $h(t)$  above the ground. That is,

$$V_f(t) = \omega.h(t) \quad (29.1)$$

where the constant of proportionality,  $\omega$ , is the angular velocity of the image in radians per second. We have also seen, from figure 29.6c, d, that the descent speed,  $V_d(t)$ , is proportional to the instantaneous forward speed. That is,

$$V_d(t) = -\frac{dh(t)}{dt} = B.V_f(t) \quad (29.2)$$

where  $B$  is a constant of proportionality corresponding to the ratio of  $V_d$  to  $V_f$ .

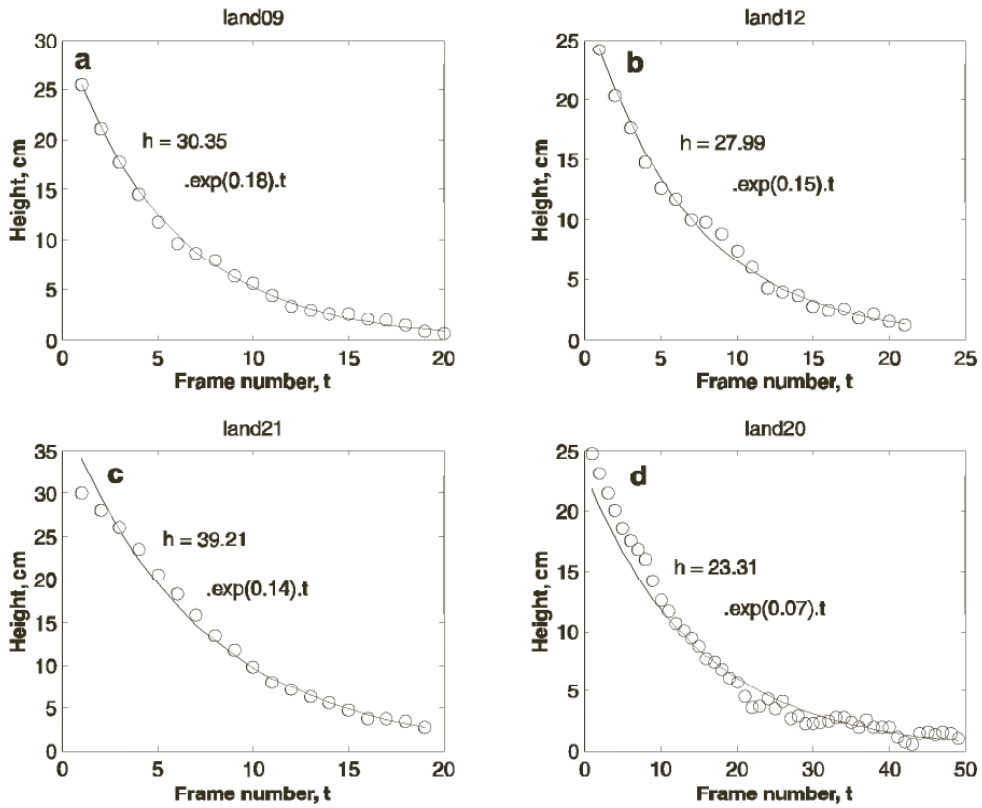
Substituting equation 29.1 into equation 29.2, we obtain

$$B.\omega.h(t) + \frac{dh(t)}{dt} = 0 \quad (29.3)$$

This differential equation can be solved for  $h(t)$  to yield

$$h(t) = h(t_0).e^{-\omega.B.(t-t_0)} \quad (29.4)$$

where  $h(t_0)$  is the height at the initial time  $t = t_0$ .



**Figure 29.7** Variation of height ( $h$ ) with time ( $t$ , denoting frame number) for four different landing trajectories. The interframe interval is 40 msec. The curves show least-squares fits of exponential functions as specified by equation 29.4.

Equation 29.4 predicts that, during landing, the height should decrease exponentially with time. The four examples shown in figure 29.7 reveal that the variation of height with time indeed approximates an exponential function very closely, thus reinforcing the model's validity.

The evolution of the forward speed,  $V_f(t)$ , during landing is obtained by inserting equation 29.4 into equation 29.1:

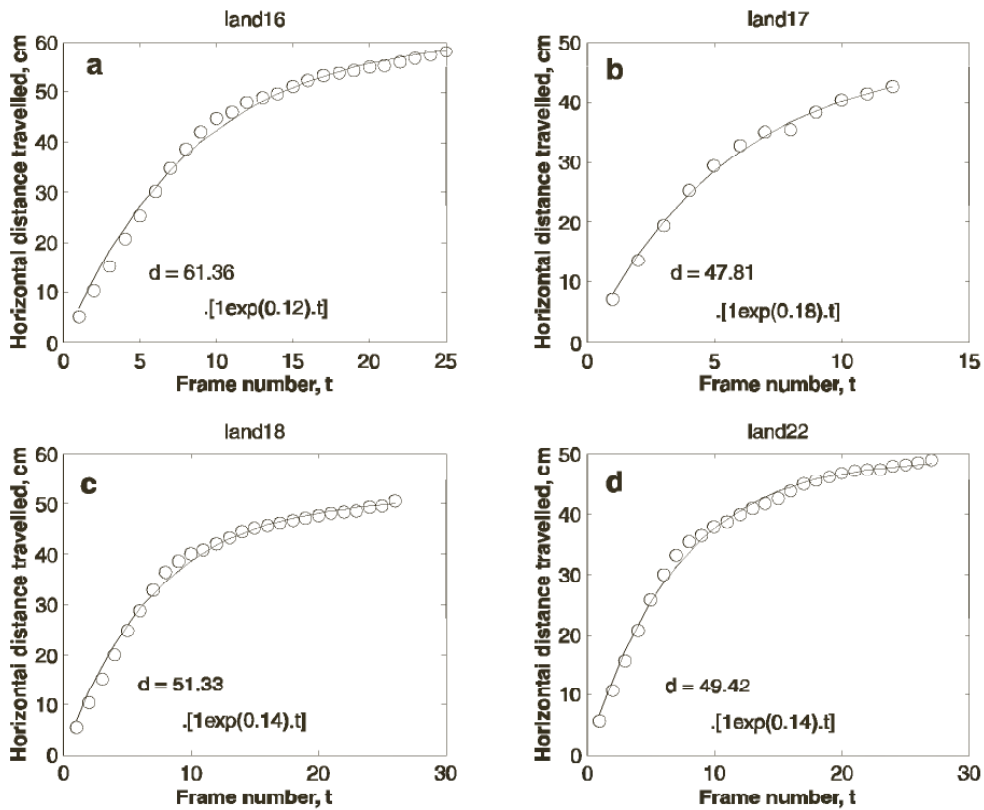
$$V_f(t) = \omega.h(t_0).e^{\omega.B(t-t_0)} \quad (29.5)$$

This implies that the bee's forward flight speed should decrease exponentially with time, a prediction that is confirmed by the data (not shown here). We can go one step further and predict the way in which the cumulative horizontal distance traveled by the bee should increase with time. This distance,  $H$ , is obtained by integrating over time the expression for the forward flight speed,  $V_f(t)$ , from equation 29.5:

$$H = \int_{t_0}^t V_f(t).dt = \int_{t_0}^t \omega.h(t_0).e^{-\omega.B.(t-t_0)}.dt \quad (29.6)$$

Integration yields

$$H = \frac{h(t_0)}{B}.[1 - e^{-\omega.B(t-t_0)}] \quad (29.7)$$



**Figure 29.8** Variation of cumulative horizontal distance travelled ( $H$ ) with time ( $t$ , denoting frame number) for four different landing trajectories. The interframe interval is 40 msec. The curves show least-squares fits of saturating exponential functions as specified by equation 29.7.

Equation 29.7 predicts that the cumulative horizontal distance travelled should increase as a saturating exponential function of time. The four examples shown in figure 29.8 reveal that the variation of horizontal distance traveled with time indeed approximates a saturating exponential function very closely. This finding further reinforces the model's validity.

In summary, landing bees regulate the forward speed of their flight by holding constant the velocity of the image of the ground in the eye. As a consequence, the forward speed is tightly linked to the height above the ground and is proportional to it. Furthermore, when a bee is in "landing mode," the rate of descent is coupled to the forward speed and is proportional to it. The coupling factor  $B$  determines the angle of descent. This landing strategy ensures that both the forward speed and the descent speed are reduced progressively as the bee approaches the ground, and are both zero at touchdown. Such a model predicts that, during landing, the height above the ground should decrease exponentially with time. It also predicts that the cumulative horizontal distance traveled should increase as a saturating exponential function. We have been able to verify both of these predictions experimentally.

These findings reveal what appears to be a surprisingly simple and effective strategy for making grazing landings on flat surfaces. A safe, smooth landing is ensured by following two simple rules: (a) adjusting the speed of forward flight to hold constant the angular velocity of the image of the surface as seen by the eye, and (b) making the speed of descent proportional to the forward speed. This produces landing trajectories in which the forward speed as well as the descent speed decrease progressively as the surface is approached, both approaching zero at touchdown.

What are the advantages, if any, of using this landing strategy? We can think of three attractive features. First, the strategy is very simple because it does not require explicit knowledge of instantaneous height or flight speed. Second, the control of forward and descent speeds is achieved by holding the image velocity *constant*. This is advantageous because the image velocity can then be maintained at a level at which the visual system is most sensitive to deviations from the “set” velocity, thereby ensuring that the control of flight is as precise as possible. An alternative strategy, for instance, might be to approach the surface at constant flight speed, decelerating only toward the end. Such a constant-speed approach, however, would cause the image velocity to increase rapidly as the surface is approached, and to reach levels at which the image velocity measurements may no longer be precise enough for adequate flight control. This situation would be avoided by the bee’s landing strategy, which holds the image velocity constant. Third, an interesting by-product of the bee’s landing strategy is that the projected time to touchdown is constant throughout the landing process (details in Srinivasan et al., 2000). In other words, if, at any time during the landing process, the bee were to stop decelerating and continue downward at constant velocity, the time to contact the ground would be the same, regardless of where this occurs in the landing trajectory. From the landing trajectories, one calculates a projected time to touchdown of about 0.22 sec. Thus, it appears that landing bees allow themselves a “safety margin” of a fifth of a second to prepare for touchdown if they were to abandon the prescribed landing strategy at any point, for whatever reason, and proceed toward the ground in the same direction without further deceleration.

The tendency exhibited by bees to adjust their forward speed according to the speed of the image in the eye is not restricted to landing situations. As we saw earlier, even cruising bees use the same technique to regulate their flight speed. Thus, the use of image speed to control flight speed appears to be a general, more or less “hard-wired” feature of insect flight. The only additional control mechanism that is brought into play during landing is that which couples the descent speed to the forward speed—to ensure a smooth touchdown.

A little reflection will reveal that the landing strategy described here can be used by an aerial vehicle to dock with any surface, regardless of its orientation: horizontal, vertical, or oblique. All that is required is that the vehicle approach the surface in a straight line, and hold the image velocity constant during the approach. This will automatically ensure that the vehicle's speed decreases as the surface is approached, ensuring smooth docking. In the special case in which the surface is approached perpendicularly, the image velocity will be zero in the "straight-ahead" direction: The flow field has a pole there. However, the strategy can still be implemented by holding constant the image speed in an annular region surrounding the pole, or in a large region centered on the pole.

The feasibility of the landing strategy described above has been tested by implementation in a computer-controlled gantry robot carrying a visual system. Vision is provided by a video camera mounted on the gantry head, which can be translated in three dimensions ( $x$ ,  $y$ , and  $z$ ). For the purpose of implementing the landing strategy, translatory motion of the camera is restricted to the forward ( $x$ ) and downward ( $-z$ ) directions. There is no rotary motion about the  $z$ -axis.

The system is placed under closed-loop control by using a computer to analyze the motion in the image sequence captured by the camera, and to control the motion of the gantry. A view of the gantry and camera is shown in figure 29.9. The floor, defined to be the landing surface, is covered with a visual texture consisting of a black-and-white Julesz pattern. The camera faces downward and views the floor. The velocity of image motion is measured by using an image interpolation algorithm (details of which are given in Srinivasan, 1994).

Landing is controlled as follows. The system is required to maintain a constant descent angle ( $\tan^{-1} B$ ) and a constant image angular velocity,  $\omega_{set}$ , as it descends. In the first time step, the gantry moves the camera head along the direction of descent at an arbitrarily chosen initial speed. The image velocity is measured during this step, using the image interpolation algorithm. Let us denote the measured image velocity by  $\omega_{meas}$ . In the next step, the speed of motion of the head is increased or decreased, depending upon whether the the measured image velocity is lower or greater than the set image velocity. Specifically, the forward speed  $V_f(i+1)$  of the camera during the next step is related to the current speed  $V_f(i)$  by

$$V_f(i+1) = V_f(i) \cdot \frac{\omega_{set}}{\omega_{meas}} \quad (29.8)$$

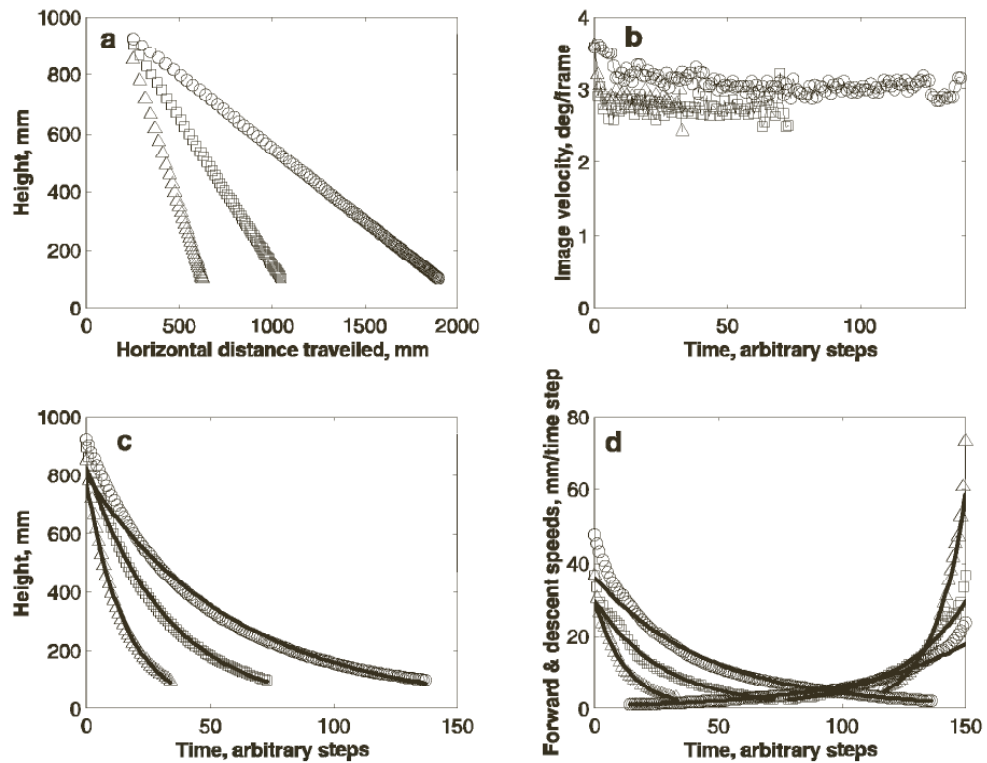
The speed of descent is also corrected by the same factor, because the forward and descent speeds are proportional to each other and linked by the desired angle of descent. This speed correction ensures that the image velocity during the next step will have the desired value  $\omega_{set}$ ,



**Figure 29.9** View of robotic gantry, showing camera head and visual texture on the floor.

provided the camera maintains its present altitude. However, because the camera continues to descend during the new step, the forward speed in the following step would have to be reduced further. Thus, both the forward and descent speeds decrease continuously as the camera descends, reaching very low values when the camera is close to the ground.

Landing trajectories generated by the gantry, using this procedure, are shown in figure 29.10a for three different descent angles:  $-26.5^\circ$ ,  $-45^\circ$ , and  $-63.5^\circ$ . The image velocities maintained during these three landings are shown in figure 29.10b. It is clear that the image velocity, though somewhat noisy, is held reasonably constant. The height of the camera decreases exponentially with time (figure 29.10c), as do the forward speed and the speed of descent. These behaviors are as expected from the proposed control strategy. Comparable results were obtained when the floor was covered with other textures, such as a newspaper,



**Figure 29.10** Landing trajectories generated by the robotic gantry. (a) Height versus distance traveled for three different descent angles:  $-26.5^\circ$  (circles),  $-45^\circ$  (squares) and  $-63.5^\circ$  (triangles). (b) Variation of image angular velocity as a function of time. The symbols in this and other panels refer to the three different descent angles, as in (a). (c) Variation of height with time. (d) Variation of forward speed and descent speed with time. For clarity the descent speed data are shown on a reversed time scale, decaying from right to left. The line curves in (c) and (d) depict least-squares fits of exponential functions to the data.

or with twigs, bark, and leaves to simulate a more natural outdoor environment (data not shown).

The results with the robotic gantry suggest that the strategy proposed is a feasible one for landing on flat surfaces. In undulating terrain, the system reduces the forward and descent speeds when the ground rises toward the camera, and increases them when the ground falls away. This is obviously a desirable feature, but has limitations in that the system cannot cope with a situation in which the ground in front rises abruptly to a level above the camera's current height.

## CONCLUSIONS

Analysis of vision in simple natural systems, such as those found in insects, can often point to novel ways of tackling tenacious problems in autonomous navigation. This is probably because insects, with their "stripped down" nervous systems, have been forced to evolve ingenious strategies to cope with visual challenges within their environ-

ment. This article has outlined three ways in which insects use motion cues to perceive their environment in three dimensions and navigate in it. The feasibility of each of these strategies has been tested by implementation in a robot. In constructing these robots, the aim is not to copy insect vision in the literal sense. Rather, the intent is to embody and test the *principles* of visual information processing that we have gleaned from the biology. The next step is to investigate whether these principles can be used to advantage in the design of visually based control systems for autonomously flying vehicles.

## ACKNOWLEDGMENTS

This research was supported partly by a grant from the Australian Defence Science and Technology Organisation, Salisbury, Grant RG 84/97 from the Human Frontiers in Science Program, and Grant N00014-99-1-0506 from the U.S. Defence Advanced Research Projects Agency and the Office of Naval Research.

## REFERENCES

- Borst, A., and Bahde, S. (1988). Visual information processing in the fly's landing system. *J. Comp. Physiol. [A]* 163: 167–173.
- Coombs, D., and Roberts, K. (1992). Bee-bot: Using peripheral optical flow to avoid obstacles. *Proc. SPIE* 1825: 714–721.
- David, C. T. (1982). Compensation for height in the control of ground speed by *Drosophila* in a new, "Barber's Pole" wind tunnel. *J. Comp. Physiol. [A]* 147: 485–493.
- Duchon, A. P., and Warren, W. H. (1994). Robot navigation from a Gibsonian viewpoint. In *IEEE International Conference on Systems, Man and Cybernetics*, San Antonio, Tex., 2–5, October. pp. 2272–2277.
- Eckert, H., and Hamdorf, K. (1980). Excitatory and inhibitory response components in the landing response of the blowfly, *Calliphora erythrocephala*. *J. Comp. Physiol. [A]* 138: 253–264.
- Gibson, J. J. (1950). *The Perception of the Visual World*. Boston: Houghton Mifflin.
- Goodman, L. J. (1960). The landing responses of insects: I. The landing response of the fly, *Lucilia sericata*, and other Calliphoridae. *J. Exp. Biol.* 37: 854–878.
- Horridge, G. A. (1977). Insects which turn and look. *Endeavour* n.s. 1: 7–17.
- Horridge, G. A. (1987). The evolution of visual processing and the construction of seeing systems. *Proc. R. Soc. Lond. B* 230: 279–292.
- Kirchner, W. H., and Srinivasan, M. V. (1989). Freely flying honeybees use image motion to estimate object distance. *Naturwissenschaften* 76: 281–282.
- Sandini, G., Gandolfo, F., Grosso, E., and Tistarelli, M. (1993). Vision during action. In Y. Aloimonos (ed.), *Active Perception*. Hillsdale, N.J.: Lawrence Erlbaum, pp. 151–190.
- Srinivasan, M. V. (1990). Generalized gradient schemes for the measurement of two-dimensional image motion. *Biol. Cybern.* 63: 421–431.



- Srinivasan, M. V. (1993). How insects infer range from visual motion. In F. A. Miles and J. Wallman (eds.), *Visual Motion and Its Role in the Stabilization of Gaze*. Amsterdam: Elsevier, pp. 139–156.
- Srinivasan, M. V. (1994). An image-interpolation technique for the computation of optic flow and egomotion. *Biol. Cybern.* 71: 401–416.
- Srinivasan, M. V. (1998). Insects as Gibsonian animals. *Ecol. Psychol.* 10: 251–270.
- Srinivasan, M. V., Chahl, J. S., Nagle, M. G., and Zhang, S. W. (1997). Embodying natural vision into machines. In M. V. Srinivasan and S. Venkatesh (eds.), *From Living Eyes to Seeing Machines*. Oxford: Oxford University Press, pp. 249–265.
- Srinivasan, M. V., and Lehrer, M. (1984). Temporal acuity of honeybee vision: Behavioural studies using moving stimuli. *J. Comp. Physiol. [A]* 155: 297–312.
- Srinivasan, M. V., Lehrer, M., Kirchner, W., and Zhang, S. W. (1991). Range perception through apparent image speed in freely flying honeybees. *Vis. Neurosci.* 6: 519–535.
- Srinivasan, M. V., and Zhang, S. W. (1997). Visual control of honeybee flight. In M. Lehrer (ed.), *Orientation and Communication in Arthropods*. Basel: Birkhäuser Verlag, pp. 67–93.
- Srinivasan, M. V., Zhang, S. W., Chahl, J. S., Barth, E., and Venkatesh, S. (2000). How honeybees make grazing landings on flat surfaces. *Biol. Cybern.* (in press).
- Srinivasan, M. V., Zhang, S. W., and Chandrashekara, K. (1993). Evidence for two distinct movement-detecting mechanisms in insect vision. *Naturwissenschaften* 80: 38–41.
- Srinivasan, M. V., Zhang, S. W., Lehrer, M., and Collett, T. S. (1996). Honeybee navigation en route to the goal: Visual flight control and odometry. *J. Exp. Biol.* 199: 237–244.
- Wagner, H. (1982). Flow-field variable trigger landing flies. *Nature (Lond)* 297: 147–148.
- Weber, K., Venkatesh, S., and Srinivasan, M. V. (1997). Insect inspired behaviours for the autonomous control of mobile robots. In M. V. Srinivasan and S. Venkatesh (eds.), *From Living Eyes to Seeing Machines*. Oxford: Oxford University Press, pp. 226–248.