


## 




$\stackrel{T}{2}$

## Introduction

This book is designed to help you teach yourself how to program with C++ under the Linux operating system. Ir just 21 days, you will learn about such fundamentals as managing I/O, loops and arrays, object-oriented programming, templates, and creating C++ applications-all in well-structured and easy-to-follow lessons. Lessons provide sample listings-complete with sample output and an analysis of the code-to illustrate the topics of the day. Syntax examples are clearly marked for handy reference.

To get you up to speed with the features, tools, and environment specific to Linux, there is a bonus week! You can learn C++ without that week, but it does build your Linux-specific skills

To help you become more proficient, each lesson ends with a set of common questions and answers, exercises, and a quiz. You can check your progress by examining the quiz and exercise answers provided in Appendix D,"Answers to Exercises and Quizzes."

## Who Should Read This Book

You do not need any previous experience in programming to learn C++ with this book. This book starts you from the beginning and teaches you both the language and the concepts involved with programming $\mathrm{C}++$. You'l find the numerous examples of syntax and detailed analysis of code an excellent guide as you begin your journey into this rewarding environment. Whether you are just beginning or already have some experience programming, you will find that this book's clear organization makes learning C++ fast and easy.

This book does not teach you how to use or install Linux, although a copy is included on the CD-ROM. There are plenty of other good books for that purpose (like the Sams Linux Unleashed series).

C++ (like the C language before it ) is a standard language, so it is the same core language under Linux as on other platforms. In other words, the $\mathrm{C}++$ you learn from this book applies to many different systems and compilers.

Two different compilers are available under Linux-gcc (GNU C/C++ Compiler) and egcs (Experimental GNU Compiler System). The CD-ROM contains egcs, whereas older distributions of Linux (and the current MS-DOS version) use gcc. From the command line, it is hard to tell the difference. However, egcs has the more advanced and current $\mathrm{C}++$ features.

Versions of the GNU compilers are available for most platforms and operating systems. So the skills you learn in this book are useful for more than just Linux.

You must be running Linux to use the included compiler.

## Conventions

Note:

These boxes highlight information that can make your C++ programming more efficient and effective.

## Frequently Asked Questions

FAQ: What do FAQs do
Answer: These Frequently Asked Questions provide greater insight into the use of the language and clarify potential areas of confusion.

## Caution:

The cautions focus your attention on problems or side effects that can occur in specific situations.


This book uses various typefaces to help you distinguish C++ code from regular English. Actual C++ code is typeset in a special monospace font. Placeholders-words or characters temporarily used to represent the real words or characters you would type in code-are typeset in italic monospace. New or important terms are typeset in italic.

In the listings in this book, each real code line is numbered. If you see an unnumbered line in a listing, you'll know that the unnumbered line is really a continuation of the preceding numbered code line (some code lines are too long for the width of the book). You will also see a line continuation character like this $f$. In this case, you should type the two lines as one; do not divide them.

The listings are also included on the CD-ROM with filenames that begin with Ist, followed by the 2 -digit lesson number, a dash, and then the 2-digit listing number within that lesson. For example, the first example in Day 1 is 1st01-01.cxx

## About the Authors

## Lead Authors

Jesse Liberty is the author of C++ From Scratch, as well as a dozen other books on C++ and Web application: development and object-oriented programming. Jesse is the president of Liberty Associates, Inc., where he provides custom programming, training, mentoring, and consulting. He is a former Vice President of electroni delivery for Citibank and a Distinguished Software Engineer at AT\&T. Jesse also serves as the series editor of Que's Programming From Scratch books. He supports his books at http://www.LibertyAssociates.com
David B. Horvath, CCP, is a Senior Consultant in the Philadelphia Pennsylvania area. He has been a consulta for over 14 years and is also a part-time Adjunct Professor at local colleges teaching topics that include C/C++ Programming, UNX, and Database Techniques. He has an M.S. degree in Organizational Dynamics fron
University of Pennsylvania (and is taking more classes as this book is being published). He has provided seminars and workshops to professional societies and corporations on an international basis.
David is the author of UNIX for the Mainframer (Prentice-Hall/PTR), contributing author to UNIX Unleashed Second Edition (with cover credit), Red Hat Linux Second Edition, Using UNIX Second Edition (Que), UNIX Unleashed Third Edition, Learn Shell Programming in 24 H
Fourth Edition, and has written numerous magazine articles.

When not at the keyboard, he can be found working in the garden, soaking in the hot tub, or engaging in vario ctivities that upset people. He has been married for over 12 years to his wife and has several dogs and cats (the number of which seems to keep going up).
David can be reached at cpplinux@cobs.com for questions related to this book, and his Web page i http://www.cobs.com/. No spam please!

## Contributing Authors

Paul Cevoli is a software engineer living in Boston, MA, where he has been a software engineer for the past 1 Paul Cevoli is a software engineer living in Boston, MA, where he has been a software engineer for the
years. He specializes in system software, device drivers, and operating system internals. Paul is currently employed in the system software group of an internet appliance company in the metro Boston area.

Jonathan Parry-McCulloch is a technical consultant for a major financial institution in the city of London, where he provides cryptographic expertise for automated trading systems for trading financial futures. He also writes cross-platform C++ GU applications, largely for fun, and is a confirmed Linux fre
Class degree in Electronic Engineering and is pleased that it bears no relevance to his job

You can reach Jonathan at jim@antipope.org, if you have questions about his portions of this book, or anything You can reach Jonathan at j @ antipope.org, if you have questions about his portions of this ba
else that crosses your mind. One word of warning: he really hates spammers. So don't do it.

Hal Moroff has been designing embedded systems for 20 years. From the East Coast originally, he worked fo veral years in Europe and Japan, and is now programming full-time in Califar Villey He write alum in Linux Magazine, and read the source whenever he can.

## Acknowledgments

I would like to thank and acknowledge the many people who made this and the previous editions of this book possible. First among them are Stacey, Robin, and Rachel Liberty. I must also thank the editors at SAMS, insluding Carol Ackerman, Tracy Dunkelberger, Holly Allender, Sean Dixon, Chris Denny, Brad Jones, Roby Thomas, Christina Smith, and Margaret Berson

Many readers contributed suggestions and improvements to this book and I thank you all. Finally, and again, thank Mrs. Kalish, who taught my sixth-grade class how to do binary arithmetic in 1965 , when neither she nor we knew why.

I want to acknowledge all the good folks at Sams (Macmillan Computer Publishing) who helped this book alo its path. Carol Ackerman (Acquisitions Editor) came to me with the idea of writing this book, Christina Smith was the Project Editor riding herd over the editing process, Margaret Berson performed the copy edit (checkins grammar, syntax, clarity, and so on), Robyn Thomas was the development editor (content and structure), and Javad Abdollahi and Sean Coughlin performed most of the technical edit along with Rory Bray. Without them
this book would not be nearly as good as it is now. I absolutely must thank my co-author, Jesse Liberty. When Carol and I were first discussing this project, I ask
for a sample of a Sams Teach Yourself in 21 Days book. She sent me a copy of Sams Teach Yourself C in 21 Days. After reviewing the book, I commented to Carol that this book would be much easier if there already existed a C++ book similar to the C book. It turned out that there already was-Jesse's. So creating this book was much easier because of the base $\mathrm{C}++$ information already written. After all, there are just so many ways to explain a while loop or polymorphisn The contributing authors, Hal, Paul, and Jon have been a huge help covering some of the very specific Linux
topics that go well beyond the C++ language itself. Deserving special attention is Jon Parry-McCulloch, who took on the GUI chapter on short notice and had to contend with the time zone differences (he lives in Englan
tor Jon is a fellow contributor to the Cult of Father Darwin Internet email discussion group (which discusses the manifestation of Darwinian theories in the daily actions and mistakes of various people). ''m not going to tell you how to find the group because we try to keep the Signal-to-Noise ratio high (if you're interested, you'll in London in 1999. There is a big difference between talking via email and talking in person. As with any large project, someone has to make a sacrifice. In my case the one who suffers most from the time
commitments a book like this takes is my wife Mary. The amazing thing is that she doesn't even complain (or least not too much

Last, and certainly not least, are the folks at DCANet (www.dca.net) who provide virtual Web site hosting for my Web page, handle my domain, POP my email, and have a large modem pool that are local calls for me. These people are very responsive; I never get busy signals, and Andrew white is a System Aumnis
without peer! They are in the Greater Philadelphia Metropolitan Area, and I've directed a number of organizations to their services. No, I'm not getting paid for this plug (but they do take good care of me),


Dedication

## This book is dedicated to the living memory of David Levine.

## - Jesse

I would like to dedicate this book to alt the good teachers I have had through the years. Those that taught that questions are not stupid, even if they may seem to be simple; that there is more to learn than exactly at is written in the book; that the best way to learn anything that requires thinking is to try it
some people call playing is often learning. Those teachers already know who they are!

I also want to dedicate this book to my students. They've asked some amazing questions. I thought I knew computing and programming before I started tecching: in having to answer their questions, they've forced me to learn even more about my topics!

## Tell Us What You Think

As the reader of this book, you are our most important critic and commentator. We value your opinion and w. to know what we're doing right, what we could do better, what areas you'd like to see us publish in, and any other words of wisdom you're willing to pass our wa

As an Associate Publisher for Sams Publishing, I welcome your comments. You can fax, email, or write $m$ c directly to let n
books stronger

Part 1

## At a Glance

As you prepare for your first week of learning how to program in $\mathrm{C}++$, you will need a few things: a compile editor, and this book. If you don't have a $\mathrm{C}++$ compiler and an editor, you can still use this book, but you won't get as much out of it as you would if you did the exercise

The best way to learn to program is by writing programs! At the end of each day you will find a workshop The best way to learn to program is by writing programs! At the end of each day you will find a workshop
containing a quiz and some exercises. Be sure to take the time to answer all the questions, and to evaluate you work as objectively as you can. The later lessons build on what you've learned in the earlier lessons, so be sure you fully understand the material before moving on.

## A Note to C Programmers

The material in the first five days will be familiar to you. Be sure to skim the material and to do the exercises to make sure you are fully up to speed before going on to Day 6, "Basic Classes." If you know C but have not

## Where You Are Going

The first week covers the material you need to get started with programming in general, and with $\mathrm{C}++$ in The first week covers the material you need to get started with programming in general, and with C++ in
particular. On Day 1, "Getting Started," and Day 2, "The Parts of a C++ Program," you will be introduced to th basic concepts of programming and program flow. On Day 3, "Variables and Constants," you will learn about
variables and constants and how to use data in your programs. On Day 4, "Expressions and Statements," you will learn how programs branch based on the data provided and the conditions encountered when the program running. On Day 5 , "Functions," you will learn what functions are and how to use them, and on Day 6 , "Basic Classes," you will learn about classes and objects. Day 7, "More Program Flow," teaches more about program flow, and by the end of the first week you will be writing real object-oriented programs

## Chapter

## Getting Started

Welcome to Sams Teach Yourself $C_{+}+$for Linux in 21 Days! Today you will get started on your way to becoming a proficient C++ programmer. You will learn:

- What the GNU compilers are and how GNU and Linux are related
- Why $\mathrm{C}++$ is the emerging standard in software development

The steps to develop a C++ program
The basics on how to use the GNU compilers
How to enter, compile, and link your first working $\mathrm{C}++$ program

## What Is GNU?

The acronym GNU stands for "GNU is Not UNIX." It is the name of a series useful software package commonly found in UNIX environments that are being distributed by the GNU project at MIT. The packages are generally free and available at various locations on the Internet (you are charged if you want a copy on a physical medium like foppy or tape). The development of the packages is a cooperative process with the work
being done by many volunteers. This effort is largely led by Richard M. Stallman (one of the developers of the being done by mat
EMACS editor).

Linux is an operating system that is very similar to UNIX (without infringing on the trademarks and copyright Linux is an operating system that is very similiar to UNX (without infringing on the trademarks and c
of commercial versions of UNIX). But it is primarily just a kernel (the core of the operating system).

Most of the commands that you use under Linux are actually part of the GNU project sponsored by the Free Software Foundation. Twenty-eight percent of a typical Linux distribution is GNU, and only $3 \%$ is truly Linux

```
Note:
Yes, the GNU acronym is self-referencing, but it was created that way on purpose. This type of naming has a 
history in products reated to, but not covered by the trademark,
```

So, in reality, you are using GNU utilities when you think you are issuing Linux commands. The folks who
work on Linux saw no reason to duplicate the work done by the GNU folks. That is part of the whole philosophy of UNIX: reuse, do not recreate

The GNU project also includes C and $\mathrm{C}++$ compilers (along with many other tools, languages, and utilities). T
 operating system such as DOS that does not allow the plus sign in filenames)

The nice thing about the GNU tools is that they exist for many platforms, not just Linux. So you can get a GN
This means that you can learn one compiler and use it many places. You can also have a personal C++ project under Linux that you work on in your own time and work on it when you are supposed to being doing "real
work" at your employer-even if you only have access to a Microsoft Windows/DOS-based machine. Of work" at your employer-even if you only ha
course, you would never do anything like that

## Getting GNU Tools

The easiest way to get the GNU compiler and other tools is to install the copy of Linux contained on the CDROM included with this book. For installation details, you will have to review the instructions on the CD-RO as this book focuses on $\mathrm{C}+$, not Linux

SAMS publishes some good Linux books also (this is more than a boldfaced plug for SAMS as I hav contributed chapters to several of those books

As mentioned before, you can also get the compiler for other platforms and operating systems.
If you already have Linux installed, there is a good chance that you already have the GNU compilers (but you may have an older version).
thest to go (besides the CD-ROM) is http://www.gnu.org. Here you can download the GNU tools and compilers for various platforms for free. You can also get information on purchasing the GNU software on CD ROM. Either way, you should consider sending them a donation. After all, it costs money to maintain Web pages and print documentation. Personally, I ordered a full set of the binaries and tools for all platforms so tha
would have them for comparison purposes. I do not get any money when you purchase from or would have them for comparison purposes. (I do not get any money when you purchase from or donate to the
Free Software Foundation; I just believe in what they are doing.)

Developing code takes time. And it takes time to get CD-ROM disks cut, duplicated, and delivered. As a result there are different versions for different platforms and different versions between distribution media. You can download the latest version or purchase a slightly older version hook was written). The current version on the CD-ROM available from the Free Software Foundation is 2.7.2. In general, the newer compiler is bette

Examples in this book will compile under both versions unless otherwise noted.
The newer versions of the compiler are part of the EGCS (Experimental GNU Compiler System)-pronounced "eggs". Much of the work of creating the new compiler was done by the folks at Cygnus
(http//www.cygnus.com). Cygnus has been named as the official maintainer of the GNU compilers (under the oversight of a steering committee)

A Brief History of C+
Computer languages have undergone dramatic evolution since the first electronic computers were built to assis in artillery trajectory calculations during World War II. Early on, programmers worked with the most primitiv computer instructions: machine language. These instructions were represented by long strings of ones and ees. Soon, assemblers were invented to map machine instructions to human-readable and -manageable mnemonics, such as ADD and MOV

In time, higher-level languages evolved, such as BASIC and COBOL. These languages let people work with something approximating words and sentences, such as Let $\mathrm{I}=100$. These instructions were translated back machine language by interpreters and compilers. An interpreter translates a program as it reads it, turning the program instructions, or code, directly into actions. A compiler translates the code into an intermediary forn This step is called compiling, and it produces an object file. The compiler then invokes a linker, which turns th object file into an executable program.

Another advantage of many compiled languages such as $\mathrm{C}++$ is that you can distribute the executable program to people who do not have the compiler. With an interpretive language, you must have the interpreter to run the program.

Some languages, such as Visual Basic, call the interpreter the runtime library. Java calls its runtime interpreter Virtual Machine (VM), but in this case the VM is provided by the browser (such as Internet Explorer or Netscape)

For many years, the principal goal of computer programmers was to write short pieces of code that would execute quickly. The program needed to be small, because memory was expensive, and it also needed to be fas execute quickly. The program needed to be small, because memory was expensive, and it also needed to be fas
because processing power was expensive. As computers have become smaller, cheaper, and faster, and as the because processing power was expensive. As computers have become smaller, cheaper, and faster, and as the
cost of memory has fallen, these priorities have changed. Today the cost of a programmer's time far outweighs cost of memory has fallen, these priorities have changed. Today the cost of a programmer's time far outweig
the cost of most of the computers in use by businesses. Well-written, easy-to-maintain code is at a premium. "Easy to maintain" means that as business requirements change, the program can be extended and enhanced without great expense.

## Programs

The word program is used in two ways: to describe individual instructions (or source code) created by the programmer, and to describe an entire piece of executable software. This distinction can cause enormous confusion, so we will try to distinguish between the source code on one hand, and the executable on the other.

A program can be defined as either a set of written instructions created by a programmer or an executable piece of software.

Source code can be turned into an executable program in two ways: Interpreters translate the source code into computer instructions, and the computer acts on those instructions immediately. Alternatively, compilers translate source code into a program, which you can run at a later time. Although interpreters are easier to work with, most serious programming is done with compilers because compiled code runs much faster. $\mathrm{C}++$ is compiled language.

## Solving Problems

The problems programmers are asked to solve have been changing. Twenty years ago, programs were created t manage large amounts of raw data. The people writing the code and the people using the program were all computer professionals. Today, computers are in use by far more people, and most know very little about how computers and programs work. Computers are tools used by people who are more interested in solving their business problems than in struggling with the computer.

Ironically, to become easier to use for this new audience, programs have become far more sophisticated. Gone are the days when users typed cryptic commands at esoteric prompts, only to see a stream of raw data. Today's programs use sophisticated "user-friendly interfaces" involving multiple windows, menus, dialog boxes, and th myriad of metaphors with which we have become familiar. The programs written to support this new approach are far more complex than those written only ten years ago.

With the development of the Web, computers have entered a new era of market penetration; more people are using computers than ever before and their expectations are very high. In the last few years, programs have become larger and more complex, and the need for object-oriented programming techniques to manage this complexity has become manifest.

As programming requirements have changed, both languages and the techniques used for writing programs hav As programming requirements have changed, both languages and the techniques used for writing programs hav
evolved. Although the complete history is fascinating, this book focuses on the transformation from procedural programming to object-oriented programming

## Procedural, Structured, and Object-Oriented Programming

Until recently, programs were thought of as a series of procedures that acted upon data. A procedure, or function, is a set of specific instructions executed one after the other. The data was quite separate from the procedures, and the trick in programming was to keep track of which functions called which other functions, and what data was changed. To make sense of this potentially confusing situation, structured programming was created.

The principal idea behind structured programming is as simple as the idea of divide and conquer. A computer program can be thought of as consisting of a set of tasks. Any task that is too complex to be described simply would be broken down into a set of smaller component tasks, until the tasks were sufficiently small and selfcontained that they were easily understood.

As an example, computing the average salary of every employee of a company is a rather complex task. You can, however, break it down into the following subtasks:

1. Find out what each person earns.
2. Count how many people you have
3. Total all the salaries
4. Divide the total by the number of people you have

Totaling the salaries can be broken down into the following steps:

1. Get each employee's record
2. Access the salary
3. Add the salary to the running total
4. Get the next employee's record.

In turn, obtaining each employee's record can be broken down into the following

1. Open the file of employees
2. Go to the correct record.
3. Read the data from disk.

Structured programming remains an enormously successful approach for dealing with complex problems. By the late 1980s, however, some of the deficiencies of structured programming had become all too clear.

First, a natural desire is to think of data (employee records, for example) and what you can do with data (sort, edit, and so on) as a single idea. Procedural programming worked against this, separating data structures from functions that manipulated that data.

Second, programmers found themselves constantly reinventing new solutions to old problems. This is often called "reinventing the wheel," which is the opposite of reusability. The idea behind reusability is to build components that have known properties, and then to be able to plug them into your program as you need them. This is modeled after the hardware world-when an engineer needs a new transistor, she usually does not inver a new one; she goes to the big bin of transistors and finds one that works the way she needs it to, or perhaps modifies it. No similar option existed for a software engineer manipulate that data.

The essence of object-oriented programming is to treat data and the procedures that act upon the data as a single object-a self-contained entity with an identity and certain characteristics of its own.

## $\mathrm{C}_{++}$and Object-Oriented Programming

C++ fully supports object-oriented programming, including the three pillars of object-oriented development: encapsulation, inheritance, and polymorphism.

## Encapsulation

When an engineer needs to add a resistor to the device she is creating, she typically does not build a new one from scratch. She walks over to a bin of resistors, examines the colored bands that indicate the properties, and picks the one she needs. The resistor is a "black box" as far as the engineer is concerned-she does not care much about how it does its work as long as it conforms to her specifications; she does not need to look inside the box to use it in her design.
he property of being a self-contained unit is called encapsulation. With encapsulation, we can accomplish dat hiding. Data hiding is the highly valued characteristic that an object can be used without the user knowing or caring how it works internally. Just as you can use a refrigerator without knowing how the compressor works, you can use a well-designed object without knowing about its internal data members.

Similarly, when the engineer uses the resistor, she does not need to know anything about the internal state of th resistor All the properties of the resistor are encapsulated in the resistor object; they are not spread out through the circuitry. It is not necessary to understand how the resistor works to use it effectively. Its data is hidden inside the resistor's casing.

C++ supports the properties of encapsulation through the creation of user-defined types, called classes. You will see how to create classes on Day 6, "Basic Classes." After you create a well-defined class, it acts as a fully encapsulated entity-it is used as a whole unit. The actual inner workings of the class should be hidden. Users of a well-defined class do not need to know how the class works; they just need to know how to use it.

## Inheritance and Reuse

When the engineers at Acme Motors want to build a new car, they have two choices: They can start from scratch, or they can modify an existing model. Perhaps their Star model is nearly perfect, but they want to add a turbocharger and a six-speed transmission. The chief engineer would prefer not to start from the ground up, but would rather say, "Let's build another Star, but let's add these additional capabilities. We'll call the new model a Quasar." A Quasar is a kind of Star, but a specialized one with new features.

C++ supports inheritance. A new type, which is an extension of an existing type, can be declared. This new subclass is said to derive from the existing type and is sometimes called a derived type. The Quasar is derived from the Star and thus inherits all its qualities but can add to them as needed. Inheritance and its application in C++ are discussed on Day 11, "Inheritance," and Day 15, "Advanced Inheritance."

## Polymorphism

The new Quasar might respond differently than a Star does when you press down on the accelerator. The Quasar might engage fuel injection and a turbocharger, whereas the Star would simply let gasoline into its Quasar might engage fuel injection and a turbocharger, whereas the Star would simply let gasoline into its
carburetor. A user, however, does not have to know about these differences. She can just "floor it," and the righ carburetor. A user, however, does not have to know about
thing will happen, depending on which car she is driving.

C++ supports the idea that different objects do "the right thing" through what is called function polymorphism and class polymorphism. Poly means many, and morph means form. Polymorphism refers to the same name taking many forms, and it is discussed on Day 10, "Advanced Functions," and Day 13, "Polymorphism."

## How C++ Evolved

As object-oriented analysis, design, and programming began to catch on, Bjarne Stroustrup took the most popular language for commercial software development, C , and extended it to provide the features needed to facilitate object-oriented programming.

Although it is true that $\mathrm{C}++$ is a superset of C and that virtually any legal C program is a legal C++ program, the leap from C to $\mathrm{C}++$ is very significant. $\mathrm{C}++$ benefited from its relationship to C for many years because C programmers could ease into their use of $\mathrm{C}++$. To really get the full benefit of $\mathrm{C}++$, however, many programmers found they had to unlearn much of what they knew and learn a new way of conceptualizing and programmers found they had to
solving programming problems.

## Should I Learn C First?

The question inevitably arises: "Because C ++ is a superset of C, should you learn C first?" Stroustrup and mos other $\mathrm{C}++$ programmers agree that not only is it unnecessary to learn C first, it may be advantageous not to do so.

This book does not assume you have any prior programming experience. If you are a C programmer, howeve the first five chapters of this book will largely be review. But it will tell you how to use the GNU compilers. Starting in Day 6 we begin the real work of object-oriented software development

## C++ and Java

$\mathrm{C}++$ is now the overwhelmingly predominant language for the development of commercial software. In recent years, Java has challenged that dominance, but the pendulum swings back, and many of the programmers who left C++ for Java have recently begun to return. In any case, the two languages are so similar that to learn one to learn 90 percent of the other

## The ANSI Standard

The Accredited Standards Committee, operating under the procedures of the American National Standards Institute (ANSI), has created an international standard for $\mathrm{C}++$

The C++ Standard is now also referred to as ISO (International Standards Organization) Standard, the NCITS (National Committee for Information Technology Standards) Standard, the X3 (the old name for NCITS Standard, and the ANSI/ISO Standard. This book will continue to refer to ANSI standard code because that is the more commonly used term.

## Note:

ANSI is usually pronounced "antsy" with a silent "t."

The ANSI standard is an attempt to ensure that $\mathrm{C}++$ is portable-ensuring, for example, that ANSI-standard compliant code you write for the GNU compilers will compile without errors, using a compiler from any othe vendor (like Microsoft). Further, because the code in this book is ANSI compliant, it should compile without errors on a Mac, a Windows box, or an Alpha

## Preparing to Program

C++, perhaps more than other languages, demands that the programmer design the program before writing it. Trivial problems, such as the ones discussed in the first few chapters of this book, do not require much design. Complex problems, however, such as the ones professional programmers are challenged with every day, do require design, and the more thorough the design, the more likely it is that the program will solve the problems it is designed to solve, on time and within budget. A good design also makes for a program that is relatively bug free and easy to maintain. It has been estimated that 90 percent of the cost of software is the combined cost of debugging and maintenance. To the extent that good design can reduce those costs, it can have a significant impact on the bottom-line cost of the project.

The first question you need to ask when preparing to design any program is, "What is the problem I'm trying to solve?" Every program should have a clear, well-articulated goal, and you will find that even the simplest programs in this book do so.

The second question every good programmer asks is, "Can this be accomplished without resorting to writing custom software?" Reusing an old program, using pen and paper (the old, original, manual, stand-by way of doing work), or buying software off the shelf is often a better solution to a problem than writing something new The programmer who can offer these alternatives will never suffer from lack of work; finding less expensive solutions to today's problems will always generate new opportunities later.

Assuming you understand the problem and that it requires writing a new program, you are ready to begin your design.

The process of fully understanding the problem (analysis) and creating a solution (design) is the necessary foundation for writing a world-class commercial application. Although these steps logically come before coding-that is, you must understand the problem and design the solution before you implement it-you are, in fact, better off learning the fundamental syntax and semantics of $\mathrm{C}++$ before learning formal analysis and desigr techniques.

## The GNU/Linux Development Environment

This book makes the assumption that you are using the GNU compilers in a text environment or similar combination. That is, your compiler has a mode in which you can write directly to the screen, without worrying about a graphical environment, such as the ones in Windows or on the Macintosh. This is known as a console mode and is standard for the GNU compilers. If you are working under a different environment, you will have to look for an option such as console or easy window or check your compiler's documentation. If you are using one of the graphical environments under Linux, your best bet is to open a terminal emulation window (like KDT) so that you can work in pure text mode

With GNU, you can use EMACS, vi, or any other text editor that you prefer. If you are using a different compiler, you will have different choices-your compiler may have its own built-in text editor, or you can use a commercial text editor to produce text files. The important thing is that whatever you write your program in, it must save simple, plain-text files, with no word processing commands embedded in the text. Examples of safe editors include Windows Notepad, the DOS edit command, Brief, Epsilon, EMACS, and vi. Many commercial word processors, such as WordPerfect, Word, and dozens of others, embed special characters but also offer a method for saving simple text files-so be sure how you are saving your file.

The files you create with your editor are called source files, and for $\mathrm{C}++$ they typically are named with the extension .cpp, .cp, or .c. In this book, we named all the source code files with the .cxx extension because the GNU compilers accept this as C++ source code on different platforms. If you are using something different, check your compiler for what it needs.

Note:
The GNU compilers treat the file extension as being significant. You should use .cxx or .c++.
Many other C++ compilers do not treat the extension as significant. If you don't specify otherwise, many will use .cpp by default. Be careful, however, some compilers treat .c files as C code and .cpp files as C++ code. Again, please check your documentation.

Do Don't

DO use a simple text editor (like vi, EMACS, or even DON'T use a word processor that saves special the DOS edit command) to create your source code, or formatting characters. If you do use a word processor, use the built-in editor that comes with your compiler. save the file as ASCII text.
DO save your files with the .cxx or .c++ extension. DO check the man pages for the GNU compiler for specifics about how it and the linker work to ensure that you know how to compile and link your programs.

## Compiling the Source Code

Although the source code in your file is somewhat cryptic, and anyone who does not know C++ will struggle to understand what it is for, it is still in what we call human-readable form. Your source code file is not a program, and it cannot be executed, or run, as a program can.

To turn your source code into a program, you use a compiler. The simplest way to invoke the g++ compiler looks like this:
g++ file.c++ -o file
Or if you are on a platform that does not support the plus sign $(+)$ in a filename, like MS-DOS, you would use the following:
gxx file.cxx -o file.ex
In both cases, you replace file in the command with the name you selected. It will produce an executable with the name file or file.exe. If you leave off the -o option, it will produce a file called a.out (Linux) or a.exe (MSDOS) respectively.

If you are using a different compiler, you will have to check its documentation to determine how to invoke it and how to tell it where to find your source code-this varies from compiler to compiler.

## The Development Cycle


 how trivial. can and will have errors, or buss. Some bugs will cause the compile to fail, some will cause the lii
to fail, and some will show up only when you run the program. Whatever type of bug you find, you musst fix it and that involves editing your source code, recompiling and
relinking and then rerunning the program. This cycle is sepresented in Figure 1.1, which diagrams the steps in relinking, and then rex rel
the develoment cyl


## Fin

## hello.cxx-Your First C++ Progran


#### Abstract

Traditional programming books begin by teaching you how to write the wo variation on that statement. This time-honored tradition is carried on here.


 Type the first program directly into your editor, exactly as shown in Listing 1.1. After you are certain it is worry too much about how it works, this example is really yust to get you comfortable with he developmecycle. Every aspect of this program will be covered over he next couple of dayt worry too much about how it works, this exampli is really just to get you comfo
cycle. Every aspect of this program will be covered over the next couple of days

The first listing of the book (Lisising 1.1 ) is ste file suol-0.o.ex,
Not all listing files will compile. When this is the case, it it son onoed in the e ex.

Caution:
The following isising. ike the rest of the isisings in this book. conainan sine numbers on the ceff. These number Listing 1.1, you should enerer:

INPUT Listing 1.1 hellocx, The Hello World Program

```
include <iostream.
```

int main()
cout << "Hello World! $\backslash n "$ ";
return 0;

This should work on all compilers and has few disadvanatages. If you prefer to use the new standard librarie
however, simply change your code to -
using namespace std;

## Using the $\mathrm{g}_{++}$Compiler

All the programs in this book were tested with the GNU $\mathrm{g}++$ compiler version 2.9 .5 ; many of them were also
tested with version 2.7.2. In theory, because this is ANSI-compliant code, all the programs in this book should tested with version 2.7.2. In theory, because this is ANSI-con
run well on any ANSI-compliant compiler from any vendor.
In theory, theory and practice are the same. In practice, they never are
To get you started, this section briefly introduces how to edit, compile, link, and run a program using the GNU compiler. If you are using a different compiler, the details of each step may be somewhat different. Even if yo
are using the GNU version 2.7 .2 or 2.9 .5 compiler, please check your documentation for details on how to proced from here.

## Building the Hello World Project

To create and test the Hello World program, follow these steps:

1. Pick an editor and create a file.
2. Enter the code as shown in Listing 1.1 (or you can use the editor to open the lst01-01.cxx file),
3. Save the file and exit your editor.
4. Save the file and exit your edito
5. Enter the compile command.
6. Enter the executable name to run the program.

## Compile Errors

Compile-time errors can occur for any number of reasons. Usually they are a result of a typo or other nadvertent minor error. Good compilers (like GNU) not only teIf you what you did wrong, they point you to
exact place in your code where you made the mistake. The great ones will even suggest a remedy! You can see this by intentionally putting an error into your program. If the Hello World program ran sn You can see this by intentionally putting an error into your program. If the Hello World program
edit it now and remove the closing brace on line 7 . Your program will now look like Listing 1.2 .
$\qquad$
include <iostream h
int main()
cout << "Hello World! ${ }^{\text {n" }}$ "
return 0;
return 0;

Recompile your program and GNU will show you the following error
.$/ 1$ st01-02.cxx: In function 'int main()',
.$/ 1$ st01-02.cxx:7: parse error at end of input
Other compilers will give you an error message like:
Hello.cpp, line 7: Compound statement missing terminating; in
f function main().
or an error, which looks like this
F: \Mcp\Tycpp21d\Testing\List0101.cpp (8) : fatal error C1004: unexpected end of file found Error executing cl.exe.

This error tells you the file and line number of the problem and what the problem is (although I admit it is somewhat cryptic). Note that the error message points you to line 7 . The compiler was not sure if you intende to put the closing brace before or after the cout statement on line 6 . Sometimes the errors just get you to the
general vicinity of the problem. If a compiler could perfectly identify every problem, it would fix the code its

## Summary

After completing this day, you should have a good understanding of how $\mathrm{C}++$ evolved and what problems it After completing this day, you should have a good understancing ond how
designed to solve. You should feel confident that learning $\mathrm{C}++$ is the right choice for anyone interested in programming in the next decade. C++ provides the tools of object-oriented programming and the performance of a systems-level language, which makes $\mathrm{C}++$ the development language of choice.
Today you learned how to enter, compile, link, and run your first C++ program, and what the normal development cycle is. You also learned a little of what object-oriented programming is all about. You will
return to these topics during the next three weeks. During the bonus week, you will learn advanced topics of working with the GNU tool set and Linux programming.

## Q\&A

Q What is the difference between a text editor and a word processor?
A A text editor produces files with plain text in them. No formatting commands or other special symb
are required by a particular word processor. Text files do not have automatic word wrapping, bold or are required by a particular wo
italic formatting, and so forth.
Q If my compiler has a built-in editor, must I use it?
A The GNU compilers do not come with built-in text editors. Linux comes with vi and EMACS. Almo all other compilers will compile code produced by any text editor. The advantages of using the built-in
text editor, however, might include the capability to quickly move back and forth between the edit and compile steps of the development cycle. Sophisticated compilers include a fully integrated developme environment, enabling the programmer to access help files, edit and compile the code in place, and to
resolve compile and link errors without ever leaving the environment. Q Can I ignore warning messages from my compiler?
A Many books hedge on this one, but ''ll stake myself to this position: No! Get into the habit, from day one, of treating warning messages as errors. $\mathrm{C}++$ uses the compile to warn you when you are doing
something you may not intend. Heed those warnings and do what is required to make them go away.

Warnings mean that the compiler is able to create an executable from your source code but doesn't Warnings mean that the compiler is able to creat
believe you really want to do what you coded!

## Q What is compile time?

A Compile time is the time when you run your compiler, in contrast to link time (when you run the linker) or runtime (when running the program). This is just programmer shorthand to identify the three
times when errors usually surface. Because the GNU compiler invokes the linker if the compilation is successful, the link time is easy to miss (but believe me, it is there).

## Workshop

The Workshop provides quiz questions to help you solidify your understanding of the material covered and exercises to provide you with experience in using what you learned. Try to answer the quiz and exercise questions before checking the answers in Appendix $D$, $A$.
understand the answers before continuing to the next day.
Quiz

1. What is the difference between an interpreter and a compiler?
2. How do you compile the source code with your compiler?
3. What are the steps in the normal development cycle?

## Exercises

1. Look at the following program and try to guess what it does without running it

1: \#include <iostream.h>
3: int main()
5: $\quad$ int $x=5 ;$


9: re
10:子
2. Type the prog
you guessed?
1: include <iostream.h>
1: include <i
2: int main()
4: cout << "Hello World $\backslash n$ "
5: return 0;
4. Fix the error in the program in Exercise 3 and recompile, link, and run it. What does it do.

## Chapter 2

## The Parts of a $\mathrm{C}_{++}$Program

C++ programs consist of objects, functions, variables, and other component parts. Most of this book is devotec
to explaining these parts in depth, but to get a sense of how a program fits together, you must see a complete to explaining these parts in depth, but to
working program. Today you will learn:

- The parts of a C+t program
- What a function is and what it does


## A Simple Program

Even the simple program HELLO.cpp from Day 1, "Getting Started," had many interesting parts. This section
reviews this program in more detail. Listing 2.1 reproduces the original version of HELLO.cp for your reviews this program in more detail. Listing 2.1 reproduces the original version of HELLO.cpp for your

INPUT Listing 2.1 HeLLo.cpp Demonstrates the Parts of a C++ Program
\#include <iostream.h>
int main()
cout << "Hello World! !n"
feturn 0 ;

## Output

$$
110 \text { worid }
$$

ANALYSIS On line 1, the file iostream.h is included into the current file.
Here is how the preprocessor works: the first character is the \# (pound) symbol, which is a signal to the
preprocessor. Another name for the pound sign or symbol is octothorpe, which has its history in the darke ages of the telephone company. Each time you start your compiler, the preprocessor is run. The peprocessor ages of the telephone company. Each time you start your compilier, the preprocessor is run. The preprocessor
reads through your source code, looking for lines that eegin with the pound symbol (\#), and acts on those line,
befores before the compiler runs. The preprocessor is discussed in detail on Day 21 , "What's Next.
Include is a preprocessor instruction that says, "What follows is a filename. Find that file and read it in right here." The angle brackets around the filename tell the preprocessor to look in all the usual places for this file.
your compiler is set up correctly, the angle brackets will cause the preprocessor to look for the file iostream. in your compiler is set up correctly, the angle brackets will cause the preprocessor to look for the file iostream.h in
the directory that holds all the .h files for your compiler. The file iostream.h (input-output-stream) is used by cou which assists with writing to the screen. The effect of line 1 is to include the file iostream.h into this program as
you had typed it in yourself. The preprocessorruns before your compiler each time the compiler is invoked. Th you had typed it in yourself. The preprocessorruns before Your compilier each time the compier is invoked. IT
preprocessor translates any line that begins with a pound symbol (\#) into a special command, getting your cod
file ready for the conile. preprocessor trans ates any
file ready for the compiler.
Line 3 begins the actual program with a function named maino. Every $\mathrm{C}_{+}+$program has a maino function. A function is a block of code that performs one or more actions. Usually functions are invoked or called by othe functions,
system).
maino, like all functions, musts state what kind of value it will return. The return value type for maino in HELLO.cpp is int, which means that this function will return an integer to the operating system when it HELLO.cpp is int, which means that this function will return an integer to the operating system when it
completes. In this case, it returns the integer value 0 , as shown on line 6 . Returning a value to the operatin, system is a relatively unimportant and little-used feature, but the $\mathrm{C}++$ standard does require that maino be
declared as shown. Note: GNU and some ohere compilers will let you declare maino to return voi. This is no longer legal
you should not get into bad habits. Have maino return int and simply reutrn 0 as the last line in main0.

Note: Some operating systems (like Linux) enable you totest the value returned by a program. The convention is
or reurun 0 to indicate that the program ended normally If you are using the assh or p pdsks shell, the 8 ? environment or eturn 0 to indiciate that the prog
variable contains this return value.

All functions begin with an opening brace (i) and end with a closing brace (\}). The braces for the main $)$ funct are on lines 4 and 7 . Everything between the opening and closing braces is considered a part of the function. The main processing of this program is on line 5 .

The object cout is used to print a message to the screen. This book covers objects in general on Day 6, "Basic Classes," and cout and its related object cin in detail on Day 16, "Streams." These two objects, cin and cout, are used in C++ to handle input (for example, from the keyboard) and output (for example, to the screen), respectively
cout is used this way: Type the word cout, followed by the output redirection operator (<<). Whatever follows output redirection operator is written to the screen. If you want a string of characters written, be sure to enclo:

A text string is a series of printable characters
The final two characters, पn, tell cout to put a new line after the wordd "Hello World!" This special code is
explained in detail when cout is discussed on Day 17 , "Namespaces."

## The maino function ends on line 7 with the closing brace.

## Brief Look at cout

On Day 16 , you will see how to use cout to print data to the screen. For now, you can use cout without fully
understanding how it works. To print a value to the screen, write the word couts followed by the insertion derstanding how it works. Thint a value to the screen, write the word cout, followed by the insertion operator ((ऽ), which you create by t
characters, $\mathrm{C}++$ treats them as one.

```
lout<<"Hello there.\n";
cout<< "Here is 5:"\<< " << "\n";
cout << "Here is a very big number:\t" << 70000<< endl;
cout<< "Here is the sum of 8 and 5:\t"<< 8+5<< endl;
<< "And a very very big num
< (double
"Don't forget to replace Jesse Liberty
return 0;
```


## Output

Hello there.
Here is 5: 5
The manipulator endl writes a new line to the screen
is a very big number: 70000
$\begin{array}{ll}\text { 's a fraction: } & 13 \\ \text { a very very big number: } & 0.625 \\ 4.2\end{array}$
And a very very big number: $4.9 e+07$
Don't forget to replace Jesse Liberty with your name
Liberty is a C++ programmer.

## Analysis

ANALYsIS On line 2, the statement \#include <iostream.h> causes the iostream.h file to be added to your source code. This is required if you use cout and its related functions.

On line 5 is the simplest use of cout, printing a string or series of characters. The symbol ln is a special formatting character. It tells cout to print a new line character to the screen; it is pronounced "backslash-n" or "new line.'

Three values are passed to cout on line 6, and each value is separated by the insertion operator. The first value the string "Here is 5:". Note the space after the colon. The space is part of the string. Next, the value 5 is passed to the insertion operator and the newline character (always in double quotes or single quotes). This causes the line
Here is 5:
to be printed to the screen. Because no newline character is present after the first string, the next value is print immediately afterward. This is called concatenating the two values
On line 7 , an informative message is printed, and then the manipulator endl is used. The purpose of endl is to write a new line to the screen. (Other uses for endl are discussed on Day 16 .)

Note: endl stands for $e n d$ line and is end-ell rather than end-one. It is commonly pronounced "end-ell."

On line 10 , a new formatting character, Lt , is introduced. This inserts a tab character and is used on lines 10 to to line up the output. Line 9 shows that not only integers, but long integers as well, can be printed. Line 11 demonstrates that cout will do simple addition. The value of $8+5$ is passed to cout, but 13 is printed
On line 12 , the value $5 / 8$ is inserted into cout. The term (float) tells cout that you want this value evaluated as a decimal equivalent, and so a fraction is printed. On line 14 , the value $7000 * 7000$ is given to cout, and the term 3 "Vaind Cols discussed Day 3, "Variables and Constants," when data types are discussed

On line 16, you substituted your name, and the output confirmed that you are indeed a $\mathrm{C}++$ programmer. It m be true, because the computer said so

## Comments

When you are writing a program, it is always clear and self-evident what you are trying to do. Funny thing, though-a month later, when you return to the program, it can be quite confusing and unclear. I really am not sure how that confusion creeps into your program, but it always does.

To fight the onset of bafflement, and to help others understand your code, you should use comments. Commen are text that is ignored by the compiler, but that may inform the reader of what you are doing at any particular point in your program

## Types of Comments

C++ comments come in two flavors: the double-slash (//) comment, and the slash-star (/*) comment. The double slash comment, which is referred to as a C++-style comment, tells the compiler to ignore everything that follow Slash comment, which is referred to a

The slash-star comment tells the compiler to ignore everything that follows until it finds a star-slash (*/) comment mark. These marks are referred to as C-style comments. Every $/ *$ must be matched with a closing *. As you might guess, C-style
of the official definition of

Many C++ programmers use the C++-style comment most of the time and reserve C-style comments for blocking out large blocks of a program. You can include C++-style comments within a block "commented out by C-style comments; everything, including the C++-style comments, is ignored between the C-style comment marks.

## Using Comments

As a general rule, the overall program should have comments at the beginning, telling you what the program Ades. Each function should also have comments explaining what the function does and what values it return does. Each function should also have comments explaining what the function does and what values it retu
And these should be updated when the program is changed. At the very least, a change history should be maintained here.

Functions should be named so that little ambiguity exists about what they do, and confusing and obscure bits of code should be redesigned and rewritten so as to be self-evident. As often as not, comments are a lazy programmer's excuse for obscurity

This is not to suggest that comments ought never be used-only that they should not be relied upon to clarify obscure code; instead, fix the code. In short, write your code well, and use comments to supplement understanding.
Listing 2.3 demonstrates the use of comments, showing that they do not affect the processing of the program o its output.

INPUT. Listing 2.3 HELP.cpp Demonstrates Comments

```
#include <iostream.h>
```

2:
3: in

3: int main(
/* this is a comment
and it extends until the closing
star-slash comment mark
cout << "Hello World! 1 "
// this comment ends at the end of the line
cout << "That comment ended! $\backslash \mathrm{n}$ ";
// double slash comments can be alone on a line

* as can slash-star comments */
return 0;


## Output

Hello World! That comment ended!

## Analysis

 lines 9,12 , and 13. The comment on line 9 ended with the end of the line, however, but the comments on lines and 13 required a closing comment mark
## A Final Word of Caution About Comments

 you might stop drawing, go sharren the pencil, and then return to what you were doing. When a program neea service performed, it can call a function to perform the service and then pick up where it left off when the a service performed, it can call a function to perform the service a
function is finished running. Listing 2.4 demonstrates this idea.

InPut Listing 2.4 Demonstrating a Call to a Function
\#include <iostream.h>
// function Demonstration Function
// prints out a useful message
prints out a useful messag
void DemonstrationFunction ()
cout << "In Demonstration Function $\backslash n "$;
// function main - prints out a message, then
a second message.
$t$ main()
cout << "In main $\backslash$ " ;
Demonstrat ionFunction();
cout $\ll$ "Back in main $\$ n"
return

In main
In Demonstration Function
Analysis to the screen and then returns.
Line 13 is the beginning of the actual program. On line 15 maino prints out a message saying it is in maino
After rinting the message, , ine 16 calls DemonstrationFunctiono. This call causes the commands in Demonstrationfunctiono to execute. In this case, the entire function consists of the code on line 7 , which


## Using Functions

Functions either return a value or the return void, meaning they return nothing. A function that adds two integers might return the sum, and thus would be defined to return
message has nothing to return and would be declared to return void

Functions consist of a header and a body. The header consists of the return type, the function name, and the
parameters to that function. The parameters to a function enable values to te peassed int the function. Thus, parameters to that function. The parameters to a function enable values to be passed into the function. Thus, ,
the function were to add two numbers, the numbers would be the parameters to the function. A typical functi he anction were t
header looks like:
int Sum(int a, int b)
A parameter is a declaration of what type of value will be passed in; the actual value passed in by the callin function is called the argument. Many programmers use these two terms, parameters and arguments, as
synonyms. Others are careful about the technical distinction. This book uses the terms interchangeably The body of a function consists of an opening brace, zero or more statements, and a closing brace. The
statements constitute the work of the function. A function may return a value, sing a retur statement. T
 automatically return void (no value) at the end of the function. The value returned must be of the type declare

Listing 2.5 demonstrates a function that takes two integer parameters and returns an integer value. Do not worr about the syntax or
in detail on Day 3 .

InPut Listing 2.5 Func.cpp Demonstrates a Simple Function
\#include <iostream.h>
int Add (int $x$, int $y)$
cout << "In Adal)
return ( $x+y$ );


# More About the GNU Compile 

 Whenever you are in doubt about how to do something, you should check the documentation for your compile,The GNU compiler g+t is no exception. Included with he distribution are manual pages and info files.
$\qquad$

Compiler options are specified using command-line switches or arguments. The most commonly used compiler options are shown in Table 2.1.

Table 2.1Compiler Options

| Option | Meaning |
| :--- | :--- |
| (no options used) | Compile and link program on command line and produce executable with <br> default name (a.out for Linux, a.exe for MS-DOS). <br> -c |
| Compile but do not link; create file with .o extension. |  |
| -Dmacro=value | Define macro inside program as value. |
| -E | Preprocess, do not compile or link; output is sent to the screen. |
| -g Include debugging information in executable. |  |
| -Idir | Include dir as directory to search for include files (where the name is within <br> <>); multiple directories are separated by semicolons (;) <br> Include dir as directory to search for library files (used by linker to resolve <br> external references); multiple directories are separated by semicolons (;). <br> -Ldir <br> Include library when linking. <br> -lilibrary |
| Optimize. |  |
| -O | Save the executable as file. |
| -Wall | Enable warnings for all code that should be avoided. |
| -w | Disable all warning messages |

You can use multiple options on each compilation command line. You should insert a space before the beginning of each option (the minus sign). You should not have a space between the option and any arguments (such as -o and file).

## GCC Compiler Tips

It is a good idea when creating an executable to use the -o option so you can specify the executable name (otherwise, the next compilation will overlay it with the default name).

By default, any error messages goes to your screen. You may want to redirect them to a file so you can review them at length using a text editor. This is especially handy when you have a lot of errors. You redirect the output to a file:
g++ lst02-01.g++ -olst02-01 > lst02-01.lst
Of course, the selection of filename is entirely up to you. I like to use a suffix of .lst.
As you build function libraries (more about functions in Day 5), you can compile them once and reuse them. If you do not change the source, there is no need to recompile. The -c option will create the intermediate compiled formatted object file. When you want to use the compiled code, you merely include its name on the command line.

As always, the biggest tip with any compiler or Linux tool is to check the manual page (or info file). Review the options, try them out, and see how they behave. They might help you!

## Summary

The difficulty in learning a complex subject, such as programming, is that so much of what you learn depends on everything else there is to learn. This chapter introduced the basic parts of a simple C++ program. It also introduced the development cycle and a number of important new terms

Q\&A

Q What does \#include do?
A This is a directive to the preprocessor, which runs when you call your compiler. This specific directive causes the file named after the word include to be read, as if it were typed in at that location in your source code.

## Q What is the difference between // comments and /* style comments?

A The double-slash comments (//) "expire" at the end of the line. Slash-star (/*) comments are in effect until a closing comment $(* /)$. Remember, not even the end of the function terminates a slash-star comment; you must put in the closing comment mark, or you will get a compile-time error.
Q What differentiates a good comment from a bad comment?
A A good comment tells the reader why this particular code is doing whatever it is doing or explains what a section of code is about to do. A bad comment restates what a particular line of code is doing. Lines of code should be written so that they speak for themselves. Reading the line of code should tell you what it is doing without needing a comment.

## Workshop

The Workshop provides quiz questions to help you solidify your understanding of the material covered and exercises to provide you with experience in using what you learned today. Try to answer the quiz and exercise questions before checking the answers in Appendix D, "Answers to Exercises and Quizzes," and make sure you understand the answers before continuing to the next day.

## Quiz

1. What is the difference between the compiler and the preprocessor?
2. Why is the function main() special?
3. What are the two types of comments, and how do they differ?
4. Can comments be nested?
5. Can comments be longer than one line?

## Exercises

1. Write a program that writes "I love C++" to the screen.
2. Write the smallest program that can be compiled, linked, and run.
3. BUG BUSTER: Enter this program and compile it. Why does it fail? How can you fix it?
\#include <iostream.h>
2: int main()
3: \{
: cout << Is there a bug here?";
return 0;
6: \}
4. Fix the bug in Exercise 3 and recompile, link, and run it

## Chapter 3

## Variables and Constants

Programs need a way to store the data they use. Variables and constants offer various ways to represent and manipulate that data.

## Today you will learn

- How to declare and define variables and constants
- How to assign values to variables and manipulate those values
- How to assign values to variables and manipulate


## What Is a Variable?

In C++, a variable is a place to store information. A variable is a location in your computer's memory in which you can store a value and from which you can later retrieve that value.

位 computer's memory can be viewed as a series of cubbyholes. Each cubbyhole is one of many, many suc holes all lined up. Each cubbyhole-or memory location-is numbered sequentially. These numbers are know as memory addresses. A variable reserves one or more cubbyholes in which you may store a value.
Your variable's name (for example, myVariable) is a label on one of these cubbyholes so that you can find it easily without knowing its actual memory address. Figure 3.1 is a schematic representation of this idea. As y take up one or more memory addresses.

## $=\square\| \| \|$

Figure 3.1 A schematic representation of memory.

## Note:

RAM is random access memory. When you run your program, it is loaded into RAM from the disk file. All variables are also created in RAM. When programmers talk about memory, it is usually RAM to which they are referring.

Variables are also known as lvalues because they can be used on the left side of an assignment operator. Th assignment operator is the equal sign (more about the assignment operator in today's lesson, in the section "Assigning Values to Your Variables," and in Day 4, "Expressions and Statements," in the "Expressions" section). Variables can also be used on the right side of the assignment operator

## Setting Aside Memor

When you define a variable in $\mathrm{C}_{++}$, you must tell the compiler what kind of variable it is: an integer, character, and so forth. This information tells the compiler how much room to set aside and what kind of value you want to store in your variable. Each cubbyhole is 1 byte large. If the type of variable you create is 4 bytes in size, it needs 4 bytes of memory
or four cubbyholes. The type of the variable (for example, integer) tells the compiler how much memory (how many cubbyholes) to set aside for the variable

Because computers use bits and bytes to represent values, and because memory is measured in bytes, it is important that you understand and are comfortable with these concepts. For a full review of this topic, please read Appendix C, "Binary, Octal, and Hexadecimal Numbers and an ASCII Chart:"

## Determining the Size of Integers and Other Data Types

On any one computer, each variable type takes up a single, unchanging amount of room. That is, an intege might be 2 bytes on one machine and four on another, but on either computer it is always the same, day in an day out

A char variable (used to hold characters) is most often 1 byte long

Note:
There is endless debate about how to pronounce char. Some say it as "car." some say it as char as in charcoal, and others say it as "care." Feel free to say it however you like, but some people do make an issue about it-as if there were one true way.
short integer is 2 bytes on most computers, a long integer is usually 4 bytes, and an integer (without the keyword short or long) can be 2 or 4 bytes. The size of an integer is determined by the computer (16-bit or 32and the compiler you use. On modern 32 -bit (Pentium) computers using modern compilers (for example, gec Listing 3.1 should help you determine the exact size of these types on your computer

A character is a single letter, number, or symbol that takes up 1 byte of memory
INPUT Listing 3.1 Determining the Size of Variable Types on Your Compute
\#include <iostream.h>
int main(
cout << "The size of an int is: \t\t" << sizeof(int)
" bytes. $\backslash n$ ";
cout << "The size of a short int is:\t" << sizeof(short)
cout << "The size of a long int is: $\backslash \mathrm{t}$ " << sizeof(long)
cout << "The size of a char is:\t\t" << sizeof(char)
" bytes. $\backslash \mathrm{n}$ ";
cout $\ll$ "The size of a float is: $\backslash t \backslash t$ " $\ll$ sizeof(float) " bytes. \n";
cout << "The size of a double is:\t" << sizeof(double)
cytes << "The size of a bool is:\t" << sizeof(bool)
"out << "The

```
return 0;
```


## OUtrut

The size of an int is: The size of a short int is The size of a long int is The size of a char is: The size of a float is: The size of a bool is:

4 bytes.
2 bytes
4 bytes.
1 bytes
4 bytes
4 bytes
8 bytes
8 bytes
1 bytes

On your comper

ANAIYSIS Most of Listing 3.1 should be pretty familiar. The one new feature is the use of the sizeof() operato in lines 5 through 11. sizeoff 0 is provided by your compiler, and it tells you the size of the object you pass in as parameter. On line 5, for example, the keyword int is passed into sizeof(). Using sizeoff), I was able to determin an in which is 4 bytes.

## Using Signed and Unsigned Integers

In addition, all integer types come in two varieties: signed and unsigned. The idea here is that sometimes you need negative numbers and sometimes you don't. Integers (short and long) without the word "unsigned" are assumed to be signed. Signed integers are either negative or positive. Unsigned integers are always positive.

Because you have the same number of bytes for both signed and unsigned integers, the largest number you can store in an unsigned integer is twice as big as the largest positive number you can store in a signed integer. An unsigned short integer can handle numbers from 0 to 65,535 . Half the numbers represented by a signed short are negative, thus a signed short can only represent numbers from $-32,768$ to 32,767 . If this is confusing, be sure to read Appendix C.

## Understanding Fundamental Variable Types

Several other variable types are built into $\mathrm{C}++$. They can be conveniently divided into integer variables (the type discussed so far), floating-point variables, and character variables.

Floating-point variables have values that can be expressed as fractions-that is, they are real numbers Character variables hold a single byte and are used for holding the 256 characters and symbols of the ASCII and extended ASCII character sets.

The ASCII character set is the set of characters standardized for use on computers. ASCII is an acronym for American Standard Code for Information Interchange. Nearly every computer operating system supports ASCII, although many support other international character sets as well.

The types of variables used in $\mathrm{C}++$ programs are described in Table 3.1. This table shows the variable type, how much room this book assumes it takes in memory, and what kinds of values can be stored in these variables. The values that can be stored are determined by the size of the variable types, so check your output against what you see in Listing 3.1.

Table 3.1Variable Types

| Type | Size | Values |
| :--- | :--- | :--- |
| bool | 1 byte | True or false |
| unsigned short int | 2 bytes | 0 to 65,535 |
| short int | 2 bytes | $-32,768$ to 32,767 |
| unsigned long int | 4 bytes | 0 to $4,294,967,295$ |
| long int | 4 bytes | $-2,147,483,648$ to $2,147,483,647$ |
| int (16-bit) | 2 bytes | $-32,768$ to 32,767 |
| int (32-bit) | 4 bytes | $-2,147,483,648$ to $2,147,483,647$ |
| unsigned int (16-bit) | 2 bytes | 0 to 65,535 |
| unsigned int (32-bit) | 4 bytes | 0 to $4,294,967,295$ |
| char | 1 byte | 256 character values, if signed, then -128 to $127 ;$ if |
|  |  | unsigned, then 0 to 255 |
| float | 4 bytes | $-1.2 \mathrm{e}-38$ to 3.4 e 38 |
| double | 8 bytes | $-2.2 \mathrm{e}-308$ to 1.8 e 308 |

Note:

The sizes of variables might be different from those shown in Table 3.1, depending on the compiler and the computer you are using. If your computer had the same output as presented in Listing 3.1, Table 3.1 should apply to your compiler. If your output was different from that shown in Listing 3.1, you should consult the manual pages for the values that your variable types can hold on your system.

Another place you can look is in the header file limits.h.

## Defining a Variable

You create or define a variable by stating its type, followed by one or more spaces, followed by the variable name and a semicolon. The variable name must begin with a letter or an underscore and can contain virtually any combination of letters, numbers, and underscores, but it cannot contain spaces. Legal variable names include x, J23qrsnf, my_Age, and myAge. Good variable names tell you what the variables are for; using good names makes it easier to understand the flow of your program. The following statement defines an integer variable called myAge:
int myAge;

Note:

When you define or declare a variable, memory is allocated (set aside) for that variable. The value of the variable will be whatever happened to be in that memory at that time. You will see in a moment how to assign a new value to that memory.

Structures and object classes behave a little differently than regular variables. You will learn that difference when you get to that day's lesson. You can define or declare an object class but not use any memory. It is when you actually create the object that space is allocated.

As a general programming practice, avoid such horrific names as J23qrsnf, and restrict single-letter variable names (such as $x$ or i) to variables that are used only very briefly. Try to use expressive names such as myAge or howMany. Such names are easier to understand three weeks later when you are scratching your head trying to figure out what you meant when you wrote that line of code.

Try this experiment. Guess what these pieces of programs do, based on the first few lines of code:

```
unsigned short x;
    unsigned short y;
    unsigned short z;
    z=x * y;
```

Example 2
int main (
unsigned short Width;
unsigned short Length;
unsigned short Area;
Area = Width * Length;
return 0;
Note:
If you compile this program, your compiler will warn you that these values are not initialized. You'll see how to
solve this problem shortly.
Clearly, the purpose of the second program is easier to guess, and the inconvenience of having to type the
onger variable names is more than made up for by how much easier it is to maintain the second program.

## Case Sensitivity

C++ is case sensitive. In other words, uppercase and lowercase letters are considered to be different. A variable named age is different from Age, which is different from AGE

Note:
Some compilers allow you to turn case sensitivity off. The GNU compilers do not, so do not be tempted to do this

Various conventions exist for naming variables, and although it does not matter much which method you adopt it is important to be consistent throughout your program

Many programmers prefer to use all lowercase letters for their variable names. If the name requires two word (for example, my car), two popular conventions are used: my_car or myCar. The latter form is called camel notation because the capitalization looks something like a camel's hump

Some people find the underscore character (my_car) to be easier to read, but others prefer to avoid the解 subsequent words are capitalized: myCar, theQuickBrownFox, and so forth.

Note:
Many advanced programmers employ a notation style that is often referred to as Hungarian notation. The ide Manind Huncarian notation is toprefix every variable with a set of characters that describes its type. Integer
beariables might begin with a lowercase letter $i$, and longs might begin with a lowercase 1. Other notations indicate variables might begin with a lowercase letter $i$, and longs might begin with a lowercases 1. Other notations indicate
constants, globals, pointers, and so forth. Most of this is much more important in $C$ programming because $C++$ constants, globals, pointers, and so forth. Most of this is much more important in C Programming because C-
supports the creation of user-defined types see Day 6 , "Basic Classese") and because C++ is strongly typed.

## Keywords

Some words are reserved by C+t, and you may not use them as variable names. Another name for keywords reserved words. These are keywords used by the compiler to control your program. Keywords include if, while,
for, and main. Your compiler manual should provide a complete list, but generally, any reasonable name for a ariable is I Your compiler manual should provide a complete list, but generally, any reasonable nam Keywords."

## Creating More Than One Variable at a Time

You can create more than one variable of the same type in one statement by writing the type and then the variable names, separated by commas. For example:

$$
\begin{aligned}
& \text { unsigned int myAge, myWeight; // two unsigned int variable } \\
& \text { long int area, width, length; } \\
& \text { // three long integers }
\end{aligned}
$$

As you can see, myAge and myWeight are each declared as unsigned integer variables. The second line declares three individual long variables named area, width, and length. The type (long) is assigned to all the variables, so you cannot mix types in one definition statement.

## Assigning Values to Your Variables

You assign a value to a variable by using the assignment operator ( $($ ). Thus, you would assign 5 to width by writing
unsigned sh
Width $=5$;
Note:
long is a shorthand version of long int, and short is a shorthand version of short int.

You can combine these steps and initialize Width when you declare it by writins
unsigned short Width $=5$
Initialization looks very much like assignment, and with integer variables, the difference is minor. Later, when constants are covered, you will see that some values must be initialized because they cannobe assigned to. The essential difference is that initialization takes place at the moment you create the variable.
Just as you can define more than one variable at a time, you can initialize more than one variable at the same time. For example:
create two long variables and initialize them
with = 5, length = 7;
This example initializes the long integer variable width to the value 5 and the long integer variable length to the value 7 . You can even mix definitions and initializations:
int myAge $=39$, yourAge, hisAge $=40$
This example creates three type int variables, and it initializes the first and third
Listing 3.2 shows a complete program, ready to compile, that computes the area of a rectangle and writes the answer to the screen

```
// Demonstration of variables
#include <iostream.h>
int main()
unsigned short int Width = 5, Length;
M
// create an unsigned short and initialize with result
// of multiplying Width by Length 
cout << "Width:" << Width << "\n";
cout << "Length:" << Length << endl;
return 0;
```


## Output

width: 5
Length: the program.

On line 6 , Width is defined as an unsigned short integer, and its value is initialized to 5 . Another unsigned sho integer, Length, is also defined, but it is not initialized. On line 7 , the value 10 is assigned to Length.
On line 11, an unsigned short integer, Area, is defined, and it is initialized with the value obtained b multiplying Width times Length. On lines 13 to 15 , the values of the variables are printed to the screen Note the special word end creates a new line.

## Using typedef

It can become tedious, repetitious, and, most important, error-prone to keep writing unsigned short int. $\mathrm{C}++$
enables you to create an alias for this phrase by using the keyword typedef, which stands for type definitio In effect, you are creating a synonym, and it is important to distinguish this from creating a new type (which you will do on Day 6) because typedef does not create a new data type
thedef is used by writing the keyword typedef, followed by the existing type, then the new name, and endin with a semicolon. For example,
typedef unsigned short int USHORT;
creates the new name USHORT that you can use anywhere you might have written unsigned short int. Listing 3 a replay of Listing 3.2, using the type definition USHORT rather than unsigned short int.

INPUT Listing 3.3 A Demonstration of typedef

```
#/ Demonstrates typedef keyword
```

typedef unsigned short int USHORT; //typedef defined
int main()

```
USHORT Width =
USHORT Length;
MSHORT Length;
USHoRT Area = Width * Length
lol
cout << "Anea:" << Area <<endl;
return 0;
```


## Output

Width:5
Length:
Area: 50

```
ANalysIS
```

2ne 5 , USHORT is typedefian the program is very much like Listing 3.2, and the output is the same

## When to Use short and When to Use long

One source of confusion for new C+t programmers is when to declare a variable to be type long and when t $t$ e
declare it to be type short. The rule, when understood, is fairly straightforward: If any chance exists that the declare it to be type short. The rule, when understood, is fairly straightforward: If any chance exists that the value you will want to put into your variable will be too big for its type, use a larger type As shown in Table 3.1, unsigned short integers, assuming that they are 2 bytes, can hold a value only up to . 335 Signed short integers split their values between positive and negative numbers, and thus their maxim value is only half that of the unsigned

Although unsigned long integers can hold an extremely large number ( $4,294,967,295$ ), that is still quite finite you need a larger number, you'll have to go to float or double, and then you lose some precision. Floats and doubles can hold extremely large numbers, but only the first 7 or 19 digits are significant on most computers. That means that the number is rounded off after that many digits

Shorter variables use up less memory. These days, memory is cheap and life is short. Feel free to use int, whicl will probably be 4 bytes on your machine

## Wrapping Around an unsigned Integer

The fact that unsigned long integers have a limit to the values they can hold is only rarely a problem, but wha happens if you do run out of room?
When an unsigned integer reaches its maximum value, it wraps around and starts over much as a car odomet might. Listing 3.4 shows what happens if you try to put too large a value into a short integer

INPUT

```
#include < < 
```


cout \ll

## Output

## Wrapping Around a signed Integer

A signed integer is different from an unsigned integer, in that half of the values you can represent are negative Instead of picturing a traditional car odometer, you might picture one that rotates up for positive numbers and Instead of picturing a traditional car odometer, you might picture one that rotates up for positive numbers and
down for negative numbers. One mile from 0 is either 1 or -1 . When you run out of positive numbers, you run right into the most negative numbers (largest absolute value) and then count back down to 0 . Listing 3.5 shows what happens when you add 1 to the maximum positive number in a short integer
\#include <iostream.h>
int main()
short int smallNumber
smallinumber = 32767;
cout << "small number:" << smallNumber << endl;
cout << "small number:" << smallNumber $\ll$ endl;
cout << "small number:" << smallNumber << endl;
return 0;
2: \}

## Output

small number:32767 small number:-32768

ANALYSIS On line 4, smallNumber is declared this time to be a signed short integer (if you do not explicitly say that it is unsigned, it is assumed to be signed). The program proceeds much as the preceding one, but the outpu is quite different. To fully understand this output, you must be comfortable with how signed numbers are represented as bits in a 2 -byte integer.

The bottom line, however, is that just like an unsigned integer, the signed integer wraps around from its highes positive value to its most negative value (highest absolute value).

## Using Character Variables

Character variables (type char) are typically 1 byte, enough to hold 256 values (see Appendix C). A char can b interpreted as a small number ( -128 to 127 or if unsigned, 0 to 255 ) or as a member of the ASCII set. ASCII stands for the American Standard Code for Information Interchange. The ASCII character set and its ISO (International Standards Organization) equivalent are a way to encode all the letters, numerals, and punctuation marks.

Note:
Compuers do not know about letters, punctuation, or sentences. All they understand is numbers. In fact, all they really know about is whether a sufficieient amount of electricicty is at a particular junction of wires. If so, it is
represented internally as ar if if not, it is represented as a 0 . By rouping ones and cros, the computer is able represented internally as a 1 ; if not, it is represented as a 0 . By grouping ones and zeros, the computer is able to
generate patterns that can be interpreted as numbers, and these, in turn, can be assigned to tetters and punctuation.

Your program is stored in memory as a bunch of ones and zeros also; the computer knows how to interpret them
properly. properly.

In the ASCII code, the lowercase letter "a" is assigned the value 97. All the lower- and uppercase letters, all the numerals, and all the punctuation marks are assigned values between 0 and 127. An additional 128 marks and symbols are reserved for use by the computer maker, although the IBM extended character set has become something of a standard.

Note:
ASCII is usually pronounced "Ask-ee."

Note:
Linux is ASCII-based; older operating systems (IBM mainframe and some others) use the EBCDIC character set. EBCDIC stands for Extended Binary Coded Decimal Interchange Code and is related to the way holes are punched in Hollerith cards (a data entry method I hope you never get to see, except in a museum).

## Characters as Numbers

When you put a character, for example, "a," into a char variable, what really is there is a number between and 127. Some systems use unsigned characters, which are numbers between 0 and 255 . The compiler know位 letter, numeral, or punctuation mark, followed by a closing single quotation mark) and one of the ASCII values
The value/letter relationship is arbitrary; there is no particular reason that the lowercase " a " is assigned the 97. A log 2ite hower, that a big difference exists between the value 5 and the character ' 5 ', The latter is actually valued at 53 , much as the letter " $a$ " is valued at 97

Inrut
INPUT.

```
include <iostream.h>
int main()
for (int i = 32; i<128; i++
    cout << (char) i;
    eturn 0;
```

Output
("\#\#\%' ()*+,./0123456789:; <>?@ABCDEFGHIJKLMNOP
QRSTUVWXYZ[\]^'abcdefghijklmnopqrstuvwxyz \{|\}

ANALYSIS This simple program prints the character values for the integers 32 through 127
Day 12, "Arrays, C Strings, and Linked Lists," provides more information on combining characters together in arrays and strings to build things like words
es some special characters for formatting. Table 3.2 shows the most common ones. You put these into your code lled the escape character), followed by the character. Thus, to put a tab character into your code, you would ente ackslash, the letter t , and then a closing single quotation mark:
r variable (tabCharacter) and initializes it with the character value $t$, which is recognized as a tab. The specia when printing either to the screen or to a file or other output device means the letter n , but when it is preceded by the escape character ( $($ ), it means new line

Table 3.2The Escape Characters

| Character | What It Means |
| :--- | :--- |
| In | New line |
| It | Tab |
| lb | Backspace |
| $"$ | Double quote |
| $Y$ | Single quote |
| $?$ | Question mark |
| $I$ | Backslash |

## Using Constants

Like variables, constants are data storage locations. Unlike variables, and as the name implies, constants don't change. You must initialize a constant when you create it, and you cannot assign a new value later.

## Literal Constants

C++ has two types of constants: literal and symbolic.
A literal constant is a value typed directly into your program wherever it is needed, for example:

```
int myAge = 39;
```

myAge is a variable of type int; 39 is a literal constant. You are not allowed to assign a value to 39 , and its value can't be changed. It is an rvalue because it can only appear on the right side of an assignment statement

## Symbolic Constants

A symbolic constant is a constant that is represented by a name, just as a variable is represented. Unlike a variable, however, after a constant is initialized, its value can't be changed.

If your program has one integer variable named students and another named classes, you could compute how many students you have, given a known number of classes, if you knew each class consisted of 15 students:
students = classes * 15;

## Note:

* indicates multiplication

In this example, 15 is a literal constant. Your code would be easier to read, and easier to maintain, if you substituted a symbolic constant for this value:
students = classes * studentsPerClass

If you later decided to change the number of students in each class, you could do so where you define the constant studentsPerClass without having to make a change every place you used that value.

Two ways exist to declare a symbolic constant in C++. The old, traditional, and now obsolete way is with : preprocessor directive, \#define and the const keyword.

## Defining Constants with \#define

To define a constant the traditional way, you would enter this:
\#define studentsPerClass 15
Note that studentsPerClass is of no particular type (int, char, and so on). \#define does a simple text substitution Every time the preprocessor sees the word studentsPerClass, it puts in the number 15 .

Because the preprocessor runs before the compiler, your compiler never sees your constant; it sees the number 15.

Just because it is old and obsolete does not mean you don't have to understand it-many programmers grew up with this directive, and there are many lines of code that use it

## Defining Constants with const

Although \#define works, a new, much better way exists to define constants in C++
const unsigned short int studentsPerClass = 15,
This example also declares a symbolic constant named studentsPerClass, but this time studentsPerClass is typed as an unsigned short int. This method has several advantages in making your code easier to maintain and in preventing bugs. The biggest difference is that this constant has a type, and the compiler can enforce that it i used according to its type

Note:
Constants cannot be changed while the program is running. If you need to change studentsPerClass, for example, you need to change the code and recompile.

DO watch for numbers overrunning the size of the
DO watch for numbers overrunning the size 0
Don't
DON'T use the term int. Use short and lons to make it lear which size number you intended
DON'T use keywords as variable names. D

## Using Enumerated Constants

Enumerated constants enable you to create new types and then to define variables of those types whose value are restricted to a set of possible values. For example, you can declare COLOR to be an enumeration, and you can define five values for COLOR: RED, BLUE, GREEN, WHITE, and BLACK.

The syntax for enumerated constants is to write the keyword enum, followed by the type name, an open brace, each of the legal values separated by a comma, and finally, a closing brace and a semicolon. Here's an exampl enum CoLor \{ Red, BLUE, GREEN, White, BLACK \};

This statement performs two tasks

1. It makes COLOR the name of an enumeration, that is, a new type

It makes RED a symbolic constant with the value 0 , BLUE a symbolic constant with the value 1 , GREE
a symbolic constant with the value 2 , and so forth.

Every enumerated constant has an integer value. If you don't specify otherwise, the first constant will have the value 0 , and the rest will count up from there. Any one of the constants can be initialized with a particular valu however, and those that are not initialized will count upward from the ones before them. Thus, if you write
enum Color \{ RED=100, BLUE, GREEN=500, WHITE, BLACK=700 \};
then RED will have the value 100 ; BLUE, the value 101 ; GREEN, the value 500 ; WHITE, the value 501 ; and BLACK the value 700 .

You can define variables of type COLOR, but they can be assigned only one of the enumerated values (in this case, RED, BLUE, GREEN,WHITE, or BLACK, or else $100,101,500,501$, or 700 ). You can assign any color value $t$ your COLOR variable. In fact, you can assign any integer value, even if it is not a legal color, although a good compiler will issue a warning if you do. It is important to realize that enumerator variables actually are of type unsigned int, and that the enumerated constants equate to integer variables. It is, however, very convenient to be able to name these values when working with colors, days of the week, or similar sets of values

With enumerated variables, you do not create a new data type-you just hide the implementation details. As a programmer, you use words like RED or BLUE, but the compiler substitutes numbers behind the scenes. This process makes creating and understanding the program a lot easier

Listing 3.7 presents a program that uses an enumerated type
Listing 3.7 A Demonstration of Enumerated Constants

```
1: #include <iostream.h>
2: int main()
3: {
enum Days { Sunday, Monday, Tuesday,
    Wednesday, Thursday, Friday, Saturday };
int choice;
cout << "Enter a day (0-6): ";
cin >> choice
    (choice == Sunday || choice == Saturday)
        cout << "\nYou're already off on weekends!\n";
    else
    else}\mathrm{ cout << "\nOkay, I'll put in the vacation day.\n";
return 0;
```


## Outpu

Enter a day (0-6):
You're already off on weekends
integer, counting upward from 0 ; thus, Tuesday's value is 2 .

The user is prompted for a value between 0 and 6 . The user cannot type the word "Sunday" when prompted fo day; the program does not know how to translate the characters in Sunday into one of the enumerated values. You can, however, test the value the user enters against one or more of the enumerated constants, as shown in line 9 . Here the use of enumerated constants makes the intent of the comparison more explicit. This same effe could have been accomplished by using constant integers, as shown in Listing 3.8, but with a little more programming effort.

## Note:

For this and all the small programs in this book, all the code you would normally write to deal with what happens hen the user types inappropriate data has been left out. For example, this program does not check, as it would i eal program, to make sure that the user types a number between

In real programs, this data validation should be included. You never know what a user (or some cracker trying to
break into your system) is going to do. The generic term for this is defensive programming because you defend the onsistency and proper behavior of your code and data.

```
1: #include <iostream.h
2: int main()
const int Sunday = 0;
const int Monday = 1;
const int Tuesday = 2;
const int Thursday = 4;
const int Friday = 5;
const int saturday =
const int Saturday = 6;
int choice;
cout << "Enter a day (0-6): ";
cin >> choice;
if (choice == Sunday || choice == Saturday)
cout << "\nYou're already off on weekends!\n";
cout << "\nOkay, I'll put in the vacation day.\n";
:
return 0;
```


## OUTPU

Enter a day (0-6): 6 You're already off on weekends

ANALYSIS The output of this listing is identical to that shown in Listing 3.7. Here, each of the constants (Sunday, Monday, and so on) was explicitly defined, and no enumerated Days type exists. Enumerated constant have the advantage of being self-documenting-the intent of the Days enumerated type is immediately clear.

## Summary

Today's lesson has discussed numeric and character variables and constants, which are used by $\mathrm{C}++$ to store data during the execution of your program. Numeric variables are either integral (char, int, short int, and long int they are floating point (float and double). Numeric variables can also be signed or unsigned. Although all the types can be of various sizes among different computers, the type specifies an exact size on any given computer.

You must declare a variable before it can be used, and then you must store the type of data that you've declarec as correct for that variable. If you put too large a number into an integral variable, it wraps around and produce an incorrect result.

This lesson also reviewed literal and symbolic constants as well as enumerated constants, and it showed two ways to declare a symbolic constant: using \#define and using the keyword const.

## Q If a short int can run out of room and wrap around, why not always use long integers?

A Both short integers and long integers will run out of room and wrap around, but a long integer will do so with a much larger number. For example, an unsigned short int will wrap around after 65,535 , where as an unsigned long int will not wrap around until 4,294,967,295. However, on most machines, a long integer takes up twice as much memory every time you declare one ( 4 bytes versus 2 bytes), and a program with 100 such variables will consume an extra 200 bytes of RAM. Frankly, this is less of a problem than it used to be because most personal computers now come with many thousands (if not millions) of bytes of memory.

## Q What happens if I assign a number with a decimal point to an integer rather than to a float? Consider the following line of code:

int aNumber = 5.4;

A A good compiler will issue a warning, but the assignment is completely legal. The number you've assigned will be truncated into an integer. Thus, if you assign 5.4 to an integer variable, that variable will have the value 5 . Information will be lost, however, and if you then try to assign the value in that integer variable to a float variable, the float variable will have only 5.

## Q Why not use literal constants; why go to the bother of using symbolic constants?

A If you use the value in many places throughout your program, a symbolic constant allows all the values to change just by changing the one definition of the constant. Symbolic constants also speak for themselves. It might be hard to understand why a number is being multiplied by 360 , but it's much easie to understand what's going on if the number is being multiplied by degreesInACircle.
Q What happens if I assign a negative number to an unsigned variable? Consider the following line of code:
unsigned int aPositiveNumber = -1;
A A good compiler will warn, but the assignment is legal. The negative number will be assessed as a bit pattern and assigned to the variable. The value of that variable will then be interpreted as an unsigned number. Thus, 1 , whose bit pattern is 1111111111111111 ( 0 xFF in hex), will be assessed as the unsigned value 65,535 . If this information confuses you, refer to Appendix C.

## Q Can I work with $\mathbf{C + +}$ without understanding bit patterns, binary arithmetic, and hexadecimal?

 A Yes, but not as effectively as if you do understand these topics. C++ does not do as good a job as some languages at "protecting" you from what the computer is really doing. This is actually a benefit because it provides you with tremendous power that other languages don't. As with any power tool, however, to get the most out of $\mathrm{C}++$, you must understand how it works. Programmers who try to program in $\mathrm{C}++$ without understanding the fundamentals of the binary system often are confused by their results.
## Workshop

The Workshop provides quiz questions to help you solidify your understanding of the material covered and exercises to provide you with experience in using what you've learned. Try to answer the quiz and exercise questions before checking the answers in Appendix D, "Answers to Exercises and Quizzes," and make sure that you understand the answers before continuing to the next day.

## Quiz

1. What is the difference between an integer variable and a floating-point variable?
2. What are the differences between an unsigned short int and a long int?
3. What are the advantages of using a symbolic constant rather than a literal constant?
4. What are the advantages of using the const keyword rather than \#define?
5. What makes for a good or bad variable name?
6. Given this enum, what is the value of BLUE?
enum COLOR \{ WHITE, BLACK = 100, RED, BLUE, GREEN = 300 \};
7. Which of the following variable names are good, which are bad, and which are invalid?
a. Age
b. !ex
c. R79J
d. TotalIncome
e. __Invalid

## Exercises

1. What would be the correct variable type in which to store the following information?
a. Your age.
b. The area of your back yard.
c. The number of stars in the galaxy.
d. The average rainfall for the month of January.
2. Create good variable names for the information in question 1.
3. Declare a constant for $p i$ as 3.14159 .
4. Declare a float variable and initialize it using your pi constant.

# Chapter 4 

Expressions and Statements
At its heart, a program is a set of commands executed in sequence. The power in a program comes from its
capability to execute one or another set of commands, based on whether a particular condition is true or false capability to execute on
Today you will learn:

- What statements are
- What block ser
- What expescions are

What expressions are
How to brand your code based on conditions
What truth is, and how to act on it

## Statements

In C++, a statement controls the sequence of execution, evaluates an expression, or does nothing (he null
statement). All C+t statements end with a semicolon, even the null statement, which is is iust he semicolon and


Unike in algebra, this statement does not mean that X equals a+b. This is read, "Assign the value of the sum . has one semicolon. is on the left si

Whitespace (which you create by using tabs, spaces, and newline c
The asigigment statement previously discussed could be written as $x=a+b$;
號

Although his last variation is perfectly legal, it it also perfectly foolish. Whitespace can be used to make you programs morr readable and easier to maintain, or it can be used to create horifici and indecipherable code. Whitespace characters sppaces,
see only the white of the paper.

## Blocks and Compound Statements

Any ylace you can put a single statement, you can put a compound statement, also called a block. A block
begins with an opening brace $(1$ and ends with c clo begins with an opening brace (t) and ends with a closing brace ( $($ ). Although every statement in the block
lock of code acts as one statement and swaps the values in the variables $a$ and

| Do |
| :--- |
| DO use a closing brace any time you have an opening |
| brace |
| Do end your statements with a semicolon. |
| Do sue e hhitespace judiciously to make your code |
| cleare. |

Compound statements are very important, as you will find when you get to the sections about iffelse, for, whii
and dolvwile statements. Thoses statements can only have one statement after them (as part of the true or fals
 of those, you have to use a compound statement (which counts as the one statement) to contain as man.
statements as you need. You will learn more about this in the section "The if Statement" in this lesson.

## Expressions

 The myriad pieces of code that qualify as expressions might surprise you. Here are three examples

```
// int const that returns 60
```

Assuming that $P$ is a constant equal to 3.14 and ScoondserMinute is a constant equal to 60 , all three of these

not only adds $s$ and $b$ and assigns the result to $x$, but reutrns the value of that assignment (the value of $x$ as
well. Thus this statement is also an experession. Because it is an experession, it can be on the right side of an
well. Thus this stat
assignment operator:
This line is evaluated in the following order

Add to b.
Assign the resull of the expression $a+b$ to $x$.
If a $, \mathrm{b}, \mathrm{x}$, and y are all integers, and if $a$ has the value 2 and $b$ has the value 5 , both $x$ and $y$ will be assigned the
value 7 . Listing 4.1 shows how $C_{++}$evaluates complex expressions
Input Listing 4.1 Evaluating Complex Expressions


Output

ANAIYSIS On line 4, the four variables are declared and initialized. Their values are printed on lines 5 and On line 7, a is assigned the value 9 . On line $8, b$ is assigned the value 7 . On line 9 , the values of $a$ and $d$ are
en value is, in turn, assigned to oy.

## Operators

An operator is asymbol that causes the compiler to take an action. Operators act on operands, and in $\mathrm{C}++$ al
operands are expressions. In $\mathrm{C}++$ several categerories of operators exist. Two of these cateoroies ser - Assignment operators

## Assignment Operato

The assigment operator $(=$ ) causes the operand on the left side of the assigmment operator to have its value
changed to the value on the right side of the assignment operator. The following expression $\mathrm{x}=\mathrm{a}+\mathrm{b}$;
esur of adidng a and to he operanax.
An operand that can legally be on the efft side of an assisnment operatar is called an $L$-value (Ivalue). That
which can be on the rights side is called (you suessed it an $r$-value ruvalue).
$\qquad$
$\mathrm{x}=35$;
// ok

This is important for you to remember. soit gets repeated: An l -value is an operand that can be on the left s s.
of an experession. An r -value is an operand that can be on the righ side of an expression. Note that all l-valt are r -values, but not arl r -values are $\mathrm{-}$-values. A . you can write $x=5$; but you cannot write $5=x$; ( $x$ can be an 1 -value or an $r$-value, 5 can only be an $r$-value
int main()



Outrut
Difference is: 50
Now difference is: 4294967248
Analussis $_{\text {The subtraction operator is is inoked on line } 10 \text {, and dhe resultis sprinted on line } 11 \text {, much as you }}$ mightexpect. The subtraction popataor is called again on ine 12 , but this ime a laggo ins.ged number is
 Integer Division and Modulu:



 until he next mutuple
"More Progam Flow:

Frequently Asked Questions

Answer: If you divide one integer by another, you get an integer as a result.
Thus $5 / 3$ will be 1
To get a factional (with decinal) reurn value, you must use floats.
5.013 .0 will give you a fractional (with decimal) answer 1.6667 .
 In this paricular case, you wan
am doing This is really a float:
Two ways existo do dot tis cass: You can use the o
operator L Lsting 4.3. liustrates cating toa float.
Invot
\#include <iostream.h>




OUtrut
e: $1 ., 666$
. 12, the integers are cast of ofot and assi
float on line: 13 and primed on on ine 15 .

## Combining the Assignment and Mathematical Operator




This meatod, however, is teribly convolued and wasefu
of the asisignent operator thus the preceding becomes

Even simplert to write, but perhaps a bit harder to read is myage $+=2$;
 the value 4 to start, $i$ t would have 6 affer this statement.

## Self-asisigned subracion $(-)$,



// start with c and increment it.
dis stalement is equivalent to the more vertoses statemen
$=c+$
$\mathrm{c}+=1$;

efix and Posttix



```
// List ing 4.4 - demonstrates use of
#/ decrement poemators
\begin{tabular}{|c|c|c|}
\hline \multicolumn{3}{|r|}{} \\
\hline \multicolumn{3}{|l|}{\multirow[t]{2}{*}{cout <
nylaet + \%}} \\
\hline & & \\
\hline \multicolumn{3}{|l|}{\multirow[b]{2}{*}{cout << "I am: " << myage <<" " years old. \(\backslash n^{\prime \prime}\);}} \\
\hline & & \\
\hline \multicolumn{3}{|l|}{cout < "You are: " < yourage <<" years oldhn} \\
\hline \multicolumn{3}{|l|}{\multirow[t]{2}{*}{cout}} \\
\hline \multicolumn{3}{|l|}{\multirow[t]{2}{*}{cout}} \\
\hline \multicolumn{3}{|l|}{\multirow[t]{2}{*}{}} \\
\hline & & \\
\hline \multicolumn{3}{|l|}{} \\
\hline
\end{tabular}
```


## Output



Avalvsis ${ }_{\text {On lines }} 7$ and 8, two integ
value are printed on lines 9 and 10 .
On line 11 , myAge is incremented using the postixix increment operator, and on line 12 , yourage is incrementec On line 17 , my Age is incremented as part of the printing statement, using the postfix increment operator.
Because it is postixi, the increment happens after the print, and so the value 40 is printed again. In contrast, o line 18 , yourfeg is incremented using the prefix increment operator. Thus it is incremented before being print


## Operator Precedence

In the following complex statement
$x=5+3$ * $8 ;$
which is performed first, the addition or the multiplication? If the addition is
8 , or 64 . If the multiplication is performed first, the answer is $5+24$, or 29 .
Every operator has a precedence value, and the complete list is shown in Appendix A. Multiplication has higt
precedence than additition; thus the value of the expression is 29 . When two mathematical operators have the same precedence, they are performed in left-to-right order. Thus $x=5+3+8 * 9+6 * 4 ;$
is evaluated multiplication first, eff to right. Thus $8^{*} 9=72$, and $6^{*} 4=24$. Now the expression is essentially $\mathrm{x}=5+3+72+24 ;$
in right-to-left order! In any case wh the precedence order does not meet your needs? Consider the expression
TotalSeconds $=$ NunMinutesToThink + NumMinutesToType $* 60$
In this expression, you do not want to multiply the NumMinuesToType variable by 60 and then add it to
NumMinuessoThink. You want to add the two variables to get the total number of minutes, and then $Y$ O NumMinuesTTThin. You want to add the two variab
multiply that number by 60 to get the total seconds.

In this case, you use parentheses to change the precedenc
precedence than any of the mathematical operators. Thus
TotalSeconds $=($ (NumMinutesToThink + NumMinutesToType) $* 60$

## will accomplish what you want

## Nesting Parentheses

For complex expressions, you uight need to nest parentheses one within another. For example, you might nee
to compute the toal seconds and then compute the total number of people who are involved before multiplyin to compute the total se
seconds times people:
TotalPersonseconds $=(($ (NunMinutesToThin
(PeopleInTheoffice + PeopleonVacation) $)$
This complicated expression is read from the inside out. First, NumMinutesToThink is added to NumMinumesToT,
because these are in the innermost parentheses. Then this sum is multiplied by 60 . Next. Peoplentheoffec is because these are in the innermost parentheses. Then this sum is multipied by 60 . Next. PeoplenTheorfic,
added to Peopleonvacation. Finaly, the total number of people found is multiplied by the toal number of seconds.
This example raises an important related issue. This expression is easy for a computer to understand, but ver
difficult for a human to read, understand, or modify. Here is the same expression rewriten using some fficult for a human to read

TotalMinutes $=$ NunMinutesToThink + NunMinutesToType
TotalMinutes $=$ NunMinutesTorhink
TotalSeconds $=$ TotalMinutes $* 60$ TotalPeople = PeopleInTheoffice + Peopleonvacatior
TotalPersonseconds $=$ TotalPeople
This example takes longer to write and uses more temporary variables than the preceding example, but ti is fal
easier to understand. Add a comment at the top to explain what this coode does and change the 60 to a symboli easier to understand. Add a comment at the top to explain what this code do d.
constant. You then will have code that is easy to understand and maintain.

## The Nature of Truth

In previous versions of $\mathrm{C}+$, all trut and falsity were represented by integers,
introduced a new type: bool. This new type has two possible values: false or tue
Every expression can be evaluated for is truth or falsity. Expressions that evaluate mathematically to zero wi



The relational operators are used to determine whether two numbers are equal or one is greater or less than th
other. Every relational statement evaluates to e either true or fasse. The relational operators are presented in Tab
 equal to $( \rangle=)$ ，and not equals $(l=)$ ．Table e 4.1 shows each realational operator，

| Name | operator | Sample | Evaluates |
| :---: | :---: | :---: | :---: |
| Equals | ＝ | $\begin{aligned} & 100=50 ; \\ & 50=50 ; \\ & \hline 0 \end{aligned}$ | $\begin{gathered} \text { false } \\ \text { tue } \end{gathered}$ |
| Not Equals | ： | $\begin{aligned} & 100!=50 ; \\ & 50!=50 ; \end{aligned}$ | $\underset{\substack{\text { true } \\ \text { false }}}{\text { to }}$ |
| Greater Than | ＞ | $\begin{aligned} & 100>50 ; \\ & 50>50 ; \end{aligned}$ | true false |
| Greater Than or Equals | ＞ |  | $\underset{\substack{\text { tue } \\ \text { tue }}}{\text { tue }}$ |
| Less Than | ＜ | $\begin{aligned} & 100<50 ; \\ & 50<50 ; \end{aligned}$ | false false |
| Less Than or Equals | ＜ | $\begin{aligned} & 100<=50 ; \\ & 50<=50 ; \end{aligned}$ | false <br> true | equals relational operator $(=\Leftrightarrow)$ ．Thisis one of the mos

common $\mathrm{C}++$ programming mistakes - be on gard
for it

The if Statemen

Normally，your program flows along line by line in the order in which it appears in your source code．The if if statement enables you to test tor condititu
parts of your oode，depending on the resul
The simplest form of an if statement is the following
（expression
statement
The expression in the parentheses can be any expression at all，but it susully contains one of the relational
expressions．If the expression has the value false，the satatenent is skipece．If it evaluates true，the statement i executed．Consider the following exampl
（bigNumber＞smal 1 Number $)$
bigNumber $=$ sma11Number
This code compares bigNumber and smallNumber．If bieNvumber is larger，the second line sets its value to the valt
of smallumber． Because a block of statements surrounded by braces is equivalent to a single statement，the following type o branch can be quite large and powerful：
（expression）
statement 1；
statement 2；
statement 3 ；
（bigNumber＞smal1Number）


This time，if bievNumber is larger than small Sumber，not only is it set to the value of smallNumber but an
informational message is printed．Listing 4.5 shows a more detailed example of branching based on relational operators

Input

／／Listing 4．5－demonstrates if statement
tinclud with relation ＜iostream．h＞
int Redsoxscore，Yankeesscore；
cout $\ll$ Enter the ser scre for the Red Sox：
cin $\gg$ RedSoxScore；
－＜＜＂
cout 《＜＂nnEnter the
cin 》》 Yankeesscore；
cout＜＜＂\n＂；

Redsoxscore＜YankeesScore
cout＜＜＂Go Yankees！！n＂；
cout＜＜＂Happy days in New York！$\backslash n " ;$
f（RedSoxScore＝＝Yankeesscore）
cout＜＜＂A tie？Naah，can＇t be．\n＂；
cout＜＂Give me the real score for the Yanks
cin＞＞Yankeesscore；
if（RedSoxScore＞Yankeesscore）
cout＜＜＂Knew it！Go sox！＂，
if（Yankeesscore $>$ RedSoxScore）
cout＜＜＂Knew it！Go Yanks！
f（Yankeesscore $==$ Redsoxscore）
cout $\ll$＂Wow，it really was
cout＜＜＂ $\begin{gathered}\text { nn } \\ \text { return } 0 ;\end{gathered}$

Output
nter the score for the Red sox： 10
Enter the score for the Yankees： 10
A tie？Naah，can＇t be
Give me the real score
Knew it！ Go Sox！
Thanks for telling me
ANalYsIs This program asks for user input of scores for two baseball teams；th．
variables．The variables are compared in the if ftatement on lines 15 ， 18 ，and 24 ．
If one score is higher than the other，an informational message is printed．If the scores are equal，the block of
code that begins on line 25 and ends on line 38 is entered．The second score is requested again，and then the
scores are compared again．

 | on lines 20 and 21 would be invoked．Then the if statement on line 24 would be tested and this would be false |
| :--- |
| on |

OUTPut
Please enter a big number: 10
eease enter a smaller number:
os. The second is bigger!




## Advanced if Statements



## if (expression



${ }_{\text {statement } 3 \text {; }}^{\text {else }}$
${ }^{\text {else }}{ }_{\text {statement4; }}$



## vevt





int first Number, seconalumber;
cout << NEnter two
cint


if (firstNumber >= secondNumber)
if (firstris seconamumber) $=$


$$
{ }^{\text {else }} \text { cout << "They are evenly divisible! ! } n^{\prime \prime} \text {; }
$$

cout << "They are not evenly divisisibe! \n";
cout << "Hey! The second one is larger! $\backslash \mathrm{n}$ ";

Outrut

| Enter tuwo numbers. |
| :--- |
| First: |
| 10 |

Second: 2
ney are evenly divisisible

 number are eitere venly divisu
appropiate message e ititer vay INPUT Listing 4.8 A Demonstration of Why Braces Help Clarify Which else Statement Goes with Which
// Listing 4.8- demonstrates why braces
(/) are mportant in nested if statements
\#include <iostream.h>

cout "< "Enter a number less than 10 or greater than 100: "

if $\left.\begin{array}{c}\text { if } \gg 10) \\ \text { if } \\ \text { ( } x>100) \\ \text { cout } \lll\end{array}\right)$

return 0;

Enter a number less than 10 or greater than 100: 20
Less than 10, Thanks
ANalvsis The programmer intended to ask for a number less than 10 or greater than 100 , check for the
correctulue, and then print a thank-you note. If the if statement on line 11 evaluates true, the following statement (line 12 ) is executed. In this case, line 12 execules when the number entered is greater than 10 . Line 12 contains an it statement also. This is statement
evaluets
13 isuri it the numbere entered is sreater than 100 . If the number is reater than 100 , the statement on line recate
If he number entered is less than 10 , the if statement on line 11 evaluates fase. Program control goes to the ne
line following the if statements in this case line 16 . If you enter a number less than 10 , the output is as ofollows The else claus on ine 14 was clearly intended to be attached to the if statement on line 11 , and thus is inden
accordingly Unfortuately, the esse statement is seally attached to the it statement on line 12 , and thus this accordingly. nitorunate
program has s subtle oug
 the iniended reselt. Firuth
long and a umber that is
example of a logic error:
$\qquad$
// Listing 4.9- - demonstrates proper use of braces
\#include
int main()
in


if $(x>10$

else

```
cout <<
```

OUvirur
Ner less than 10 or greater than 100 :
ANalvsIs The braces on lines 12 and 15 make
on line 16 applies to the if on line 11 , as intended
The user typed 20 , so the if statement on line 11 is true; however, the if statement on line 13 is fate, so nothin
is prined. It would be betere if the programmer put another clse clause after line 14 so that errors would be t and a message printed.

emember to program defensively!

## Logical Operators


 it is NOT a holiday or ifitis a weekend, then call the police $e$.
this kind of evaluation. Thess operators are isted in Table 4.2

Table 4.2The Logical Operator

| operator | Symbol | Example |
| :---: | :---: | :---: |
| and | $8 \times$ | expession 184. expession2 |
|  |  |  |
| or not | ! | lexpresion |

## Logical AND

A ogical AND statement evaluates two expressions, and if oht expressions are true, the logical AND statemn
is true as well. If it is true that you are hungry, and it it strue that you have money, hen it it strue that you can bu lunch. Thu
would evaluate ruse if both $x$ and y are equal to 5 , and it would
that booth sides must be true for the entire expression to be true.
Note that the logical ANI,
Day 21, "Whars SNext."
ogical OR


$$
\text { evaluates tuu if either or or is equal to } 5 \text {, or if both are equal to } 5 \text {. }
$$

te that the logical OR is two | |symbols. A single |symbol is a different operator, which is discussed on Day

A logical Not statement evaluates tru if the expression being tested is false. Again, if the expression being
tested is false, the value of the test is true! Thus if ( $!(x==5)$ )

## is true only if $x$ is not equal to 5 . This is the same as writi

## Short Circuit Evaluation

 the compiler will Nor 8 g
requires both to be tuue.

## Relational Precedence

Relational operators and logical operators, because they are C++ expressions, each return a value: true or false. Relational operators and logical operators, because they are C++ expressions, each return a value: true or
Like all expressions, they have a precedence order (see Appendix A) that determines which relations are Like all expressions, they have a precedence order (see Appendix A) that determines which rela
evaluated first. This fact is important when determining the value of the statement, for example:

```
if ( }x>5&& y>5 || z>5
```

It might be that the programmer wanted this expression to evaluate true if both $x$ and $y$ are greater than 5 or if $z$ It might be that the programmer wanted this expression to evaluate true if both $x$ and $y$ are greater than 5 or if
is greater than 5 On the other hand the programmer might have wanted this expression to evaluate true only $i$ is greater than 5 and if it is also true that either $y$ is greater than 5 or $z$ is greater than 5 .
If $x$ is 3 , and $y$ and $z$ are both 10 , the first interpretation will be true ( $z$ is greater than 5 , so ignore $x$ and $y$ ), bu the second in be false (it is not true that $x$ is greater than 5 and thus it does not matter what is on the right si of the (\&\&) symbol because both sides must be true Although precedence will d
make the statement clearer: $(x>5) \& \&(y>5| | z>5)$

Using the values from earlier, this statement is false. Because it is not true that x is greater than 5 , the left side the AND statement fails, and thus the entire statement is false. Remember that an AND statement requires
both sides be true-something is not both "good tasting" AND "good for you" if it is not good tasting.

Note:
It is often a good idea to use extra parentheses to clarify what
programs that work and that are easy to read and to understand.
Whenever you are in doubt about the order of operations, put the parentheses in. It
some small number of milliseconds and probably result in the same execution time.

## More About Truth and Falsehood

In $\mathrm{C}++$, zero evaluates false, and all other values evaluate true. Because an expression always has a value, many C++ programmers take advantage of this feature in their if statements. A statement such as
f (x)

```
/ if x is true (nonzero)
```

$\mathrm{x}=$
can be read as "If x has a nonzero value, set i to 0 ." This is a bit of a cheat; it would be clearer if written $x!=0$
$x=0 ;$

Both statements are legal, but the latter is clearer. It is good programming practice to reserve the former meth for true tests of logic, rather than for testing for nonzero values.

```
if (!x) (/ if lif is false (zero)
```

The second statement, however, is somewhat easier to understand and is more explicit if you are testing for th The second statement, however, is somewhat easier to
mathematical value of x rather than for its logical state

## Conditional (Ternary) Operato

The conditional operator (?:) is C++'s only ternary operator; that is, it is the only operator to take three terms.
The conditional operator takes three expressions and returns a value
(expression1) ? (expression2) : (expression3)
This line is read as "If expression1 is true, return the value of expression2; otherwise, return the value of expression3." Typically, this value would be assigned to a variable

```
// Listing 4.10 - demonstrates the conditional operator
#include <iostream.h>
int main()
int x, y, z;
    out<< "First.";
    cin>> x;
    cin >> y;
    if (x>y)
    z = x;
    z=
    cout << "z:""><
    = (x>y) ? x : y 
    cout << "z:""
    eturn 0;
```


## OUTPUT

Enter two
First: 5
Second
E: :
Analysis
2eger variables are created: $x, y$, and $z$. The first two are given values by the user. The if 19.

The conditional operator on line 22 makes the same test and assigns z the larger value. It is read like this: "If is greater than $y$, return the value of $x$; otherwise, return the value of $y$." The value returned is assigned to $z$. That value is printed on line 24 . As you can see, the conditional statement is a shorter equivalent to the if..els statemer

## Summary

## Q Why use unnecessary parentheses when precedence will determine which operators are acted on

 first?A Although it is true that the compiler will know the precedence and that a programmer can look up the precedence order, code that is easy to understand is easier to maintain.
Q If the relational operators always return true or false, why is any nonzero value considered true? A The relational operators return true or false, but every expression returns a value, and those values can also be evaluated in an if statement. Here is an example:

$$
\text { if }(x=a+b)==35)
$$

This is a perfectly legal $C++$ statement. It evaluates to $a$ value even if the sum of $a$ and $b$ is not equal to 35. Also note that $x$ is assigned the value that is the sum of $a$ and $b$ in any case.

## Q What effect do tabs, spaces, and newline characters have on the program?

A Tabs, spaces, and newline characters (known as whitespace) have no effect on the program, although judicious use of whitespace can make the program easier to read.

## Q Are negative numbers true or false

A All nonzero numbers, positive and negative, are true.

## Workshop

The Workshop provides quiz questions to help you solidify your understanding of the material covered and exercises to provide you with experience in using what you have learned. Try to answer the quiz and exercise questions before checking the answers in Appendix D, "Answers to Exercises and Quizzes," and make sure tha you understand the answers before continuing to the next day

## Quiz

1. What is an expression?
2. Is $x=5+7$ an expression? What is its value?
3. What is the value of $201 / 4$ ?
4. What is the value of $201 \% 4$ ?
5. If myAge, $a$, and $b$ are all int variables, what are their values after
```
myAge = 39;
```

a = myAge++;
b = ++myAge;
6. What is the value of $8+2 * 3$ ?
7. What is the difference between $\mathrm{x}=3$ and $\mathrm{x}==3$ ?
8. Do the following values evaluate true or false?
a. 0
b. 1
c. -1
d. $x=0$
e. $x==0 / /$ assume that $x$ has the value of 0

## Exercises

1. Write a single if statement that examines two integer variables and changes the larger to the smaller, using only one else clause.
2. Examine the following program. Imagine entering three numbers, and write what output you expect
```
#include <iostream.h>
int main()
{
int a, b, c;
cout << "Please enter three numbers\n";
    cout << "a: ";
    cin >> a;
    cout << "\nb: ";
    cin >> b;
    cout << "\nc: ";
    cin >> c;
    if (c = (a-b))
    {cout << "a: ";
    cout << a;
    cout << "minus b: ";
    cout << b;
    cout << "equals c: ";
    cout << c << endl;}
    else
    cout << "a-b does not equal c: " << endl;
    return 0;
: }
```

3. Enter the program from Exercise 2; compile, link, and run it. Enter the numbers 20, 10, and 50. Did you get the output you expected? Why not?
4. Examine this program and anticipate the output:
```
#include <iostream.h>
int main()
{
    int a = 1, b = 1, c
    if (c = (a-b))
    cout << "The value of c is: " << c;
    return 0;
}
```

5. Enter, compile, link, and run the program from Exercise 4. What was the output? Why?

## Chapter 5

## Functions

Although object-oriented programming has shifted attention from functions and toward objects, functions nonetheless remain a central component of any program. Today you will learn:

- What a function is and what its parts are

How to declare and define functions
How to pass parameters into functions
How to return a value from a function

- How to build and use function libraries
-What the standard libraries are and contain


## What Is a Function?

A function is, in effect, a subprogram that can act on data and return a value. Every C++ program has at least on function- main(). When your program starts, main() is called automatically. main) might call other functions, some of which might call still others.
Each function has its own name, and when that name is encountered, the execution of the program branches to the body of that function. This is called calling the function. When the function returns, execution resumes on the next line of the calling function. This flow is illustrated in Figure 5.1


Figure 5.1 When a program calls a function, execution switches to the function and then resumes at the lin after the function call

Well-designed functions perform a specific and easily understood task. Complicated tasks should be broken down into multiple functions, and then each can be called in turn.

Functions come in two varieties: user-defined and built-in. Built-in functions are part of your compiler package-they are supplied by the manufacturer for your use. User-defined functions are the functions you write yourself.

## Return Values, Parameters, and Arguments

Functions can return a value. When you call a function, it can do work and then send back a value as a result o that work. This is called its return value, and the type of that return value must be declared. Thus if you write int myFunction();
you are declaring that myFunction will return an integer value
You can also send values into the function. The description of the values you send is called a parameter list.
int myFunction(int someValue, float someFloat);
This declaration indicates that myFunction will not only return an integer, it will take an integer value and a float as parameters. A parameter describes the type of the value that will be passed into and the variable name used in
when the function is called. The actual values you pass into the function are called the arguments.
int theValueReturned $=$ myFunction $(5,6.7)$;
Here you see that an integer variable theValueReturned is initialized with the value returned by myFunction, and that the values 5 and 6.7 are passed in as arguments. The type of the arguments must match the declared parameter types.

## Declaring and Defining Functions

Using functions in your program requires that you first declare the function and that you then define the function. The declaration tells the compiler the name, return type, and parameters of the function. The definiti tells the compiler how the function works. No function should be called from any other function if it has not first been declared. A declaration of a function is called a prototype.

## Declaring the Function

Three ways exist to declare a function:

- Write your prototype into a file, and then use the \#include directive to include it in your program. - Write the prototype into the file in which your function is used
ny other function. When you do this, the definition acts as own declaration.

Although you can define the function before using it and thus avoid the necessity of creating a function prototype, this is not good programming practice for three reasons.

First, it is a bad idea to require that functions appear in a file in a particular order. Doing so makes it hard to maintain the program as requirements change.

Second, it is possible that function A() needs to be able to call function B() , but function B() also needs to be abl to call function A() under some circumstances. It is not possible to define function A) before you define function B() and also to define function $\mathrm{B}($ before you define function A() , so at least one of them must be declared in any case.

Third, function prototypes are a good and powerful debugging technique. If your prototype declares that your function takes a particular set of parameters or that it returns a particular type of value, and then your function does not match the prototype, the compiler can flag your error instead of waiting for it to show itself when you run the program.

## Using Function Prototypes

Many of the built-in functions you use will have their function prototypes already written in the files you include in your program by using \#include. For functions you write yourself, you must include the prototype

The function prototype is a statement, which means it ends with a semicolon. It consists of the function's returr type and signature. A function signature is its name and parameter list

The parameter list is a list of all the parameters and their types, separated by commas. Figure 5.2 illustrates the parts of the function prototype.
$\square$

Figure 5.2 Parts of a function prototype
The function prototype and the function definition must agree exactly about the return type and signature. If they do not agree, you will get a compile-time error. Note, however, that the function prototype does not need contain the names of the parameters - just their types. A prototype that looks like this is perfectly legal.
long Area(int, int);
This prototype declares a function named Area() that returns a long and that has two parameters, both integers. Although this is legal, it is not a good idea. Adding parameter names makes your prototype clearer. The same function with named parameters might be
long Area(int length, int width);

It is now obvious what this function does and what the parameters are.
The compiler does not require the variable names and essentially ignores them in the prototype. They are there for us humans.

#  <br> $n$ scuare fe return 0 ； 

int Area（int yardiength，int yardwidth）
return yardLength＊yardwidth；

## ototype Syntax：

## eturn＿type fux

eturn＿type function＿name（［type parametervame］．．）
statements；

 within the body of the function must b be tern
semicolon（i）．It ends with a closising rare．
If the function returns a value，it should end wit
appear anywhere in he body of the function．
Every function has reteurn type．If one is not explicitly designated，the return type will be in．Be sure to gil Function Prototype Examples

f／l returns void，has one parameter
Int returns void，has one paramete
BadFunction
paraneters
on Definition Examples：
ong FindArea（1ong 1，long w
return 1 ＊w；
void PrintMessage（int whichmsg）
（



## Execution of Functions

 section laeter in this day）．
Local Variables
Not only can you pass in variables tothe function，but you also can declare variables within the body of the
function．This is sone using local variables，so named because they exis only locally within the function itel
when areturs，the local varaibles are no olongear available．

## \＃include＜iostream．h＞

float Convert（float），
int main（）
float Temprer；
float Tempeel；

loat Convert（float Temprer）


Outrut
Here＇s the temperature in celsius： 100
Please enter the temperature in Fahrenheit： 32
Here＇s the temperature in celsius： 0
Please enter the temperature in Fahrenheit： 85

## Analysis

 On lines 6 and 7 two float variables are declared, one to hold the temperature in Fahrenheit and on hold the temperature in degrees Celsius. The user is prompted to enter a Fahrenheit temperature on line 9, that value is passed to the function Convert().Execution jumps to the first line of the function Convert() on line 19, where a local variable, also named Temp is declared. Note that this local variable is not the same as the variable TempCel on line 7. This variable exis only within the function Convert0. The value passed as a parameter, TempFer, is also just a local copy of the variable passed in by maino

This function could have named the parameter FerTemp and the local variable CelTemp, and the program would work equally well. You can enter these names again and recompile the program to see this work.

The local function variable TempCel is assigned the value that results from subtracting 32 from the parameter TempFer, multiplying by 5 , and then dividing by 9 . This value is then returned as the return value of the func and on line 11 it is assigned to the variable TempCel in the maino function. The value is printed on line 13 The program is run three different times. The first time, the value 212 is passed in to ensure that the boiling poin
of water in degrees Fahrenheit ( 212 ) generates the correct answer in degrees Celsius (100). The second test is the freezing point of water. The third test is a random number chosen to generate a fractional result.
As an exercise, try entering the program again with other variable names as illustrated here

```
#include <iostream.h>
```

int main()
int main()
float TempFer;
cout << "Please enter the temperature in Fahrenheit: "
cin >> TempFer;
TempCel $=$ Convert (TempFer)
out << "\nHere's the temperature in Celsius:";
cout << TempCel << endl;
return 0;
float Convert (float Fer)
float Cel;
el $=($ (Fer -32$) * 5) / 9$
return Cel;
You should get the same results.

A variable has scope, which determines how long it is available to your program and where it can be accessed. Variables declared wien at beck out of existence" wher
your program.

## Normally scop loops on Day 7

None of this matters very much if you are careful not to reuse your variable names within any given function.

## Global Variables

Variables defined outside of any function have global scope and thus are available from any function in the program, including maino.

Local variables with the same name as global variables do not change the global variables. A local variable wit the same name as a global variable hides the global variable, however. If a function has a variable with the sam name as a global variable, the name refers to the local variable-not the global-when used within the function

## Input

Listing 5.3 Demonstrating Global and Local Variables

```
#include <iostream.h>
void myFunction(); // prototype
int x=5,y = 7;
// global variables
cout << "x from main:" << x << "\n";
    myFunction();
    cout << "Back from myFunction!\n\n";
    cout << "x from main: " << x << "\n"
    return 0;
void myFunction()
    int y = 10;
    cout << "x from myFunction:" << x << "\n";
```


## OUTput

$x$ from main: 5
y from main: 7
x from myFunction:
from myFunction: 10
Back from myFunction!
$x$ from main: 5

ANALYSIS This simple program illustrates a few key, and potentially confusing, points about local and globa variables. On line 4 , two global variables, $x$ and $y$, are declared. The global variable x is initialized with the
value 5 , and the global variable y is initialized with the value 7 .

On lines 8 and 9 in the function main $($, these values are printed to the screen. Note that the function main $($ defines neither variable; because they are global, they are already available to main 0 .

When myFunction 0 is called on line 10 , program execution passes to line 18 , and a local variable, $y$, is defined and initialized with the value 10 . On line 21 , myFunctiono prints the value of the variable $x$, and the global variable $x$ is used, just as it was in main). On line 22 , however, when the variable name $y$ is used, the local variable $y$ is used, hiding the global variable with the same name

The function call ends, and control returns to maino, which again prints the values in the global varibles. Note that the global variable $y$ was totally unaffected by the value assigned to myFunctiono's local y variable

## Global Variables: A Word of Caution

In $\mathrm{C}++$, global variables are legal, but they are almost never used. $\mathrm{C}++$ grew out of C , and in C global variables are a dangerous but necessary tool. They are necessary because there are times when the programmer needs to make data available to many functions, and he does not want to pass that data as a parameter from function to function is invisible to another function. This can and does create bugs that are very difficult to find.

## More on Local Variables

ariables declared within the function are said to have local scope. That means, as discussed, that they are visible and usable only within the function in which they are defined. In fact, in C++ you can define variable anywhere within the function, not just at its top. The scope of the variable is the block in which it is defined. Thus, if you define a variable inside a set of braces within the function, that variable is available only within that block. Listing 5.4 illustrates this idea

INPUT Listing 5.4 Variables Scoped Within a Block

```
// Listing 5.4 - demonstrates variables
// scoped within a block
#include <iostream.h>
void myFunc();
int main()
    int x = 5;
    cout << "\nIn main x is: " << x;
    myFunc();
    cout << "\nBack in main, x is: " << x << endl
        return 0;
void myFunc()
int x = 8;
    cout << "\nIn myFunc, local x:" << x << endl;
    {
        cout << "\nIn block in myFunc, x is: " << x;
        int x = 9;
        cout << "\nVery local x: " << x;
    cout << "\nOut of block, in myFunc, x:" << x << endl;
```


## Output

In main x is: 5
In myFunc, local $x$ :
In block in myFunc, $x$ is: 8
Very local x: 9
out block, in myFunc, $x$ : 8

ANAIYSIS This program begins with the initialization of a local variable, x , on line 10 , in main(). The printout on line 11 verifies that x was initialized with the value 5

MyFunc() is called, and a local variable, also named x , is initialized with the value 8 on line 22. Its value i printed on line 23 .

A block is started on line 25 , and the variable $x$ from the function is printed again on line 26. A new variable also named $x$, but local to the block, is created on line 28 and initialized with the value 9

The value of the newest variable $x$ is printed on line 30 . The local block ends on line 31 , and the variable crea on line 28 goes "out of scope" and is no longer visible
When x is printed on line 33 , it is the x that was declared on line 22 . This x was unaffected by the x that was defined on line 28 ; its value is still 8 .

On line 34, MyFunc() goes out of scope, and its local variable x becomes unavailable. Execution returns to line 15, and the value of the local variable $x$, which was created on line 10 , is printed. It was unaffected by either of the variables defined in

Needless to say, this program would be far less confusing if these three variables were given unique names!

## Function Statements

Virtually no limit exists to the number or types of statements that can be placed in the body of a function. Although you cannot define another function from within a function, you can call a function, and of course, maino does just that in nearly every C++ program. Functions can even call themselves, which is discussed in th section on recursion.

Although no limit exists to the size of a function in C++, well-designed functions tend to be small. Many programmers advise keeping your functions short enough to fit on a single screen so that you can see the enture function at one time. This is a rule of thumb, often broken by very good programmers, but a smaller function is easier to understand and maintain.
Each function should carry out a single, easily understood task. If your functions start getting large, look for places where you can divide them into component tasks

## More About Function Arguments

Function arguments do not all have to be of the same type. It is perfectly reasonable to write a function that Function arguments do not all have to be of the same type. I

Any valid $\mathrm{C}++$ expression can be a function argument, including constants, mathematical and logical expressions, and other functions that return a value

## Using Functions as Parameters to Functions

Although it is legal for one function to take as a parameter a second function that returns a value, it can make for code that is hard to read and hard to debug.

## As alue. You could write vale

Answer $=($ double (triple $($ square $($ cube $($ myValue $))))$;
This statement takes a variable, myValue, and passes it as an argument to the function cube (), whose return valu is passed as an argument to the function square), whose return value is in turn passed to triple(), and that return value is passed to double). The return value of this doubled, tripled, squared, and cubed number is now passed Answe
It is difficult to be certain what this code does (was the value tripled before or after it was squared?), and if the answer is wrong, it will be hard to figure out which function failed.

An alternative is to assign each step to its own intermediate variable

## Parameters Are Local Variables

The arguments passed into the function are local ot the finction. Changes made to the arguments do not affec
the values in the calling function. This is hown as pasing by value, which means alocal copy of ach
 INPUT Listing 5.5 A Demonstration of Passing by Value
// Listing 5.5 - demonstrates passing by value \#include <iostream.h>
void swap (int x, int y);
int main()
int $\mathrm{x}=5, \mathrm{y}=10$,
cout << Main. Before swa
swap $x, y)$, Mi.
sout $<$ MMain. After swa
ut <<"Main
return 0 ;
void swap (int x, int y)
int temp;
cout << "Swap. Before swap, x : " << x <<" y:" << y << " $1 \mathrm{n} "$; temp $=x ;$
$x=y ;$
$y=t e r i m$ temp;

## More About Return Values

ons return a value or return void. Void is a signal to the compiler that no value will be returned.
To return a value from a function, write the keyword reurun followed by the value you want to return. The vall
might itseff be an expression that returns vavalue. For example: return $\begin{aligned} & \text { return }(x>5) ; \\ & \text { return (MyPunction () }\end{aligned}$,


 $\substack{\text { rogramex } \\ \text { execuled. }}$


```
INPUT Listing 5.6 A Demonstration of Multiple reum Satement:
```

// Listing 5.6 - demonstrates multiple return
/statements
\#include <iostream.h>
int Doubler (int AmountroDouble),
${ }^{\text {int }}$ main()
int result
int input;
cout 《"Enter
cin 》input;
cout $\ll$ "nnefore doubler is called....";
cout $\ll "$ "ninput: $"><$ input $\ll "$ doubled: $" \ll$ result

## esult $=$ Doubler (input) ;

cout << "nBack from Doubler....nn";
cout << "\ninput: $" \ll$ input $\ll "$. doubled: " << result
return 0;
if (oricinal <= 10000

The one exception to this rule is if the function prototype declares a default value for the parameter. A default ave is a value to use if none is supplied. The preceding declaration could be rewritten as long myFunction (int $\mathrm{x}=50$ );
This prototype says, "myFunctiono returns a long and takes an integer parameter. If an argument is not supplied the default value of 50 ." Because parameter names are not required in function prototypes, this declaration could have been written a
long myFunction (int $=50$ );
The function definition is not changed by declaring a default parameter. The function definition header for this function would be

## long myFunction (int x

If the calling function did not include a argument, the compiler would fill $x$ with the default value of 50 . The name of the default parameter in the prototype need not be the same as the name in the function header; the default value is assigned by position, not name.
Any or all of the function's parameters can be assigned default values. The one restriction is this: If any of the parameters does not have a default value, no previous parameter may have a default value.
f the function prototype looks lik
long myFunction (int Param1, int Param2, int Param3);
you can assign a default value to Param2 only if you have assigned a default value to Param3. You can assign a default value to Param1 only if you assigned default values to both Param2 and Param3. Listing 5.7 demonstrate the use of default values.

```
// of default parameter values
```

\#include <iostream.h>
int VolumeCube(int length, int width $=25$, int height $=1$ );
int main()
int length $=100$;
int width $=50$
int volume;
volume = VolumeCube (length, width, height);
cout << "First volume equals: " << volume << " $\backslash n "$
volume $=$ VolumeCube (length, width)
cout << "Second

volume = VolumeCube (length)
out << "Th time volume equals: " << volume << "\n"
return 0;
volumeCube (int length, int width, int height)
return (length * width * height);

## Output

First volume equals: 10000 Second time volume equals: 5000

ANalysIS On line 6 , the VolumeCubeo prototype specifi rameters. The last two have default values.

This function computes the volume of the cube whose dimensions are passed in. If no width is passed in, a This function computes the volume of the cube whose dimensions are passed in. If no width is passed in, a
width of 25 is used and a height of 1 is used. If the width but not the height is passed in, a height of 1 is used is not possible to pass in the height without passing in a width.
On lines 10 to 12 , the dimensions length, height, and width are initialized, and they are passed to the VolumeCube function on line 15. The values are computed, and the result is printed on line 1

Execution returns to line 18 , where VolumeCube() is called again, but with no value for height. The default value is used, and again the dimensions are computed and printed.

Execution returns to line 21 , and this time neither the width nor the height is passed in. Execution branches for third time to line 27. The default values are used. The volume is computed and then printed

| Do | Don't |
| :---: | :---: |
| DO remember that function parameters act as local variables within the function. | DON'T try to create a default value for a first parameter if no default value exists for the second. This also applies to the second and other additional parameters: Do not give a parameter a default value if the parameter to the right does not have one. DON'T forget that arguments passed by value cannot affect the variables in the calling function. DON'T forget that changes to a global variable in one function change that variable for all functions. |

## Overloading Function

$\mathrm{C}_{++}$enables you to create more than one function with the same name. This is called function overloading. Th ctions must differ in their parameter list with a different type of parameter, a different number of parameter or both. Here is one example:
int myFunction (int, int);
int myFunction (long, long)
int myFunction (long,
(long) ;
myFunction() is overloaded with three parameter lists. The first and second versions differ in the types of the parameters, and the third differs in the number of parameters

The return types can be the same or different on overloaded functions. You should note that two functions wit the same name and parameter list, but different return types, generate a compiler error.

## Function overloading is also called fun polymorphic function is many-formed.

Function polymorphism refers to the capability to "overload" a function with more than one meaning. By changing the number or type of the parameters, you can give two or more functions the same function name, and the right one will be called by matching the parameters used. This enables you to create a function that ca average integers, doubles, and other values without having to create individual names for each function, such a geins 0 , AverageDoubles 0 , and so on.
Suppose you write a function that doubles whatever input you give it. You would like to be able to pass in a int a long, a float, or a double. Without function overloading, you would have to create four function names;
int DoubleInt(int); long DoubleLong (long); doat DoubleFloat (float);
double DoubleDouble (double)
With function overloading, you make this declaration:
int Double(int);
/// Listing 5.8 - demonstrates
include <iostream.h>
int Double (int);
long o oubuel (ong);
float Double ( float,
float Double (float);
double Double (double)


int
ing
float $\begin{aligned} & \text { doubledrnt; } \\ & \text { doubleding; } \\ & \text { doubleafloas }\end{aligned}$
$\begin{array}{ll}\text { 1ong } & \text { doubledong; } \\ \text { float } \\ \text { double } & \text { doubledr) ati } \\ \text { doubledpouble; }\end{array}$

doublearnt $=$ Double (myInt);
doubleadiong $=$ D ouble (myLIong)
doubledFloat $=$ Double (nyFloat);
doubleadoouble $=$ Double (nypouble),


return 0;
cout $\ll$ "In Double (int
return $2 *$ oriqinal
Iong Double (long original)
cout << "In Double (1ong) $\backslash \mathrm{n}^{\prime} ;$
return $2 \star$ original $;$

Double (float) ${ }^{n}$
$\underset{\substack{\text { cout << "In Double (f) } \\ \text { return } 2 * \\ *}}{\text { original }}$
double Double (double original)
cout << "In Double (double) $\backslash \mathrm{n}^{\prime}$ ";
return $2 *$ origina $; ~$

## Special Topics About Functions

 Function recursion is one of those wonderful, esoteric bit
cut through a thomy problem not asili solved otherwise.

When you define a function. normally the compiler creates just one set of instructions in memory. When you
call the function execution of the rergram iumps to those instructions and when the function rem jump back to the next line in inte calling function. If you call the function 10 times, ,your propram iumps to 1




 in a big gain More important, the increased size brings is own performance cost.
What is the rull of thumb? ? f yo have a small function-one or two statements-itits a candidate for inine.
When in doubt, though, leave it out. Listing 5.5 demonstrates an ininine tunction.
INPUT Listing 5.9 Demonstrates an Inline Functio
demonstrates inline functions
include <iostream.h>
line int Double(int)
main()

target = Double (target);
cout << "Targee: $"$ << target << end1,
target = Double (target);
cout <<"Target: $" \ll$ target $\ll$ endi,
target $=$ Double (target);
cout << "Target: $" \ll$ target << endl;
return $0 ;$
int Double (int target)
return $2 *$ target; later is quite funny (when it happens to someone else).

It is important to note that when a function calls itself, a new copy of that function is run. The local variables in the second version are independent of the local variables in the first, and they cannot affect one another directly any more than the local variables in main 0 can affect the local variables in any function it calls, as wa: illustrated in Listing 5.4

To illustrate solving a problem using recursion, consider the Fibonacci series:
1,1,2,3,5,8,13,21,34...
Each number, after the second, is the sum of the two numbers before it. A Fibonacci problem might be to dermine what the twelfth number in the series is

One way to solve this problem is to examine the series carefully. The first two numbers are 1. Each subsequen number is the sum of the previous two numbers. Thus the seventh number is the sum of the sixth and fifth numbers. More generally, the nth number is the sum of $n-2$ and $n-1$, as long as $n>2$
Recursive functions need a stop condition. Something must happen to cause the program to stop recursing, or i will never end. In the Fibonacci series, $\mathrm{n}<3$ is a stop condition.

The algorithm to use is the following:

1. Ask the user for a position in the series
2. Call the fibo function with that position, passing in the value the user entered.
3. The fib) function examines the argument ( $n$ ). If $\mathrm{n}<3$, it returns 1 ; otherwise, fibo calls itsel
(recursively) passing in $\mathrm{n}-2$, calls itself again passing in $\mathrm{n}-1$, and returns the sum
( you call fib(1), it returns 1. If you call fib(2), it returns 1. If you call fib(3), it returns the sum of calling fib(2) and fib(1). Because fib(2) returns 1 and fib(1) returns 1 , fib(3) will return 2 .
If you call fib(4), it returns the sum of calling fib(3) and fib(2). We already established that fib(3) returns 2 (by calling fib(2) and fib(1)) and that fib(2) returns 1, so fib(4) will sum these numbers and return 3, which is the four number in the series.

Taking this one more step, if you call fib(5), it will return the sum of fib(4) and fib(3). We established that fib(4) returns 3 and fib(3) returns 2 , so the sum returned will be 5
This method is not the most efficient way to solve this problem (in fib(20) the fibo function is called 13,529 times!!, but it does work. Be careful-if you feed in too large a number, you will run out of memory. Every
time $f$ bbo is called, memory is set aside. When it returns, memory is freed With recursion memory time fib( is called, memory is set aside. When it returns, memory is freed. With recursion, memory continues
be set aside before it is freed, and this system can eat memory very quickly. Listing 5.10 implements the fib( function

Caution: When you run Listing 5.10 , use a small number (less than 15 ). Because this
each call of the function, it produces a lot of output and can consume a lot of memory.
INPUT Listing 5.10 Demonstrates Recursion Using the Fibonacci Series

```
#include <iostream.h>
int fib (int n);
int main()
int n, answer;
in>> n;
    cout << "\n\n";
    answer = fib(n);
    cout << answer << " is the " << n << "th Fibonacci number\n";
    return 0;
t fib (int n)
    cout << MPro
        cout << "Return 1!\n"
        return (1);
    else
    cout << "Call fib(" << n-2 <<") and fib(" << n-1 << ").\n";
    return( fib(n-2) + fib(n-1));
```


## OUTrut



Note: Some compilers (other than GNU have difficulty with the use of operators in a cout statement. If you
receive a warning on line 28 , place parentheses around the subtraction operation so that line 28 becomes


The program asks for a number to find on line 9 and assigns that number to n . It then calls fibo with n . Executi banches to the fib() function, where, on line 20 , it prints its argument.

The argument n is tested to see whether it is less than 3 on line 21 ;
returns the sum of the values returned by calling fib() on $\mathrm{n}-2$ and $\mathrm{n}-1$
It cannot return these values until the call (to fib()) is resolved. Thus, you can picture the program diving into fii repeatedly until it hits a call to fib, which returns a value. The only calls that return a value immediately are the calls to fib(2) and fib(1). These are then passed up to the waiting callers, which, in turn, add the return value to their own, and then they return. Figures 5.4 and 5.5 illustrate this recursion into fib)

That means that a call is made to fib(4) [because $n=6$, fib( $n-2$ ) is the same as fib(4)] and another call is made to fib(5) [ fib(n-1)], and then the function you are in [fib(6)] waits until these calls return a value. When these retur a value, this function can return the result of summing those two values

Because fib(5) passes in an argument that is not less than 3 , fib() will be called again, this time with 4 and 3 . fib( will in turn call fib(3) and fib(2)

The output traces these calls and the return values. Compile, link, and run this program, entering first 1 , then 2 then 3 , building up to 6 , and watch the output carefully
This would be a great time to start experimenting with your debugger. Put a break point on line 20 and then trace into each call to fib, keeping track of the value of $n$ as you work your way into each recursive call to fib.
Recursion is not used often in C++ programming, but it can be a powerful and elegant tool for certain needs.

> Note: Recursion is a tricky part of advanced programming. It is presented here because it can be useful to understand the fundamentals of how it works, but do not worry too much if you do not fully understand all ti details.

## How Functions Work—A Look Under the Hood

When you call a function, the code branches to the called function, parameters are passed in, and the body of function is executed. When the function completes, a value is returned (unless the function returns void), and control returns to the calling function.

How this task accomplished. How does the code know where to branch? Where are the variables kept when they are passed in? What happens to variables that are declared in the body of the function? How is the return value passed back out? How does the code know where to resume?
introductory books do not try to answer these questions, but without understanding this information, you will find that programming remains a fuzzy mystery. The explanation requires a brief tangent into a discussio of computer memory

## Levels of Abstraction

One of the principal hurdles for new programmers is grappling with the many layers of intellectual abstraction Computers, of course, are only electronic machines. They do not know about windows and menus, they do n know about programs or instructions, and they do not even know about ones and zeros. All that is really going on is that voltage is being measured at various places on an integrated circuit. Even this is an abstraction: Electricity itself is just an intellectual concept representing the behavior of subatomic particles

Few programmers bother with any level of detail below the idea of values in RAM. After all, you do not need understand particle physics to drive a car, make toast, or hit a baseball, and you do not need to understand the
und understand particle physics to drive a car, make toast, or hit a baseball, and you do not need to understand th electronics of a computer to program one

You do need to understand how memory is organized, however. Without a reasonably strong mental picture of where your variables are when they are created and how values are passed among functions, it will all remain unmanageable mystery

## Partitioning RAM

When you begin your program, Linux sets up various areas of memory based on the requirements of the compiled program. As a C++ programmer, you will often be concerned with the global name space, the free store, the registers, the code space, and the stack

Global variables are in global name space. We will talk more about global name space and the free store ir coming days, but for now we will focus on the registers, code space, and stack

Registers are a special area of memory built right into the Central Processing Unit (or CPU). They take care of internal housekeeping. A lot of what goes on in the registers is beyond the scope of this book, but what we are encerned with is the set of registers responsible for pointing, at any given moment, to the next line of code. W.
 line of code is to be executed next.
code itself is in code space, which is that pat of memory set aside to hold the binary form of the instructions you created in your program. Each line of source code is translated into a series of instructions, and each of these instructions is at a particular address in memory. The instruction pointer has the address of the next instruction to execute. Figure 5.6 illustrates this idea.


Figure 5.6 The instruction pointer.
The stack is a special area of memory allocated for your program to hold the data required by each of the functions in your program. It is called a stack because it is a last-in, first-out queue, much like a stack of dish at a cafeteria, as shown in Figure 5.7.
 like a line at a theater: The first one in line is the first one off. A stack is more like a stack of coins: If you stacl 10 pennies on a tabletop and then take some back, the last three you put on will be the first three you take off.

When data is pushed onto the stack, the stack grows; as data is popped off the stack, the stack shrinks. It is not possible to pop a dish off the stack without first popping off all the dishes placed on after that dish

A stack of dishes is the common analogy. It is fine as far as it goes, but it is wrong in a fundamental way. A more accurate mental picture is of a series of cubbyholes aligned top to bottom. The top of the stack is whatev cubby the stack pointer (which is another register) happens to be pointing to.

Each of the cubbies has a sequential address, and one of those addresses is kept in the stack pointer register above the top of the stack is considered to be off the stack and invalid. Figure 5.8 illustrates this idea
$\square$
Figure 5.8 The stack pointe

## The Stack and Functions

## The following is what happens when a program, running on most systems, branches to a function:

1. The address in the instruction pointer is incremented to the next instruction past the function call. That address is then placed on the stack, and it will be the return address when the function returns. 2. Room is made on the stack for the return type you declared. On a system with 2 -byte integers, if the
return type is declared to be int, another 2 bytes are added to the stack, but no value is placed in these bytes.
2. The
3. The address of the called function, which is kept in a special area of memory set aside for that
purpose, is loaded into the in purpose, is
4. The current top of the stack is now noted and is held in a special pointer called the stack frame.
Everything added to the stack from now until the function returns will be considered "local" to the
function. function.
5. The instruction now in the instruction pointer is executed, thus executing the first instruction in the function.
6. Local variables are pushed onto the stack as they are defined.

When a function is ready to return, the following occurs:

1. The return value is placed in the area of the stack reserved at step 2 above
2. The stack is then popped all the way up to the stack frame pointer, which effectively throws away all the local variables and the arguments to the function.
3. The return value is popped off the stack and assigned as the value of the function call itself
4. The current top of the stack is now noted and is held in
5. The current top of the stack is now noted and is held in a special pointer called the stack frame.
Everything added to the stack from now until the function returns will be consided "local" to the Everything added to the stack from now until the function returns will be considered "local" to the function.
6. The a 5. The address stashed away in step 1 of the call is retrieved and put into the instruction pointer. The
program thus resumes immediately after the function call, with the value of the function retrieved.保 Some of the details of this process change from compiler to compiler, or between computers, but the esse
ideas are consistent across environments. In general, when you call a function, the return address and the ideas are consistent across environments. In general, when you call a function, the return address and the parameters are put on the stack. During the life of the function, local variables are added to the stack. When th function returns, these are all removed by popping the stac

## the life of the function

Multiple Source File Programs (Programmer-Created Function Libraries)
One of the greatest strengths of C and $\mathrm{C}++$ is the ability to reuse code. Now all languages support this ability.
However for many of them the mechanism is one of cut and paste (also known as cloning): when you want to However, for many of them, the mechanism is one of cut and paste (also known as cloning): when you want to use existing code, you copy the source into the new program and go from there. The problem with this approach
is the maintenance nightmares. When the business rule embodied in the code changes (and they all will at some time), all copies of the code have to change. But each version is a little different and has to be carefully reviewed before it can be changed.

C and $\mathrm{C}++$ provide a better mechanism for code reuse: function and object libraries. These are collections of code that is already compled only thjer (.o or .obe). Wh proper assiphe, there is one versio only that all programm
recompiled or relinked

In addition to the object file, you should create a header file that contains the prototypes for those functions.
Linux also supports a "library" file in the form of an archive file. An archive file contains multiple object files The linker knows how to read files in plain object format or extract them from an archive file. The ar comman is used to manipulate archives. Check the manual page or info file for more information.

In addition, you can create project files (known as make files, after the command that uses them-make). Thes project files control compilation and linking based on rules that you create. In the simplest form, your program will recompile and link if any of its parts change-main source, header files, and libraries.
The next section shows you how to create and use simple libraries of functions; the section after that introduces the GNU make command (gmake)

## How to Create and Use Function Libraries with g+-

A simple function library consists of three parts: function source (the functions themselves), function prototypes in a header file, and the compiled functions.
You can break the function out of the Fibonacci series program shown in Listing 5.10. Listing 5.11 contains header file and should be saved as fib

INPUT Listing 5.11 Heading file (fib.h) to Demonstrate Function Libraries

```
1: // fib.h header fil
2:#ifndef -FIB_H
3:#define -FIB_H
5:#endif
l
Input
```

```
#include <iostream.h>
```

\#include <iostream.h>
Listing 5.12
Listing 5.12
int fib (int n)
int fib (int n)
cout << "Processing fib(" << n << ")..."; 23
cout << "Processing fib(" << n << ")..."; 23
if ( }\textrm{n}<3
if ( }\textrm{n}<3
cout << "Return
cout << "Return
}
}
cout << "Call fib(" << n-2 <<") and fib(" << n-1 << ").\n";
cout << "Call fib(" << n-2 <<") and fib(" << n-1 << ").\n";
return(fib(n-2) + fib(n-1));
return(fib(n-2) + fib(n-1));
Listing 5.13 contains the main program that calls the fibo function. This is the same as the program in Listing .

``` INPUT. Listing 5.13 Demonstrates Recursion Using the Fibonacci Series-Using Libraries
```

1: \#include <iostream.h>
2: Hincluae fib.h
int main()
int n, answer;
cout << "Enter number to find:";
cin >> n; "\n\n";
answer = fib(n);
cout << answer << " is the " << n << "th Fibonacci number\n";
return 0;

```
17:

\section*{Or you can use two separate commands:}
g++ -c lst05-12.c++
g++ lst05-13.c++ lst05-12.0 -o lst05-13
The advantage of the second method is that you can compile 1st05-12.c++ into an object file once and provide the object and header file to other programmers. That allows the programmers to use your function without being able to change it (because they do not have the source code).s

\section*{Creating Project Files (for make)}

In the simplest form, the make command under Linux allows you to set up a project and let those rules control what gets compiled when. Only when needed (when something changes) will code actually be changed.

These rules are stored in a file known as a makefile (which is also the default name for the file), although you can give it any name you want

The general format of a makefile consists of two different kinds of lines. The first is the target/dependency, which defines the target to be made along with any files that it depends on. The other lines are the commands required to produce that target. The makefile for Listings 5.11 through 5.13 is shown in Listing 5.14.

\section*{INPUT.}

Listing 5.14 Makefile for Use with Listings 5.11 through 5.13 (1st05-14.mak)
```

1: lst05-12.o: lst05-12.c++ fib.h
2: g++ -c lst05-12.c++
3: \#
4: lst05-13: lst05-13.c++ lst05-12.0 fib.h
5: g++ lst05-13.c++ lst05-12.0 -o lst05-13

```

The first line shows that 1st05-12.o (the function library, our target) depends on 1st05-12.c++ and fib.h. Line 2 is the command to produce that target. Line 3 is a comment.

The line 4 shows that 1st05-13 (the executable, our target of this step) depends on 1st05-13.c++, lib05-12.o, and fib.h. Line 5 is the command to produce that target.
\[
\begin{aligned}
& \text { Note: Lines } 2 \text { and } 5 \text { have a tab character before the g++ command. This must be a tab character. if you use any } \\
& \text { other character here (including spaces), you will get the following error message from make: } \\
& \text { yourfile.mak:2: *** missing separator. Stop. }
\end{aligned}
\]

When fib.h changes and you try to make the executable using the following command, it builds all the parts needed:
make -f lst05-14.mak lst05-13

\section*{Output}
g++ -c lst05-12.c++
gcc: -lgpp: linker input file unused since linking not done gcc: -lstdcx: linker input file unused since linking not done gcc: -lm: linker input file unused since linking not done g++ lst05-13.c++ lst05-12.0 -o lst05-13

\section*{Analysis} added an extra space at the end of a line) before issuing the make command. This command is smart enough t analyze the rules you give it to determine whether there are any other dependencies. It determined that because 1st05-13 needs 1st05-12.o and fib.h changed since the last time 1st05-12.o was created, 1st05-12.c++ needed to be recompiled before 1st05-13.c++

The make command is much more sophisticated than this simple example shows. Take a look at the info page for more information; make supports a programming language of its own

You can create simple project files like this one for your programs. Then you do not need to remember what the dependencies are.

\section*{Standard C++ Library Functions (libg++)}

All compilers that subscribe to the ANSI standard come with a series of functions for you to use. These functions are said to be built-in because they come with the compiler; you do not write them yourself. The GNI compiler is no exception. There is a standard library of functions for C programs and a superset library of functions and objects for C++ programmers.

You will learn about many of the standard objects in later days. You can always use the info command to get more information:
info libg++
You can get details on the C standard functions through the info command
info libc
The standard functions for C and \(\mathrm{C}++\) are very important. You can use them instead of writing your own. This makes program creation easier (less to do) and greatly simplifies maintenance because the compiler vendors (GNU Project contributors/Free Software Foundation) worry about maintaining those functions. They are categorized in the following sections by Math, Character, and General functions.

\section*{Math Functions}

The standard library math functions allow you to perform common (and some not-so-common) mathematical operations. This allows you to worry about how to use the results instead of worrying about how to calculate the results. For example, there is a function to calculate the square root of a positive number. As a programmer, you do not need to know how to code this-another programmer has done so. You just need to know what to do with the results (or why you would want to get the square root of a number).

Some of the most commonly used math functions are shown in Table 5.1.
Table 5.1Common Mathematical Functions
\begin{tabular}{|c|c|}
\hline Function & Usage \\
\hline ceil(f) & Rounds f to the smallest integer not less than \(\mathrm{f}: \operatorname{ceil}(3.3)\) is 4 , ceil( -3.3 ) is -3 . \\
\hline \(\cos (\mathrm{f})\) & Calculates the cosine of f (where f is in radians): \(\cos (0.0)\) is \(1, \cos (3.1415926)\) is 1. \\
\hline \(\exp (\mathrm{f})\) & Exponential function in the form of \(\mathrm{e}^{\mathrm{f}}: \exp (0.0)\) is \(1, \exp (1.0)\) is 2.718 (e). \\
\hline fabs(f) & Floating absolute value of \(\mathrm{f}: \mathrm{fabs}(-1.0)\) is \(1, \mathrm{fabs}(0.0)\) is 0 , fabs(1.0) is 1 . \\
\hline floor(f) & Rounds f to the largest integer that is not greater than f : floor(3.3) is 3 , floor(3.3 ) is -4 ; compare with the ceil() function. \\
\hline fmod (f, g) & Calculates the remainder of \(\mathrm{f} / \mathrm{g}\) as a float. \\
\hline \(\log\) (f) & Natural logarithm of f using base \(\mathrm{e}: \log (1.0)\) is \(0.0, \log (2.718)\) is 1.0 . \\
\hline \(\log 10\) (f) & Base \(10 \log\) arithm of \(\mathrm{f}: \log 10(0.0)\) is undefined, \(\log 10(1.0)\) is \(0.0, \log 10(10.0) 1\), \(\log 10(1000.0)\) is 3 . \\
\hline pow (f, g) & Raises f to the g power ( fg ) : \(\operatorname{pow}(2.0,3.0)\) is 8.0 , pow ( \(3.0,2.0)\) is 9.0 . \\
\hline \(\sin (\mathrm{f})\) & Calculates the sine of f (where f is in radians): \(\sin (0.0)\) is 0 , and \(\sin (3.1415926\) / 2.0) is 1. \\
\hline sqrt (f) & Square root of f : sqrt (4.0) is 2 , sqrt (9.0) is 3.0. \\
\hline \(\tan\) (f) & Calculates the tangent of f (where f is in radians): \(\tan (0.0)\) is 0 , and \(\tan\) (3.1415926/4.0) is 1 . \\
\hline acos, asin, atan & Arc (anti or inverse) cosine, sine, and tangent. \\
\hline acosh, asinh, atanh & Arc (anti or inverse) hyperbolic cosine, sine, and tangent. \\
\hline cosh, sinh, tanh & Hyperbolic cosine, sine, and tangent. \\
\hline
\end{tabular}

For all these functions, you must include the math.h header file (in order to get the functional prototypes).

\section*{Character and String Functions}

Character and string manipulation is a common task in most programs. As a result, a large library of functions has built up over time and have become part of the standard.

As with the math functions, these functions exist to make your life easier!
Some of the most commonly used string and character functions are shown in Table 5.2.

Table 5.2Common String and Character Functions
\begin{tabular}{|c|c|}
\hline Function & Usage \\
\hline strcpy (s1, s2) & Copies s2 into s1 character by character until the null terminator is copied. \\
\hline strncpy (s1, s2, n) & Copies s2 into s1 character by character until the null terminator is copied or \(n\) characters are moved; \(s 1\) is not null terminated if the count terminated the copy. \\
\hline strcat (s1, s2) & Copies s2 onto the end of s1. \\
\hline strncat (s1, s2, n) & Copies, at most, n characters of s 2 onto the end of s 1 . \\
\hline strcmp (s1, s2) & Compares s1 with s2 character by character and returns the difference. If the strings are the same, 0 is returned; otherwise, a positive or negative value shows the difference between the two. \\
\hline strncmp (s1, s2, n) & Compares, at most, \(n\) characters of \(s 1\) with \(s 2\) character by character and returns the difference. If the strings are the same, 0 is returned; otherwise, a positive or negative value shows the difference between the two. \\
\hline strlen (s1) & Returns the length (in characters, not counting the null terminator) of s1. \\
\hline strlwr (s1) & Converts s1 to lowercase. \\
\hline strtok (s1, s2) & Breaks s1 into tokens based on values in s2. This parses s1 into pieces delimited by characters in s 2 . The first time strtok is called, s 2 is provided; each subsequent time, use NULL instead of s 2 . \\
\hline strtod (s1), strtol(s1) & Converts s1 into an integer or long numeric value. \\
\hline strupr(s1) & Converts s1 to uppercase. \\
\hline
\end{tabular}

For all these functions, you must include the string.h header file (in order to get the functional prototypes).

\section*{General Functions}

There are many other functions that are part of the standard library. Some of the more common functions that are not math- or character-based are shown in Table 5.3.

Table 5.3Other Commonly Used Functions

Function

\section*{Usage}
system (s1)
malloc (n)
free (ptr)
realloc(ptr, n)
abs (n)
asctime ( t )
bsearch ()
qsort ()
isalpha (c)
isdigit (c)
islower (c)
isupper (c)

Issues s1 as a command to the operating system or command shell.
Allocates \(n\) bytes of memory and returns a pointer to that address. Returns NULL on failure.
Frees memory allocated via malloc or realloc.
Allocates n bytes growing or shrinking space previously allocated (malloc or realloc). Allows implementation of "dynamic" arrays.
Returns integer absolute value of n . abs ( \((-1)\) is 1 , abs ( 0 ) is 0 , and abs ( 1 ) is 1 .
Converts internal format time \(t\) to printable time.
Performs binary search on sorted data structure.
Sorts data structure.
Returns true if character c is alphabetic.
Returns true if character c is numeric digit.
Returns true if character c is lowercase alphabetic.
Returns true if character c is uppercase alphabetic.

There are many, many, many more functions available in the standard library. Check the info or manual pages for more information about these and other functions. Some of the functions are very complex (such as bsearch and qsort) and require much more explanation, which can be best found in the documentation.

\section*{Much More}

Remember, code reuse is one of the basics of C and \(\mathrm{C}++\). It is also one of the basics of software engineering. Any code that you can reuse is code that does not have to be written new or maintained. It can be much more robust and tested more strenuously than that written by individual programmers. And if changes are required, they can be changed one place (in the library) instead of having to track down all the code.

Check the standard libraries ( C and \(\mathrm{C}++\) ) before writing functions of your own-one that solves your problem may already exist. It may take several standard library functions to perform your desired task, but it will save you time and effort.

I remember one project during which I was out of town attending a conference. There was a junior consultant on the project. He was a good coder and understood C well, but he was not all that familiar with the library. He needed to sort some data within his program (not a lot, so the sort algorithm was not important). Because very few people memorize how to write sort routines, he had to dig up his notes from college, code, test, and tweak the code. This took him several hours over a couple of days. I returned from the conference, and on hearing about his coding, pulled down the volume with the standard library, turned to the qsort page, and showed it to him.

In an hour or two, in the same day, using one standard function (and a special function to determine data order), he could have had his sort. Instead he took longer, had to work harder, and created more code to maintain. I used the masculine gender in this description because the person I worked with was male. The same mistake can be made no matter what your gender.

Remember this lesson!

This chapter introduced functions. A function is, in effect, a subprogram into which you can pass parameters and from which you can return a value. Every C++ program starts in the main() function, and main(), in turn, can call other functions.

A function is declared with a function prototype, which describes the return value, the function name, and its parameter types. A function can optionally be declared inline. A function prototype can also declare default variables for one or more of the parameters.

The function definition must match the function prototype in return type, name, and parameter list. Function names can be overloaded by changing the number or type of parameters; the compiler finds the right function based on the argument list.

Local function variables, and the arguments passed into the function, are local to the block in which they are declared. Parameters passed by value are copies and cannot affect the value of variables in the calling function

\section*{Q Why not make all variables global?}

A At one time, this was exactly how programming was done. As programs became more complex, however, it became very difficult to find bugs in programs because data could be corrupted by any of the functions-global data can be changed anywhere in the program. Years of experience have convinced programmers that data should be kept as local as possible, and access to changing that data should be narrowly defined.
Q When should the keyword inline be used in a function prototype?
A If the function is very small, no more than a line or two, and won't be called from many places in your program, it is a candidate for inlining.

\section*{Q Why aren't changes to the value of function arguments reflected in the calling function?}

A Arguments passed to a function are passed by value. That means that the argument in the function is actually a copy of the original value. This concept is explained in depth in the section "How Functions Work-A Look Under the Hood.'
Q If arguments are passed by value, what do I do if I need to reflect the changes back in the calling function?
A On Day 8, pointers will be discussed. Use of pointers will solve this problem, as well as provide a way around the limitation of returning only a single value from a function.
\(\mathbf{Q}\) What happens if \(\mathbf{I}\) have the following two functions:
int Area (int width, int length = 1); int Area (int size);
Will these overload? A different number of parameters exist, but the first one has a default value.
A The declarations will compile, but if you invoke Area with one parameter, you will receive a compiletime error: ambiguity between Area(int, int) and Area(int).

\section*{Workshop}

The Workshop provides quiz questions to help you solidify your understanding of the material covered and exercises to provide you with experience in using what you learned today. Try to answer the quiz and exercise questions before checking the answers in Appendix D, "Answers to Exercises and Quizzes," and make sure tha you understand the answers before continuing to the next day.

\section*{Quiz}
1. What are the differences between the function prototype and the function definition?
2. Do the names of parameters have to agree in the prototype, definition, and call to the function?
3. If a function does not return a value, how do you declare the function?
4. If you do not declare a return value, what type of return value is assumed?
5. What is a local variable ?
6. What is scope?
7. What is recursion?
8. When should you use global variables?
9. What is function overloading?
10. What is polymorphism?

\section*{Exercises}
1. Write the prototype for a function named Perimeter(), which returns an unsigned long int and takes two parameters, both unsigned short ints.
2. Write the definition of the function Perimeter() as described in Exercise 1. The two parameters represent the length and width of a rectangle. Have the function return the perimeter (twice the length plus twice the width).
3. BUG BUSTER: What is wrong with the function in the following code?
\#include <iostream.h>
void myFunc(unsigned short int \(x\) )
int main()
\{
unsigned short int \(x, y\);
\(y=\) myFunc (int);
cout << "x: " << x << " y: " << y << " \(\backslash n "\);
\}
void myFunc(unsigned short int \(x\) )
\{
return ( \(4{ }^{*} \mathrm{x}\) ) ;
\}
4. BUG BUSTER: What is wrong with the function in the following code?
\#include <iostream.h>
int myFunc(unsigned short int \(x\) );
int main()
\{
unsigned short int \(x, y\);
\(y=\) myFunc (x);
cout << "x: " << x << " y: " << y << "\n".
cout
\(\}\)
int myFunc(unsigned short int \(x\) );
\({ }_{i}^{\text {int }}\)
return (4*x);
\}
5. Write a function that takes two unsigned short integer arguments and returns the result of dividing the first by the second. Do not do the division if the second number is zero, but do return -1 .
6. Write a program that asks the user for two numbers and calls the function you wrote in Exercise 5. Print the answer, or print an error message if you get -1 .
7. Write a program that asks for a number and a power. Write a recursive function that takes the number to the power. Thus, if the number is 2 and the power is 4 , the function will return 16 .

\section*{Chapter 6}

\section*{Basic Classes}

Classes extend the built-in capabilities of \(\mathrm{C}++\) to assist you in representing and solving complex, real-world problems. Today you will learn:
- What classes and objects are
- How to define a new class and create objects of that class
- What member functions and member data are
- What constructors are and how to use them

\section*{Creating New Types}

You have already learned about a number of variable types, including unsigned integers and characters. The type of a variable tells you quite a bit about it. For example, if you declare Height and Width to be unsigned shor integers, you know that each one can hold a number between 0 and 65,535 , assuming an unsigned short integer is 2 bytes. That is the meaning of saying they are unsigned integers; trying to hold anything else in these variables causes an error. You cannot store your name in an unsigned short integer, and you should not even try

Just by declaring these variables to be unsigned short integers, you know that it is possible to add Height to Widt and to assign that number to another variable.

The type of these variables tells you
- Their size in memory
- What information they can hold
- What actions can be performed on them

More generally, a type is a category. Familiar types include car, house, person, fruit, and shape. In \(\mathrm{C}++\), the programmer can create any type needed, and each of these new types can have all the functionality and power o the built-in types.

\section*{Why Create a New Type?}

Programs are usually written to solve real-world problems, such as keeping track of employee records or simulating the workings of a heating system. Although it is possible to solve complex problems by using programs written with only integers and characters, it is far easier to grapple with large complex problems if you can create representations of the objects that you are talking about.

In other words, simulating the workings of a heating system is easier if you can create variables that represent rooms, heat sensors, thermostats, and boilers. The more closely these variables correspond to reality, the easier it is to write the program.

\section*{Introducing Classes and Members}

You make a new type by declaring a class. A class is just a collection of variables-often of different types-combined with a set of related functions.

One way to think about a car is as a collection of wheels, doors, seats, windows, and so forth. Another way is th think about what a car can do: It can move, speed up, slow down, stop, park, and so on. A class enables you to encapsulate, or bundle, these various parts and various functions into one collection, which is called an object.

Encapsulating everything you know about a car into one class has a number of advantages for a programmer Everything is in one place, which makes it easy to refer to, copy, and manipulate the data. Likewise, clients of your class - that is, the parts of the program that use your class - can use your object without worrying about what is in it or how it works

A class can consist of any combination of the variable types and also other class types. The variables in the clas are referred to as the member variables or data members. A Car class might have member variables representin: the seats, radio type, tires, and so forth

Member variables, also known as data members, are the variables in your class. Member variables are part of your class, just as the wheels and engine are part of your car

The functions in the class typically manipulate the member variables. They are referred to as member functions or methods of the class. Methods of the Car class might include Start() and Brake(). A Cat class might have data members that represent age and weight; its methods might include Sleep(), Meow(), and ChaseMice()

Member functions, also known as methods, are the functions in your class. Member functions are as much a pa of your class as the member variables. They determine what your class can do.

\section*{Declaring a Class}

To declare a class, use the class keyword followed by an opening brace, and then list the data members and methods of that class. End the declaration with a closing brace and a semicolon. The declaration of a class called Cat looks like this:
```

class Cat

```
cl
\{
    unsigned int itsAge;
    unsigned int itsWeight;
    void Meow();
\};

Declaring this class does not allocate memory for a Cat. It just tells the compiler what a Cat is, what data it contains (itsAge and itsWeight), and what it can do (Meow()). It also tells the compiler how big a Cat is-that is, how much room the compiler must set aside for each cat that you create. In this example, if an integer is 4 byte a Cat is 8 bytes big: itsAge is 4 bytes, and itsWeight is another 4 bytes. Meow() takes up no room because no storage space is set aside for member functions (methods).

\section*{A Word on Naming Conventions}

As a programmer, you must name all your member variables, member functions, and classes. As you learned o Day 3, "Variables and Constants," these should be easily understood and meaningful names. Cat, Rectangle, and Employee are good class names. Meow(), ChaseMice(), and StopEngine() are good function names because they tell you what the functions do. Many programmers name the member variables with the prefix its, as in itsAge, itsWeight, and itsSpeed. This helps to distinguish member variables from nonmember variables.

C++ is case sensitive, and all class names should follow the same pattern. That way, you never have to check how to spell your class name; was it Rectangle, rectangle, or RECTANGLE? Some programmers like to prefi every class name with a particular letter-for example, cCat or cPerson-whereas others put the name in all uppercase or all lowercase. The convention used in this book is to name all classes with initial capitalization, a in Cat and Person

Similarly, many programmers begin all functions with capital letters and all variables with lowercase. Words are usually separated with an underscore-as in Chase Mice-or by capitalizing each word-for example, ChaseMice or DrawCircle

As mentioned a few days ago when you learned about variables, the important idea is that you should pick on style and stay with it through each program. Over time, your style will evolve to include not only namins conventions, but also indentation, alignment of braces, and commenting style.

Note: It is very common for development companies to have house standards for many style issues. This ensures that all developers can easily read one another's code

Even where there are not formal standards, groups develop their preferred methods.

\section*{Defining an Objec}

You define an object of your new type the same as you define an integer variable:
unsigned int GrossWeight;
// define an unsigned integer
// define a Cat

Cat Frisky;
This code defines a variable called GrossWeight, whose type is an unsigned integer. It also defines Frisky, which is an object whose class (or type) is Cat.

\section*{Comparing Classes and Object}

You never pet the definition of a cat; you pet individual cats. You draw a distinction between the idea of a cat and the particular cat who right now is shedding all over your living room. In the same way, C++ differentiate and the particular cat who right now is shedding all over your living room. In the same way, \(\mathrm{C}++\) differentiate between the class Cat, which is the idea of a cat, and each individual Cat object. Thus Frisky is an object of type别

An object is an individual instance of a class. The technical term for creating an object is instantiation
Note: Sometimes terminology gets reused and misused. The meaning of individual terms or which term should
be used where is often a matter of discussion and convention. As a result, the correct word today can be the
incorrect word tomorrow.
In this book, the words declare and define are used interchangably to mean you are specifying or defining
something. When you describe the members (data and functions) in an object class, you are declaring that class.
When working with normal variables (like int or float), there is no difference between the process of describibng and
creating. With object classes there is.
You describe the object class (define or declare) at one place in the program (often in a happ file) and then create a
specific object later on.
With object classes, when you actually allocate memory, you are creating a specific object. Some authors use the
word specify for this purpose.

\section*{Accessing Class Members}

After you define an actual Cat object-for example, Frisky-you use the dot operator (.) to access the members of that object. Therefore, to assign 50 to Frisky's Weight member variable, you would write

Frisky.itsWeight = 50;
In the same way, to call the Meow() function, you would write
Frisky.Meow();
When you use a class method, you call the method. In this example, you are calling Meow() on Frisky

\section*{Assign to Objects, Not to Classes}

In C++ you do not assign values to types; you assign values to variables. For example, you would never write
int \(=5 ; \quad / /\) wrong

The compiler would flag this as an error because you cannot assign 5 to an integer. Rather, you must define ar integer variable and assign 5 to that variable. For example
```

int x;

```

This is a shorthand way of saying, "Assign 5 to the variable x , which is of type int." In the same way, you woul not write

Cat.itsAge=5;
// wrong

The compiler would flag this as an error because you cannot assign 5 to the age part of a Cat. Rather, you mus define a Cat object and assign 5 to that object, for example
```

Cat Frisky;
Frisky.itsAge = 5; // just like int x;

```

\section*{If You Don't Declare It, Your Class Won't Have It}

Try this experiment: Walk up to a three-year-old and show her a cat. Then say, "This is Frisky. Frisky knows a trick. Frisky, bark." The child will giggle and say, "No, silly, cats can't bark."

\section*{If you wrote}
\begin{tabular}{ll} 
Cat Frisky; & // make a Cat named Frisky \\
Frisky.Bark() & // tell Frisky to bark
\end{tabular}
the compiler would say, "No, silly, Cats can't bark." Actually, the GNU compiler will tell you:
example.cxx: In function 'int main()':
example.cxx:25: no member function 'Cat::Bark()' define

The compiler knows that Frisky cannot bark because the Cat class does not have a Bark() function. The compile would not even let Frisky meow if you did not define a Meow() function
\begin{tabular}{|l|l|}
\hline Do & Don'T \\
\hline \begin{tabular}{l} 
DO use the keyword class to declare a class. \\
DO use the dot operator (.) to access class members \\
and functions.
\end{tabular} & \begin{tabular}{l} 
DON'T confuse a declaration with a definition. A \\
declaration says what a class is. A definition sets aside \\
memory for an object. \\
DON'T confuse a class with an object. \\
DON'T assign values to a class. Assign values to the \\
data members of an object.
\end{tabular} \\
\hline
\end{tabular}

Defining Scope as Private Versus Public

\title{
int itsAge;
int itsweight;
}
int main()
Cat Frisky;
Frisky-itsAg
Frisky.itsAge \(=5 ; \quad / /\) assign to the member variable
cout \(\ll\) "rrisky is a cat who is ";
cout << "Frisky is a cat who is ";
cout \(<\) Frisky.itsAge << " years old. \({ }^{\prime}\) n";
return 0;

\section*{Output}

ANAIYIS Ame of the new class comes after the keyword class. In this case, it is Cal

The body of the declaration begins with the opening brace in line 7 and ends with a closing brace and a The body of the declaration begins with the opening brace in line 7 and ends with a closing brace and a
semicolon in line 1 L. Line 8 contains the keyword public, which indicates that everything that follows is public
until the keyword private or the end of the class declaration.

Lines 9 and 10 contain the declarations of the class members itsAge and isweight.
Line 14 begins the main function of the program. Fisky is defined in line 16 as an instance of a Cat-that is, as
Cat object In Ine 17 . Firsy's ase is set to 5 In lines 18 and 19 . the istise member varible is used to print out Cat object. In line 17, Frisky's age is set to 5 . In lines 18 and 19 , the itsAge member variable is used to print out

Note: Try commenting out line 8 and try to recompile. You will receive an error on lines 17 and 19 because itsege
will on ologer have public access. The default tor classes is private access. The GNU compiler will tell you this as
follows.


\section*{Make Member Data Private}

As a general rule of design, you should keep the member data of a class private. Therefore, you must create
public funct public functions known as accessor methods to set and get the private member variables. These accessor
methods are the member functions that other parts of your program call to get and set your private member variables

A public accessor method is a class member function used either to read the value of a private class member
variable or to set its value.
Why bother with this extra level of indirect access? After all, it is simpler and easier to use the data instead of rking through aceessor functions.
Accessor functions enable you to separate the details of how the data is stored from how it is used. This enabl you to change how the data is stored without having to rewrite functions that use the data.

If a function that needs to know a Cat's age accesses itsAge directly, that function would need to be rewritten i
you, as the author of the Cat class, decided to change how that data is stored. By having the function call GetA eeo, your Cat class can easily return the right value no matter how you arive at the age The clling Getageo, your Cat class can easily return the right value no matter how you arrive at the age. The calling
function does not need to know whether you are storing it as an unsigned integer or a long, or whether you ar computing it as needed.
This technique makes your program easier to maintain. It gives your code a longer life because design chang do not make your program obso
hides the details from the caller.
Listing 6.2 shows the Cat class modified to include private member data and public accessor methods. Note th Listing 6.2 shows the Cat class
this is not an executable listing.

Input Listing 6.2 A Class with Accessor Methods
// Cat class declaration
// Data members are private, public accessor methods
// mediate setting and getting the values of the private data
lass Ca
public:
ansigned int GetAge ()
oid SetAge (unsigned int Age)
unsigned int GetWeight ()
void SetWeight (unsigned
public member functions
// private member data
\(\begin{array}{ll}\text { unsigned int } & \text { itsAge; } \\ \text { unsigned int } & \text { itsweight; }\end{array}\)

\section*{Distinguishing Privacy and Securit}

Declaring methods or data private enables the compiler to find programming mistakes before they become bugs. Any programmer worth his consulting fees can find a way around privacy if he wants to. Stroustrup, the inventor of \(\mathrm{C}+\mathrm{+}\),
said, "The C++ access control mechanisms provide protection against accident-not against fraud." (ARM, 1990)

\section*{The class Keyword}

Syntax for the class keyword is as follows
class class_name
// access control keywords here
// class variables and methods
// class variables and methods declared here
You use the class keyword to declare new types. A class is a collection of class member data, which are variables o various types, inclucing other classes. The class also containes class functions-or
used to manipulate the data in the class and to perform other services for the class.
\(\left\lvert\, \begin{aligned} & \text { You define objects of the new type in much the same way in which you define any variable. State the type (cla } \\ & \text { then the variable name (the object). You access the class members and functions by using the dot (.) operator. }\end{aligned}\right.\)
You use access control keywords to declare sections of the class as public or private. The defaul for access control
is private. Each keyword changes the access control from that point on to the end of the class or until the next acces. is private. Each keyword changes the access control from that point on to the e
control keyword. Class declarations end with a closing brace and a semicolon.
Example 1
class Cat
public:
unsigned int Age;
unsignt Weight
void
void Meow()
Cat Frisky;
Frisky.Age \(=8\);
Frisk. Weight
Frisky-Meow();
class Car
public:
void Start ();
void Accelerate
void Accelerat
void Brake();
void SetYear(int year)
int Getyear)
private:
int
int Year;
Char Model [255];
Car oldFaithfuli \(\quad / /\) end of class declaration
int bought;
OldFaithful. Set Year (84) ;
bought = OldFaithful.
OldFaithful.Start ();

\section*{Implementing Class Methods}

As you have seen, an accessor function provides a public interface to the private member data of the class. Eacl accessor function, along with any other class m
implementation is called the function definition.
```

// Demonstrates declaration of

```
```

\#include <iostream.h> // for

```
\begin{tabular}{ll} 
class Cat & // begin declaration of \\
public: & // begin public section \\
int GetAge (); & // accessor function \\
void SetAge (int age); & // accessor function \\
void Meow(); & // general function \\
private: \\
int itsAge; & begin private section
\end{tabular}
    public:
int GetAge
(
    void Meow();
    private:
int itsAge;
                begin private section
// GetAge, Public accessor function
returns value of itsAge member
    ,
definition of SetAge, public
    accessor function
void Cat: : Setage (int age)
    // set member variable itsAge to
// value passed in by parameter age
    itsAge \(=\) age;
definition of Meow method
    returns: void
// parameters: None
    cout << "Meow. \n";
create a cat, set its age, have it
// meow, tell us its age, then meow again
nt main()
    Cat Frisky;
Frisky.SetA
    Frisk. SetAge (
Frisky.Meow();
cout \(\ll\) "Frisky
    cout << "Frisky is a cat who is ".
cout \(\ll\) Frisky;
    cout << Frisky.GetAge() << " years old. \(\backslash n^{\prime \prime}\);
Frisky.Meow ()
    Frisky.Meow
return 0;

Lines 6-14 contain the definition of the Cat class. Line 8 contains the keyword public, which tells the compiler that what follows is a set of public members. Line 9 has the declaration of the public accessor method GetAge(). GetAge() provides access to the private member variable itsAge, which is declared in line 13. Line 10 GetAge(). GetAge() provides access to the private member variable itsAge, which is declared in line 13. Line 10
has the public accessor function SetAge(). SetAge() takes an integer as an argument and sets itsAge to the value of that argument

Line 11 has the declaration of the class method Meow(). Meow() is not an accessor function. Here it is a general method that prints the word Meow to the screen.

Line 12 begins the private section, which includes only the declaration in line 13 of the private member variable itsAge. The class declaration ends with a closing brace and semicolon in line 14

Lines 18-21 contain the definition of the member function GetAge(). This method takes no parameters; it returns an integer. Note that class methods include the class name followed by two colons and the function name (Line 18). This syntax tells the compiler that the GetAge() function that you are defining here is the one that you declared in the Cat class. With the exception of this header line, the GetAge() function is created the same as any other function.

The GetAge() function takes only one line; it returns the value in itsAge. Note that the main() function cannot access itsAge because itsAge is private to the Cat class. The main() function has access to the public method GetAge(). Because GetAge() is a member function of the Cat class, it has full access to the itsAge variable. This access enables GetAge() to return the value of itsAge to main().

Line 26 contains the definition of the SetAge() member function. It takes an integer parameter and sets the value of itsAge to the value of that parameter in line 30. Because it is a member of the Cat class, SetAge() has direct access to the member variable itsAge.

Line 37 begins the definition, or implementation, of the Meow() method of the Cat class. It is a one-line function that prints the word Meow to the screen, followed by a new line. Remember that the \(\backslash n\) character prints a new lin to the screen.

Line 44 begins the body of the program with the familiar main() function. In this case, it takes no arguments. In line 46, main() declares a Cat named Frisky. In line 47, the value 5 is assigned to the itsAge member variable by way of the SetAge() accessor method. Note that the method is called by using the object name (Frisky) followed by the member operator (.) and the method name (SetAge()). In this same way, you can call any of the othe methods in a class.

Line 48 calls the Meow() member function, lines 49 and 50 print a message using the GetAge() accessor. Line 51 calls Meow() again.

\section*{Understanding Constructors and Destructors}

Two ways exist to define an integer variable. You can define the variable and then assign a value to it later in the program, for example
```

int Weight; // define a variable
... // other code here
Weight = 7; // assign it a value

```

Or you can define the integer and immediately initialize it. For example,
int Weight = 7; // define and initialize to 7

Initialization combines the definition of the variable with its initial assignment. Nothing stops you from changing that value later. Initialization ensures that your variable is never without a meaningful value.

Classes have a special member function called a constructor, which initializes the member data of that class. The constructor can take parameters as needed, but it cannot have a return value-not even void. The constructor is a class method with the same name as the class itself.

Whenever you declare a constructor, you should also declare a destructor. Just as constructors create and initialize objects of your class, destructors clean up after your object and free any memory you might have allocated. A destructor always has the name of the class, preceded by a tilde ( \(\sim\). Destructors take no arguments and have no return value. Therefore, the Cat declaration includes
~Cat ();

Note: You can have multiple constructors using function overloading. The compiler selects which constructor function to use by the parameters you provide in the call or creation. But be very careful when combining multiple constructors and default arguments-the compiler needs to be able to pick the constructor you want.

\section*{Default Constructors and Destructors}

If you do not declare a constructor or a destructor, the compiler makes one for you. The default constructor and destructor take no arguments and do nothing

\section*{Frequently Asked Questions}

FAQ: Is it called the default constructor because it has no arguments or because it is the constructor providec by the compiler if I don't declare one?

Answer: A constructor that takes no arguments is called a default constructor-whether the compiler creates it for you, or you create it yourself! You receive a default constructor by default.

Confusingly, the default destructor is the destructor provided by the compiler. Because all destructors take no parameters, what distinguishes the default destructor is that it takes no action-it has an empty function body.

\section*{Using the Default Constructor}

What good is a constructor that does nothing? In part, it is a matter of form. All objects must be constructed an destructed, and these do-nothing functions are called at the right time. However, to declare an object without passing in parameters, such as
Cat Rags; // Rags gets no parameters

The constructor took one parameter, you would write
Cat Frisky (3);
In the event that the constructor takes no parameters at all (that is, that it is a default constructor), you leave of the parentheses and write
Cat Frisky;
This is an exception to the rule that states all functions require parentheses, even if they take no parameters This is why you are able to write
```

Cat Frisky;

```
This is interpreted as a call to the default constructor. It provides no parameters, and it leaves off the a function body in which you might initialize the class.

\section*{InPUT. Listing 6.4 Using Constructors and Destructors}
// Demonstrates declaration of a constructor and
// destructor for the Cat class
\begin{tabular}{|c|c|}
\hline \#include <iostream.h> & // for cout \\
\hline class Cat & // begin declaration of the class \\
\hline & \\
\hline public: & // begin public section \\
\hline Cat(int initialAge) ; & // constructor \\
\hline \(\sim\) Cat (); & // destructor \\
\hline int GetAge() ; & // accessor function \\
\hline void SetAge(int age); & // accessor function \\
\hline void Meow();
private: & // begin private section \\
\hline int itsAge; & // member variable \\
\hline ); & \\
\hline // constructor of Cat, & \\
\hline Cat: : Cat(int initialage) & \\
\hline 1 & \\
\hline itsAge = initialage; & \\
\hline & \\
\hline Cat: :~Cat() & // destructor, takes no action \\
\hline
\end{tabular}

\title{
/ GetAge, Public accessor function
returns value of itsAge member
}
int Cat: :GetAge()
return itsAge;
Definition of SetAge, public
/ accessor function
void Cat::SetAge(int age)
// set member variable itsAge to
// value passed in by parameter age
l/ value passed
itsge \(=\) age
definition of Meow method
// returns: void
// action: Prints "meow" to screen
void Cat::Meow ()
cout << "Meow.\n";
// create a cat, set its age, have it
// meow, tell us its age, then meow again.
int main()
Cat Frisky (5);
Frisky.Meow();

cout << Frisky
Frisky. Meow ()
Frisky. MetAge (7);
cout <<"Now Frisky is "; ; " years old.\n";
\[
\begin{aligned}
& \text { cout << Fris? } \\
& \text { return } 0 \text {; }
\end{aligned}
\]

\section*{Output}

ANALYSIS Listing 6.4 is similar to 6.3 , except that line 9 adds a constructor that takes an integer. Line 10 declares the destructor, which takes no parameters. Destructors never take parameters, and neither constructo nor destructors return a value-not even void

\section*{Lines 19 to 22 show the implementation of
accessor function. There is no return value.}

Lines 24 to 26 show the implementation of the destructor \(\sim\) Cat0. This function ine tefinito of the fuction if you declare it in the class declaration

Line 58 contains the definition of a Cat object, Frisky. The value 5 is passed in to Frisk's constructor. No need exists to call SetAgeo because Frisk was created with the value 5 in its member variab
61. In line 63 , Frisky's isAge variable is reassigned to 7 . Line 65 prints the new value.
\begin{tabular}{|l|l|}
\hline Do & \begin{tabular}{l} 
Don'T \\
\hline DO use constructors to initialize your objects.
\end{tabular} \\
\begin{tabular}{l} 
DON'T give constructors or destructors a return \\
value. \\
vON'T give destructors parameters.
\end{tabular} \\
\hline
\end{tabular}

\section*{Using const Member Functions}

If you declare a class method const, you are promising that the method will not change the value of any of the emicolon The of the censtant member fir the semicolon. The d
void. It looks like this
d SomeFunction() const

\section*{Accessor function
accessor functions}
void SetAge (in
int GetAge () ;
SetA geo cannot be const because it changes the member variable itsAge. GetA ge0, on the other hand, can and
should be con member variable itsAge. Therefore, the declaration of these functions should be written like this:

\section*{Distinguishing Interface from Implementation}

As you have learned, clients are the parts of the program that create and use objects of your class. You can th lic interface to your class declat a contract with these clients. The contract tell how your class will behave

In the Cat class declaration, for example, you create a contract that every Cat's age can be initialized in it. constructor, assigned to by its SetAge) accessor function, and read by its GetAge() accessor. You also promise that every Cat will know how to Meowo. Note that you say nothing in the public interface about the member variable itsAge; that is an implementation detail that is not part of your contract. You will provide an age (GetAge() and you will set an age (SetAge()), but the mechanism (itsAge) is invisible If you make GetAge) a const
the Cat on which it is called.

C+is绪 contracts.

Caution: Listing 6.5 does not compile!

INPUT Listing 6.5 A Demonstration of Violations of the Interface
```

// Demonstrates compiler errors

# 

\#include <iostream.h> // for cout
class Cat
public:
Cat(int initialAge);
Cat (); (
void Meow();
private:
int itsAge;
// constructor of Cat,
itsAge = initialAge
cout << "Cat Constructor\n";
Cat::~Cat()
// destructor, takes no action
cout << "Cat Destructor\n"
/ GetAge, const function
/ but we violate const!
int Cat::GetAge() const
return (itsAge++); // violates const
definition of SetAge, public
accessor function
void Cat::SetAge(int age)
// set member variable itsAge to
value passed in by parameter age
itsAge = age;
// definition of Meow method
returns: void
action: Prints "meow" to screen
void Cat::Meow()
out << "Meow.\n"
demonstrate various violations of the
interface, and resulting compiler errors
Cat Frisky;
// doesn't match declaration
Frisky.Meow();
// No, silly, cats can't bark

```
    return 0;
Amaver

Line 11 declares GetAge() to be a const accessor function-as it should be. In the body of GetAge(), however, it variable itsAge is incremented. Because this method is declared to be const, it must not change the value of issAge. Therefore, it is flagged as an error when the program is compiled.

In line 13 , Meow( is not declared const. Although this is not an error, it is bad programming practice. A better design takes into account that this method does not change the member variables of Cat. Therefore, Meow should be const.

Line 58 shows the definition of a Cat object, Frisky. Cat now has a constructor, which takes an integer as parameter. This means that you must pass in a parameter. Because no parameter exists in line 58 , it is flagged
an error. an error

Tip: If you provide any constructor, the compile will not provide one at al. Thus, if you create a constructor that
takes a parameter, you will then have no default constructor unless you write your own.

Line 60 shows a call to a class method, Bark(). Bark) was never declared. Therefore, it is illega
Line 61 shows itsAge being assigned the value 7 . Because itsAge is a private data member, it is flagged as an
Why Use the Compiler to Catch Errors?
Although it would be wonderful to write 100 percent bug-free code, few programmers have been able to d so. However, many programmers have developed a system to help minimize bugs by catching and fixin so. However, many prog
them early in the process

Although compiler errors are infuriating and are the bane of a programmer's existence, they are far better t the alternative. A weakly typed language enables you to violate your contracts without a peep from the compiler, but your program will crash at runtime-when, for example, your boss is watching
Compile-time errors-that is, errors found while you are compiling-are far better than runtime errors-that \(\left\lvert\, \begin{aligned} & \text { is, errors found while you are executing the program. This is because compile-time errors can be found much } \\ & \text { more reliably. It is possible to run a program many times without going down every possible code }\end{aligned}\right.\) more reliably. It is possible to run a program many times without going down every possible code path. Thu
a runtime error can hide for quite a while. Compile-time errors are found every time you compile. Thus, the are easier to identify and fix. It is the goal of quality programming to ensure that the code has no runtime bugs. One tried-and-true technique to accomplish this is to use the compiler to catch your mistakes early in the development process.

\section*{Where to Put Class Declarations and Method Definitions}

Each function that you declare for your class must have a definition. The definition is also called the function
implementation. Like other functions, the definition of a class method has a function header and a function implem.
body.
The definition must be in a file that the compiler can find. Most \(\mathrm{C}++\) compilers want that file to end with .c or
. cpp. GNU compilers will handle extensions of C. .cpp. .cxx, or \(. c++\). The listing files use. cxx so that you can .cpp. GNU compilers will handle extensions of .C, .cpp,. cxx, or c.c+. The listing files use .cxx so that you can use them under Linux or a Microsoff Windows/DOS FAT file system. If you are using a different compiler, yo
will have to determine which it prefers. Be careful with case sensitivity under Linux-.CXX is not the same as .cxx.

Note: GNU compilers assume that files ending with. care \(C\) programs, and that \(\mathrm{C}++\) program files end with .cpp
You can use any extension, but .cpp will minimize confusion.
are free to put the declaration in this file as well, but that is not good programming practice. The You are free to put the declaration in this file as well, but that is not good programming practice. The
convention that most programmers adopt is to put the declaration into what is called a header file, usually with the same name but ending in .h., hp, or .hpp. This book names the header files with .hpp, but check your
compiler to see what it prefers.

For example, you put the declaration of the Cat class into a file named cat.hpp, and you put the definition of the For example, you put the declaration of the Cat class into a file named cat.hpp, and you put the definition of
class methods into file called cat.cxx. You then attach the header file to the \(\mathrm{C}++\) source file by putting the following code at the top of cat.cxx:
\#include "Cat.hpp"
This tells the compiler to read cathipp into the file, the same as if you had typed in its contents at this point. Not Some compilers insist that the capitalization agree between your \#\#include statement and your file system.
Why bother separating the contents of your .hpp file and your .cxx file if you are just going to read the .hpp file back into the .cxx file? Most of the time, clients of your class do not care about the implementation specifics.
Reading the header file tells them everything they need to know; they can ignore the implementation files. In Reading the header file tells them everything they need to know; they can ingore the imy
addition, you might very well end up including the .hpp file into more than one .cpp file
Using include (header) files is just good software engineering-it is easier to reuse your code
Note: The declaration of a class tells the compilier what the class is, what data it holds, and what functions it has.
The declaration ot the class is s called its interface because it tells the wser how to interact with the class. The The declaration of the class is called its interface because it tells the user how
intertace is usually stred in an .hpp file, which is refered to as a header file.

The function definition tells the compiler how the function works. The function definition is called the Implementation of the class method, and it it sepp in a. cpp file. The inplementation details of the class are of
concern only to the author of the class Clients of the classs that is, he parts of the program that use the concern only to the author of the class. Clients of the class-t hat is, the parts of the program that use the
class-neither need nor care to know how the functions are implemented.

\section*{Applying Inline Implementation}

Just as you can ask the compiler to make a regular function inline, you can make class methods inline. The keyword inline appears before the return type. The inline implementation of the GetWeight) function, for example, looks like this
```

inline int Cat::GetWeight(

```
return itsWeight; // return the Weight data member

You can also put the definition of a function into the declaration of the class, which automatically makes that function inline, for example:
class Cat
public:
int GetWeight () \(\{\) return itsweight;

Note the syntax of the GetWeight) definition. The body of the inline function begins immediately after the declaration of the class method; no semicolon is used after the parentheses. Like any function, the definition begins with an opening brace and ends with a closing brace. As usual, whitespace does not matter; you could have written the declaration as
```

class Ca

```
public
int GetWeight() const
return itsWeight;
void SetWeight (int aweight)
Listings 6.6 and 6.7 re-create the Cat class but they put the declaration in cat her and the implementation of the
functions in catcxx. Listing 6.7 also changes the accessor functions and the Meowo function to inline
```

INPUT. Listing 6.6 Cat Class Declaration in cat.hp

```
\#include <iostream.h>
class Cat
class Cat
3: f public:
    \(\underset{\sim}{\text { Cat (int }} \underset{\sim}{\text { Cat () }}\)
    ~Cat (); \(\quad\) (/ inline
int GetAge () const \(\{\) return itsAge; \}
void SetAge (int age) \(\{\) itsAge \(=\) age; \(\} \quad / /\) inline!
        void SetAge (int age) \(\begin{aligned} & \text { itsAge }=\text { age; }\} \\ & \text { void Meow() const } \\ & \left.\text { ic cout << "Meow. } \backslash \mathrm{n}^{\prime \prime} ;\right\}\end{aligned} \quad\) // inline!
            int itsAge;
```

// and inclusion inline functions
\#include "cat.hpp" // be sure to include the header files!
Cat::Cat(int initialAge) //constructor
itsAge = initialAge;
Cat::~Cat()
// Create a cat, set its age, have it
int main()
Cat Frisky (5)
Frisky.Meow()
lol
Frisky.Meow();
Frisky.SetAge (7);
Cout << "Now Frisky is ";
cout << Frisky.GetAge() <<"",
return 0;

``` .

\author{
Lines 18 to 29 repeat the m .
}

\section*{Using Classes with Other Classes as Member Data}

It is not uncommon to build up a complex class by declaring simpler classes and including them in the declaration of the more complicated class. For example, you might declare a wheel class, a motor class, a declaration of the more compicated class. For example, you might declare a whee class, a motor class, a
transmion class, and so forth, and then combine them into a car class. This declares a relationship. A car transmission class, and so forth, and then combii
motor, it has wheels, and it has a transmission.

Consider a second example. A rectangle is composed of lines. A line is defined by two points. A point is Consider a second example. A rectangle is composed of lines. A line is defined by two points. A point is
definied by an x coordinate and a y coordinate. Listing 6.8 shows a complete declaration of a Rectangle class, defined by an \(x\) coordinate and a a coordinate. Listing 6.8 shows a complete declaration of a Rectangle class,
might appear in rect.pp short of recangle.hpp). Because a rectangle is defined as four lines conecting four
points, and each point refers to a coordinate on a a graph, you first decclare a Pooint class to hold the \(x, y\) coordina of each point. Listing 6.8 shows a complete declaration of both classes.

\section*{INPUT Listing 6.8 Declaring a Complete Class}
```

\#include <iostrea
class Point // holds x,y coordinates
public:
void
|\mp@code{void SetX(int x) { itsX = x;}
int GetX()const { return itsX;
private:
int itsX;
int itsY;

```
    class Rectangle
    public:
        Rectangle (int top, int left, int bottom, int right);
~Rectangle ()

        int GetBottom() const \{ return itsBottom,


        \(\begin{array}{ll}\text { void SetUpperLeft (Point Location) } & \text { \{itsUpperLeft }=\text { Location; \}} \\ \text { void SetLowerLeft (Point Location) } & \text { (itsLowerLeft }=\text { Location; }\end{array}\)
        void SetUpperRight (Point Location)
void SetLioweright (Point Location)
(istoperRight \(=\) Location;
        void SetTop (int top) \(\left\{\begin{array}{l}\text { itsTop }=\text { top; }\end{array}\right\}\)
void SetIeft (int left) \(\{\) itsLeft \(=\) left;
        void SetBottom (int bottom) \& itsBottom \(=\) bottom;
void SetRight (int right)
itsRight \(=\) right;
        int GetArea() const
    private:
        itsupperLeft,
        Point itsupperRight,
        Point itsLowerLertight;
        \(\begin{array}{ll}\text { int } & \text { itsTop; } \\ \text { int } \\ \text { itsLeft }\end{array}\)
        int
int
itsBottom;
itsRight;
    // end rect.hpp

\section*{INPUT. Listing 6.9 rect.cpI}
```

// Begin rect.cpp
Rectangle::Rectangle(int top, int left, int bottom, int right)

```
    itsTop \(=\) top;
itsLeft \(=1\) left
    itssottom \(=\) bottom;
    itsUpperLeft, Setx(left);
    itsUpperRight.SetX(right),
    itsLowerLeft. SetX(left);
itsLowerLeft. SetY(botto
    itsLowerRight.SetX(right); ;
itsLowerRight.SetY (bottom);
// compute area of the rectangle by finding sides,
Rectangle:: GetArea() const
    int Width \(=\) itsRight-itsLeft;
int Height \(=\) itsTop - itsBottom;
    return (Width * Height);
nt main(
    //initialize a local Rectangle variable
Rectangle MyRectangle \((100,20,50,80)\)
    int Area \(=\) MyRectangle.GetArea()
    cout << "Area: ">< Area <<" " \(\backslash\) ";
cout << "Upper Left X Coordinate
    cout << MyRectangle.GetUpperLeft ().Getx () << endl;

ANALYsIS Lines 3 to 14 in Listing 6.8 declare the class Point, which is used to hold a specific \(x, y\) coordinate o graph. As written, this program does not use Points much. However, other drawing methods require Points.

Line 17 begins the declaration of a Rectangle class. A Rectangle consists of four points that represent the corners of the Rectangle.

The constructor for the Rectangle (line 20) takes four integers, known as top, left, bottom, and right. The four parameters to the constructor are copied into four member variables (Listing 6.9), and then the four Points are established.

In addition to the usual accessor functions, Rectangle has a function GetArea() declared in line 43. Instead of storing the area as a variable, the GetArea() function computes the area on lines 28 through 30 of Listing 6.9. To do this, it computes the width and the height of the rectangle, and then it multiplies these two values.

Getting the x coordinate of the upper-left corner of the rectangle requires that you access the UpperLeft point and ask that point for its \(X\) value. Because GetUpperLeft() is a method of Rectangle, it can directly access the private data of Rectangle, including itsUpperLeft. Because itsUpperLeft is a Point and Point's itsX value is private, GetUpperLeft() cannot directly access this data. Rather, it must use the public accessor function GetX() to obtain that value.

Line 33 of Listing 6.9 is the beginning of the body of the actual program. Until line 36, no memory has been allocated, and nothing has really happened. The only thing you did is tell the compiler how to make a point anc how to make a rectangle, in case one is ever needed.

In line 36, you define a Rectangle by passing in values for Top, Left, Bottom, and Right.
In line 38, you make a local variable, Area, of type int. This variable holds the area of the Rectangle that you created. You initialize Area with the value returned by Rectangle's GetArea() function.

A client of Rectangle could create a Rectangle object and get its area without ever looking at the implementation o GetArea().
rect.hpp is shown in Listing 6.8. Just by looking at the header file, which contains the declaration of the Rectangl class, the programmer knows that GetArea() returns an int. How GetArea() does its magic is not of concern to the user of class Rectangle. In fact, the author of Rectangle could change GetArea() without affecting the programs that use the Rectangle class.

\section*{Frequently Asked Questions}

FAQ: What is the difference between declaring and defining?
Answer: A declaration introduces a name of something but does not allocate memory. A definition allocates memory.

With a few exceptions, all declarations are also definitions. The most important exceptions are the declaration of a global function (a prototype) and the declaration of a class (usually in a header file).

\section*{Using Structures}

A very close cousin to the class keyword is the keyword struct, which is used to declare a structure. In C++, a structure is the same as a class, except that its members are public by default. You can declare a structure exactly as you declare a class, and you can give it the same data members and functions. In fact, if you follow the good programming practice of always explicitly declaring the private and public sections of your class, no difference will exist whatsoever.

Try re-entering Listing 6.8 with these changes:
- In line 3, change class Point to struct Point.
- In line 17, change class Rectangle to struct Rectangle.

Now run the program again and compare the output. No change should have occurred.

\section*{Why Two Keywords Do the Same Thing}

You probably are wondering why two keywords do the same thing. This is an accident of history. When C++ was developed, it was built as an extension of the C language. C has structures, although C structures do not have class methods. Bjarne Stroustrup, the creator of \(\mathrm{C}++\), built upon structs, but he changed the name to class to represent the new expanded functionality.

Do
DO put your class declaration in an .hpp file and your
member functions in a .cpp file.
DO use const whenever you can.
DO understand classes before you move on.

\section*{Summary}

Today you learned how to create new data types called classes. You learned how to define variables of these new types, which are called objects.

A class has data members, which are variables of various types, including other classes. A class also includes member functions-also known as methods. You use these member functions to manipulate the member data and to perform other services.

Class members, both data and functions, can be public or private. Public members are accessible to any part of your program. Private members are accessible only to the member functions of the class.

It is good programming practice to isolate the interface, or declaration, of the class in a header file. You usually do this in a file with an .hpp extension. The implementation of the class methods is written in a file with a .cpp extension.

Class constructors initialize objects. Class destructors destroy objects and are often used to free memory allocated by methods of the class.

Q How big is a class object?
A A class object's size in memory is determined by the sum of the sizes of its member variables. Class methods don't take up room as part of the memory set aside for the object.

Some compilers align variables in memory in such a way that 2-byte variables actually consume somewhat more than 2 bytes. Check your compiler manual to be sure, but at this point you do not need to be concerned with these details.

Q If I declare a class Cat with a private member itsAge and then define two Cat objects, Frisky and Boots, can Boots access Frisky'S itsAge member variable?
A Yes. Private data is available to the member functions of a class, and different instances of the class can access each other's data. In other words, if Frisky and Boots are both instances of Cat, Frisky's membe functions can access Frisky's data and also Boots's data.
Q Why shouldn't I make all the member data public?
A Making member data private enables the client of the class to use the data without worrying about how it is stored or computed. For example, if the Cat class has a method GetAge(), clients of the Cat class can ask for the cat's age without knowing or caring whether the cat stores its age in a member variable or computes its age on-the-fly.
\(Q\) If using a const function to change the class causes a compiler error, why shouldn't I just leave out the word const and be sure to avoid errors?
A If your member function logically shouldn't change the class, using the keyword const is a good way to enlist the compiler in helping you find silly mistakes. For example, GetAge() might have no reason to change the Cat class, but your implementation might have this line:
if (itsAge \(=100\) ) cout << "Hey! You're 100 years old\n";
Declaring GetAge() to be const causes this code to be flagged as an error. You meant to check whether itsAge is equal to 100 , but instead you inadvertently assigned 100 to itsAge. Because this assignment changes the class-and you said this method would not change the class - the compiler is able to find the error.

This kind of mistake can be hard to find just by scanning the code. The eye often sees only what it expects to see. More importantly, the program might appear to run correctly, but itsAge has now been set to a bogus number. This will cause problems sooner or later

\section*{Q Is there ever a reason to use a structure in a C++ program?}

A Many C++ programmers reserve the struct keyword for classes that have no functions. This is a throwback to the old C structures, which could not have functions. Frankly, I find it confusing and poor programming practice. Today's methodless structure might need methods tomorrow. Then you'll be forced either to change the type to class or to break your rule and end up with a structure with methods

\section*{Workshop}

The Workshop provides quiz questions to help you solidify your understanding of the material covered and exercises to provide you with experience in using what you learned. Try to answer the quiz and exercise questions before checking the answers in Appendix D, "Answers to Exercises and Quizzes," and make sure you understand the answers before continuing to the next day's lesson

\section*{Quiz}
1. What is the dot operator, and what is it used for?
2. Which sets aside memory-declaration or creation?
3. Is the declaration of a class its interface or its implementation?
4. What is the difference between public and private data members?
5. Can member functions be private?
6. Can member data be public?
7. If you declare two Cat objects, can they have different values in their itsAge member data?
8. Do class declarations end with a semicolon? Do class method definitions?
9. What would the header be for a Cat function, Meow, that takes no parameters and returns void?
10. What function is called to initialize a class?

\section*{Exercises}
1. Write the code that declares a class called Employee with these data members: age, yearsOfService, and Salary.
2. Rewrite the Employee class to make the data members private, and provide public accessor methods to get and set each of the data members.
3. Write a program with the Employee class that makes two Employees; sets their age, YearsOfService, and Salary; and prints their values.
4. Continuing from Exercise 3, provide a method of Employee that reports how many thousands of dollars the employee earns, rounded to the nearest 1,000 .
5. Change the Employee class so that you can initialize age, YearsOfService, and Salary when you create the employee.
6. BUG BUSTER: What is wrong with the following declaration?
class Square
\({ }^{\text {cla }}\)
public:
int Side;
\}
7. BUG BUSTER: Why isn't the following class declaration very useful?
class Cat
\{
int GetAge()const;
private:
int itsAge;
\};
8. BUG BUSTER: What three bugs in this code will the compiler find?
```

class TV
{
public:
void SetStation(int Station);
int GetStation() const;
private:
int itsStation
};
int main()
int
TV myTV;
myTV.itsStation = 9
TV.SetStation(10);
TV myOtherTv(2);
return 0;
Looping
Many programming problems are solved by repeatedly acting on the same datat. Two ways to do this are
reeursion (discuscsed on Deay 5 ."Functionsi) and iteration. Iteration means doing the same thing again an
recursion (discussed on Day 5 , "Functions") and ito
again. The principal method of iteration is the loop.
The Roots of Looping: The goto Statement
In the primitive dayy of early computer science, programs were nasty, brutish, and short. Loops consisted of
label, some statemens, and a jump.

Cation:
Good progaramers limit their use of this stalement!

## INPUT Listing 7.1 Looping with the Keyword

| // Listing 7.1 <br> // Looping with goto |  |  |
| :---: | :---: | :---: |
| \#include <iostream.h> |  |  |
| int main() |  |  |
|  |  |  |
|  | int counter $=0$; | // initialize counter |
|  | counter + +; | // top of the loop |
|  | cout << "counter: |  |
|  | if goto loop; | // jump to the top |
|  | cout << "Complete return 0; | Counter: " << counter << ". $\backslash \mathrm{nn}^{\prime \prime}$; |

## while Loops

A wilie loop causes your program to repaeat asequence of statements as long as the starting condition remains
true. In the example of gov in Listing 7.1 , counter was incermented until it was equal to 5 . Listing 7.2 show $t$ Input

```
/// Listing 7.2
clude <iostream.h>
main()
    int counter = 0;
        Counter++; // body of the loop
    cout<<"co
```

Eample



tincluade ciostrean.h>
int main()




 sma11++;
1arge-=2;


Outrut
Enter a small number: ${ }^{2}$
Enter a
Enarge number: 100000

 On lines 12.1015 , the numbers are entered. Line 20 sets upa a while loop, which will continue only as long as
threc onditions are met:

 When any of the three conditions in the wiil loop fails, the loop ends, and execulion of the program coninuus
after he while loop sclosing brace on line 29 .
out \ll

if (smal1 \% skip == 0) // skip the decrement?
cout <
continue;
if (1arge $==$ target) $\quad / /$ exact match for the target? cout << "Target reached!
break;
$\qquad$

Once again, $\mathrm{C}+$ + gives you several ways to accomplish the same thing. No experienced $\mathrm{C}++$ programmer would use
a for loop in this way, but it does illustrate the flexibility of the for statement. In fact, it is possible, using break and
continue, to create a for loop with none of the three statements. Listing 7.12 illustrates how.
Listing 7.12 Illustrating an Empty for Loop Statement
/Listing 7.12 illustrating
//empty for loop statement
$\qquad$
$\qquad$
$\qquad$
$\qquad$
$\qquad$
$\qquad$
$\qquad$
$\qquad$
$\qquad$
$\qquad$
$\square$
$\qquad$
$\qquad$
$\qquad$
$\qquad$
today.
Empty for Loops
Because so much can be done in the header of a for statement, at times you will not need the body to do anything at
all. In that case, be sure to put a null statement $(\underset{\sim}{)}$ as the body of the lopo. The semicolon can be on the same line as
the header, but this is easy to overlook. Listing 7.13 illustrates how to use a null body in a for loop.
INPUT Listing 7.13 Illustrates the Null Statement in a for Loop
$\qquad$
$\qquad$
for (int $\qquad$
$\qquad$
$\qquad$
Nothing is lef
for loop: the a
$\qquad$
Nested Loops
Loops may be nested with one loop sitting in the body of another. The inner loop will be executed in full for every
execution of the outer loop. From experience, it seems that for loops are the type most commonly nested. All three o
the loop types can be nested, but the while and do...while are just less common. Listing 7.14 illustrates writing marks
into a matrix using nested for loops.
$\qquad$
$\qquad$ \#include <iostream.h>
int main()
$\qquad$
$\qquad$
$\qquad$
$\qquad$
$\qquad$ \#include <i \#include
int main()
// i
$\qquad$
$\qquad$
The GNU compile
warning looks like:
$\qquad$
$\qquad$
$\qquad$ \#include <iostream.h>
int main()
$\qquad$
$\stackrel{\text { for (in }}{1}$
$\qquad$
$\qquad$

Summary of Loops
On Day 5 you learned how to solve the Fibonacci series problem using recursion. To review briefly, a Fibonacci
series starts with $1,1,2,3$, and all subsequent numbers are the sum of the previous two: 1,1,2,3,5, , ,13,21,34...

\#include <iostream.h>
$\underset{\substack{\text { int fib (int } \\ \text { int main) }}}{\text { position) }}$
int answer, position
cout
cout
cin

answer $=$ fib (position);
cout
cout $<$ answer $\lll$ is
ceut
return
rosition
cout < $<$ position <<"th Fibonacci
return $0 ;$
int fib(int n)
int minusTwo=1, minusone $=1$, answer=2,
return 1;
for ( $\mathrm{n}-\mathrm{=} 3 ; \mathrm{n} ; \mathrm{n}-\mathrm{-})$

return answer;

## 



 the followwing alogotith:

a. Puting hhe value currenly in minusone into minusTwo.
b. Puting the value currenty in in anser into miniusone.
c. Adding minuson e and
c. Derenenting
3. When n reaches 0 , return the answer.

This is exactly how you would solve this problem with pencil and paper. If you were asked for the fith
Fibonaccinumber, you would write 1,1,2,

 been writen
$\qquad$


 | ieration. Mirroo |
| :--- |
| blazingly fast. |

switch Statements

 switch (expression

```
Case valueone: statement;
    valueTwo: statementi,
case valuen: statement;
```


 defalut statement
statenent ends.

## $\overline{\text { Note: }}$




illustrates use of the swich statement.
Invot Listing 7.16 Demonstrating the swich Statemen
//Listing 7.16
\#include <iostream.h>

UTput

cin $\gg$ nunber;
switch (number)


```
cout<<"\n\n"
```

Output
Enter a num
Bxaclelent
Mastefun!
Incredible! $!$
Running the program a second time:
Enter a number between 1 and 5 :
Too large!



$\qquad$
Listing 7.17 returns to the for(;) loop discussed earlier. These loops are also called forever loops, as they will
loop forever if arraak in son encountere. The orever loop is used to put up amenus solicita chocice erom the
user, act on the choice, and then return to the menu. This will continue unil the user chooses to exit
$\qquad$
$\qquad$
$\qquad$
A forever loop is a loop that does not have an exit con
used. Forever loops are also known as eternal loops.
INPUT. Listing 7.17 Demonstrating a forever Loop
$\qquad$
$\qquad$
$\qquad$
$\qquad$
$\qquad$
$\qquad$
$\qquad$
$\qquad$
$\qquad$
$\qquad$
$\qquad$
$\qquad$
$\qquad$
$\qquad$
$\qquad$
$\qquad$
$\qquad$
$\qquad$
$\qquad$
$\qquad$
$\qquad$
$\qquad$
$\qquad$
$\qquad$
$\qquad$
$\qquad$
$\qquad$
$\qquad$
(1) Choice one.
(2) Choice two.
(3) Choice three.
(4) Redisplay menu.
(5) Quit.
$\qquad$
$\qquad$
$\qquad$
$\qquad$
$\qquad$
$\qquad$
$\qquad$
Q\&A
$\qquad$
The Workshop provides quiz questions to help you solidify your understanding of the material covered as well
as exercises to provide you with experience in using what you learned today. Try to answer the quiz and
eexecisis questions before checking the answers in Appendix D. "Answers to Exercisises and Quizzes," and mak
sure ye undersand the aswers befoe contin the
Quiz
$\square$

# Part 2 

## At a Glance

You have finished the first week of learning how to program in C+t. By now you should feel comfortable
entering programs, using your compiler, and thinking about objects, classes, and program flow.

## Where You Are Going

Week 2 begins with pointers. Pointers are traditionally a difficult subject for new C++ programmers, but you
 discusses pointers, and Day 9, "References," "teaches references, which are a close cousin to pointers. On Day
10, "Advanced Functions," you will see how to overload functions, and Day 11, "Inheritance,", introduces 10 "Advanced Functions, you will see how to overload funcions, and ay 11, "Aheritance," "infoduces Lists," you will learn how to work with arrays, strings, and collections. Day 13, "Polymorphism,", extends the lessons of Day 11 to discuss polymorphism, and Day 14, "Special Classes and Functions," ends the week with
discussin

## Chapter 8

Pointers
One of the most powerful tools available to a C

- What pointers are
- How to declare and use pointers
What the free store is and how
ters present two special challenges when you are learning $\mathrm{C}++$ : They can be somewhat confusing, and it is not immediately obvious why they are needed. This lesson explains how pointers work, step
fully understand the need for pointers, as you continue to work through the rest of the book.


## What Is a Pointer?

A pointer is a variable that holds a memory address.
To understand pointers, you must know a little about computer memory. Computer memory is divided into
sequentially numbered memory locations. Each variable is located at a unique location in memory, known ddress. Figure 81 shows a schematic representation of the storage of an unsigned long integer variable at
$\square$

```
Listing 8.1 Demonstrates address of operator
and addresses of local variables
```

\#include <iostream.h>
int main()
unsigned short shortvar=5;
unsigned long longVar=65535;
long sVar $=-65535$;
cout << "shortVar: $\backslash$ t" $\ll$ shortVar;
cout $\ll$ M
cout << " Address of shortVar
cout << \&shortvar $\ll " \$ n"
cout << "longVar: $\backslash t "$ " longVar;
cout <<" Address of longvar: 1 t"
cout << "Address of longVar
cout $\ll$ dlongVar $\ll " \backslash n " ; ~$

return 0;

Output

| shortVar: | A | Address of shortvar: | b39£d2 |
| :--- | :--- | :--- | :--- |
| longVar: | 65535 | Address of | longVar: |
| longfcc |  |  |  |
| sVar: | -65535 | Address of sVar: | b39fc8 |

(Your printout may look different depending on the configuration of your system; some GNU compilers will show all the addre
happens to you.)

ANAIYsIS Three variables are declared and initialized: an unsigned short in line 8 , an unsigned long in line 9 , and a long in line 10 . Their values and addresses are printed in lines $12-16$ by using the address of the operator (\&).
The value of shortVar is 5 , as expected, and its address is b39fd2. This complicated address is specific to the
computer. te computer, the compiler, and the operating system, and it may change slightly each time the program is run.
Your results will be Your results will be different. What does not change, however, is that the difference in the first two addresses
2 bytes if your computer uses 2 -byte short integers. The difference between the second and third is 4 breser your computer uses 4 -byte long integers. Figure 8.2 illustrates how the variables in this program would be your computer use
stored in memory.

You do not need to know the actual numeric value of the address of each variable. What you care about is that each one has an address and that the right amount of memory is set aside. Youtelel the compilier how much
memory to allow for your variables by declaring the variable type; the compiler automatically assigns an memory to allow for your variables by declaring the variable type; the compiler automatically assigns an
address for it. A long integer is typically 4 bytes, for example, meaning that the variable has an addess to address for it. A I
bytes of memory. about the size a pointer to that long int each being 4 bytes is just a coincidence. You should make no assumption store an address in an unsigned long integer)

## toring the Address in a Pointe

Every variable has an address. Even without knowing the specific address of a given variable, you can store th
address in a pointer.
an int

Note that page is a variable the same as any of the variables. When you declare an integer variable (type int), i pAge is just a different yype of variable.
In this example, pAge is initialized to the address constant NULL. A pointer whose value is NULL is called a $n u$ pointer. All pointers, when they are created, should be initialized to something. If you do not know what you
want to assign to the pointer, assign NuLL. A pointer that is not intitialized is called widd are very dangerous.

The first line creates a variable--howOId, whose type is unsigned short int-and initializes it with the value 50 . The second line declares pAge to be a pointer to type unsigned short int and initializes it to NULL. You know the
pAge is a pointer because of the asterisk (*) after the variable type and before the variable name.

The third and final line assigns the address of howOld to the pointer pAge. You can tell that the address of how is being assigned because of the address of the operator ( $\&$ ). If the address of the operator had not been used, value of howold would have been assigned. That might, or might not, have been a valid addres
At this point, pAge has as its value the address of howOld. howOld, in turn, has the value 50 . You could hav accomplished this with one less step, as in
unsigned short int howOld $=50$; $\quad$ / make a variable
pAge is a pointer that now contains the address of the howold variable. Using page, you can actually determine the value of howOld, which in this case is 50 . Accessing howOId by using the pointer pAge is called indirection because you are indirectiy a
to access a variable's value.

Indirection means accessing the value at the address held by a pointer. The pointer provides an indirect way t get the value held at that address.

## Selecting Pointer Names

Pointers can have any name that is legal for other variables. Many programmers follow the convention o naming all pointers with an initial $p$, as in pAge or pNumber

## Using the Indirection Operator

The indirection operator (*) is also called the dereference operator. When a pointer is dereferenced, the value the address stored by the pointer is retrieved

Normal variables provide direct access to their own values. If you create a new variable of type unsigned shor int called yourAge, and you want to assign the value in howold to that new variable, you would write

```
\begin{array} { l } { \text { unsigned short in } } \\ { \text { yourAge = howold;} } \end{array}
```

pointer provides indirect access to the value of the variable whose address it stores. To assign the value in hold to the new variable yoursage by way of the pointer pAge, you would write

```
unsigned short int yourAge;
```

yourAge = *pAge;

The indirection operator (*) in front of the variable page means "the value stored at." This assignment say "Take the value stored at the address in pAge and assign it to yourAge

Note: The indirection operator ${ }^{* *}$ is used in two distinct ways with pointers: declaration and dereference. When a
pointer is declared, the star indicates that it is a pointer, not a normal variable. For example,
unsigned short * pAge = NULL; // make a pointer to an unsigned short
When the pointer is dereferenced, the indirection operate
the pointer is to be accessed, rather than the address itsel
*pAge $=5$; // assign 5 to the value at pAge
Also not that his same
call, based on context.

## Understanding Pointers, Addresses, and Variable

 held by the pointer. This is the source of much of the confusion about pointers.theVariable $=5$,
int $*$ pPointer $=\$$ theVariable
the variable is declared to be an integer variable initialized with the value 5 . pPointer is declared to be a pointer an ineger; it is initialized with the address of the Variable. pPointer is the pointer. The address that pPointer the address of theVariable. The value at th
representation of theVariable and pPointer.

## $\square$

## Figure 8.3 A schematic representation of memory

## Manipulating Data by Using Pointers

After a pointer is assigned the address of a variable, you can use that pointer to access the data in that variab
Listing 8.2 demonstrates how the address of a local variable is assigned to a pointer and how the pointer manipulates the values in that variable

INPUT Listing 8.2 Manipulating Data by Using Pointers

```
/ Listing 8.2 Using pointers
```

\#include <iostream.h>
typedef unsigned short int USHORT
nt main()
USHORT myAge; $\quad$ // a variable
USHORT * pAge $=$ nULL; // a pointer
myAge $=5$;





cout <<"myyAge = " < $_{\text {ln" }}$ "
cout $\ll "$ mp
myAge $=9 ;$
cout << "myAge:" << myAge << "\n";
return 0

## Examining the Address

Pointers enable you to manipulate addresses without ever knowing their real value. After today, you can take it fait that when you assign the address of a variable to a pointer, it really has the address of that variable as on faith that when you assign the address of a variable to a pointer, it really has the ad
value. But just this once, why not check to make sure? Listing 8.3 illustrates this idea.

INPUT Listing 8.3 Finding Out What Is Stored in Pointer

```
// Listing 8.3 What is stored in a pointer.
#include <iostream.h>
int main()
    unsigned short int myAge = 5, yourAge = 10;
        unsigned short int * pAge = &myAge; // a pointer
        unsigned short int * pAge &myAge; N a poin
        cout << "&myAge:\t" << &myAge << "\t&yourAge:\t"
        rAge <<"\n";
        cout << "pAge:\t" << pAge << "\n";
        pAge = &yourAge; // reassign the pointer
        cout << "myAge:\t">< myAge << "\tyourAge:\t"
        cout << "&myAge:\t" << &myAge << "\t&yourAge:\t"
        cout << "pAge:\t" << pAge << "\n";
        cout << "*pAge:\t" << *pAge << "\n";
    cout << 0;
```


## Output

| myAge: | 5 | yourAge: | 10 |
| :--- | :--- | :--- | :--- |
| \& myAge: | $0 \times 355 \mathrm{C}$ | \& yourAge: $0 \times 355 \mathrm{E}$ |  |
| pAge: | $0 \times 355 \mathrm{C}$ |  |  |
| *pAge: | 5 |  |  |
| myAge: | 5 | yourAge: | 10 |
| \&myAge: | $0 \times 355 \mathrm{C}$ | \& yourAge: $0 \times 355 \mathrm{E}$ |  |
| pAge: | $0 \times 355 \mathrm{E}$ |  |  |
| *pAge: | 10 |  |  |
| \&pAge: | $0 \times 355 \mathrm{~A}$ |  |  |

(Your output may look different.)

## Analysis

 Lines 10 and 11 print the values and the addresses of myAge and yourAge. Line 12 prints the contents of pAge, which is the address of myAge. Line 13 prints the result of dereferencing pAge, which prints the value at pAge-the value in myAge, or 5This is the essence of pointers. Line 12 shows that pAge stores the address of myAge, and line 13 shows how to get the value stored in myAge by dereferencing the pointer pAge. Make sure that you understand this fully befo get the value stored in myAge by dereferencing th
you go on. Study the code and look at the output.

In line 14, pAge is reassigned to point to the address of yourAge. The values and addresses are printed again. T output shows that pAge now has the address of the variable yourAge and that dereferencing obtains the value in yourAge.
Line 19 prints the address of pAge itself. Like any variable, it has an address, and that address can be stored in pointer. (Assigning the address of a pointer to another pointer will be discussed shortly.)

| Do |
| :--- |
| DO use the indirection operator (*) to access the data |
| stored at the address in a pointer. |

DO initialize all pointers either to a valid address or to
DO remember the difference between the address in a pointer and the value at that address.

## Using Pointers

To declare a pointer, write the type of the variable or object whose address will be stored in the pointe followed by the pointer operator (*) and the name of the pointer. For example
unsigned short int * pPointer = NULL;

To assign or initialize a pointer, prepend the name of the variable whose address is being assigned with the address of operator ( () . For example,
unsigned short int theVariable $=5$;
unsigned short int * pPointer $=\&$ theVariable;
To dereference a pointer, prepend the pointer name with the dereference operator (*). For example,
unsigned short int theValue $=$ *pPointer

## Why Would You Use Pointers?

So far you have seen step-by-step details of assigning a variable's address to a pointer. In practice, though, yo would never do this. After all, why bother with a pointer when you already have a variable with access to tha value? The onl
pointers work.

Now that you are comfortable with the syntax of pointers, you can put them to good use. Pointers are used, mo Now that you are con
often, for three tasks:

- Managing data on the free store
- Accessing class member data and functions

The rest of this lesson focuses on managing data on the free store and accessing class member data anc functions. Tomorrow you will learn about passing variables by reference.

## The Stack and the Free Stor

On Day 5, Functions, in the section "How Functions Work-A Look Under the Hood, five areas of memo are mentioned:

- Global name space
- The free st
- Code space
- The stack

Local variables are on the stack, along with function parameters. Code is in code space, of course, and global
variables are in global name space. The registers are used for internal housekeeping functions variables are in global name space. The registers are used for internal housekeeping functions, such as keepin track of the top of the stack and the instruction pointer. Just about all remaining memory is given over to the free store, which is sometimes referred to as the heap.

// Listing 8.4 Allocating and deleting a pointer
tine lided


```
*)
```



```
cout << "*pHe
```

Output

Analusis adaress of the local variable. Line 9 declares another pointer but initializes it with her result obtained from calling for a new int. This allocates space on the free store for an int. Line 10 verifies that memory was allo
and the pointer is valid (not null by bsing it If no memory can be allocated, he pointer is null and and error ,
 for a null pointer before initializizing the value. You could replace line 10 in Listing 8.4 w with he the following: if (pheap $==$ NuLL)
//handle the error here:
Cases the next steps
*pHeap $=7$;
 In line 14 , the memory allocated in line 9 is returned to the free strer by y call to delee. This frees the memory
and sisassociates the poiner from that memory. pHeap is now free to point to other memory. It is reassigned i and disassociates the pointer from that memory. p Heap is inow free to point to oher memory
lines 15 and 16 , and line 17 prints he result. Line 18 restores that memory 0 the fre store.
Although hine 18 is redundant (the end of the program would have returned that memory), it is a good idea to
free this memory explicitl. If the program changes or is extended, having already taken care of this step will
benficil

Another way you might inadvertenly create a memory leak is by reassigning your pointer before deleteing it
memory to which it points. Consider thi code fragment:
unsigned short int *ppointer $=$ new unsigned short int
*ppointer
fpor
ppointer $=$ new
tppointer $=84 ;$
Line 1 creates ppoiner and assigns it the address of an area on the free store. Line 2 stores the value 72 in that
area of memory. Line 3 reassigs ppoiner to another area of memory. Line 4 places the value 84 in that area
 been reassigne
program ends.
The code should have been written like this




 Cat tpcat =new Cat;

Deleting Object:

INvUT Listing 8.5 Crating and Deleting Ojicets on the Free Store
ITisting 8.5
Ccreating objects on the free store
\#include <iostream.
class simplecat

int itsAge;

cout <<"Destructor called.ln";


```
M
Simen
\iclol
```

Ounan


 Accessing Data Members
 (*pRags) .Getage (1);

[^0]include <iostream.h>
class simplecat
public:
Simplecat (1)
sin
SimpleCat ();
it Getage ()const (return *itsAge;
/other methods
/ /other methods
private:
int $\star_{\text {itsAge }}$
int ${ }^{\text {itssweight }} ;$
SimpleCat: : SimpleCat()
itsAge $=$ new int (2);
itsweight $=$ new int (5) ;
SimpleCat::~~SimpleCat ()
delete itsAge;
delete its itswight;
nt main()
Simplecat Frisky;
cout <<"Frisky is
Frisky.GetAge ( $) \ll$
a years
Frisky.SetAge (5);
cout <<"Frisky is
return 0 ;
Answer: What is on the stack is the local variable Frisky. That variable has two pointers, each of which is
taking up some stack space and holding the address of an integer allocated on the heap. Thus, in the exampl
taking up some stack space and holding the address of an integer allocated on $n$th
8 byyes are on the stack (assuming 4 -byte pointers) and 8 bytes are on the heap.

Now this would be prety silly in a real program unless a good reason existed for the Cat object to hold its
members by reference. In this case, there is no good reason, but in other cases, this would make a lot of sense This begs the question: What are you trying to accomplish? Understand, too, that you must start with design. This begs the question: What re you trying to accomplish? Understand, too, that you must start with design
what you designed is an object that refers to another boject, but the second bbect may come int existence
before the first obiect and continue after the first bieject is onene, the first object must contain the second by what you desi:
before the firs
reference.

For example, the first object might be a window and the second object might be a document. The window need access to the document, but it does not control the lifetime of the document. Thus, he window needs to hold document by reference

## Special Pointers and More issues

There are still many things to learn about using and pointers. The following sections will cover:

> - The this Pointer
> - Sonsty, Woild,or, ors Danging Pointers

## The this Pointel

Every class member function has a hidden parameter: the this pointer. this points to the individual object.
Therefore in each call to Geitese or sets ge0. the this pointer for the object is included as a hidden pera ereore, in each call to Getageo) or set geo), the this pointer for the object is included as a hidden paramett It is possible to use the this pointer explicitly, as Listing 8.8 illustrates.
$/$ Listing 8.8
$/$ Using the th
\#include <iostream.h>
class Rectangle public: Rectangle ();
$\sim$ Rectangle ();
 void SetWidth(int width) (itsWidth =width;
int GetWidth () const (return itsWidth; )

```
l
```

Rectangle: :Rectangle()
itswidth $=5 ;$
itsLength $=10$
Rectangle: : ~Rectangle ()
int main()
Rectangle theRect;
cout <<"theRect is
cout <<"theRect is" <<<therect.GetLength
<<"theRect is" <<theRect.Getwidth()
therect. SetLength (20);
theRect.Setwidth (10);
cout <<"therect is " <<theRect.GetLength()
long. 1 "n
feet wide. $\backslash n$
return $0 ;$
Output
therect is 10 feet long.
theRect is 5 feet wide.
theRect is 20 feet long.
ANALYSIS The Seete ength) and Geite ength0 accessor functions explicitly use the this pointer to access the member variables of the Rectangle object. The Setwidh anc If that were all there was to the this pointer, there would be little point in bothering you with it. The chis pointe
however, is a pointer, itstores the memory address of an object. As such, it can be a powerful tool. You will see a practical use for the this pointer on Day 10 , "Advanced Functions," when operator overloadin
discussed. For now, your goal is to know about the his pointer and to understand what it is: a pointer to the
. discusse. ${ }^{\text {objo }}$
object iself.

## Stray, Wild, or Dangling Pointers

One source of bugs that are nasty and difficicut to find is stray pointers. A stray pointer (also called a wild or dangling pointer) is created when you call delete on a pointer-thereby freeing the memory that it points to
then you do not set it to NuLL. If you then try to use that pointer again without reassigning it, the result is unpredictable and, if you are lucky, your program will crash.

It is as though the Acme Mail Order company moved away, but you still pressed the programmed button on
your phone It it possible that nothing terityle happens-a telephone rings in adeeserted warehouse Perhaps your phone. It is possible that nothing terrible happens-a telephone rings in a deserted warehouse. Perhaps
telephone number has been reassigned to munitions factory and your call detonates an explosive and blow telephone number $h$
up your whole city!
In short, be careful not to use a pointer atier you have called delete on it. The pointer still points to the old arc
memory, but the compiler is free to put othere data theree using the pointer can cause your proeqrat memory, buu the compiler is free to put other datat here, using the pointer can cause your program to crash.
Worse, your program might proced merily on its way and crash several minutes later. This is called a time bomb, and it is no fun. To be safe, after you delete a pointer, set it to NulL. This disarms the pointe

INPUT. Listing 8.9 Creating a Stray Pointer

```
//Listing 8.9
/Demonstrates a stray pointer
#include <iostream.h>
int main()
    USHORT *pInt =new USHORT;
    *pInt =10;
    delete pInt:" <<*pInt <<endl;
    delete pInt;
    long *pLong =new long;
    *pLong =90000;
    cout <<"*pLong:" <<*pLong <<endl,
    *pInt =20; //uh oh,this was deleted!
    cout <<"*pInt:"><*pInt <<endl;
    cout <<"*pLong:" <<*pLong <<endl
    delete pLong;
}
```


## Output

*pInt:10
*pLong:90000
*pInt:20
*pInt:20
*pLong: 65556
(Your output may look different.)

ANALYSIS Line 8 declares pint to be a pointer to USHORT, and plnt is pointed to newly allocated memory. Lin 9 puts the value 10 in that memory, and line 10 prints its value. After the value is printed, delete is called on the pointer. plnt is now a stray, or dangling, pointe

Line 13 declares a new pointer, pLong, which is pointed at the memory allocated by new. Line 14 assigns the Line 13 declares a new pointer, pLong, which is
value 90000 to pLong, and line 15 prints its value.

Line 17 assigns the value 20 to the memory that plnt points to, but plnt no longer points anywhere that is valid The memory that plnt points to was freed by the call to delete, so assigning a value to that memory is certain disaster.
Line 19 prints the value at plnt. Sure enough, it is 20 . Line 20 prints 20 , the value at pLong; it has suddenly been changed to 65556 . Two questions arise:

1. How could pLong's value change, given that pLong was not touched?
2. Where did the 20 go when pInt was used in line 17 ?

As you might guess, these are related questions. When a value was placed at pInt in line 17 , the compiler happi placed the value 20 at the memory location that plnt previously pointed to. However, because that memory was freed in line 11 , the compiler was free to reassign it. When pLong was created in line 13 , it was given plnt's old
memory location. (On some computers this may not happen, depending on where in memory these values are stored.) When the value 20 was assigned to the location that pnt previously pointed to, it wrote over the value pointed to by pLong. This is called "stomping on a pointer." It is often the unfortunate outcome of using a stray pointer.
This is a particularly nasty bug because the value that changed was not associated with the stray pointer. The This is a particularry nasty bug because the value that changed was not associated with the stray pointer. Th
change to the value at pLong was a side effect of the misuse of pInt. In a large program, this would be very change to the value at
difficult to track dowr

Just for fun, here are the details of how 65,556 got into that memory address

1. plnt was pointed at a particular memory location, and the value 10 was assigned.
2. delete was called on pInt, which told the compiler that it could put something else at that location. The
pLong was assigned the same memory location. pLong was assigned the same memory location.
3. The value 90000 was assigned to
4. The value 90000 was assigned to "pLong. The particular computer used in this example stored the 4
byte value of $90,000(00015 \mathrm{FF} 90$ ) in byte-swapped order. Therefore, it was stored as 5 F 9000 ol 4. pInt was assigned the value 20 - or 0014 in hexadecimal notation. Because pInt still pointed to the same address, the first two bytes of pLong were overwritten, leaving 00140001 .
5. The value at pLong was printed, reversing the bytes back to their correct order of 00010014 , which
was translated into the value of 65556 . was translated into the value of 65556 .

## Frequently Asked Questions

FAQ: What is the difference between a null pointer and a stray pointer?
Answer: When you delete a pointer, you tell the compiler to free the memory, but the pointer itself continues to exist. It is now a stray pointer.

Hen you then write myPtr =NULL;, you change it from being a stray pointer to being a null pointer.
Normally, if you delete a pointer and then delete it again, your program is undefined. That is, anything migh
happen-if you are lucky, the program will crashh. If you delete a null pointer, nothing happens; it is safe.
Using a stray or a null pointer (for example, writing myPtr $=5$;) is illegal and it might crash. If the pointer is null, it will crash, another benefit of null over stray. The authors prefer predictable crashes because they are easier to debug.

## const Pointers

You can use the keyword const for pointers before the type, after the type, or in both places. For example, all th You can use the keyword const fc
following are legal declarations:
const int *pone;
const int *const pThree;
pOne is a pointer to a constant integer. The value that is pointed to cannot be changed.
pTwo is a constant pointer to an integer. The integer can be changed, but pTwo cannot point to anything else,
pThree is a constant pointer to a constant integer. The value that is pointed to cannot be changed, and pThree cannot be changed to point to anything else

The trick to keeping this straight is to look to the right of the keyword const to find out what is being declared constant. If the type is to the right of the keyword, it is the value that is constant. If the variable is to the right the keyword const, it is the pointer variable itself that is constant.
const int *p1; // the int pointed to is constan

## const Pointers and const Member Functions

On Day 6, "Basic Classes," you learned that you can apply the keyword const to a member function. When a function is declared const, the compiler flags as an error any attempt to change data in the object from within that function

Listing 8
/Using poi
rs with const methods
Lincluad <iostream.h>
ubiic:



$\underset{\substack{\text { private: } \\ \text { int } \\ \text { int tenentht } \\ \text { itswidith; } ;}}{ }$
itswidth $=5 ;$
itstengnt $=10$
${ }^{\text {Rectangle: : -Rectang1e (1) }}$
int main()

 ,





Outrut



 Lines $36-38$ print their values

## Pointer Arithmetic


Invot




hile (Getword (buffer, word, wordoff set))
cout <<"Got this word: " <<word <<enal










return true;

## Summary

Pointers provide a powerful way to access data by indirection. Every variable has an address, which can be obtained using the address of operator (\&). The address can be stored in a pointer

Pointers are declared by writing the type of object that they point to, followed by the indirection operator ${ }^{*}$ ) and the name of the pointer. Pointers should be initialized to point to an object or to NULL.

You access the value at the address stored in a pointer by using the indirection operator $(*)$. You can declare const pointers, which cannot be reassigned to point to other objects, and pointers to const objects, which cannot be used to change the objects to which they point

To create new objects on the free store, you use the new keyword and assign the address that is returned to a pointer. You free that memory by calling the delete keyword on the pointer. delete frees the memory, but it does not destroy the pointer. Therefore, you must reassign the pointer after its memory has been freed.

## Q\&A

## Q Why are pointers so important?

A Today you saw how pointers are used to hold the address of objects on the free store and how they are used to pass arguments by reference. In addition, on Day 13, "Polymorphism," you will see how pointers are used in class polymorphism.

## Q Why should I bother to declare anything on the free store?

A Objects on the free store persist after the return of a function. Additionally, the capability to store objects on the free store enables you to decide at runtime how many objects you need, instead of having to declare this in advance. This is explored in greater depth tomorrow.

## Q Why should I declare an object const if it limits what I can do with it?

A As a programmer, you want to enlist the compiler in helping you find bugs. One serious bug that is difficult to find is a function that changes an object in ways that are not obvious to the calling function. Declaring an object const prevents such changes.

## Workshop

The Workshop provides quiz questions to help you solidify your understanding of the material covered ano exercises to provide you with experience in using what you learned. Try to answer the quiz and exercise questions before checking the answers in Appendix D, "Answers to Exercises and Quizzes," and make sure you understand the answers before continuing to the next day

## Quiz

1. What operator is used to determine the address of a variable?
2. What operator is used to find the value stored at an address held in a pointer?
3. What is a pointer?
4. What is the difference between the address stored in a pointer and the value at that address?
5. What is the difference between the indirection operator and the address of operator?
6. What is the difference between const int *ptrOne and int * const ptrTwo?

## Exercises

1. What do these declarations do?
a. int *pOne;
b. int vTwo;
c. int *pThree $=\& v T w o$;
2. If you have an unsigned short variable named yourAge, how would you declare a pointer to manipulate yourAge?
3. Assign the value 50 to the variable yourAge by using the pointer that you declared in Exercise 2.
4. Write a small program that declares an integer and a pointer to the integer. Assign the address of the integer to the pointer. Use the pointer to set a value in the integer variable
5. BUG BUSTER: What is wrong with this code?
```
#include <iostream.h>
int main()
{
    int *pInt;
    *pInt =9;
    cout <<"The value at pInt:" <<<*pInt;
    return 0;
```

\}
6. BUG BUSTER: What is wrong with this code?

```
int main()
{
    int SomeVariable =5;
    cout <<"SomeVariable:" <<SomeVariable <<"\n "
    int *pVar =&SomeVariable;
    pVar =9;
    cout <<"SomeVariable:" <<*pVar <<"\n ";
    return 0;
}
```


## Chapter 9

## What Is a Reference?

 you dot to the reference is really done to the target.

 reference to that variable by writing the following

This is read as s somenere and isa reference to an integer that is inititilized to refer to soment. Listing 9.1 shows
how references are created and used.


INPUT Listing g.1 Creating and Using References

```
//I isting 9.1
// Denonstrat
```

<iostream.h>
int main()
$\underbrace{}_{\substack{\text { int } \\ \text { int } \\ \text { intone; } \\ \text { \&rsomeref }}}=$ intone



## Wing the Address of Operator \& on Reference


INPUT Listing 9.2 Taking the Adress of a Reference
$/ / /$ I. isting 9.2 Demonstrating the use of References
\#include <iostream.h>
int intone


return 0;

OUtrut

| rsomeref: 5 |
| :--- |
| sintone: |
| co |
| $0 \times 55$ |

Note: Your output may differ on the last wo lines

 opactis

```
iliam_Jefferson_Clinton
```

You might then declare a reference to Pessiden and initialize it with his objec
President sBill_clinton = willi iam_Jefferson_clinton


Normaly, when you use areference, you do not tse the addre.
you would use the target variale. This is shown on line 13 .
References Cannot Be Reassigned

```
INPUT Listing ,.3 Assigning to a Refrenc
```

| //Listing 9.3 |
| :--- |
| $/ /$ Reass $i$ igning |

\#include <iostream.h>
nt main()



OUTP
5
5
$0 \times 213 \mathrm{e}$
$0 \times 213 \mathrm{e}$
:
$0 \times 213 e$
$0 \times 2130$
$0 \times 213 \mathrm{e}$

## What Can Be Referenced?

Any obiect can be ereferenced, including seer-defined objects. Note hat you create a refererence to an object, bu int \& rintref = int; // wrong
You must initialize e Inkee to a particular integer, such as this

```
\begin{array} { l } { \text { int howBig =200;} } \\ { \text { int \& Intref = howBig,} } \end{array}
```

```
In the same way, you do not initialie a referene to a can
```

```
In the same way, you do not initialie a referene to a can
```

cat \& rcatref = car; $/ /$ wrong
Cat frisky
CAT $\&$ rcateef $=$ frisky;

normal class member access operal
object. Listing 9.4 illustrates this.
Input Listing 9,4 References to objects
// Listing 9.4
finclude <iostream.h>
class simplecat


inte
int itsAge
itsNueight,
itsage $=$ age;
itsweight $=$ weight;
Simplecat $\left.\begin{array}{c}\text { Frisky }(5,8) ; \\ \text { Simplecat } \\ \&\end{array}\right)$ rcat $=$ Frisky;

" $"$ pounds. $\backslash n^{\prime \prime} ;$
Luma




ample 1:



## Using Null Pointers and Null References





## Passing Function Arguments by Reference

On Day 5 ."Functions" " you learned that functions have two limitations: Arguments are passed by value, and the
reums satement can return only one value.
 using a poiner, or you pass by reference using a reference.
Passing an object by reference e enables the funcion to change the object being referred to.
int main()

## 



oid swap (int x, int y)
int temp;


OUtrut


Anairsis $_{\text {This program initializes two variables in maino and then passes them to the swap function, which appears to }}$ swap them. When they are examined again in maino, they are unchanged!
The problem here is that x and y are being passed to swap 0 by value. That is, local copies were made in the function. What you want is to pass $x$ and $y$ by reference.

Two ways to solve this problem are possible in $\mathrm{C}++$ : You can make the parameters of swap) pointers to the original value Two ways to solve this problem are possible in C+
or you can pass in references to the original values

## Making swap() Work with Pointers

When you pass in a pointer, you pass in the address of the object, and thus the function can manipulate the value at that address. To make swapo change the actual values using pointers, the function, swap0, should be declared to accept two $i$ pointers. Then, by dereferencing the pointers, the values of xand y will, in fact, be swapped. Listing 9.6 demonstrates th
ind

INPUT. Listing 9.6 Simulating Pass by Reference Using Pointers

```
// Demonstrates simulated pass by
/ reference Using pointers
#include <iostream.h>
void swap(int *x, int *y);
int main()
int x = 5, y = 10;
    cout << "Main. Before swap, x:" << x
    y: "<< Y << "\n"
    Swap(&x,&y); ;
"y:"<<
void swap (int *px, int *py)
    int temp;
    Cout << "Swap. Before swap, *px:
    temp = *px;
    *px = *py;
    cout << "Swap. After swap, *px:" << *px
<< *py << "\n";
```


## Output

Main. Before swap, x: 5 y: 10
Swap. Before swap, *px: $5 *_{\text {py }}: 10$
Swap. After swap, ${ }^{\text {ppx: }}: 10{ }^{\text {ppy }}: 5$
Swap. After swap, *px: $10{ }^{*}{ }^{\text {py }}$
Main. After swap, x: 10 y: 5

## Analysis

On line 5 , the prototype of swap0 is changed to indicate that its two parameters will be pointers to in rather than int variables. When swap 0 is called on line 12 , the addresses of x and y are passed as the arguments.
On line 19, a local variable, temp, is declared in the swapo function. temp need not be a pointer, ;it will just hold the value o $* \mathrm{px}$ (that is, the value of x in the calling function) for the life of the function. After the function returns, temp will no longe "px (that is,
be needed.

On line 23 , temp is assigned the value a px. On line 24 , the val pxis assigned to the value at $p$ y. On line 25 , the valu stashed in temp (that is, the original value at px) is put into p .
The net effect of this is that the values in the calling function, whose address was passed to swap), are, in fact, swapped.

## Implementing swap() with Reference

The preceding program works, but the syntax of the swapo function is cumbersome in two ways. First, the repeated need dereference the pointers within the swapo function makes it error-prone and hard to read. Second, the need to $p$
address of the variables in the calling function makes the inner workings of swapo overly apparent to its users. It is a goal of $\mathrm{C}+$ to prevent the user of a function from worrying about how it works. Passing by pointers puts the burde
on the calling function rather than where it belongs-on the called function. Listing 9.7 rewrites the swapo function, usin: references.

INPUT. Listing 9.7 swap0 Rewritten with References

```
/Listing 9.7 Demonstrates passing by reference
    #include <iostream.h>
    void swap(int &x, int &y);
    int main()
        int x = 5, y = 10;
        cout << "Main. Before swap, x:"<< x<<
        swap (x,y);
        cout<< "Main. After swap,
        return 0
        void swap (int &rx, int &ry)
        int temp
        cout << "Swap. Before swap, rx:" << rx <<
            temp = rx;
            rx = ry;
            cout << "Swap. After swap,
```


## Output

Main. Before swap, $x: 5 \mathrm{y}: 10$
. After swap, rx: 10 ry: 5
Mai 12. On line 13 , the function swap0 is called, but note that $x$ and $y$, not their addresses, are passed. The calling functio

When swapo is called, program execution jumps to line 18 , where the variables are identified as references. Their values : printed on line 22 , but note that no special operators are required. These are aliases for the original values and can be used

On lines 24-26, the values are swapped, and then printed on line 28 . Program execution jumps back to the calling functic and on line 14 , the values are printed in maino. Because the parameters to swapo are declared to be references, the value from maino are passed by reference, and thus are changed in main0 as well. Another successful method.

#  

Class or function. The actual implementation is hidden from the client. This enables hheWhen Colonel John Roebling designed the Brooklyn Bridge, he worried in detail about how the concrete was
poured and how the wire for the bridge was manufactured. He wa in intimately involved in the mechanical and
 time by using well-understood building materials, without regard to how their manufacturer produced them It is the goal of $C+$ to enable programmers to rely on well-understood classes and functions without regard to
theiri internal workings. These "component parts' can be assembled to produce a program, much the same way wires. pipes. clamps, and other ratts are assembled to produce buildings and brid

## Returning Multiple Values Using Pointers

 objects, this effectively enables the function to reiutn two pieces of information. This approach by yasses
return value of the function, which can then be eeserved for feporting errors.

```
Nues: two as pointer prameters and one as the return value of the function
```

```
Nues: two as pointer prameters and one as the return value of the function
```

```
Nues: two as pointer prameters and one as the return value of the function
```

```
Nues: two as pointer prameters and one as the return value of the function
```

```
INPUT Listing 9.8 Returning Values with Pointe
```

//Listing 9.8
// Returning
// using pointers
short Factor (int n, int* psquared, int* pCubed);
int main()
int number, squared, cubed,
short error;
cout $\ll$ "Enter
cin $\gg$ number;
error $=$ Factor (number, \&squared, \&cubed)
if (!error

else
cout
return 0 ;
short Factor(int n, int *psquared, int *pcubed


eturn Value;
/ Listing 9.
// Returnin
Returning multip
using references
nclude <iostream.h>
typedef unsigned short USHort;
enum ERR_CooE ( success, ERROR
RR CODE Factor (USHort, Ushorta, USHORT $\varangle$ )
int main()
USHORT number, squared, cubed;
ERR_CODE result;
out 《< "Enter a namber (0)
f (result $==$ SUCCESS)

$\substack{\text { else } \\ \text { cout } \\ \text { eturn } \\ \text {; }}$
位
if ( $\mathrm{n}>20$ )


## Passing by Reference for Efficiency

Each time you pass an object into a function by value, a copy of the object is made. Each time you return an object from a function by value, another copy is made.

On Day 5 , you learned that these objects are copied onto the stack. Doing so takes time and memory. For smal objects, such as the built-in integer values, this is a trivial cost.
However, with larger, user-created objects, the cost is greater. The size of a user-created object on the stack is the sum of each of its member variables. These, in turn, can each be user-created objects, and passing such a massive structure by copying it onto the stack can be very expensive in performance and memory consumptio

Another cost occurs as well. With the classes you create, each of these temporary copies is created when the compiler calls a special constructor: the copy constructor. Tomorrow you will learn how copy constructors we
and how you can make your own, but for now it is enough to know that the copy constructor is called each tin and how you can make your own, but for now it is enough to know that the copy constructor is called each tin atemporary copy of the object is put on the stac

When the temporary object is destroyed, which happens when the function returns, the object's destructor i called. If an object is returned by the function by value, a copy of that object must be made and destroyed a called.
well.
With large objects, these constructor and destructor calls can be expensive in speed and use of memory. To illu larger ar
called.

Listing 9.10 creates the Simplecat object and then calls two functions. The first function receives the Cat by and then returns it by value. The second one receives a pointer to the object, rather than the object itself, and returns a pointer to the object.

InPUT. Listing 9.10 Passing Objects by Reference Using Pointers

```
//Listing 9.10
// Passing pointers to objects
#include <iostream.h>
class SimpleCat
public:
SimpleCat ();
SimpleCat (Simp
SimpleCat::SimpleCat()
SimpleCat::SimpleCat(SimpleCat&)
    cout << "Simple Cat Copy Constructor...\n";
SimpleCat::~SimpleCat()
    cout << "Simple Cat Destructor.
SimpleCat Functionone (SimpleCat theCat);
int main()
    cout << "Making a'
        ImpleCat Frisky;
        FunctionOne(Frisky); ;
        cout << "Calling FunctionTwo.
            FunctionTwo(&Frisky);
        keturn 0;
    // FunctionOne, passes by value
    SimpleCat FunctionOne(SimpleCat theCat)
            cout << "Function One. Returning...\n";
            return theCat;
    // functionTwo, passes by reference 
            cout << "Function Two. Returning.
            return theCat;
```


## Outrut

Making a cat
Simple Cat Construct
Simple Cat Copy Constru
Function One. Returning
Simple Cat Copy Construc
Simple Cat Destructor
Simple Cat Destructor
Calling FunctionTwo...
: Function Two. Returning
Simple Cat Destructor copy constructor and producing line 6 .
The return value from FunctionOne() is not assigned to any object, so the temporary value created for the returr thrown away, calling the destructor, which produces output line 7 . Because FunctionOne) has ended, its local copy goes out of scope and is destroyed, calling the destructor and producing line 8 .

Program execution returns to main(), and FunctionTwo( is called, but the parameter is passed by reference. No copy is produced, so there is no output. FunctionTwo( prints the message that appears as output line 10 and the
\#include <iostream.h>
${ }^{\text {class }}$ simplecat
${ }^{\text {pubIIc: }} \underset{\substack{\text { Simplecat (); } \\ \text { simpleceat (si }}}{\text { Sin }}$


${ }^{\text {private: }}{ }^{\text {int }}$ itshge;
${ }^{\text {Simplecat: : Simplecat() }}$
cout $\ll$ "si;
itstge
$=1 ;$
Simplecat: : Simplecat (simplecat $\alpha$
cout << "simple cat copy constructor...\n
simplecat: : simplecat()
cout << "simple cat Destructor....n";
int main()

|  |
| :---: |
|  |  |
|  |  |
|  |  |
|  |  |
|  |  |
|  |  |
|  |  |
|  |  |
|  |  |
|  |  |





Outrut

$\qquad$
$\qquad$
cout << "Simple cat Destructor....n";
int main()

out <<"Frisky is" " < Frisky. Getage() <<" years olddn",


// functionTwo. passes a ref to a const object
const Simplecat \& FunctionTwo (const simplecat

 FunctionTwo() now takes and returns a reference to a constant object. Again, working with references is somewhat simpler than working with pointers, and the same savings and efficiency are achieved, as well as the safety provided by using const

## const References

C++ programmers do not usually differentiate between "constant reference to a SimpleCat object" and "reference to a constant SimpleCat object." References themselves can never be reassigned to refer to another object, and so they are always constant. If the keyword const is applied to a reference, it is to make the object referred to be constant

## Understanding When to Use References and When to Use Pointers

C++ programmers strongly prefer references to pointers. References are cleaner and easier to use, and they do a better job of hiding information, as you saw in the previous example.

References cannot be reassigned, however. If you need to point first to one object and then to another, you must use a pointer. References cannot be null, so if any chance exists that the object in question may be null, you must not use a reference. You must use a pointer.

An example of the latter concern is the operator new. If new cannot allocate memory on the free store, it returns a null pointer. Because a reference cannot be null, you must not initialize a reference to this memory until you check that it is not null. The following example shows how to handle this:

```
int *pInt = new int;
if (pInt != NULL)
int &rInt = *pInt;
```

In this example a pointer to int, pInt, is declared and initialized with the memory returned by the operator new. The address in pInt is tested, and if it is not null, pInt is dereferenced. The result of dereferencing an int variable is an int object, and rInt is initialized to refer to that object. Thus, rInt becomes an alias to the int returned by the operator new.

| Do | Don't |
| :--- | :--- |
| DO pass parameters by reference whenever possible. <br> DO return by reference whenever possible. <br> DO use const to protect references and pointers whenever <br> possible. | DON'T use pointers if references will work <br> DON'T return a reference to a local object. |

## Mixing References and Pointers

It is perfectly legal to declare both pointers and references in the same function parameter list, along with objects passed by value. Here is an example:

CAT * SomeFunction (Person \&theOwner, House *theHouse, int age);
This declaration says that SomeFunction takes three parameters. The first is a reference to a Person object, the second is a pointer to a House object, and the third is an integer. It returns a pointer to a CAT object.

The question of where to put the reference (\&) or indirection (*) operator when declaring these variables is a great controversy. You may legally write any of the following:

1: CAT\& rFrisky;
2: CAT \& rFrisky;
3: CAT \&rFrisky;

Note: Whitespace is completely ignored, so anywhere you see a space here you may put as many spaces, tabs, and new lines as you like

Setting aside freedom of expression issues, which is best? Here are arguments for all three:

The argument for case 1 is that rFrisky is a variable whose name is rFrisky and whose type can be thought of as "reference to CAT object." Thus, this argument goes, the \& should be with the type

The counter-argument is that the type is CAT. The \& is part of the "declarator," which includes the variable name and the ampersand. More important, having the \& near the CAT can lead to the following bug:

CAT\& rFrisky, rBoots;
Casual examination of this line would lead you to think that both rFrisky and rBoots are references to CAT objects, but you would be wrong. This really says that rFrisky is a reference to a CAT, and rBoots (despite its name) is not a reference but a plain old CAT variable. This should be rewritten as follows:

CAT \&rFrisky, rBoots;

The answer to this objection is that declarations of references and variables should never be combined like this. The right answer is

CAT\& rFrisky;
CAT boots;
Finally, many programmers opt out of the argument and go with the middle position, that of putting the \& in the middle of the two, as illustrated in case 2 .

Of course, everything said so far about the reference operator ( $\&$ ) applies equally well to the indirection operator (*). The important thing is to recognize that reasonable people differ in their perceptions of the one true way. $^{\text {a }}$ Choose a style that works for you, and be consistent within any one program; clarity is, and remains, the goal

Many programmers like the following conventions for declaring references and pointers

```
1. Put the ampersand and asterisk in the middle, with a space on either side.
```

2. Never declare references, pointers, and variables all on the same line.
```
// Listing 9.13
    // Returning a reference to an object
    include <iostream.h>
    class SimpleCat
public:
        SimpleCat (int age, int weight);
        ~simplecat ()
        Mint GetAge() { return itsAge; },
    My.
        int itsAge;
    SimpleCat::simpleCat(int age, int weight)
        itsAge = age; ;
    SimpleCat &TheFunction()
    int main()
        SimpleCat &rCat = TheFunction();
        SimpleCat &rCat = TheFunctin
    return 0
    SimpleCat &TheFunction()
        SimpleCat Frisky (5,9)
```

Outpu
.$/ 1$ st09-13.cxx: In function ‘class simplecat $\&$ TheFunction (1)':
.$/ 1$ sto9-13.cxx:38: warning: reference to local variable 'Frisky' returned
Caution: The GNU compiless will not compile this program. It also will not compile on the Borland compiler,
but i will on Microsoft compilers.

Reurning a reference to ilocal varable is a bad codng practice

```
$/ Listing 9.14 _
class SimpleCat
class S
        Mimplecat (in
        Mimplecat() () weight)
        \mathrm{ int GetAge() {return itsAge; },}
        l
    SimpleCat::SimpleCat(int age, int weight)
        itsAge = age; ;
    SimpleCat & TheFunction();
    int main()
        SimpleCat & rCat = TheFunction();
        col
        cout<<"&rcat:"<< &rcat<< endl;
        Simplecat * pCat = &rCat;
    return 0
    SimpleCat &TheFunction()
    Simplecat * pFrisky = new SimpleCat (5,9);
    cout<<"pFrisky;
```

Ourver


Three solutions exist to this problem. The first is to declare a SimpleCat object on line 28 and to return that cat from TheFunction by value. The second is to go ahead and declare the SimpleCat on the free store in TheFunction(), but have TheFunction() return a pointer to that memory. Then the calling function can delete the pointer when it i done

The third workable solution, and the right one, is to declare the object in the calling function and then to pass it to TheFunction() by reference.

## Pointer, Pointer, Who Has the Pointer?

When your program allocates memory on the free store, a pointer is returned. It is imperative that you keep a pointer to that memory because after the pointer is lost, the memory cannot be deleted and becomes a memor leak.

As you pass this block of memory between functions, someone will "own" the pointer. Typically, the value in the block will be passed using references, and the function that created the memory is the one that deletes it. Bi this is a general rule, not an ironclad one

It is dangerous for one function to create memory and another to free it, however. Ambiguity about who owns the pointer can lead to one of two problems: forgetting to delete a pointer or deleting it twice. Either one can cause serious problems in your program. It is safer to build your functions so that they delete the memory they create.

If you are writing a function that needs to create memory and then pass it back to the calling function, consider changing your interface. Have the calling function allocate the memory and then pass it into your function by解 delete it

Do Don't
DO pass parameters by value when you must. DON'T pass by reference if the item referred to may go out of DO return by value when you must. scope.

## Summary

Today you learned what references are and how they compare to pointers. You saw that references must be initialized to refer to an existing object and cannot be reassigned to refer to anything else. Any action taken on reference is in fact taken on the reference's target object. Proof of this is that taking the address of a reference returns the address of the target
You saw that passing objects by reference can be more efficient than passing by value. Passing by reference also allows the called function to change the value in the arguments back in the calling function.
You saw that arguments to functions and values returned from functions can be passed by reference, and that this can be implemented with pointers or with references.
You saw how to use pointers to constant objects and constant references to pass values between functions safely while achieving the efficiency of passing by reference

## Q\&A

Q Why have references if pointers can do everything references can?
A References are easier to use and to understand. The indirection is hidden, and no need exists to repeatedly dereference the variable.
Q Why have pointers if references are easier?
A References cannot be null, and they cannot be reassigned. Pointers offer greater flexibility but are slightly more difficult to use
Q Why would you ever return by value from a function.
Af the object being returned is local, you must return by value or you will be returning a reference to a non-existent object.
Q Given the danger in returning by reference, why not always return by value?
Far greater efficiency is achieved in returning by reference. Memory is saved and the program runs faster.

## Workshop

The Workshop contains quiz questions to help solidify your understanding of the material covered and exercise to provide you with experience in using what you learned. Try to answer the quiz and exercise questions befor hecking the answers in Appendix D, "Answers to Exercises and Quizzes," and make sure you understand the answers before continuing tomorrow.

## Quiz

1. What is the difference between a reference and a pointer?

When must you use a pointer rather than a reference?
What does new return if there is insufficient memory to make your new object
4. What is a constant reference?
5. What is the difference between passing by reference and passing a reference?

## Exercises

. Write a program that declares an int, a reference to an int, and a pointer to an int. Use the pointer and he reference to manipulate the value in the int.
2. Write a program that declares a constant pointer to a constant integer. Initialize the pointer to an
iteger variable, varOne. Assign 6 to varOne. Use the pointer to assign 7 to varOne. Create a second integer variable, varTwo. Reassign the pointer to varTwo. Do not compile this exercise yet.
Now compile the program in Exercise 2. What produces errors? What produces warnings?
Write a program that produces a stray pointer.
5. Fix the program from Exercise 4.

Write a program that produces a memory leak.
. Fix the program from Exercise 6 .
8. BUG BUSTER: What is wrong with this program?

```
#include <iostream.h>
class CAT
i
    public:
            CAT(int age) { itsAge = age; }
            ~CAT() {}
            int GetAge() const { return itsAge;}
            private:
            int itsAge;
};
CAT & MakeCat(int age)
int main()
int age = 7;
CAT Boots = MakeCat(age);
cout << "Boots is " << Boots.GetAge() << " years old\n";
return 0;
}
CAT & MakeCat(int age)
    CAT * pCat = new CAT(age)
    return *pCat;
}
```

9. Fix the program from Exercise 8.

## Overloading Member Functions








/ Const ructor inple enentat ion
Rectangle: : Rectangle (int width, int height)


Prawshape ( itswidth, itssei ight)

for (int $\mathrm{i}=0 ; \mathrm{i}$ Cheight; $\mathrm{i}+\mathrm{+}$ )
for (int $j=0 ; j<$ width; $j+$ )
cout < " "x
/ Oriver program to demonstrate overloaded functions
int main ()



$\qquad$
$\qquad$
$\qquad$

|  |
| :---: |
|  |  |
|  |  |
|  |  |

Outpur

## 

## Choosing Between Default Values and Overloaded Functions






## erstanding Default Constructor



 Overloading Constructors

$/ / / \begin{aligned} & \text { List ing } 10.3 \\ & \text { Overioading }\end{aligned}$
*include <iostream.h>
class Rectangle
public:
Rectangle ();
Rectangle (int width, int length),
Recenget

int itswidth;
int itstength;
Rectangle: : Rectang1e()
itswidth $=5 ;$
itsLength $=10 ;$
int main()
 int awidth, alength;
cout
cin
cinterner a widt cin
cout < N NnEtee
cin $\gg$ atength;


Ineut

ectangle: : Rectangle
it swidth ( width),

Ourw
 as much as possible.

## Using the Copy Constructor

In add dition to providing a defaut constructor and destructor the com
The copy constructor is called every ine a coppy of a $n$ obiect is made

 cat (const cat \& thecat);

Here the car construcor t,
is to make a copy of thecat




## 




$$
\begin{aligned}
& x_{\text {it tsWeight }}=9 ; \\
& \text { CAT (const CAT } \& \text { rh }
\end{aligned}
$$


$\qquad$


Output



# // Listing 10.12 

$\underset{\substack{\text { class } \\ \text { publii } \\ \text { Co } \\ \text { Con }}}{\text { a }}$
Cuic:
Counter (1)
counter

const counter operator+ (1) $i=x ; / /$ prefix
const counter operatort+ (int); $/ /$ postfix
private
int
itsvali

Counterغ Counter: : operator++()
$\underset{\substack{++i t \text { svali } \\ \text { return } x \text { this }}}{ }$
const Counter Counter:: operator++(int X)
Counter temp (*this)
++ itsvali
$\xrightarrow{+ \text { teturvali }}$ temp;






## $\mathrm{int}_{\text {main(1) }}$




## return 0 ;



$)^{\text {return o; }}$
Outrut


varnhree
than osy:
varThree = varone. Add (varTwo),
Nota big change, but enough to make the progam easier to use and undessand.

${ }^{\text {Operator Overioading: Binary Operators }}$
 xample 1

唔
Counter Counter: : operator-(const Counter \& rhs);

## 




son Operator Overloadin

 What to overioad










 Invot
// Listing 10.15
\#include <iostream.h>
${ }^{\text {class }}$ Cat
public:
cars
copy
construct



cat: :carl)

if (this $==$ erhs)
return $n$ this $s$,


Outrut

## - Corsion

 user defined
Conere objec
$\qquad$
Invot

```
| Listing 10.16't Inmis code won't compile!
include <iostream.h>
    ass counter
    c
```



```
    l
main()
```




 out, uness you tell it, that given an int, t, ithould assign nat hav value to the member variable isval.

| rivate |
| :---: |
| int $i$ itsval |

Counter: : Counter ():
itsval ( 0 :



Outrut

It is a tundamental aspect of human intelligence to seek out, recognize, and create relationsthips among
concerss. We build hierarchies, matrices, newworks, and other interelationships oexplain and understanc concepts. We build hierarchies, matrices, networks, and other interereationsthips to explain and understand dha
ways in which things interact. C + antempst to cappure
this in in ineritance hierarchies. Today you will learn:

- What inheritance is
- How to derive one class from another

What protected access is and how to use it
What virtual functions are

## What Is Inheritance?

What is a dog? When you look at your pet, what do you see? I see four legs in service to a mouth. A biologis network of interacting organs, a physicist sees atoms and forces at work, and a taxonomist sees a representative of the species Canis domesticu.

It is that last assessment that interests us at the moment. A dog is a kind of canine, a canine is a kind of nammal, and so forth. Taxonomists divide the world of living things into Kingdom, Phylum, Class, Order mammal, and so orth. Taxo
Family, Genus, and Species. This hierarchy establishes an is-a relationship. A member of the species Homo sapiens is $a$ kind of primate. Yt
see this relationship everywhere: A station wagon is $a$ kind of car, which is $a$ kind of vehicle. A sundae is $a$ kir of dessert, which is a kind of food.

## When we say something is ak car is a special kind of vehicle

Any object that is is-a has some characteristics of the higher order description. A car is a vehicle. So is a A car is not a truck. But they both have certain similar characterisics that they share with the common

## Inheritance and Derivation

The concept dog inherits-that is, it automatically gets-all the features of a mammal. Because it is a mamma we know that it moves and that it breathes air. All mammals, by definition, move and breathe air. The concept of a dog adds the idea of barking, wagging its tail,
done, barking when I'm trying to sleep, and so on.

We can divide dogs into working dogs, sporting dogs, and terriers; and we can divide sporting dogs into retrievers, spaniels, and so forth. Finally
subdivided into Labradors and Goldens.
A Golden is a kind of retriever, which is a sporting dog, which is a dog, and thus a kind of mammal, which is kind of animal, and therefore, a kind of living thing. This hierarchy is represented in Figure 11.11 using the
Unified Modeling Language (UML) Arrows point from more specialized types to mere C++; attempts to represent these relationships by enabling you to define classes that derive from one anoth Derivation is a way of expressing the is-a relationship. You derive a new class, Dog,
do not have to state explicitly that dogs move because they inherit that from Mamma

A class that adds new functionality to an existing class is said to derive from that original class. The origin class is said to be the new class's base class
If the Dog class derives from the Mammal class, Mammal is a base class of Dog. Derived clases erest their base classes. Just as dog adds $c$ e
methods or data to the Mammal class.

| $\stackrel{\square}{\square}$ |
| :---: |
| - |
| 㟋 |

t

## Building Classes that Represent Animals

To facilitate the discussion of derivation and inheritance, this lesson will focus on the relationships among, number of classes representin,

In time you will develop a whole set of farm animals, including horses, cows, dogs, cats, sheep, and so forth You will create methods for these classes so that they can act in ways he child might expect, but for nc you can stub-out each method with a simple print statement

Stubbing-out a function means that you write only enough to show that the function was called, leaving the details for later when you have more time.
enable the animals to act more real istically.

## The Syntax of Derivatio

When you declare a class, you can indicate what class it derives from by writing a colon after the class name the type of derivation (public or otherwise), and the class from which it derives. The following is an exampl class Dog : public Mamma

The type of derivation will be discussed later in this lesson. For now, always use public. The class from whic you derive must have been declared earlier, or you wiil get a compiler error. Listing 11.1 illustrates how to
declare a Dog class that is derived from a INPUT Listing 11.1 Simple Inheritance
//Listing 11.1 Simple inheritance
include <iostream.h>
class Mammal
public:
// constructors
Mammal() ;
${ }_{\sim}^{\text {Manmal () }}$;
//accessors
void SetAge (int);
int GetWeight () const
//other methods
void Speak () const;
void Sleep () const;
protected:
int itsAge;
int itsWeight;
lass Dog : public Mammal
public:
// Const
Dog();
$\xrightarrow[\sim \operatorname{Dog}(1) ;]{ }$
// Accessors
BREED GetBreed) const
void SetBreed (BREED)
// other methods
WagTail();
BegForFood
rotected:
BREED itsBreed;

## Output

The hierarchy has to begin somewhere; this program begins with Mammal. Because of this decision, some member variables that might properly belong in a higher base class are now represented here. Certainly al animals have an age and weight, for example, if Mammal is derived from Animal, we might expect to inheri those attributes. As it is, the attributes appear in the Mammal class

To keep the program reasonably simple and manageable, only six methods have been put in the Mamm class-four accessor methods, Speak 0 , and Sleep $($.

The Dog class inherits from Mammal, as indicated on line 29. Every Dog object will have three member variables itsAge, itsWeight, and itsBreed. Note that the class declaration for Dog does not include the member variables its A and itsWeight. Dog objects inherit these variables from the Mammal class, along with all Mammal's methods except the copy operator and the constructors and destructor

## Private Versus Protected

You may have noticed that a new access keyword, protected, was introduced on lines 24 and 45 of Listing 11.1. Previously, class data had been declared private. However, private members are not available to derived classe You could make itsAge and itsWeight public, but that is not desirable. You do not want other classes accessing these data members directly

What you want is a designation that says, "Make these visible to this class and to classes that derive from this class." That designation is protected. Protected data members and functions are fully visible to derived classes but are otherwise private

In total, three access specifiers exist: public, protected, and private. If a function has an object of your class, it can access all the public member data and functions. The member functions, in turn, can access all private dat members and functions of their own class, and all protected data members and functions of any class from which they derive

Thus, the function Dog::WagTailo can access the private data itsBreed and can access the protected data in the Mammal class.

Even if other classes are layered between Mammal and Dog (for example, DomesticAnimals), the Dog class will still be able to access the protected members of Mammal, assuming that these other classes all use public inheritance. be able to access the protected members of Mammal, assuming that ti"
Private inheritance is discussed on Day 15, "Advanced Inheritance."

Listing 11.2 demonstrates how to create objects of type Dog and access the data and functions of that type.

```
INPUT. Listing 11.2 Using a Derived Objec
```

```
/Listing 11.2 Using a derived object
enum BREED { GOLDEN, CAIRN, DANDIE, SHETLAND, DOBERMAN, LAB },
class Mammal
public:
    // constructors
    e(2), itsWeight (5){
    Mammal() {}
    /accessors
    int GetAge()const { return itsAge; }
    oid SetAge(int age) { itsAge = age; }
    int GetWeight() const { return itsWeight;
    void SetWeight(int weight) { itsWeight = weight; }
    //Other methods
    void Speak()const { cout << "Mammal sound!\n"; }
protected:
    int itsAge;
        int itsWeight
};
class Dog : public Mammal
public:
    // Constructors
    Dog():itsBreed (GOLDEN) {}
    ~Dog() {}
    // Accessors
    BREED GetBreed() const { return itsBreed;
    void SetBreed(BREED breed) { itsBreed = breed; }
    // Other methods
    void WagTail() const { cout << "Tail wagging...\n"; }
private
    BREED itsBreed;
};
int main()
Dog fido;
    fido.Speak();
    fido.WagTail()
    out << "Fido is " << fido.GetAge() << " years old\n";
return 0
```

OUtrut
Mammal sound!
Tail wagging...

Analysis
2-47, the On lines 6-27, the Mammal class is declared (all its functions are inline to save space here). On line 29-47, the Dog class is declared as a derived class of Mammal. Thus, by these declarations, all Dogs have an age, a weight, and a breed

On line 51, a Dog is defined: Fido. Fido inherits all the attributes of a Mammal, as well as all the attributes of a Dog. Thus, Fido knows how to WagTail(), but he also knows how to Speak() and Sleep(0.

## Constructors and Destructors

Dog objects are Mammal objects. This is the essence of the is-a relationship. When Fido is created, his base constructor is called first, creating a Mammal. Then the Dog constructor is called, completing the construction of the Dog object. Because we gave Fido no parameters, the default constructor was called in each case. Fido does constructed. Thus, both constructors must be called. When Fido is destroyed, first the Dog destructor will be called and then the destructor for the Mammal part of Fid
Each destructor is given an opportunity to clean up after its own part of Fido. Remember to clean up after your Dog! Listing 11.3 demonstrates this
cout < Memamal


Sraed (conomen)






'


1: varimur and initializes itsWeight with the value 5 .

Dog has overloaded five constructors on lines 35-39. The first is the default constructor. The second takes the Dog has overloaded five constructors on lines $35-39$. The first is the default constructor. The second takes the
age, which is the same parameter that the Mammal constructor takes. The third constructor takes both the age ar the weight, the fourth takes the age and the breed, and the fifth takes the age, the weight, and the breed.

Note that on line 74, Dog's default constructor calls Mammal's default constructor. Although it is not strictly necessary to do this, it serves as documentation that you intended to call the base constructor, which takes no parameters. The base constructor would be called in any case, but actually doing so makes your intentions explicit.

The implementation for the Dog constructor, which takes an integer, is on lines 80-85. In its initialization phase (lines 81-82), Dog initializes its base class, passing in the parameter, and then it initializes its breed.

Another Dog constructor is on lines $87-93$. This one takes two parameters. Once again it initializes its base clas by calling the appropriate constructor, but this time it also assigns weight to its base class's variable itsWeight. Note that you cannot assign to the base class variable in the initialization phase. Because Mammal does not hav a constructor that takes this parameter, you must do this within the body of the Dog's constructor

Walk through the remaining constructors to make sure you are comfortable with how they work. Note what is initialized and what must wait for the body of the constructor
The output has been numbered so that each line can be referred to in this analysis. The first two lines of output represent the instantiation of Fido, using the default constructor

In the output, lines 3 and 4 represent the creation of rover. Lines 5 and 6 represent buster. Note that the Mamma nstructor that was called is the constructor that takes one integer, but the Dog constructor is the constructor that takes two integers.

After all the objects are created, they are used and then go out of scope. As each object is destroyed, first the Dog destructor and then the Mammal destructor is called, five of each in total.

## Overriding Functions

A Dog object has access to all the member functions in class Mammal, as well as to any member functions, such as WagTail), that the declaration of the Dog class might add. It can also override a base class function. Overriding a function means changing the implementation of a base class function in a derived class. When you make an object of the derived class, the correct function is called

When a derived class creates a function with the same return type and signature as a member function in the base class, but with a new implementation, it is said to be overriding that method.

When you override a function, it must agree in return type and in signature with the function in the base class The signature is the function prototype other than the return type: that is, the name, the parameter list, and the keyword const, if used.

Listing 11.5 illustrates what happens if the Dog class overrides the Speak() method in Mammal. To save room, the accessor functions have been left out of these classes
isting 11.5 Overriding a Base Class Method in a Derived Class

```
//Listing 11.5 Overriding a base class method in a derived class
#include <iostream.h>
M BREED { GOLDEN, CAIRN, DANDIE, SHETLAND, DOBERMAN, LAB }
class Mammal
{
        // constructors
        Mammal() { cout << "Mammal constructor...\n"; }
        //Other methods
        void Speak()const { cout << "Mammal sound!\n"; }
        void Sleep() const { cout << "shhh. I'm sleeping.\n"; }
protected:
        nt itsAge;
    };
    class Dog : public Mammal
    public:
        // Constructors
        ~Dog(){ cout << "Dog destructor...\n"; ; }
        // Other methods
        void WagTail() const { cout << "Tail wagging...\n"; }
        void BegForFood() const { cout << "Begging for food...\n"; }
        void Speak() const { cout << "Woof!\n"; }
    private:
        BREED itsBreed;
    int main()
        Mammal bigAnimal;
        Dog fido;
        MAnimal.Speak()
        fido.Speak
    return 0;
```


## Output

```
/Listing 11.6 Hiding methods
```

finclude <iostream.h>
lass Man

cout << "Mannal move ";
cout << distance <<" steps. $\backslash$ n
int itsAge;
解 itsAge;
ass Dog : public Mamma
public:
other compilers may give a warning that you are hiding a function!
void Move () const
nt main()
Marnal bi
Dog fido
bigAnimal
big
bigAnimal. Move ();
bigAnimal. Mov
fidoo Movel);
/I fido. Move (10);
ceturn $0 ;$

## Output

Mannal move one ste
Mammal move 2 steps

Analusis
Ass declares the verarloaded mods and data have been removed from these classes. On Iines 8 and 9 ,

Line 33 . however, is commented out because it causes a compile-time error. Althoush the Doe class could ha
called the Movecint method if it had not veeridden the version of Moveo without parameters, now that t has done so, it must verride both if it wants to use both. Otherwise, it will hide the method that it does not veerrid
This is reminisent of the rule that if you supply any constructor, the compiler will no longer supply a default enstructor. The GNU compilers give the following message if you uncomment line 33:

It is a common mistake to hide a base class method when you intend to override it, by forgeting to include


Overriding Versus Hidin
In the next section, virtual methods are described. Overriding a virtual method support

## Calling the Base Metho

If you have overridden the base method, it is still possible to call it by fully qualifying the name of the method
You do this by writing the base name, followed by two colons and then the method name, for example: You do this by wri
Mammal:Move).
ido. Marmal: : Move (10) ;
/Listing 11.7 Calling base method from overridden method.
include <iostream.h>
ass Mannal
void Move () const $f$ cout <<"
void Move (int distance) const cout << "Manmal move"" << distance,
cout $\ll$ " staps. ${ }^{\text {n"; }}$;
ected:
int itsAge;
ass Dog : public Manma
ublic:
void Dog: :Move() cons
cout << "In dog move
Mammal: : Move (3);
$t$ main
Mannal biganimal
Dog fido;
bigAnimal
IgAnimal. Move (2);
fido. Mammal: $:$ Move ( 6 )

Output
Mannal move 2 steps
Mamnal move 6 steps
ANALYSIS On line 35 , a Mammal, bieanimal, is created, and on line 36 , a Dos, fido, is created. The method call The programmer wanted to invoke Movecint on the Dog object, but had a problem. Dog overides the Move0
method, but does not overload it and does not provide a version that takes an in. This is solved by the explici call to the base class Move(int) method on line 38 .

When a pointer to a Mammal is used, the vpru continues to point to the correct function, depending on the "real
type of the object. Thus, when Speak) is invokede, the correct function is invoked.

## You Cannot Get There from Here

 one in a derived class. Normally, you would want to execute ?
The term virtual function magic is often applied to this proces



```
//Listing 11.10 Data slicing when passing by value
```

include <iostream.h>

```
class Mammal
```

$\underset{\substack{\text { public: } \\ \text { Marma } \\ \text { virtu }}}{\substack{\text { (1ass }}}$

protected:
int itsAge $;$
Class Dog : public Mammal
public:
class Cat : public Mammal
public:
void Speak () const ( cout << "Meow! \n";
void ValueFunction
void PtrFunction
(Mannal);
(Mannal*) ;
vid RefFunction
int main()
Manmal $*$ ptr=NULL;
int choice;
int choice
while (1)
bool feuit = false;
cout <<"(1) dog (2) cat (0) Quit: "
cin $\gg$ choice;
switch (choice
break;
case 1:
break;
case
break;
bref
defal:
default: $\mathrm{ptr}=$ new Mamnal;
reak;
if (fQuit)

ValueFunction (*ptr),
return 0 ;
id ValueFunction (Mamnal Manmalvalue)
MannalValue.Speak();
oid PtrFunction (Manmal * pMarmal)
pManmal->Speak ();
id RefFunction (Mammal \& rMammal)
rMammal.Speak ();


```
//Listing 11.11 Virtual copy constructor
```

\#include <iostream.h>
class Mammal
public:


virtual void Speak () const i cout << "Manmal speak! \n"; \}
irtual Mammal* Clone () $\{$ return new
int GetAge() const \{ return itsAge; $\}$
rotected:
int itsAge

cout << "Mammal Copy Constructor...\n";
class Dog : public Manmal
public:

virtual (const Dog \& rhs)
Dog
oid
Did Speak ()const $\{$ cout << "Woof! \n"; \}
virtual Mammal* Clone() \{return new Dog (*this); \}
Dog::Dog(const Dog \& rhs):
Mammal(rhs)
cout << "Dog copy constructor
lass Cat : public Marmal
public:



Cat: :Cat (const Cat \& rhs):
cout << "Cat copy constructor
enum Animals \{ MAMMAL, DOG, CAT\};
const int NumAnimaltypes $=3$;
const int
int $\operatorname{main}()$
Mammal *theArray [NumAnimalTypes];
Manmal* ptr;
int choice,
int choice, $i ;$
for $(i=0 ; i<N u m A n i m a l$ Types; $i++)$
cout <<"(1)dog (2) cat (3) Mammal: "
cin >> choice;
switch (choice
case DoG: ptr = new Dog;
break;
case CAT:
btr $=$ new Cat;
case cAT: ptr = new Cat;
break;
default ptr = new Mamma
default:
break; ptr = new Mammal;
theArray[i] = ptr
Marmal *OtherArray [NumAnimalTypes];
theArray[i]->
theArray [i]->Speak (1);
OtherArray $[$ i] $=$ theAr
or ( $\mathrm{i}=0 ; i \operatorname{i}$ NumAnimalTypes;
OtherArray $[\mathrm{i}]$->Speak();
eturn 0 ;

ANALYsIS Listing 11.11 is very similar to the previous two listings, except that a new virtual method has bee added to the Mammal class: Cloneo. This method returns a pointer to a new Mammal object by calling the copy structor, passing in itself ("this) as a const reference.

Dog and Cat both override the Clone/ method, initializing their data and passing in copies of themselves to their
own copy cons own copy constructors. Because Clone) is virtual, this will effectively create a virtual copy constructor, as shown on line 81

## The user is prompted to choose dogs, ca choice is stored in an array on line 75 .

As the program iterates over the array, each object has its Speako and its Cloneo methods called, in turn, on line
80 and 8 . The result of the Cloneo call is a pointer to a copy of the object, which is the stored in a seond 80 and 8
on line 8

On line 1 of the output, the user is prompted and responds with 1 , choosing to create $a$ dog. The,
constructors are invoked. This is repeated for Cat and for Mammal on lines 4 and 7 of the output.
Line 9 of the output represents the call to Speak0 on the first object, the Dog. The virtual Speak0 method is callec and the correct version of Speak) is invoked. The Cloneo function is then called, and because this is also virtua Dog's Cloneo method is invoked, causing the Mammal constructor and the Dog copy constructor to be called. The same is repeated for Cat on lines 12-14, and then for
iterated, and each of the new objects has Speako invoked.

## The Cost of Virtual Methods

Because objects with virtual methods must maintain a v-table, some overhead occurs in having virtual methods If you have a very small class from which you do not expect to derive other classes, there may be no reason to have any virtual methods at all

When you declare any methods virtual, you have already paid most of the price of the v-table (although each entry does add a small memory overhead). At that point, you want the destructor to be virtual, and the assumption will be that all other methods probably will be virtual as well. Take a long, hard look at any nonvirtual methods, and be certain you understand why they are not virtual.

DO use virtual methods when you expect to derive from a class
DON'T mark the constructor as virtual DO use a virtual destructor if any methods are virtual

## Summar

Today you learned how derived classes inherit from base classes. This lesson discussed public inheritance and virtual functions. Classes inherit all the public and protected data and functions from their base classes

Protected access is public to derived classes and private to all other objects. Even derived classes cannot access private data or functions in their base classes

Constructors can be initialized before the body of the constructor. At that time, the base constructors are invoked and parameters can be passed to the base class

Functions in the base class can be overridden in the derived class. If the base class functions are virtual, and if the object is accessed by pointer or reference, the derived class's functions will be invoked, based on the runtime type of the object pointed to

Methods in the base class can be invoked by explicitly naming the function with the prefix of the base class name and two colons. For example, if Dog inherits from Mammal, Mammal's walk() method can be called with Mammal::walk().

In classes with virtual methods, the destructor should almost always be made virtual. A virtual destructor ensures that the derived part of the object will be freed when delete is called on the pointer. Constructors cannot be virtual. Virtual copy constructors can be effectively created by making a virtual member function that calls the copy constructor

## Q\&A

Q Are inherited members and functions passed along to subsequent generations? If Dog derive from Mammal, and Mammal derives from Animal, does Dog inherit Animal's functions and data?
A Yes. As derivation continues, derived classes inherit the sum of all the functions and data in all their base classes.

## Q If, in the preceding example, Mammal overrides a function in Animal, which does Dog get, the

 original or the overridden function?A If Dog inherits from Mammal, it gets the function in the state Mammal has it: the overridden function. Q Can a derived class make a public base function private?
A Yes, the derived class can override the method and make it private. It then remains private for all subsequent derivation.
Q Why not make all class functions virtual?
A Overhead occurs with the first virtual function in the creation of a v-table. After that, the overhead is trivial. Many C++ programmers feel that if one function is virtual, all others should be. Other programmers disagree, feeling that there should always be a reason for what you do.
Q If a function (SomeFunc()) is virtual in a base class and is also overloaded, so as to take either an integer or two integers, and the derived class overrides the form taking one integer, what is called when a pointer to a derived object calls the two-integer form?
A The overriding of the one-int form hides the entire base class function, and thus you will get a compile error complaining that that function requires only one int.

## Workshop

The Workshop provides quiz questions to help you solidify your understanding of the material that was covere and exercises to provide you with experience in using what you learned today. Try to answer the quiz and exercise questions before checking the answers in Appendix D, "Answers to Exercises and Quizzes," and mak sure you understand the answers before continuing to tomorrow's lesson.

Quiz

1. What is a $v$-table?
2. What is a virtual destructor?
3. How do you show the declaration of a virtual constructor?
4. How can you create a virtual copy constructor?
5. How do you invoke a base member function from a derived class in which you've overridden that function?
6. How do you invoke a base member function from a derived class in which you have not overridden that function?
7. If a base class declares a function to be virtual, and a derived class does not use the term virtual when overriding that class, is it still virtual when inherited by a third-generation class?
8. What is the protected keyword used for?

## Exercises

1. Show the declaration of a virtual function that takes an integer parameter and returns void.
2. Show the declaration of a class Square, which derives from Rectangle, which in turn derives from Shape 3. If, in Exercise 2, Shape takes no parameters, Rectangle takes two (length and width), but Square takes only one (length), show the constructor initialization for Square.
3. Write a virtual copy constructor for the class Square (in Exercise 3)
4. BUG BUSTER: What is wrong with this code snippet?
void SomeFunction (Shape);
Shape * pRect = new Rectangle;
SomeFunction(*pRect);
5. BUG BUSTER: What is wrong with this code snippet?
```
class Shape()
{
public:
    Shape();
    virtual ~Shape();
    virtual Shape(const Shape&);
```

\};

Chapter 12
Arrays, C Strings, and Linked Lists
In previous lessons, you declared a single int, char, or other object. You often want to declare a collection of
objects, such as 20 ins or a liter of c AAss. Today, you will learn:

- What arrays are and how to declare them
What strings are and how to use character

What strings are and how to use character arrays to make the
The raloioship between arrays and pointers
How to use
The ereationship betwen arrays and pointers
How to use pointer arithmetic with arrays

## What Is an Array?

An array is a collection of data storage locations, each of which holds the same type of data. Each storage
location is called an elemento f the array
 represens the distance fom



## Accessing Array Element

 second, and so forth.

This can be somewhat confusing. The aray Some Arayy[3] has three elements. They are Somearray[0]
Someatray 11 , and SomeAray $[2] 1$ ough Some A tray $n-1]$.

INPUT Listing 12.1 Using an Integer Array

```
#/Listing 12.1 - Arrays
int mainn
    int myArray[5];
        i=0; i<5; i++) // 0-4
        cout << "Value for myArray[" << i << "]:";
    for (1)=0,1<5, 1+ (
    *)
```

ANalvsIS Line declares an array called myArray, which holds five integer variables. Line 8 establishes a lo that counts from 0 through 4 , which is the proper set of offsets for a five-element array. The user is prompled
for a value, and that value is saved at the correct offset into the array.
The first value is saved at myA Aray[0], the second at my yaray 11 , and so forth. The second for loop prints each
value to the screen.


## Writing Past the End of an Array

When you write a valu to an element in an array, the compier computes where to tore the value based on the
size of each element and the subscript. Suppose that you ask to write over the value at Long toryy 5 , which is the


If you ask to write at Long trayy 50 I), the compiler ignores the fact that no such element exists. It computes how
faar past the first element it should look (200 byyess and then writes over whatever is at that location. This can far past he first element titshould Look (200 bytes) and then write enperitable resulst that If ocation. This cant
virtually any data, and writing your new value your program will crash immediately. If not you will eventully get strange results much later in your program, an
you will have a difficult time figuring out what went wrong. In a robust operating system like Linux, a program crash will no affect the operating system itself. In a non
robust environment flike MS-DOS robust enviroment (like MS-DOS), a program can overwitt portions of the operating system, causing your
machine to foeeze. These problems are particularly difficult to debugg In rare cases, a rogue program can affec
m. Linux-primarily when it truns with root privilieges.




Input Listing 12.2 Writing Past the End of an Array
$/$ Listing 12.2
Demonstrates
of an array
$\underset{\substack{\text { *include } \\ \text { int main() } \\ \text { (iostream.h> }}}{ }$
/ sentinels
long sentinel
long sentinelone [3],
long TargetArray [25

int $i ; \quad i \quad i=0, i<3 ; i++$
for $(i=0$,
sentinelone $[i]$



${ }^{\prime}$ cout << "sentinelo



cout $\ll$ "sent inelone $[$ "
cout $\ll$ sentinelone $[$ i

$\square$

## Fence Post Errors

## Initializing Arrays

 For xample,
Int Integerarray [5] $=\{10,20,30,40,50)$;

cout << "The value at Tuesday is: " << ArrayWeek (Tue)
return 0;

OUtput

Analvsis Laysinhweek is cequal to C .
Line 11 uses the e enumerated constant 7 at a a an offset into the array.
element of the array, Aray Weekl2), is returned and printed in line 11 .
$\longdiv { \text { Arrays } }$
To declare an array, write the type of oject stored, followed by the name of the array and a subscript with til
number of obiects ot obe held in the earay
Example
int MyIntegerarray 900 ;
Example:
1ong * ArrayofpointersToLongs [100];
To access n .
int theNinethinteger $=$ MyIntegerarray $[8]$
long * pLong $=$ ArrayOfPointersToLongs [8];
Arrays count from zero. An array of $n$ items is numbered from $0-n-1$.
 com much room is needed for each object based on the class declaration. The class must have a defau
how how much rom is needed or each object based on the class declaration. The class must have a.
constructor that takes no arguments so that the obiects san be created when the array is defined.


```
Misting 12.4 - An array of objects
```

include <iostream.h>
${ }_{1}{ }^{\text {class }}$ CAT
$\underset{\substack{\text { public: } \\ \text { CAT) } \\ \sim \text { CAT) } \\ \text { ( } \\ \text { (its }}}{\text { it }}$

ivate:
int itsfge,
int itsweight;
int main()
Car Litter [5];
int
for
in
int
for $(i, i=0 ; i<5 ; i+1)$
Liteer $[i]$. Setage $(2 * i+1)$,
for ( $\mathrm{i}=0 ; i<5 ; i++$ )

${ }_{r}^{\text {return } 0 ;}$
Output


Amwselinis.

The first of lop (lines 23 and 24 sest the age of each of the five CATs in the array. The second tor loop (line
$26-30$ ) accessess ach member of fthe aray and acals ceets geo.

```
ndividual CAT's Geateseo method is callea
```


oard and a two-dimensional array.
SQUARE Board [8] [8];
You could also represent the same data with a one-dimensional. 64 scgure a aray. For example
square board [64];

corresponds to
Board [0] [3]

Initializing Multidimensional Arrays
You can initiaize multidimensional arrays. Assign hhe lis of values to array elements in order. with hhe last
array subscript changing while each of the former holds steady. Therefore, if you have an array
int thearray [5] [3];

You initialize this array by writing
Int thearray $[5] 3]=\{1,2,3,4,5,6,7,8,9,10,11,12,13,14$,
For the sake of clarity, you could group the intializations with braces. For example:

$(11,11,12\})$
$(13,14,15)\} ;$
The compiler ignores the inner braces, which make it easier to understand how the numbers are distributed
Each value must be seppated by a comma, without regard to the braces. The enitie initialization set must be
within braces, and it must end with semicololo.
Listing 12.5 creates a two-dimensional aray. The first dimension is the
dimension consiss of the double of each valuc in the first dimension.
\#include <iostream.h>
int main()



OUtrut

#   

ivate:
int itsage;
int
itsweight,

CAT $*$ Fami 1 y $(5001$;



```
cout<<"cat #"><i+1<<"";";
```

pointer-is added to the array.

## Using a Pointer to an Array Versus an Array of Pointers

Examine the following three declaratic

Fanily On is an array of 500 CATs. FamilyTwo is an array of 500 pointers to CATs. FanilyThee is a pointer to an
arayy of 500 cats.
 This raises the thorny issue of how pointers relatet to arrays. In the third case, Family Three is a pointer to an ar
That is, the address in Family Thee is the address of the first item in that aray. This is sexactly the case for

## Using Pointers with Array Names

(
CAT Fanily $[50]$;
It i s legal to use array names as constant pointers, and vice versa. Therefore, Family +4 is a legitimate way of
accessing the data a accessing hie data a ar famly
 Family 4 is 16 byles past the start of the array. If each object is a CAT that has 4 long member variables of 4 the satro t t the array.

Input Listing 12.7 Creating an Array by Using neen
// Listing 12.7-An array on the free store
\#include <iostream.h>
${ }_{1}^{\text {class Cat }}$
$\underset{\substack{\text { public: } \\ \text { CAT) } \\ \sim \text { CAT() }}}{\text { Cist }}$

rivate:
int itsage;
int itsweight
cat : : ~cat ()
cout << "Destructor callea! ! \n";

CAT * Family $=$ new CAT [500) ;
for ( $\mathrm{i}=0 ; 1$ < $500 ; \mathrm{i}^{\text {+ }}$ )
Family [i] .SetAge (2 ${ }^{\star}{ }_{i}+1$ );
for (i=0; i<500; i+t)


ANAIYsIS
Line 26 declares the array Family, which holdd 500 CAT objects. The entire array is created on the
free store with the call to new CATI[500].

## Deleting Arrays on the Free Store

What happens to the memory allocated for these CAT objects when the array is destroyed? Is there a chance of memory leak? Deleting Family automatically returns all the memory set aside for the array if you use the delete j operator, remembering to include the square brackets. The compiler is smart enough to destroy each object in the array and to return its memory to the free store.
To see this, change the size of the array from 500 to 10 in lines 26,29 , and 34 . Then uncomment the cout so see this, change the size of nhe array from an the array is destroyed, each CAT object destructor is called.
statement in line 21 . When line 43 is reached and the When you create an item on the heap by using new, you always deletet that item and free its memory with delete.
Similarly, when you create an array by using new <class> \sizel you delete that array and free all its memory with ilarly, when you create an array by using new <class> (sizel], you delete that array and free all its memory wit delete[ I. The brackets signal the compiler that this array is being deleted

If you leave the brackets off, only the first object in the array will be deleted. You can prove this to yourself b removing the bracket on line 38. If you edited line 21 so that the destructor prints, you should now see only CAT object destroyed. Congratulations! You just created a memory leak

| Do |
| :--- |
| Do remember that an array of $n$ items is numbere |
| from zero through $n-1$. | an array.

## What are char Arrays?

A string is a series of characters. The only strings you have seen until now have been unnamed string constant used in cout statements, such as

In $\mathrm{C}_{++}$a string is an array of chars ending with a null character. You can declare and initialize a string the sam as you would any other array. For example

The last character, ${ }^{40}$, is the null character, which many $\mathrm{C}++$ functions recognize as the terminator for a strin , Although this character-by-character approach works, it is difficult to type and it adt
for error. $\mathrm{C}++$ enables you to use a shorthand form of the previous line of code. It is
char Greeting[] = "Hello World";

## You should note two things about this syntax:

quoted string, no commas, and no braces.
quoted string, no commas, and no braces.
The string Hello World is 12 bytes. Hello is 5 bytes, the space 1 , World 5 , and the null character
You can also create uninitialized character arrays. As with all arrays, it is important to ensure that you do no
put more into the buffer than there is room for

## INPUT. Listing 12.8 Filling an Arra

$$
\text { //Listing } 12.8 \text { char array buffers }
$$

\#include <iostream.h>
int main()
char buffer [80],
cout << "Enter the string: ";
cout << "Here's the buffer: " << buffer << endl
return 0 ;

$$
\text { eeturn } 0 \text {; }
$$

## Ourver

## ANALYSIS $\begin{aligned} & \text { On line } 7 \text {, a buffer is declared to hol } \\ & \text { string and a terminating null char }\end{aligned}$

On line 8, the user is prompted to enter a string, which is entered into buffer on line 9. It is the syntax of cin write a terminating null to buffer after it writes the strin,
Two problems occur with the program in Listing 12.8. First, if the user enters more than 79 characters, cii writes past the end of the buffer. Second, if the user enters a space, cin thinks that it is the end of the string it stops writing to the buffer.

- The buffer to fill

The maximum number of characters to ge
The delimiter to terminate input
The default delimiter is newline. Listing 12.9 illustrates its use
INPUT. Listing 12.9 Filling an Arras

```
/Listing 12.9 using cin.get()
Include <iostream.h>
nt main()
    Char buffer [80];
    cin.get (buffer, 79); ; %ing: "; get up to 79 or newlin
    cout return 0;
```

```
1: // listing 12.10 - strcpy, strncpy, strlen, and strcat
    #include <iostream.h>
    int main()
        char String1[] = "No man is an island";
        char String2[80] = "";
        char String3[80] = "";
        cout << "String1:" << String1 << endl;
        strcpy(String2,String1);
        cout << "String2:" << String2 << endl;
        strncpy (String3, String1, 5); // not entire string
        String3[5] = `\0'; % // you need a null terminator
    cout << "String3 after strncpy:" << String3 << endl;
        out << "String1 is" << strlen (String1) <<
        "bytes long, \n String2 is "}<
        strlen (String2)
        "bytes long, \n and String3 is " <<
        strlen (String3) << " bytes long" << endl;
    strcat(String3, String1);
    cout << "String3 after strcat:" << String3 << endl;
    cout << "String1 is still" << strlen (String1) <<
        " bytes long, \n String2 is still " <<
        strlen (String2)
        " bytes long,\n and String3 is now"<<
        strlen (String3) << " bytes long" << endl;
    return 0;
    }
```


## Output

String1: No man is an island String2: No man is an islano String3 after strncpy: No ma String1 is 19 bytes long,
String2 is 19 bytes long, and String3 is 5 bytes long
String3 after strcat: String3 after strcat: No maNo man is an island
String1 is still 19 bytes long,
String2 is still 19 bytes lon
and String3 is now 24 bytes long
Anatrs
ANALYSIS The header file string.h is included in line 3. This file contains the prototypes of the string function Lines 6,7 , and 8 allocate three character arrays; the first is initialized to "No man is an island," and the other two are empty.
On line 10 , the strcpy) function takes two character arrays-a destination followed by a source. If the source were larger than the destination, strcpy) would overwrite past the end of the buffer.

To protect against this, the standard library also includes strncpy) as shown on line 12. This variation takes a maximum number of characters to copy. strncpyO copies up to the first null character or the maximum numbe terminator is not included in the destination string So you have to put one there yourself as shown on line 1
strien() takes one character array argument and returns the length of that string up to, but not including, the nul terminator. This should be at least one less than the size of the array containing the string.

On line 20 , the strcat) function takes two character arrays-a destination followed by a source. The source arra will be appended to the end of the string stored in the destination array. If the size of source plus the size of th original destination were larger than the destination array can contain, strcat) would overwrite past the end of the buffer (strncat) can be used to prevent this).
Of course, there are many other string functions available

## Using Strings and Pointer:

As you learned earlier today, when you use the array name alone (without a subscript contained in squar brackets), you are referencing the address of the array. This variable is behaving like a pointer. It is a special pointer-a constant pointer. It is a constant pointer because you are not allowed to change where it points to The reason behind that is: arrays do not move; if you changed the pointer itself, you would not be able to reference the memory allocated for that array.

位e that in Listing 12.10, it is the array name that is used as arguments to he string functions. You can also use pointers to strings in place of arrays. The only difference is that you are not allowed to change a string in place (you are not allowed to concatenate another string to it as was done to the String3 array). The languag andard does not prohibit this but the compiler or operating system may not allow you (treating strings as be read-only). Even if you are not prevented from doing so, you should not change a string in place. Trust me on this one-it can cause all kinds of side effects and is not portable

```
1: // listing 12.11 - Strings and Pointers
    #include <iostream.h>
    include <string.h>
    nt main()
        char *String1 = "No man is an island"
        Char String2[80]=
        char String3i801="; << String1 << endl;
        strcpy(String2,String1);
        cout << "String2: " << String2 << endl;
        strncpy (String3, String1, 5); // not entire string
        String3[5] = `\0'; // you need a null terminator
        cout << "String3 after strncpy: " << String3 << endl;
        cout << "String1 is" << strlen (String1) <<
        " bytes long, \n String2 is" <<
        trlen (String2)
        bytes long, \n and String3 is " <<
        trlen (String3) << " bytes long" << endl;
        Strcat (String3, String1); 
        cout << "String1 is still" << strlen (String1) <<
            bytes long, \n String2 is still" <<
            trlen (String2) <<
            bytes long, \n and String3 is now "
            ",
        String1 = "Now is the time for all good people
        cout << "String1: " << String1 << endl;
        *ring3, String1)
        cout << "String3 after strcat2: " << String3 << endl;
        <<<"String1 is now" << strlen (String1) <<
        strlen (String2)
        bytes long, \n and String3 is now " <<
            strlen (String3) << " bytes long" << endl;
        return 0
    }
```

```
OutPut
String1: No man is an island
String2: No man is an island
String3 after strncpy: No ma
String1 is 19 bytes long,
    String2 is 19 bytes long,
    and String3 is 5 bytes long
String3 after strcat: No maNo man is an island
String1 is still }19\mathrm{ bytes long,
    String2 is still 19 bytes long,
    and String3 is now 24 bytes long
String1: Now is the time for all good people...
String3 after strcat2: [ccc:ic]
No maNo man is an islandNow is the time for all good people...
String1 is now 38 bytes long,
    String2 is still }19\mathrm{ bytes long,
    and String3 is now 62 bytes long
```

ANALYSIS The header file string.h is included in line 3. This file contains the prototypes of the string functions. Line 6 allocates a pointer to a character string ("No man is an island"), and lines 7 and 8 allocate two character arrays.

Through line 26, Listing 12.11 is exactly the same as Listing 12.10 (see that listing if you need an explanation of those lines).

Line 27 changes String1. It does not modify the original contents ("No man is an island"); it causes String1 to point to a different string in memory ("Now is the time for all good people..."). This is a subtle distinction but very important because you should not use a string pointer (like String1) as the destination of strcpy(), strncpy(), strcat(), or similar functions.

You also cannot code something like the following:

String2 = "this will not compile";
This does not work because the variable String2 was allocated as an array and you are not allowed to change the address constant that is the beginning of the array.

Line 29 is the same as line 20 . The output of line 30 will be different from that of line 21 because String 1 is different and String 3 retains its contents from line 20.

The final version of String3 is not very readable because there are no separators between the strings concatenated into it. A good technique is to strcat() another character into the destination string (like a blank) as follows:
strcpy(result_string, original_string);
strcat (result_string, " ");
strcat(result_string, next_string);

A space will appear between the contents of original_string and next_string in the result_string variable. The strcat() function requires a string, which is why the single space in the first call is enclosed in double quotes-to make i a string (of length one).

You can use string pointers with any of the string functions.

## String Classes

Most C++ compilers come with a class library that includes a large set of classes for data manipulation. A standard component of a class library is a String class.
$\mathrm{C}++$ inherited the null-terminated string and the library of functions that includes strcpy() from C , but these functions are not integrated into an object-oriented framework. A String class provides an encapsulated set of data and functions for manipulating that data, as well as accessor functions so that the data itself is hidden from the clients of the String class.

If your compiler does not already provide a String class-and perhaps even if it does-you might want to write your own. GNU provides a String class. The remainder of this lesson discusses the design and partial implementation of String classes.

At a minimum, a String class should overcome the basic limitations of character arrays. Like all arrays, character arrays are static. You define how large they are. They always take up that much room in memory, even if you do not need it all. Writing past the end of the array is disastrous.

A good String class allocates only as much memory as it needs and always enough to hold whatever it is given. If it cannot allocate enough memory, it should fail gracefully.
initial test
Hello ; nice to be here!
Hello world; nice to be here
Hellx World; nice to be here!
Hellx World; nice to be here! Another string
Why does this work? string.

This String class overloads the offset operator ([ ] ), operator plus ( + ), and operator plus-equals ( $+=$ ). The offset perator is overloaded twice: once as a constant function returning a char and again as a nonconstant function returning a reference to a char.

The nonconstant version is used in statements such a
SomeString[4]= 'x';
as seen in line 161. This enables direct access to each of the characters in the string. A reference to the charac is returned so that the calling function can manipulate it
. the copy constructor (line 63). Note that rhs[i] is accessed, yet rhs is declared as a const String \&. It is not legal t access this object by using a nonconstant member function. Therefore, the offset operator must be overloaded with a constant accessor.
If the object being returned was large, you might want to declare the return value to be a constant reference However, because a char is only 1 byte, there would be no point in doing that

The default constructor is implemented in lines 33 - 39 . It creates a string whose length is 0 . It is the conventio of this String class to report its length not counting the terminating null. This default string contains only of this string clas
terminating null.

The copy constructor is implemented in lines 63-70. It sets the new string's length to that of the existing string-plus 1 for the terminating null. It copies each character from the existing string to the new string, and it null-terminates the new string.

號 53 -60 implement the constructor that takes an existing C-style string. This constructor is similar to the copy cor
strlen().

On line 28, another constructor, string(unsigned short), is declared to be a private member function. It is the inten of the designer of this class that no client class ever create a String of arbitrary length. This constructor exists only to help in the internal creation of Strings as required, for example, by operator $+=$, on line 131. This will b discussed in depth when operator+= is described later.
The String(unsigned short) constructor fills every member of its array with NULL. Therefore, the for loop checks f $\mathrm{i}<=$ len rather than $\mathrm{i}<$ len.

The destructor, implemented in lines 73-77, deletes the character string maintained by the class. Be sure to include the brackets in the call to the delete operator so that every member of the array is deleted, instead of on the first.

The assignment operator first checks whether the right side of the assignment is the same as the left side. If it is the current string is deleted, and the new string is created and copied into place. A reference is returned to facilitate assignments such as

String1 = String2 = String3
he offset operator is overloaded twice. Rudimentary bounds checking is performed both times. If the use attempts to access a character at a location beyond the end of the array, the last character-that is, len-1-is returned

Lines 117-128 implement operator plus ( + ) as a concatenation operator. It is convenient to be able to write
String3 $=$ String1 + String2
and have String 3 be the concatenation of the other two strings. To accomplish this, the operator plus ( + ) function computes the combined length of the two strings and creates a temporary string temp. This invokes the private constructor, which takes an integer, and creates a string filled with nulls. The nulls are then replaced by the contents of the two strings. The left side string (*this) is copied first, followed by the right side string (rhs)
The first for loop counts through the string on the left side and adds each character to the new string. The secon for loop counts through the right side. Note that $i$ continues to count the place for the new string, even as $j$ coun for loop counts th
into the rhs string.

Operator plus returns the temp string by value, which is assigned to the string on the left side of the assignment (string1). Operator $+=$ operates on the existing string-that is, the left side of the statement string $1+=$ string2. It (string) . works same as operator plus, except that the temp value is assigned to the current string (*this = temp) in line 142.

The main() function (lines 145-175) acts as a test driver program for this class. Line 147 creates a String object b using the constructor that takes a null-terminated C-style string. Line 148 prints its contents by using the accessor function GetString(). Line 150 creates another C-style string. Line 151 tests the assignment operator, a line 152 prints the results.

Line 154 creates a third C-style string, tempTwo. Line 155 invokes strcpy to fill the buffer with the characters nice to be here! Line 156 invokes operator $+=$ and concatenates tempTwo onto the existing string s1. Line 158 prin the results.
In line 160 , the fifth character in s1 is accessed and printed. It is assigned a new value in line 161. This invoke the nonconstant offset operator ([]). Line 162 prints the result, which shows that the actual value has, in fact, the nonconstar
been changed.

Line 164 attempts to access a character beyond the end of the array. The last character of the array is returned as designed.

Lines 166 and 167 create two more String objects, and line 168 calls the addition operator Line 169 prints the results.

Line 171 creates a new String object, s4. Line 172 invokes the assignment operator. Leference in line 21 , but here the program passes in a C-style string. Why is this legal?
resting

The answer is that the compiler expects a String, but it is given a character array. Therefore, it checks whether it can create a String from what it is given. In line 12, you declared a constructor that creates Strings from character arrays. The compiler creates a temporary String from the character array and passes it to the assignment operato This is known as implicit casting, or promotion. If you had not declared-and provided the implementation for-the constructor that takes a character array, this assignment would have generated a compiler error
Linked Lists and Other Structures
Arrays are much like Tupperware. They are great containers, but they are of a fixed size. If you pick a containe that is too large, you waste space in your storage area. If you pick one that is too small, its contents spill all ove and you have a big mess

$\underset{\substack{\text { furioss: } \\ \text { norsss: }}}{\text { Demonostrate } 1 \text { Inted } 1 \text { 1sot }}$


..............
$1 / 2$
0












$0=$

${ }^{\text {suritank }}$ besul




// Tail nodes is just a sent






$=$

$1,{ }^{\text {Private: }}$ node $\boldsymbol{m y n y e x t}$


$\underset{\substack{\text { mynext } \\ \text { roturn mins }}}{\substack{\text { mise }}}$
CII get all the credit and do none of the work


${ }^{*}$ myjeadi

myyead - neen neasuloder


${ }_{\text {data }}^{t}$ maino $)$






## What Have You Learned?

Dorothy said, "If I ever go looking for my heart's desire again, I won't look any further than my own backyard." Although it is true there is no place like home, it is also true that this is nothing like procedura programming. In procedural programming, a controlling method would examine data and invoke functions.

In this object-oriented approach, each individual object is given a narrow and well-defined set o responsibilities. The linked list is responsible for maintaining the head node. The head node immediately passe the new data on to whatever it points to, without regard to what that might be

Internal nodes are marginally more complicated; they ask their existing object to compare itself with the nev object. Depending on the result, they then insert or they just pass it along.

Note that the InternalNode has no idea how to do the comparison; that is properly left to the object itself. All th InternalNode knows is to ask the objects to compare themselves and to expect one of three possible answers. Given one answer, it inserts; otherwise, it just passes it along, not knowing or caring where it will end up.

So which object is in charge? In a well-designed object-oriented program, no one is in charge. Each object doe its own little job, and the net effect is a well-running machine

## Using Array Classes Instead of Built-in Arrays

Writing your own array class has many advantages over using the built-in arrays. For starters, you can preven array overruns. You might also consider making your array class dynamically sized: At creation it might have only one member, growing as needed during the course of the program.

You might want to sort or otherwise order the members of the array. You might consider a number of powerfu array variants. Among the most popular ar

- Ordered collection: Each member is in sorted order.
- Set: No member appears more than once
- Dictionary: This uses matched pairs in which one value acts as a key to retrieve the other value.
- Sparse array: Indices are permitted for a large set, but only those values actually added to the array consume memory. Thus, you can ask for SparseArray [5] or SparseArray[200], but it is possible that memor is allocated only for a small number of entries.
- Bag: An unordered collection that is added to and retrieved in random order

By overloading the index operator ([ ] ), you can turn a linked list into an ordered collection. By excluding duplicates, you can turn a collection into a set. If each object in the list has a pair of matched values, you can use a linked list to build a dictionary or a sparse array.

## Summary

Today you learned how to create arrays in $\mathrm{C}++$. An array is a fixed-size collection of objects that are all th same type

## Arrays do not do bounds checking. Therefore it is legal-even if disastrous-to read or write past the end of a

 array. Arrays count from 0 . A common mistake is to write to offset $n$ of an array of $n$ members.Arrays can be one-dimensional or multidimensional. In either case, the members of the array can be initializec as long as the array contains either built-in types, such as int, or objects of a class that has a default constructor Arrays and their contents can be on the free store or on the stack. If you delete an array on the free store remember to use the brackets in the call to delete

Array names are constant pointers to the first elements of the array. Pointers and arrays use pointer arithmetic find the next element of an array
You can create linked lists to manage collections with sizes you will not know at compile time. From linked lists, you can create any number of more complex data structures

Strings are arrays of characters, or chars. C++ provides special features for managing char arrays, including th capability to initialize them with quoted strings.
cher

Q\&A
Q What happens if $I$ write to element 25 in a 24 -member array?
A You will write to other memory, with potentially disastrous effects on your program.
Q What is in an uninitialized array element?
A Whatever happens to be in memory at a given time. The results of using this member without assigning a value are unpredictable
Yes With simple arrays.
Yes. With simple arrays you can use pointers to combine them into a new, larger array. With strings, ou can use some of the built-in functions, such as strcat, to combine strings.
號
An array must have a fixed size, whereas a linked list can be sized dynamically at runtime.
A Built-in arrays are quick and easy to use.
Q Must a string class use a char * to hold the contents of the string?
A No. It can use any memory storage the designer thinks is best.

## Workshop

The Workshop provides quiz questions to help you solidify your understanding of the material covered and exercises to provide you with experience in using what you have learned. Try to answer the quiz and exercise questions before checking the answers in Appendix D. "Answers to Exercises and Quizzes," and make sure understand the answers before continuing to Day 13 .

## Quiz

1. What are the first and last elements in SomeArray[25]?
2. How do you declare a multidimensional array?
3. Initialize the members of the array in Question 2 .
4. How many elements are in the array SomeArray[10][5][20]?

Can your maximum number of elements that you can add to a linked list?
7. What is the last character in the string "Brad is a nice guy"?

Exercises

1. Declare a two-dimensional array that represents a tic-tac-toe game board.
. Write the code that initializes all the elements in the array you created in Exercise 1 to the value 0 .
2. Write the declaration for a Node class that holds integers.
3. BUG BUSTER: What is wrong with this code fragment?

$$
\begin{aligned}
& \text { unsigned short SomeArray[5][4]; } \\
& \text { for (int } i=0 ; i<4 ; i++) \\
& \text { for (int } j=0 ; j<5 ; j++) \\
& \quad \text { SomeArray }[i][j]=i+j ;
\end{aligned}
$$

5. BUG BUSTER: What is wrong with this code fragment?
for (int $i=0 ; i<=5 ; i++$ )
for (int $j=0 ; j<=4 ; j+$

- What multiple inheritance is and how to use it
- What virtual inheritance is
- What abstract datat types are


## Problems with Single Inheritance

Suppose you were working with your animal classes for a while, and you divided the class hierarchy into Bird.
and Mammals. The Bird class includes the member function Flyo. The Mammal class has been divided in and Mammals. The Birrd class includes the member function Fly,. The Mammal class has been divided into a
number of types of Mammals, including Horse. The Horse class includes the member functions Whinnyo and $\begin{aligned} & \text { number } \\ & \text { Gallopo. }\end{aligned}$.

Suddenly, you realize you need a Pegasus object: a cross between a Horse and a Bird. A Pegasus can Fly0, it can Suddenly, you realize you need a Pegasus object: a cross between a Horse ar
WhinnyO, and it can Gallopo. With single inheritance, you are in quite a jam

## You can make Pegasus a B it will not be able to Flyo.

Your first solution is to copy the Flyo method into the Pegasus class and derive Pegasus from Horse. This works nine, a he cost remember to change the other. Of course, a developer who comes along months or years later to maintain you
code must also know to fix both places. ode must also know to fix both place
Soon, however, you have a new problem. You want to create a list of Horse objects and a list of Bird objects.
You would like to be able to add your Pegasus objects to either list, but if a Pegasus is a a list of Birds.

You have a couple of potential solutions. You can rename the Horse method Gallopp to Moveo, and then overrid Moveo in your Pegasus object to do the work of Flyo. You would then override Move0 in your other horses to 0 Pegasus::Move(long distance)

```
if (distance ) 
```

fly (distance);
else
gallop (distance)

This is a bit limiting. Perhaps one day Pegasus will want to fly a short distance or gallop a long distance. Your next solution might be to move Flyo up into Horse, as illustrated in Listing 1
cannot fly, so you have to make this method do nothing unless it is a Pegasu

Input. Listing 13.1 If Horses Could Fly...

```
// Listing 13.1. If horses could fl
```

\#include <iostream.h>
class Horse
public:
void
void Gallop() $\begin{aligned} & \text { cout << "Galloping....n"; }\} \\ & \text { virtual void Fly () (cout << "Horses can't }\end{aligned}$
private:
int itsAge;
class Pegasus : public Horse
class Pegasus : public Horse
public:

const int
int main()
Horse* Ranch [NumberHorses];
Horse phorse;
int choice, i ;
for ( $\mathrm{i}=0$; $\mathrm{i}<$ NumberHorses; $\mathrm{i}++$ )
cout << "(1) Horse (2) Pegasus: "
cin >> choice: cin >> choice;
if (choice $==2$ ) phorse
else pHorse $=$ new Horse;
Ranch $[i]=$ phorse; cout $\ll " \backslash n " ;$
for $(i=0$; $i<N$ Ranch[i]->Fly ();
delete Ranch[i];
$\qquad$

## Output

(1) Horse (2) Pegasus:
(1) Horse (2) Pegasus:
(1) Horse (2) Pegasus:
(1) Horse (2) Peqasus: (1) Horse (2) Pegasus:
(1) Horse
(2) Pegasus:
se (2) Pegasu
Horses can't fly.
I can fly! I can fly! I can fly!
can fly! I can fly! I can fly!

ANAIYsIS This program certainly works, although at the expense of the Horse class having a Flyo method. O line 10 , the method Flyy is provided to Horse. In a real-world classs you might have it issue an error, or fail
quietly. On line 18 , the Pegasus class overides the Fly method so "do the right thing," represented here by quielty. On line 18 , the Pe
printing a happy message.
The array of Horse pointers on line 24 is used to demonstrate that the correct Flyo method is called, based on th runtime binding of the Horse or Pegasus object

Note: These examples have been strirped down to their bare essentials to illustrate the e point suder
consideration. Constructors, virtual destructors, and so on have been removed to keep the code simple.

## Percolating Upward

Putting the required function higher in the class hierarchy is a common solution to this problem and results in many functions "percolating up" into the base class. The base class is then in grave danger of becoming a glo namespace for all the functions that might be used by any of the derived cl cl .
the class typing of C+t, and can create a large and cumbersome base class.

In general, you want to percolate shared functionality up the hierarchy, without migrating the interface of eac In general, you want to percolate shared functionality up the hierarchy, without migrating the interface of eac
class. This means that if two classes that share a common base class for example, Horse and Bird both share Class. This means that ind chases hat have a function in common (both birds and horses eat, for example), you would want to move tha
Animal functionality up into the base class and create a virtual function

What you want to avoid, however, is percolating an interface (such as Fly) up where it does not belong just so
that you can call thot , Casting Down

An alternative to this approach, still within single inheritance, is to keep the Fly) method within Pegasus and or call it it the pointer is actually pointing to a Pegasus object. To make this work, you will need to be able to ass
your pointer what type it is really pointing to. This is known as Runtime Type Identification (RTTI). Using your pointer what type it is really pointing to. This is k
RTTI has only recently become an official part of $\mathrm{C}+$.

Note: Beware of using RTTI in your programs. Its use may be an indication of poor design. Consider using virtual
functions, templates, or multiple inheritance instead.
To call Fly), however, you must cast the pointer, telling it that the object it is pointing to is a Pegasus object, Horse. This is called casting down because you are casting the Horse object down to a more derived type
C++ now officially, though perhaps reluctantly, supports casting down using the new dynamic_cast operator. It works this way

If you have a pointer to a base class such as Horse, and you assign to it a pointer to a derived class, such as regasus, you can use the Horse pointer polymorphically. If you then need to get at the Pegasus object, you create Pegasus pointer and use the dynamic_cast operator to make the conversion
At runtime, the base pointer will be examined. If the conversion is proper, your new Pegasus pointer will be fine If the conversion is improper, if you did not really have a Pegasus object after all, your new pointer will be null Listing 13.2 illustrates this point.

## INPUT. Listing 13.2 Casting Down

```
// Listing 13.2 Using dynamic_cast.
#include <iostream.h>
enum TYPE { HORSE, PEGASUS };
class Horse
public
    virtual void Gallop(){ cout << "Galloping...\n"; }
    private:
        int itsAge;
    class Pegasus : public Horse
    public:
        virtual void Fly() { cout <<
    };
    const int NumberHorses = 5,
    nt main()
        Horse* Ranch [NumberHorses];
        Horse* pHorse
        for (i=0; i<NumberHorses; i#43;+)
            cout << "(1)Horse (2) Regasus:"
            cin >> choice;
            pHorse = new Pegasus;
            else
            pHorse = new Horse;
            Manorse = new Ho
        cout << "\n";
        for (i=0; i<NumberHorses; i++)
            Pegasus *pPeg = dynamic_cast< Pegasus *> (Ranch[i])
            (pPeg)
            pPeg->Fly();
            cout << "Just a horse\n";
                delete Ranch[i];
    turn 0;
```

49:
$50:$

## Output

(1) Horse (2) Pegasus: 1
(1) Horse (2) Pegasus:
(1) Horse (2) Pegasus: 2
(1) Horse (2)Pegasus: (1) Horse (2) Pegasus: 1

```
Just a horse
Just a horse
```

I can fly! I can fly! I can fly!
I can fly! I can fly! I can fly!
I

## Frequently Asked Questions

FAQ: When compiling, I got an error from my GNU $\mathrm{g}++$ compiler (version 2.7 .2 or earlier)
1st13-02.cxx: In function ‘int main()':
1st13-02.cxx:42: cannot take typeid of object when -frtti is not specifie 1st13-02.cxx:42: invalid type argument
f maybe
FAQ: What should I do?
Answer: This is one of the GNU compiler's more confusing error messages. Unfortunately, the version you are using may not support RTTI even though it suggests ways around it.

You can try recompiling with the -frtti option, but if you get the warning that the option is not recognized, yo
patarom and specific yersion does not support this capability. You can try recompiing with the -frtio option, but if you get und
platform and specific version does not support this capability

Version 2.9 .5 supports RTTI by default.

ANalysis This solution also works. Fly0 is kept out of Horse, and it is not called on Horse objects. When it is called on Pegasus objects, however, they must be explicitly cast; Horse objects do not have the method Fly), so th pointer must be told it is pointing to a Pegasus object before being used.
The need for you to cast the Pegasus object is a warning that something may be wrong with your design. This program effectively undermines the virtual function polymorphism because it depends on casting the object to program effectively ur
its real runtime type.

## Adding to Two Lists

The other problem with these solutions is that because you declared Pegasus to be a type of Horse, you cannot ac a Pegass object to a list of Birds. You have paid the pice of either moving Flyo up into Horse or casting down th pointer, and yet you still do not have the full functionality you need.

One final, single inheritance solution presents itself. You can push Flyo, Whinny(), and Gallop0 all up into a common base class of both Bird and Horse: Animal. Now, instead of having a list of Birds and a list of Horses, can have one unified list of Animals. This works, but percolates more functionality up into the base classes.
Alternatively, you can leave the methods where they are but cast down Horses and Birds and Pegasus objects, but that is even worse!

## Multiple Inheritance

It is possible to derive a new class from more than one base class. This is called multiple inheritance. To deriv
from more than the base class, you separate each base class by commas in the class designation. Listing 13.3 from more than the base class, you sepanate each base class by commas in the class designation. Listing 13.3
illustrates how to declare Pegasus so that it derives from both Horeses and Birds. The program then adds Pegasus

```
// Listing 13.3. Multiple
```

include <iostream.h>
class Horse
class Hors
public:
Horse()

private:
int itsAge;
class Bird
$\underset{\substack{\text { public: } \\ \text { Bird () }}}{ }$

virtual void Chirp() const
virtual void Fly() const
cout
ate:
lass Pegasus : public Horse, public Bird
public:
void
左

Const int MagicNumber $=2 ;$
int main()
$t$ main()
Horse* Ranch [MagicNumber];
Bird** Aviary [MagicNumber];
Bird* Aviary [Ma
Horse $*$ phorse;
Bird $*$ pBird;
ind choice,
for ( $\mathrm{i}=0 ; \mathrm{i}$ i<MagicNumber; $\mathrm{i}+\mathrm{+}$ )
Cout << "nn(1) Horse (2) Pegasus: ";
cin >> choice:
cin > choice;
if (choice $==2$
(choice $=$ 2)
phorse $=$ new Pegasus,
else
pHorse $=$ new Hors
Ranch $[i]=$ pHorse;
for ( $\mathrm{i}=0$; i<MagicNumber; ${ }^{\text {i+ }}$ )

cin $\gg$ choice;
if (choice $==$
pBird = new Pegasus;
pBird $=$ new Bird;
Aviary[i] $=$ PBird;

cout << "\nRanch $["$,
Ranch $[i]->\operatorname{Whinny}() ;$
delete Ranch $[i]$;



```
return 0
```


# Bird constructor... Pegasus constructor <br> Horse destructor. Pegasus destructo 


 On lines 16-28, a Bird class is declared. In addition to its constructor and destructor, this class has two method Chipo and Flyo, both of which print idenityying messages. In a real program these might, for example, activati Finally, on lines $30-36$, the class Pegasus is declared. It derives both from Horse and fir
overrides the Chirpo method to call the Whinyo method, which it inherits from Hose. Two lists are created. a Ranch with pointers to Horse on line 41 , and an Aviary with pointers to Bird on line $42 . \mathrm{C}$
lines $46-55$, Horse and Pegasus objects are added to the Ranch. On lines 56 -65, Bird and Pegasus objects are addec Invocations of the virtual methods on both the Bird pointers and the Hose pointers do the right things for Pegas
objects. For example, on line 78 the members of the Aviary array are used to call Chirpo on the objects to whicl objects. For example, on line 78 the members of the Aviary array are used to call Chirpo on the objects to whic
they point. The Bird class declares this to be a virtual method, so the right function is called for each object.

## Declaring Multiple Inheritan

Declare an object to inherit from more than one class by listing the base classes following the colon after the Example 1

class

Example
class Schnoodle : public Schnauzer, public poodle

## The Parts of a Multiply Inherited Objec

// Listing 13.4 .
\#include <iostream.
typedef int HANDS;
typedef int HANDS;
enum Colior ( Red, Green, Blue, Yellow, white, Black, Brown )
class Horse
$\stackrel{1}{\text { public: }} \mathrm{Horse}$
Hic:
Horse (CoLor color, HANDS height) ;
virt
virtual $\sim$ Horse() f cout << "Horse destructor.
virtual void Whinny() const (cout << "Whinny

Virtual CoLor Getcolor() const ( return itsColor;


Horse: : Horse (COLOR color, HANDS height)
itscolor (color), itsHeight (height)
cout << "Horse constructor... $\backslash \mathrm{n} "$;
${ }^{\text {class }}$ Bird
public:
Bird ( (coLor color, bool migrates);
virtual $\sim$ Bird() (cout << NBird)

virtual coLor GetColor () const ifeturn itscolor; f
virtual bool GetMigration () const (return itsMigration;
private:
color itscolor;
bool itsMigration

Bird: : Bird (CoLor color, bool migrates):
itscolor (color),
cout << "Bird constructor...\n",
class Pegasus : public Horse, public Bir
ublic:
void Chirp () const \{ Whinny () ; )
Pegasus (COLOR, HANDS, bool
~Pegasus() (cout << "Pegasus destructor
~
return itsNumberBelievers;
private:
long itsNumberBelievers;
Pegasus:: Pegasus
Cotor acolor,
HANDS height,
HANDS height,
bool migrates,
long Numpeli iever
Iong Numbelieve) :
Horse (acolor, height),
Bird (acolor, migrates),
itsNumberBeli ievers (NumBelieve)
cout << "Pegasus constructor...\n"

```Pegasus \(\star_{\text {preg }}\)
peeg - FIy ()\(;\)
    MPe->Ply();
    <pPeg->GetHeight()
    cout << " hands tal1 and ",
    cout << "it does migrate.";
    cout << "it does not migrate.";
    *)
    #
```

\# Listing 13.5
Common base classes
\#include <iostream.h>
typedef int HANDS;
enum CoLOR $\{$ Red, Green, Blue, Yellow, White, Black, Brown \}
class Animal // common base to both horse and bira
public:

private:
int itsAge;
Animal::Animal(int age)
sAge (age
cout << "Animal constructor...\n";
Class Horse : public Anima
public:
Horse
l
Horse (CooLor color, HANDS height, int age);
virtual $\sim$ Horse()
( cout << "Horse destruct


HANDS itsHeight;
COLOR itscolor;
Horse: : Horse (COLOR color, HANDS height, int age)
Animal (age).
( Animal (age),
itscolor (coil
(color (color), itsHeight (height)
cout << "Horse constructor... $\backslash \mathrm{n} "$;
Class Bird : public Animal
public:
Bird (cooor color, bool migrates, int age);
virtual $\sim$ Bird) ( (cout << "Bird destructor.
virtual $\sim$ Bird ( $)$ (cout << "Bird destructor.
virtual void Chirp() const , cout <<"Chirp. virtual void Chirp () const
virtual void Fly () const
t cout << "I can fly! I can fly! I can fly!";
virtual CoLor Getcolor () const \{ return itscolor; $\}$
vitutual bool GetMigration() const \{return itsMigration;

$$
\begin{aligned}
& \text { protected: } \\
& \text { CoLOR } \mathrm{it}
\end{aligned}
$$

$$
\begin{aligned}
& \text { coLor itscolor; } \\
& \text { bool itsMigration }
\end{aligned}
$$

Bird::Bird(CoLor color, bool migrates, int age):
Animal (age),
itscolor (coior), itsMigration (migrates)
cout << "Bird constructor...\n"
class Pegasus : public Horse, public Bird ${ }_{\text {public }}$
void Chirp () const $\{$ Whinny (); ;
Pegasus (CoLor, HANDS, bool,
Pegasus (COLOR, HANDS, bool, long, int);
virtual $\sim$ Pegasus() (cout << "Pegasus des
virtual
virtual long GetNumberBelievers () const
( return itsNumberBelievers;
virtual CoLor Getcolor () const $\{$ return Horse:: itscolor;
virtual int GetAge() const $\{$ return Horse::GetAge (); )
rivate:
long itsNumberBelievers;
Pegasus:: Pegasus
COUOR acolor,
COLOR acolor,
HANDS height,
bool migrates
bool migrates,
long Numbeli ieve
Horse (acolor, height, age),
Bird (aColor, migrates, age),
itsNumberBeli ievers (NumBelieve
cout << "Pegasus constructor...\n"
$\begin{aligned} & \text { Pegasus } \star_{\text {preg }}=\text { new Pegasus (Red, 5, true, 10, 2) } \\ & \text { int age }=\text { preg->cetAge (); }\end{aligned}$
$\begin{aligned} & \text { int age }=\text { preg->GetAge (); } \\ & \text { cout } \ll \text { "This pegasus is }\end{aligned}$
$\begin{aligned} & \text { delete ppeg; } \\ & \text { return 0; }\end{aligned}$
/ Listing 13.6
Virtual inheritance
\#include <iostream.h>
typedef int hands;
enum CoLor (Red,
G
class Animal // common base to both horse and bird.
$\underset{\substack{\text { public: } \\ \text { Animal (int) }) ; ~}}{ }$


$$
\begin{gathered}
\text { private: } \\
\text { int itsAge; }
\end{gathered}
$$

Animal: : Animal (int age):
it shage (age)
cout << "Animal constructor...\n";
class Horse : virtual public Animal
$\substack{\text { class } \\ \text { public: } \\ \text { Horse }}$


 HanD ${ }^{\text {itsHR ight; }}$,
COLOR itscololor;

Horse:: Horse (CoLor color, hands height, int age)
 scolor (color), itsHei ght (height)
cout << "Horse constructor... $\mathrm{n}^{\prime \prime}$
class Bird : virtual public Animal
$\underset{\substack{\text { public: } \\ \text { Bird } \\ \text { ind }}}{\text { end }}$


 protected
Cooor itscolor;
bool itsMigration;

Bird:: Bird (Color color, bool migrates, int age): Animal (age),
itscolor (color), , itsMigration(migrates)

$$
\text { out << "Bird constructor. } \ldots \backslash n^{n}
$$

class Pe
public:

 virtual long Get Numberbeli ievers(, cons
return $i$ isNumberBel ievers;

asus: : Pegasus
COLOR acolor,
Aor acolor,
HANDS hei ight,
bool migrates,
long umpeli ieve
long Numbe 1,
int age):
Horse (acolor, hei , hht, age),
Bird (acolor, migrates, age),

out << "Pegasus const
main()
Example 1
class Horse: virtual public Animal
class Bird: virtual public Animal
class Pegasus: public Horse, public Bira
$\left\lvert\, \begin{aligned} & \text { Class Schnauzer: virtual public Dog } \\ & \text { Class Poodle: virtual public Dog }\end{aligned}\right.$
class Schnoodie : public Schnauzer, public Poodle

## xins and Capabilities Classes

 A mixin, or capability class, is a class shat adds functionality without adding much or any dataCapability classes are mixed into a derived class the same as any other class might be, by declaring the derive
casss o inherit publicy from them. The only difference between a cappobility class and any ohter chass sis stat
 noting hat at times a ty you
class

 from both.
 programmers who used to take a summer break there, especially while working with the objectectorieneed Abstract Data Types



Invot Listing 13.7 Shape Classe.
/Listing 13.7. Shape classes.
\#include 〈iostream.h>
${ }^{\text {class shape }}$
$\underset{\substack{\text { publice } \\ \text { shape (1) } \\ \text { virtual }}}{\substack{\text { and }}}$

class circle: public Shape


int
int itssadius;
itsciricumference,
oid Circle: : Draw()
cout <<"Circle drawing routine here! \nn";


Ie: :Draw ()
for (int $j=0 ; j<i$ itwidth; $j++$
for (int $=0$,
cout $\ll \times x ;$
class Square : public Rectangle
c:
square (int len)
square (int
len,


Square: : Square (int 1en)
Rectangle (len, len)
$\underset{\substack{\text { Square: : Square (int } \\ \text { Rectangle ( } 1 \text { len, }, \text { width })}}{\text { int width }):}$
f (GetLength() $:=$ Getwidth())
cout << "Error, not a square
main(

while (! fquit)

switch (choice)
case 0 : ffuit
break true

sp->Draw ();
delete spi
cout $\ll \cdots 1 n^{\prime}$,
return 0 ;

Outru
(1) Circle (2) Rectang1e (3) Square (0) Quit:

$(1)$ Circle
$\times \times x \times x \times x$
$\times \times \times \times x$
$x \times x$




## Pure Virtual Functions


virtual void Draw() $=0$;

$\qquad$

# 1ass square: public Rectangle 



## 




while (1)

sitch (choice)

1. sp - new Circle (5);
$\mathrm{sp}=$ new Rectangle $(4,6)$;
$\mathrm{sp}=$ new square
Equit $=$ true;
${ }_{\substack{\text { itenit) } \\ \text { break; }}}$

$\qquad$

## 

```
\/. Listing 13.10
Minerving ADrs from other mDTs
```

enum CoLor (Red, Green, Blue, Yellow, white, Black, Brow
class Animal // cormon base to both Mannal and Fish
public:



virtual
vate
vate
$\substack{\text { aninal: }: \text { an } \\ \text { it } t \text { ange } \\ \text { ase }}$
cout << "Animal constructor... $\mathrm{nn}^{n}$;
Class Mamaal : pulic annimal



$\underset{\substack{\text { class } \\ \text { pub }}}{ }$




${ }^{\text {Class Horse }}$ : pulic Mamai

| public: |
| :---: |
| Hors |

    Horsse (int age, coror color
    Mamal ( Iage),
itsocolor (color)




protected.
Cotor itscolor ;
class Dog : public Mamaal




$\underset{\substack{\text { protected } \\ \text { Cotor } \\ \text { itscolors; }}}{ }$
int main()

int choice;
bool tequit
while (1)
cout 《< "(1) Dog (2) Horse (3) Fish (0) Quit: ";
cin 》" choice;
switch (choice)
Case 1:
break ;
panimal $=$ new
$\operatorname{Dog}(5, B r o w n) ;$
ase 2: pAnimal $=$ new Horse (4, Black) ;
3: panimal $=$ new Fish (5):
breaki
defanult: fquit $=t r u e ; ~$
breaki
${ }_{\text {if }}^{\text {if (feuit) }}$ break;


panimall->Move (1);
panimal-s>liep
delete


```
        \({ }_{\text {return }}{ }^{\text {cout }}\)
```


## Which Types Are Abstract:

## The Observer Pattern

A very hot trend in $\mathrm{C}++$ is the creation and dissemination of design patterns. These are well-documented solutions to common problems encountered by $\mathrm{C}++$ programmers. As an example, the observer pattern solves a common problem in inheritance.

Imagine that you develop a timer class that knows how to count elapsed seconds. Such a class might have a class member itsSeconds, which is an integer, and it would have methods to set, get, and increment itsSeconds.

Now suppose we further assume that your program wants to be informed every time the timer's itsSeconds member is incremented. One obvious solution would be to put a notification method into the timer. However, notification is not an intrinsic part of timing, and the complex code for registering those classes that need to be informed whenever the clock increments does not really belong in your timer class.

More importantly, after you work out the logic of registering those who are interested in these changes and then notifying them, you would like to abstract this out into a class of its own and be able to reuse it with other classes that might be "observed" in this way.

Therefore, a better solution is to create an observer class. Make this observer an abstract data type with a pure virtual function Update().

Now create a second abstract data type, called Subject. Subject keeps an array of Observer objects and also provides two methods: register() (which adds observers to its list) and Notify(), which is called when there is something to report.

Those classes that want to be notified of your timer's changes inherit from Observer. The timer itself inherits from Subject. The Observer class registers itself with the Subject class. The Subject class calls Notify when it changes (in this case when the timer updates).

Finally, we note that not every client of timer wants to be observable, and thus we create a new class called ObservedTimer, which inherits both from timer and from Subject. This gives the ObservedTimer the timer characteristics and the capability to be observed.

## A Word About Multiple Inheritance, Abstract Data Types, and Java

Many C++ programmers are aware that Java was based in large part on C++, and yet the creators of Java chose to leave out multiple inheritance. It was their opinion that multiple inheritance introduced complexity that worked against the ease of use of Java. They felt they could meet $90 \%$ of the multiple inheritance functionality using what are called interfaces.

An interface is much like an abstract data type in that it defines a set of functions that can only be implemented in a derived class. However, with interfaces, you do not directly derive from the interface; you derive from another class and implement the interface, much like multiple inheritance. Thus, this marriage of an abstract data type and multiple inheritance gives you something akin to a capability class without the complexity or overhead of multiple inheritance. In addition, because interfaces can have neither implementations nor data members, the need for virtual inheritance is eliminated.

Whether this is a bug or a feature is in the eyes of the beholder. In either case, if you understand multiple inheritance and abstract data types in C++, you will be in a good position to move on to using some of the more advanced features of Java should you decide to learn that language as well.

The observer pattern and how it is implemented both in Java and C++ is covered in detail in Robert Martin's article "C++ and Java: A Critical Comparison," in the January 1997 issue of $C++$ Report.

## Summary

Today you learned how to overcome some of the limitations in single inheritance. You learned about the danger of percolating interfaces up the inheritance hierarchy and the risks in casting down the inheritance hierarchy. You also learned how to use multiple inheritance, what problems multiple inheritance can create, and how to solve them using virtual inheritance.

You also learned what abstract data types are and how to create abstract classes using pure virtual functions. You learned how to implement pure virtual functions and when and why you might do so. Finally, you saw how to implement the observer pattern using multiple inheritance and abstract data types.

## Q\&A

## Q What does percolating functionality upward mean?

A This refers to the idea of moving shared functionality upward into a common base class. If more than one class shares a function, it is desirable to find a common base class in which that function can be stored.

## Q Is percolating upward always a good thing?

A Yes, if you are percolating shared functionality upward. No, if all you are moving is interface. That is, if all the derived classes can't use the method, it is a mistake to move it up into a common base class. If you do, you'll have to switch on the runtime type of the object before deciding if you can invoke the function.

## Q Why is switching on the runtime type of an object bad?

A With large programs, the switch statements become big and hard to maintain. The point of virtual functions is to let the virtual table, rather than the programmer, determine the runtime type of the object. Q Why is casting bad?
A Casting isn't bad if it is done in a way that is type-safe. If a function is called that knows that the object must be of a particular type, casting to that type is fine. Casting can be used to undermine the strong type checking in $\mathrm{C}++$, and that is what you want to avoid. If you are switching on the runtime type of the object and then casting a pointer, that may be a warning sign that something is wrong with your design.
Q Why not make all functions virtual?
A Virtual functions are supported by a virtual function table, which incurs runtime overhead, both in the size of the program and in the performance of the program. The v-ptr, or virtual-function pointer, is an implementation detail of virtual functions. Each object in a class with virtual functions has a v-ptr, which points to the virtual function table for that class. If you have very small classes that you don't expect to subclass, you may not want to make any of the functions virtual.

## Q When should the destructor be made virtual?

A Any time you think the class will be subclassed, and a pointer to the base class will be used to access an object of the subclass. As a general rule of thumb, if you've made any functions in your class virtual, be sure to make the destructor virtual as well.
Q Why bother making an Abstract Data Type-why not just make it non-abstract and avoid creating any objects of that type?
A The purpose of many of the conventions in $\mathrm{C}++$ is to enlist the compiler in finding bugs, so as to avoid runtime bugs in code that you give your customers. Making a class abstract-that is, giving it pure virtual functions-causes the compiler to flag any objects created of that abstract type as errors.

## Workshop

The Workshop provides quiz questions to help you solidify your understanding of the material covered and exercises to provide you with experience in using what you have learned. Try to answer the quiz and exercise questions before checking the answers in Appendix D, "Answers to Exercises and Quizzes," and make sure you understand the answers before continuing to tomorrow's lesson.

## Quiz

1. What is a down cast?
2. What is the v-ptr?
3. If a round-rectangle has straight edges and rounded corners, and your RoundRect class inherits both from Rectangle and from Circle, and they in turn both inherit from Shape, how many Shapes are created when you create a RoundRect?
4. If Horse and Bird inherit from Animal using public virtual inheritance, do their constructors initialize the Animal constructor? If Pegasus inherits from both Horse and Bird, how does it initialize Animal's constructor?
5. Declare a class vehicle and make it an abstract data type.
6. If a base class is an ADT, and it has three pure virtual functions, how many of these must be overridden in its derived classes?

## Exercises

1. Show the declaration for a class JetPlane, which inherits from Rocket and Airplane
2. Show the declaration for 777, which inherits from the JetPlane class described in Exercise 1.
3. Write a program that derives Car and Bus from the class Vehicle. Make Vehicle be an ADT with two pure virtual functions. Make Car and Bus not be ADTs.
4. Modify the program in Exercise 3 so that Car is an ADT, and derive SportsCar, Wagon, and Coupe from Car. In the Car class, provide an implementation for one of the pure virtual functions in Vehicle and make it non-pure.

Chapter 14

## What static member data (variables and static member functions ate How to use static member variales and static member functions

How to use static member variables and static member functions
How or cata end manipulate pointers fonuctions and pointers to member function
How to work with arays of pointers to functions

## atic Member Data

Until now, you have probably thought of the data in each object as unique to that object and not shared among
obiects in a class. If you have five Cat objects, for example, each has iss own age, weight, and other data. The agg位
A times. however, you will want to keep track of a pool of data. For example, you might want to know how man
 . You can think of a static member as belonging to the class rather than to the object. Normal member data is one
obiect, but static members are one per class Listing 14.1 declares a ca objeet with a sataic data member,


Input Listing 14.1 Static Member Data
$/$ LIisting 14.1 static data members
*include <iostream.h>

$\underset{\text { private: }}{\text { int } i t \text { itsAge }}$
int Cat::HowManyCats = 0 ;
int main() const int MaxCats = 5 ;
Cat * Cathouse [Maxcats],
for


竍



```
return 0;
```

Deleting the one which is 0 years old
There are 4 cats left There are 4 cats left!
Deleting the one which
The There are 3 ane ats which is 1 is 1 years old
Deleting the one which is 2 years old There are 2 cats left!
Deleting the one which $i$ is 3 years old
There are 1 cats left!
Deleting the one which is 4 years old

The declaration of HowManyCats does not define an integer, no storage space is set aside. Unike the non-static


It is a common mistake to forget to define the static member variables of classes. Do not let this happen to you!
course, if it does, the GNU compiler will catch i with w number of pithy error messagas such a s the ofllowing:



(fach undeclared identifier is re 1st14-01.cxx:9: 'HownanyCat
1st14-01.cx:
At
undeclared (first use this function) 1st14-01.cxx: 19: 'int Cat: HowManyCats', is not a static member of 'class ce 1st14-01.cxx: In function 'int main()':
1st14-01.cxx:31: 'HowManyCats' is not a member of type 'Cat

## def ined symbol Cat: : HowManyCat

You do not need to do this for istseg because it is a non-static member variable and is defined each time you mak
a cat object, which you do here on line 26 . The constructor for Cat increments the static member variable on line 8 . The destructor decrements it on line 9 .
Thus, atany moment HowMany Cas has an a cocurate measure of how many Cat objects were created but not yet
destroyed. destroyed.

```
\begin{array} { l } { \text { The driver program on lines 21-40 instantiates five Cass and puts shem in aa} } \\ { \text { and tus HowManyCass is incemented five times from is intial value of 0.} } \end{array}
```

 Notet that HowManyCats is pulicic and is accessed directly by maino. No reason exists to expose this member varia
in this way. It is peferabbele to make it private along with the ohher member variables and provide a public aceess
 access this datad directly, without necessarily having a Cat object available, you have t .
shown in Listing 14.2, or rovidida a sataic member function, as discussed later today.
INPUT Listing 14.2 Accessing Staitic Members Without an object
/Listing 14.2 static data members
class cat
public:
Cat (int age): itsAge (age) ) HowManyCats + ;
virtual ~Cat () $\mid$ HowManyCats-;

$\underset{\substack{\text { private: } \\ \text { int } \\ \text { itsAge }}}{ }$
int Cat: : HowManyCats $=0$,
void TelepathicFunction ()
int main()
const int MaxCats $=5$; int
Cat *Cathouse [MaxCats];
for ( $\mathrm{i}=0 ; \mathrm{i}$ iMaxCats; $;$ i+t)
Cathouse[i] = new cat (i);
TelepathicFunction();

eturn 0;
void TelepathicFunction()
cout << "There are ";
cout << Cat: :HowManyCats << " cats alive! $\backslash \mathrm{nn} "$

IListing 14.3 private static data menbers
include $<$ iostream.h h
class cat
pubicic:




for ( $\mathrm{i}=0 ;$ i $\mathrm{imaxatats} ; ~ i+t)$


return ${ }^{\prime}$;



## rs to Function





$\qquad$
／／Listing 19.6 without function pointers

## 

nt main（）


[^1]
break；

break；

case 4：


$\qquad$

$\qquad$

1 1.5 $=2$ Shorthand Invoction
 Prunc ( $x$ )


NWVU

void suare (ints, in

$t$ main ()






, return 0 ;
oid Printevas (int $x$, int 8 )

oid Sguare (int axx, int $8 x y$



路滑

oid cetvals (int \& rvalone, int \& rvaritwo


Include cisostrean h.h



athen
void (tprunc) (ints, int





## return $0 ;$



id Sguare (int $8 \mathrm{xx}, \mathrm{int} \varepsilon \mathrm{rx})$

${ }^{\text {oida cube }}$ men






## Using typedef with Pointers to function:


Invort
fo Makce Pineres so Finctions Mor Reatable

I Iisting 14.9. Using typedef to make pointors to



pe pranc;






oid Square (int arx, int ary)

${ }_{i}^{\text {voia cube tmp }}$



8 (int $\&$ rvalone, int $\&$ rvaltwol


mas wam


tass

: public Mamal

class horse : public Mamanal


innt animali
ant nemad
hool fouit
bout
while efouit $=$ = false



| if tequit) |
| :---: |
| break; |




```
Woof!
```

(0) Quit (1) dog (2) ca
(1) Speak (2) Move: 1
Meow!
(0) Quit (1) dog (2) cat
(1) Speak (2) Move: 2
(1) Speak
Galloping
(0) Quit (1)

ANALYSIS On lines 4-14, the abstract data type Mammal is declared with two pure virtual methods, Speakk ar Move0. Mammal is subclassed into Dog, Cat, and Horse, each of which overrides Speak) and Move0. The driver program in maino asks the user to choose which type of animal to create, and then a new subclass 0
Animal is created on the free store and assigned to prt on lines $54-56$. The user is then prompted for which method to invoke, and that method is assigned to the pointer pfunc on lin
66 or 67. st version 2.2 compiles this without complaint; 2.9 .5 gives you the following warnings (because The user is then prompted for which method to invoke, and that method is assigned to the pointer princ on lin
66 or 6.9 gtve ersion 2.7 .2 compies this without complaint; 2.95 gives you the following warnings (because
you are passing addresses around):
.$/ 1$ st14-10.cxx: In function 'int main()' (1st14-10.cxx:67: warning: assuming \& on 'Manna1: :Speak () const on Mammal: :Move () const

On line 70, the method chosen is invoked by the object created, by using the pointer prt to access the object an pFunc to access the functior
Finally, on line 71 , delete is called on the pointer pr to return the memory set aside for the object to the free
store. Note that no reason exists to call delete on pFinc because this is a pointer to code, not to an object on thh free store. In fact, atempming to do so will generate a compile-time error
fore

## Arrays of Pointers to Member Functions

As with pointers to functions, pointers to member functions can be stored in an array. The array can be
initialized with the addresses of various member functions, and these can be invoked by offsets into the array.
init initialized with the addresses of variou
Listing 14.1 illustrates this technique.

INPUT Listing 14.11 Array of Pointers to Member Function

```
include <iostream.h>
class
    \,
    *)
    typedef void (Dog::*PDF) () const
    const int MaxFuncs = 7,
        logFunctions
        Dog::Move,
        l
        Mog* pDog=NULL;
        Method; =false;
        hile (!fquit)
        cout "(0) Quit (1) Speak (2)Move (3) Eat (4)Grow1";
        in >> Method;
        fquit = true;
        else
        MDog = new Dog;
        (pDog->*DogFunctions[Method-1])();
```

(0) Quit (1) Speak (2) Move (3) Eat (4) Grow1 (5) Whimper (6) Roll Over
(7) flay Dead:
Orowl (5) Whimper (6) Roll over
0) Quit (1) Speak
7) Play Dead: 4
(7) Play (1) Speak (2) Move (3) Eat (4) Growl (5) Whimper (6) Roll Over
this the end of Little Caesar?
(1) Quit (1) speak (2) Move (3) Eat (4) Grow1 (5) Whimper (6) Roll Over
(2) Play Dead: 0>
ANALYSIS On lines 5-15, the class Dog is created, with seven member functions all sharing the same return
type and signature. On line 17, a typedef declares PDF to be a pointer to a member function of Dog that takes no

On lines $21-28$, the array Doffunctions is declared to hold seven such member functions, and it is initialized wi
the addresses of these functions. As with Listing 14.10 , $\mathrm{g}+\mathrm{y}$ version 2.7 .2 did not complain about this code;
2.9.5 produced the following:

On lines 36 and 37 , the user is prompted to pick a method. Unless Quit is ip cked, a new Dog is created on the
heap, and then the correct method is invoked on the array on line 46
hotshot $C++$ programmers in your company; ask them what it does:
g->*DogFunctions (Method-1) () ;
What do you think would happen if an out-of-range value was entered (like -1 or 8)? Because C+t has no arra)
bound checking, you would not get a compile or runtime error. You may get completely unexpected results as
your program attempts to use the address stored in that memory as a function. Because you u idinot set the
in that memory (as was done in lindes
nume memory (as was done in lines $21-28$ ), you have no idea what is there. That location may b
numeric value that will be treated as if it was the address of a function with unpredictable results
Once again, this is a bit esoteric, but when you need a table built from member functions, it can make your
program much easier to read and understand.

## Don't

 DON'T use pointer-to-membsimpler solutions are possible.

## Summary

Today you learned how to create static member variables in your class. Each class, rather than each object, has
one instance of the static member variable. It it possible to access this member variable without an object of the class tye by fully yualifying the name a assumning that you declared the static member to have public access

Because static member functions do not have a this pointer, they also cannot be made const. const in a member function indicates that the this pointer is const.

You also learned how to declare and use pointers to functions and pointers to member functions. You saw how to create arrays of these pointers and how to pass them to functions.

Pointers to functions and pointers to member functions can be used to create tables of functions that can be selected from at runtime. This can give your program flexibility that is not easily achieved without these pointers.

## Q\&A

## Q Why use static data when you can use global data?

A Static data is scoped to the class. In this manner, static data is available only through an object of the class, through an explicit call using the class name if it is public, or by using a static member function. However, static data is typed to the class type, and the restricted access and strong typing makes static data safer than global data.
Q Why use static member functions when you can use global functions?
A Static member functions are scoped to the class and can be called only by using an object of the class or an explicit full specification (such as ClassName::FunctionName()).

## Q Is it common to use many pointers to functions and pointers to member functions?

A No, these have their special uses, but are not common constructs. Many complex and powerful programs have neither.

## Workshop

The Workshop contains quiz questions to help solidify your understanding of the material covered and exercises to provide you with experience in using what you learned today. Try to answer the quiz and exercise questions before checking the answers in Appendix D, "Answers to Exercises and Quizzes," and make sure you understand the answers before continuing tomorrow.

## Quiz

1. Can static member variables be private?
2. Show the declaration for a static member variable.
3. Show the declaration for a static function.
4. Show the declaration for a pointer to function returning long and taking an integer parameter.
5. Modify the pointer in Question 4 so that it's a pointer to member function of class Car.
6. Show the declaration for an array of 10 pointers as defined in Question 5.

## Exercises

1. Write a short program declaring a class with one member variable and one static member variable. Have the constructor initialize the member variable and increment the static member variable. Have the destructor decrement the member variable.
2. Using the program from Exercise 1, write a short driver program that makes three objects and then displays their member variables and the static member variable. Then destroy each object and show the effect on the static member variable.
3. Modify the program from Exercise 2 to use a static member function to access the static member variable. Make the static member variable private.
4. Write a pointer to member function to access the non-static member data in the program in Exercise 3 , and use that pointer to print the value of that data.
5. Add two more member variables to the class from the previous questions. Add accessor functions that get the value of these values and give all the member functions the same return values and signatures. Use the pointer to member function to access these functions.

## Week 2

## In Review

The Week in Review program for Week 2 brings together many of the skills you acquired over the past fortnight and produces a powerful program.

This demonstration of linked lists utilizes virtual functions, pure virtual functions, function overriding, polymorphism, public inheritance, function overloading, forever loops, pointers, references, and more. Note that this is a different linked list from the one shown on Day 12, "Arrays, C Strings, and Linked Lists"; in C++ there are many ways to accomplish the same thing.

# Part 3 

At a Glance
 Where You Are Going




## Containmen

 forth).

## /t default const String : : String ()


$/ /$ cout << "tDe fault
$/ /$ constructorcount + ;
$/$ private (helper) constructor, used only by
$\$ /$ chass methods for creating a new string of $1 /$ required size. Null
String:
string
(int
Ien)
 tsLen=1en;
$/ /$ cout $\ll$ String $($ int $)$
Constructor $\backslash \mathrm{n}^{\prime \prime}$
Count $+;$
// Converts a character array to a string
String: String coonst char * const ostring)
itsten $=$ strlen (cstring),
itsstring $=$ new char $[$ it sil


// copy const ructor
String: : String (const
ItsLen=rhs. Getten ();
itsstring $=$ new char $[$ itsLen+1 $]$


string: : sctrining () $)$ alloctored memory
delete $[1]$ itsstring,
itsten $=0 ;$


| 1f $\begin{array}{c}\text { (this } \\ \text { return this } \\ \text { urns }\end{array}$ |
| :--- |





Inon constant offset operator, returns
// feference to character so it can be
char \& string: :operator [] (int offset)
${ }_{\text {if }}^{\text {if (offset }>\text { itsten) }}$ return itsstring $($ itsten- 1 )
${ }_{\text {return }}$ resstring $[$ i t sLen- 1 ]


${ }_{\text {return }}{ }^{\text {itsstring }}$ (offset);




// changes current string, returns nothing
void String: :operatatort=(const Stringe
rhs)



temp totatalien
$x$ this $=$ temp;


```
tring.hp\mp@subsup{p}{}{\prime}
yee
();
e();
operator= (const Employee &);
& GetFirstName() const
Mrn itsFirstName; }
ring & GetLastName() const { return itsLastName; 
l
FirstName (const String & fName)
sFirstName = fName; } Name
LastName (const String & 1Name)
LastName = 1Name; }
Address(const String & address)
esAddress = address;}
itsFirstName;
itsLastName;
itsLastName;
itsAdaress;
mployee
ame ("")
mployee (char * firstName, char * lastName
ddress, long salary)
Name (firstName)
Mme (lastName),
ss(address)
nployee (const Employee & rhs):
Mployee (const Employee &
l
%(rhs.GetSalary())
Employee () {}
Employee::operator= (const Employee & rhs
== &rhs,
N *this;
Name = rhs.GetFirstName(),
ame = rhs.GetLastName()
* = rhs.GetAddress();
Chis;
Edie("Jane","Doe","1461 Shore Parkway", 20000);
Salary (50000)
lol
LastName (LastName);
"Name:";
<< Edie.GetLastName().GetString();
"\\< &ddress:";
MnAdaress:";
&die.GetAddress().GetString();
".\nSalary:";
```k , I included the implementation with the declaration of the class. In a real program, you would save the class
the implementation in String.cpp. You would then add String.cpp into your program (using a makefile) and have
Parkway.

\section*{the Contained Class}
ve special access to the member variables of String. If the Employee object Edie tried to access the member variat member variable, it would get a compile-time error. This is not much of a burden, however. The accessor member variabe, it would get a compie--ime error. This is not much of a burden, however. The accessor
ce for the sting class, and the employe class need not worry about the implementation details, just as it does no
r variable, itsSalary, stores its information.

\section*{tained Members}
 g that all the string accessors, such as Geffirstiame), return a constant reference. Because operatort is not

\section*{Cost of Containment}

It is important to note that the user of an Employee class pays the price of each of those string objects each time one is constructed or a copy of the Employee is made. Uncommenting the cout statements in Listing 15.1 , lines \(38,51,63,75,84\), and 99 , reveals how often these are
called. Listing 15.3 rewrites the driver program to add print statements indicating where in the program objects called. Listing 15.
are being created.

Note: To compile this listing, uncomment lines 38, 51, 63, 75, 84, and 99 in Listing 15.1.
INPUT. Listing 15.3 Contained Class Constructors
```

\#include "String.hpp"
class Employee
public:
Employee();
mployee (char
CEmployee ();
Employee (const Employee\&);
const String \& GetFirstName() const
t String \& GetFirstName ()
const String \& GetLastName() const { return itsLastName;
const String \& GetAddress() const { return itsAddress;
long GetSalary() const { return itsSalary; }
void SetFirstName(const String \& fName)
{ itsFirstName = fName; }
SetLastName(const String \& 1Name)
{ itsLastName = 1Name; }
void SetAddress(const String \& address)
{ itsAddress = address;
voi
ivate: itsFirstName;
String itsLastName;
long itsSalary;
Employee::Employee()
itsFirstName (""),
itsAddress(""),
itsSalary(0)
Employee::Employee(char * firstName, char * lastName,
char * address, long salary)
tName(firstName)
itsLastName (lastName)
itsSalary(salary)
Employee::Employee (const Employee \& rhs)
itsFirstName(rhs.GetFirstName())
itsLastName (rhs.GetLastName ())
itsSalary(rhs.GetSalary())
(})
Employee:: ~Employee () {}
Employee \& Employee::operator= (const Employee \& rhs)
if (this == \&rhs)
return *this;
itsFirstName = rhs.GetFirstName()
itsLastName = rhs.GetLastName();
itsLastName = rhs.GetLastName();
itsSalary = rhs.GetSalary();
return *this;
int main()
cout << "Creating Edie...\n";
Employee Edie("Jane","Doe","1461 Shore Parkway", 20000)
Edie.SetSalary (20000);
Edie.SetFinstName("Edythe")
cout << "Creating temporary" string LastName...N",
String "Creacing temporary string LastName
String LastName("Levine");
cout << "Name:";
cout<< Edie.GetFirstName().GetString();
cout <<"" "<< Edie.GetLastName().GetString()
cout << "\nAddress:";
"\nsalary:"
Edie.GetSalary();
cout << Edie.;
return 0;

```
Output
Creating Edie
    \(\begin{array}{ll}\text { String (char*) } & \text { constructor } \\ \text { String (char**) } & \text { constructor }\end{array}\)
    tring (char*) constructor
    Calling SetFirstName with char * \(\begin{array}{r}\text { String (chan }\end{array}\)
    String (char*) constructor
    String (char*) cons
String destructor
    Creating temporary string LastName
String (char*)
Name: String(char*) constructo
Name: Edythe Levine
Address: 1461 Shore Parkway
Salary: \(\begin{aligned} & 20000 \\ & \text { String destructor }\end{aligned}\)
    String destructor
String destructor
    String destructor
    String destructor

ANALYSIS Listing 15.3 uses the same class declarations as Listings 15.1 and 15.2. However, the cout statements have been uncommented. The output from Listing 15.3 has been numbered to make analysis easier.

On line 3 of Listing 15.3, the statement Creating Edie... is printed, as reflected on line 1 of the output. On line 73 an Employee object, Edie, is created with four parameters. The output reflects the constructor for String being called three times, as expectec
Line 75 prints an information statement, and then on line 76 is the statement Edie.SetFirstName("Edythe"). This statement causes a temporary string to be created from the character string "Edythe", as reflected on lines 6 and statement causes a temporary sting to be created from the character sting Edyhhe, as reflected on ine,
of the output. Note that the temporary string is destroyed immediately after it is used in the assignment of the outp
statement.
On line 78, a String object is created in the body of the program. Here the programmer is doing explicitly what the compiler did implicitly on the previous statement. This time you see the constructor on line 9 of the output the compiler did implicitly on the previous statement. This time you see the constructor on line 9 of the
but no destructor. This object will not be destroyed until it goes out of scope at the end of the function.
*include "String.hpp"
class Employee
public:
Employ
End

Employee (const Employee \(\&\) );
Employee \(\&\) operator= (const
const String \& GetFirstName() const Const String \& GetLastName () const (return itsLastName; const String \& GetAddress () const ( return it
long Getsalary () const ( return itssalary; )
void SetFirstName (const String \& fName)
itsFirstName
\(=\) fName

SetLastName (const tring \& 1 Name)
itstastrame \(=\) N Name ; )
id Setaddress (const string \& address)
\{ itsAddress \(=\) address; \(\}\)
void setsalary (long salary) \(\{\) itssalary \(=\) salary; \() ~\)
String
String \(\begin{aligned} & \text { itsFirstName; } \\ & \text { itstastName; }\end{aligned}\)
\(\begin{array}{ll}\text { Ctring } & \begin{array}{l}\text { itsAddrass; } \\ \text { itsSalary; }\end{array}\end{array}\)
Employee: : Employee () :
itsFirstimame ("),
itsLast Name (""),
tstaddress (""),
itsAddress ("'),
itssalary (0)
Employee: : Employee (char * first Name, char * lastName,
char * address, long salary): Har \(*\) adaress, 1 ing sal LsFirst Name (firstNan ssAddress (address),
ssalary (salary)

Employee : : Employee (const Employee \&,
itsFirstMame (rhs. GetFi rstName ()),
LsLastName (rhs. Get Last Name ()
tsAddress (rhs. GetAddress)
tstaddress (rhs. Geetaddress(),
tssalary (rhs. GetSalary ())

Employee \& Employee::operator= (const Employee \& rhs)

tsFirstName \(=\) rhs. GetFirstName (),
tsLastName \(=\) rhs. GetLastName ()
sAddress \(=\) rhs. GetAddress ()
return *this

int main()
Employee Edie ("Jane", "Doe", "1461 Shore Parkway", 20000);
die.Setsalary (20000); Edie. SetFirst Name ("Edyt String Last Name ("Levine");
cout << "Constructor count:" ";
cout << String: :Constructorcount \(\ll\) endl Print Func (Edie) ;
cout \(\lll \ll c o n s t r u c t o r\)
out << string : : Constructortorcount << endl,
cout << String:
rintFunc (EDie) ;
cout <<"Construe
cout << String: : Constructorcount << end1;
return 0

id rPrint Func (const Employee\& Edie)
\begin{tabular}{|c|c|}
\hline \multicolumn{2}{|l|}{\multirow[t]{5}{*}{\begin{tabular}{l}
"Name: " \\
Edie.GetFirstName().GetString() \\
<< Edie.GetLastName().GetStr \\
"\nAddress:
\end{tabular}}} \\
\hline & \\
\hline & \\
\hline & \\
\hline & \\
\hline
\end{tabular}
mam manam
\[
=
\]
\[
1
\]
\[
x^{2}=
\]
it imen unit icery
\[
25=
\]
\[
=v=
\]
Anse oret

Nos.

yout
크를

Tanctactases sace
dict ofter

 antury hernatur


 , minnandem
\(\qquad\)
    and
    \(=2\)
    ,
    \(=\)
    \(4=5\)
    ancmant
    max
    \(=2\)
    emem
    mase
    em

    里
    䏛
    , wes
    and
    mismex
    为
    mon
    mon
    \(==\)
    まus
    20.

    vin
    里

    \(=\)
5
    \(y=\)

    \(=4=\)
    \(=\)
    Nomen
    \({ }^{2}\)
    \(\sqrt{2}\)
    atryata
    aty
    \(=-\)
    aticim

\section*{unctions and Operator Overloading}

\title{

}












\(=\)





5





\({ }^{\text {el }}\) ereturn itestringloffeet],










nt mainn)

\(\underset{\substack{\text { string } \\ \text { strint } \\ \text { string } \\ \text { string } \\ 55 ;}}{ }\)








\section*{Overloading the Insertion Operatc}

\section*{}
cout << thestring. GetStri
```

< tnestrin

```

\title{
INPUT Listing 15.9 Overloading opratar<<0
}


\begin{abstract}






\end{abstract}
\begin{tabular}{l} 
/ default constru \\
String: :string () \\
\hline
\end{tabular}
itsstring \(=\) new char [11,
itisstring \([0]=\uparrow(0) ;\)
it


 tsten=1en;
cout \(\ll\)
" t tstring (int
(/ Converts a character array to a string
String: String const char * const cstring)


// constructorctount + ;
/ copy constructor
string: string coonst
the



delete \([1\) itsstring;
itstuen \(=0 ;\)
its




/ cout < " \(\mid\) tstring operator= \(=n_{n " ;}\)
/non constant offset operator, returns
\(/ /\) refreneac to to haracacter so so it cean be be



\({ }_{\text {return itsstring }}\) (offset];




// changes current string, returns nothing
void string: operatort+(Const strings rhs)




 thestream << thestring. it sstring;
return thestreami
main()

cout << "myAge: " << itsAge << " years.";
The implementation of this friend function is on lines 155-159. All this really does is hide the implementation details of feeding the string to the ostream, and that is just as it should be. You will see more about overloading this operator and operator>> on Day 16.

\section*{Summary}

Today you saw how to delegate functionality to a contained object. You also saw how to implement one class of another by using either containment or private inheritance. Containment is restriced in that the new functions of the contained object. Containment is simpler to use than private inheritance, and should be used when possible.

You also saw how to declare both friend functions and friend classes. Using a friend function, you saw how overload the extraction operator, to allow your new classes to use cout the same as the built-in classes do.
Remember that public inheritance expresses is-a, containment expresses has-a, and private inheritance expresses implemented in terms of. The relationship delegates to can be expressed using either containment or

Q\&A
Q Why is it so important to distinguish between is-a, has-a, and implemented in terms of
A The point of \(\mathrm{C}++\) is to implement well-designed, object-oriented programs. Keeping these modeling. Furthermore, a well-understood design will more likely be reflected in well-designed code. Q Why is containment preferred over private inheritance?
A The challenge in modern programming is to cope with complexity. The more you can use objects as black boxes, the fewer details you have to worry about and the more complexity you can manage. Contained classes hide their details; private inheritance exposes th
A Making one class a friend of another exposes the implementation details and reduces encapsulation The ideal is to keep as many of the details of each class hidden from all other classes as possible.
Q If a function is overloaded, do you need to declare each form of the function to be a friend?
A Yes, if you overload a function and declare it to be a friend of another class, you must declare friend for each form that you want to grant this access to.

\section*{Workshop}

The Workshop contains quiz questions to help solidify your understanding of the material covered and exercis
 before checking the answers in Appendix D. "Answers to Exercises and Quizzes," and make sure yo understand the answers before going to the next day.
1. How do you establish an is-a relationship?
2. How do you establish a has-a relationship?
3. What is the difference between containment and delegation?
4. What is the difference between delegation and implemented in terms of?
3. What is a friend function?
7. If Dog is a friend of Boy, is Boy a friend of Dog?
8. If Dog is a friend of Boy, and Terrier derives from Dog, is Terrier a friend of Boy?
9. If Dog is a friend of Boy and Boy is a friend of House, is Dog a friend of House.
10. Where must the declaration of a friend function appear?

\section*{Exercises}
1. Show the declaration of a class, Animal, that contains a data member that is a string object.
2. Show the declaration of a class, BoundedArray, that is an array,
3. Show the declaration of a class, Set, that is declared in terms of an array.
4. Modify Listing 15.1 to provide the String class with an extraction operator (>>)
5. BUG BUSTER: What is wrong with this program?
\#include <iostream.h>
class Animal
void setValue (Animal\& , int)
class Animal
public:
int GetWeight ()const \& return itsWeight;
int GetAge() const \{ return itsAge; \}
private:
nt itsWeight
int itsAge;
; ;
void setValue (Animald theAnimal, int theWeight
friend class Animal;
theAnimal.itsWeight \(=\) theWeight;
int main()
\{ Animal peppy; setValue (peppy, 5);
6. Fix the listing in Exercise 5 so that it compiles. 7. BUG BUSTER: What is wrong with this code?
\#include <iostream.h>
class Animal
void setvalue (Animal\& , int) ;
void setValue (Animal\&, int,int);
class Animal
friend \(v\)
private:
private:
int itsWeight
nt itsAge;
void setValue (Animal\& theAnimal, int theWeight)
theAnimal.itsWeight = theWeight;
void setValue (Animal\& theAnimal, int theWeight, int theAge)
theAnimal.itsWeight = theWeight; theAnimal.itsAge \(=\) theAge;
int main()
Animal peppy;
Animal peppy;
setValue (peppy, 7,9);
\}
8. Fix Exercise 7 so that it compiles
－What streams are and how they are used
－How to manage input and output using streams

\section*{Overview of Streams}

C＋＋does not，as part of the language，define how data is written to the screen or to a file，nor how data is rea into a program．However，these are clearly essential parts of working with \(\mathrm{C}++\) ，and the standard \(\mathrm{C}++\) library into a program．However，these are clearly essential parts of working
includes the iostream library，which facilitates input and output（I／O）．

The advantage of having the input and output kept apart from the language and handled in libraries is that it i easier to make the language＂platform－independent．＂That is，you can write C＋＋programs on a PC and then recompile them and run them on a Sun workstation．Or you can compile using the GNU compiler on the PC
（during the day at work for instance）and take that code home to use under Linux The compiler creator supp （during the day at work for instance）and take that code home to use under Linux．The compiler creator suppli the right library，and everything works．At least that is the general theory

\section*{Note：A library is a collection of object files（．o under Linux，obj on a PC）that can be linked to your program
provide additional functionality．This is the most basic form of code reuse and has been around since ancient \\ programmers punched 1 s and 0 s into the punched cards．}

Encapsulation
The iostream classes view the flow of data from your program to the screen as being a stream of data，one byte following ano the stis data variables．

One principal goal of streams is to encapsulate the problems of getting the data to and from the disk or the creen．After a stream is created，your program works with the stream and the stream sweats the details．Fig 16.1 illustrates this fundamental ide
\(\square\)
Figure 16．1 Encapsulation through stream：

\section*{Buffering}

Writing to the disk（and to a lesser extent the screen）is very＂expensive．＂It takes a long time（relatively speaking）to write data to the disk or to read data from the disk，and execution of the program is generally blocked by disk writes and reads．To solve this problem，streams provide＂buffering．＂Data is written into th stream，but it is not written back out to the disk immediately．Instead，the stream＇s buffer fills and fills，and when it is full，it writes to the disk all at once

Picture water trickling into the top of a tank and the tank filling and filling，but no water running out of the bottom．Figure 16.2 illustrates this idea．
\(\square\)
Figure 16.2 Filling the buffer
When the water（data）reaches the top，the valve opens and all the water flows out in a rush．Figure 16.3 illustrates this idea


Figure 16．3 Emptying the buffe
After the buffer is empty，the bottom valve closes，the top valve opens，and more water flows into the buffer tank，as illustrated in Figure 16．4．
\(\square\)

\section*{Figure 16．4 Refilling the buffer}

Every once in a while you need to get the water out of the tank even before it is full．This is called＂flushing the buffer．＂Figure 16.5 illustrates this idea．
\(\square\)
Figure 16.5 Flushing the buffer

\section*{Streams and Buffers}
you might expect， \(\mathrm{C}++\) takes an object－oriented view toward implementing streams and buffers．
－The streambuf class manages the buffer，and its member functions provide the capability to fill，empty， flush，and otherwise manipulate the buffer
－The ios class is the base class to the input and output stream classes．The ios class has a streambuf object The istream and ostre
am classes derive from the ios class and specialize input and output stream
－The iostream class is derived from both the istream and the ostream classes and provides input and output methods for writing to the screen．

\section*{andard I／O Objects}

When a \(\mathrm{C}++\) program that includes the iostream classes starts，four objects are created and initialized：
－cin（pronounced＂see－in＂）handles input from the standard input，the keyboard．
－cout（pronounced＂see－out＂）handles output to the standard output，the screer．
cerr（pronounced＂see－err＂）handles unbuffered output to the standard erro device，the screen．Because this is unbuffered，everything sent to cerr is written to the standard error device immediately，without waiting for the buffer to fill or for a flush command to be received．
device，the screen．It is common for this to be＂redirected＂to a log file，as dest section．

Note：The iostream class library is added automatically to your program by the compiler．All you need to do to use
these functions is put he appropriate include statement at the top of your program lisiting．

\section*{Redirection}

Each of the standard devices，input，output，and error，can be redirected to other devices．Standard error is ofter redirected to a file，and standard input and output can be piped to files using operating system command
Redirecting refers to sending output（or input）to a place different than the default The redirection operators fo Reairecing refers to sending output（or input）to a place different than the defaut．The redirec
Linux（as well as other UNIX versions and DOS）are（＜）redirect input and（ \(>\) redirect output．

Piping refers to using the output of one program as the input of another．
Linux（and other UNIX versions）provides advanced redirection capabilities：redirect output（＞），redirect inpu （く），and redirect output to the input of a second program（）

Redirection is more a function of the operating system than of the iostream libraries. C++ just provides access t four standard devices; it is up to the user to redirect the of devices to whateverer alternatives are needed

\section*{Input Using cin}

The global object cin is responsible for input and is made available to your program when you include iostream. In previous examples, you used the overloaded extraction operator (>>) to put data into your program variables. How does this work? The syntax, as you may remember, is the followins

\section*{int someVariable;}
cout << "Enter a numk
cin >> someVariable;
The global object cout is discussed later today; for now, focus on the third line, cin >> someVariable;: What can you guess about cin

Clearly, it must be a global object because you did not define it in your own code. You know from previous operator experience that cin has overloaded the extraction operator (>>) and that the effect is to write whateve data cin has in its buffer into your local variable, someVariable
What may not be immediately obvious is that cin has overloaded the extraction operator for a great variety of parameters, among them int\&, short\&, long\&, double\&, float\&, chart, char*, and so forth. When you write cin >> somel ariabe, ,the ype of son
following function is called:
```

istream \& operator>> (int \&)

```

\section*{Nete that because the parameter is passed by reference, the extraction operator is able to act on the origin} variable. Listing 16.1 illustrates the use of cin

\section*{InPUT. Listing 16.1 cin Handles Different Data Type}
```

/Listing 16.1 - character strings and cin
include <iostream.h>
int main(ef)
int myInt;
long myLong;
double myDouble;
unsigned int myUnsigned;
cout << "Long:
cin >> myLong;
cin >> "Double:
cin >> myDouble;
cin >> "Unsigned:
cout << "\n\nInt:<br>" << myInt << endl;
cout<< "Long:\t" << myLong << endl;
cout << "Float:\t" << myFloat << endl;
return 0

```

\section*{Outrut}
int: 2
Long: 70000
Double: 987654321
Float: 3.33
Unsigned: 25
Int: \(\quad 2\)
Long: \(\quad 70000\)
Doble \(9.87654 \mathrm{e}+08\)
Float: \(3.33 \quad 2\)
Unsigned: \(\quad 25\)

ANaIYsIS On lines 7-11, variables of various types are declared. On lines 13-22, the user is prompted to ent lues for these variables, and the results are printed (using cout) on lines \(24-28\).

The output reflects that the variables were put into the right "kinds" of variables, and the program works as yc might expec

\section*{Strings}
cin can also handle character pointer (char**) arguments; thus, you can create a character buffer and use cin to fi it. For example, you can write the following:
char YourName [50]
cout << "Enter your name: ";
If you enter "Jesse", the variable YourName will be filled with the characters \(\mathrm{J}, \mathrm{e}, \mathrm{s}, \mathrm{s}, \mathrm{e}, \mathrm{l}\). The last character a null character, cin automatically ends the string with a null character, and you must have enough room in th buffer to allow for the entire string plus the null. The null signals "end of string" to the standard librar
functions discussed on Day 21 , "What's Next."

Note: The null character ( \(\left(\mathrm{v}^{\circ}\right)\) is different from the NUL pointer. They may contain the same value in memory
(binary yeros) but they sere
bifferet t

String Problems
After all this success with cin, you might be surprised when you try to enter a full name into a string. cin believ that whitespace is a separator. When it sees a space or a newline character, it assumes the input for the parameter is complete,
illustrates this problem

InPut
```

/Listing 16.2 - character strings and cin
\#include <iostream.h>
int main()
char YourName[50];
cout << "Your firs
cout << "Here it is:" << YourName << endl;
cout << "Your entire name:"
cin >> YourName;
*)

```

\section*{INPUT Listing 16.3 Multiple Input}
//Listing 16.3 - character strings and cin
\#include <iostream.h>
int main()
int myInt;
long myLong;
float myFloat;
hsigned int myUnsigned;
char myWord [50];
cin >> myInt;
cout << "Long: ";
cin >> myLong;
cin >> myDouble;
cout << "Float:"
cin >> myFloat;
cout << "Word:
in >> myWord;
cin >> myUnsigned
```

cout << "\n\nInt:\t" << myInt << endl;
cout << "Long:\t" << myLong << endl;
cout << "Double:\t" << myDouble << endl,
cout << "Float:\t" << myFloat << endl;

```
cout << "Unsigned: \(\backslash t "><\) myUnsigned \(\ll\) endl
cout \(\ll " \backslash n \backslash n\) nnt, Long, Double, Float, Word, Unsigned: ";
cin \(\gg\) myInt \(\gg\) myLong \(\gg\) myDouble;
cin \(\gg\) myFloat \(\gg\) myWord \(\gg\) myUnsigned
cout \(\ll "\) "nnnInt: \(\backslash t " \lll\) myInt \(\ll\) endl
cout << "Long: 1 t" \(\ll\) myLong \(\ll\) endi;
cout << "Double: \(\backslash \mathrm{t}\) " << myDouble << endl;
cout << "Float: \(\backslash t\) " \(\ll\) myFloat << endl;
out <<<"Word: \t" << myWord << endl;
out << "Unsigned:\t" << myUnsigned << endl,
return 0;

\section*{Output}

Int: 2
Long: 30303
Long: 30303
Double: 393939397834
Float: 3.33
Word: Hello
Unsigned: 85
\(\begin{array}{ll}\text { Int: } & 2 \\ \text { Long: } & 30303\end{array}\)
Double: \(3.93939 \mathrm{e}+11\)
Float: 3.33
Word: Hello
Whsigned: 85
Int, Long, Double, Float, Word, Unsigned: 33049383938474736.66 bye Int: \(\quad 3\)
Long: 304938
Long: \(\quad 304938\)
Double: \(3.93847 \mathrm{e}+08\)
Float: 6.66
Word: bye
Word: bye
Unsigned: 4294967294
ANALYSIS Again, several variables are created, this time including a char array. The user is prompted for inpu anain, several variables are created,

On line 34 , the user is prompted for all the input at once, and then each "word" of input is assigned to the appropriate variable. It is to facilitate this kind of multiple assignment that cin must consider each word in the in appropriate variable. It is to facilitate this kind of multiple assignment that cin must consider each word in hee inh
to be the full input for each variable. If cin was to consider the entire input to be part of one variable's input, this kind of concatenated input would be impossible.

Note that on line 42 , the last object requested was an unsigned integer, but the user entered -2 . Because cin believe it is writing to an unsigned integer, the bit pattern of -2 was evaluated as an unsigned integer, and when written ou by cout, the value 4294967294 was displayed. The unsigned value 4294967294 has the exact bit pattern of the signe value -2

Later in this lesson, you will see how to enter an entire string into a buffer, including multiple words. For now, th question arises, "How does the extraction operator manage this trick of concatenation?"

\section*{operator>> Returns a Reference to an istream Objec}

The return value of cin is a reference to an istream object. Because cin itself is an istream object, the return value of extraction operation can be the input to the next extraction
int Varone, varTwo, varThree;
cout << "Enter three numbers:
> VarOne >> varTwo >> varThree;
When you write cin >> VarOne >> varTwo >> varThree;, the first extraction is evaluated (cin >> VarOne). The return value from this is \(\gg\) ther istream object and had written this
((cin >> varOne) >> varTwo) >> varThree;

You will see this technique repeated later when cout is discussed.

\section*{Other Member Functions of cin}

In addition to overloading operator>>, cin has a number of other member functions. These are used when fine control over the input is required

\section*{Single Character Input}
operator>> taking a character reference can be used to get a single character from the standard input. The membe getf can also be used to obtain a single character, and can do so in two ways: get) can be used with no

\section*{Using get() with No Parameters}

The first form of get) is without parameters. This returns the value of the character found and will return EOF (er of file, the end of the is is work:
```

// Listing 16.4-Using get () with no parameters
eam.h>
nt main()
char ch; (ch = cin.get())!= EOF)
cout << "ch:" << ch << endl;
很缺"\nDone!\n";

```

ANALYsIS \({ }_{\text {On }}\) line 6 , a local character variable is declared. The while loop assigns the input received from cingeet to th, and if it it not EOF, the string is printed out. This output is buffered until an end of line is read,
however. When EOF is encountered (when you press CrtI+D under Linux, or CrrIt on a DOS machine), the loop exits.
of ser) although hit is now part of the

\section*{Using get 0 with a Character Reference Parameter}

When a character is passed as input to geto, that character is filled with the next character in the input stream,
The return value is ani ostream object, and so this form of getol can be concatenated, as illustrated in istiting 18

Listing 16.5 - Using get () with parameters
nt main()
char a, b, c;
cout << "Enter three letters:"
cin.get (a) .get (b). .get (c);
cout \(\ll\)
return 0;
Output
Enter three letters: one

ANALYSIS On line 6 , three character variables are created. On line 10 , cin.gett is called three times
 called and the third leteter is put in in .
```

cin.get (a) >> b

```

In this form, in.get(a) evaluates to cin, so the second phrase is cin >> b;,

Output
Enter string one: Now is the time Stringone: Now is the time
Enter string two: For all good
StringTwo: For
```

// Listing 16.7 - Using getline()

```
int main()
char stringOne [256];
har strig
cout << "Enter string one:";
cin. getline (stringone, 256);
cout << "stringOne: " << stringOne << endl;
cout << "Enter string two: ";
cin >> stringTwo;
out << "stringTwo: " << stringTwo << endl
cout << "Enter string three: "
cin. getline(stringThree, 256);
cout \(\ll\) "stringThree: \(" \ll\) stringThree \(\ll\) endl;
turn 0;
\}

\section*{OUtput}

Enter string one: one two three
stringOne: one two three
Enter string two: four five six
stringTwo: four
Enter string three: stringThree: five six
ANAIYSIS This example warrants careful examination; it contains some potential surprises. On lines 6-8 hree character arrays are declared.

On line 10 , the user is prompted to enter a string, and that string is read by getline(). Like get(), getline() takes a buffer and a maximum number of characters. Unlike get() however, getline) reads the terminating newline character and throws it away. With gett), the terminating newline character is not thrown away. It is left in th input buffer.

On line 14, the user is prompted again, and this time the extraction operator is used. The user enters "four fiv six" and the first word, four, is put in stringTwo. The string "Enter string three" is then displayed, and getline() called again. Because "fiive six" is still in the input buffer, it is immediately read up to the newline character; getline() terminates and the string in stringThree is printed on line 20

The user has no chance to enter string "three" because the second getline() call is fulfilled by the string remainin in the input buffer after the call to the extraction operator on line 15. There is no way for the user to ever ent string "three" because stringTwo ends when the Enter key is pressed, but the Enter key remains in the inp buffer and immediately satisfies the second getline() call.

The extraction operator (>>) reads up to the first whitespace and puts the word into the character array
The member function get) is overloaded. In one version, it takes no parameters and returns the value of the character it receives. In the second version, it takes a single character reference and returns the istream object b reference

In the hird and final version, geto takes a character array, a number of characters to get, and a termination character (which defaults to a newline character). This version of get) reads characters into the array until it g to one fewer than its maximum number of characters or it encounters the termination character, whichever comes first. If get() encounters the termination character, it leaves that character in the input buffer and stops reading characters.

The member function getline() also takes three parameters: the buffer to fill, one more than the maximum number of characters to get, and the termination character getline() functions the same as get) does with these parameters, except getline() throws away the terminating character.

\section*{cin.ignore() for Input Cleanup}

Alimes, you want to ignore the remaining characters on a line until you hit either end of line (EOL) or end file (EOF). The member function ignore0 serves this purpose. ignore) takes two parameters: the maximum number of characters to ignore and the termination character. If you write ignore( \(\left(80\right.\), \(\mathrm{nn}^{\prime}\) '), up to 80 characters wi be thrown away until a newline character is found. The newline character is then thrown away and the ignore() statement ends. Listing 16.8 illustrates the use of ignore

\section*{InPUT. Listing 16.8 Using ignore(}
```

Histing 16.8 - Using ignore()
\#include <iostream.h>
int main()
char stringOne[255];
char stringTwo[255];
cout << "Enter string one:
in.get(stringOne,255);
out << "String one:" << stringOne << endl;
out << "Enter string two:";
in.getline(stringTwo, 255)
Out << "String two:" << stringTwo << endl;
cout << "\n\nNow try again...\n";
cout << "Enter string one: ";
cin.get(stringOne,255);
cout << "String one:" << stringOne<< endl;
cin.ignore(255,'\n');
out << "Enter string two:"
cin.getline(stringTwo,255);

```
    turn 0

\section*{Output}

Enter string one:once upon a time
String one: once upon a time
Enter string two: String two:
Now try again.
Enter string one: once upon a time
String one: once upon a time
String Two: there was a
On lines 6 and 7 , two character arrays are created. On line 9 , the user is prompted for input and types "once upon a time" and presses Enter. On line 10, get() is used to read this string. get) fills stringOne and terminates a the newline character, but leaves the newline character in the input buffe

On line 13 , the user is prompted again, but the getineo on line 14 reads the newline character that is already in
the buffer and terminates inmediatly before the user can On ine 1 , hin user is prompted again, but the eetineo on line 14 reads the
On line 19, the user is prompted again and puts in the same first line of input. This time, however, on line 23 ignore) is sused th "eat" the newline character. Thus, when the getineo call on line 26 is reached, the input buffe peek() and putback()

The input object cin has two additional methods that can come in rather handy: peeko, which looks at but does not extract the next character, and putbacko, which inserts a character into the input stream. Listing 16.9
illustrates how these might be wsed illustrates how these might be usec

INPUT Listing 16.9 Using peeko and putbacko
```

// Listing 16.9-Using peek() and putback()
int main()
coute< ch; "enter a phrase:
if (ch == !!')
else
cout << ch; ()
eturn 0;

```

\section*{enter a phrase: Now! is
NowSisthestimeforsfuns}

ANALYSIS O line 6 , a character variable, ch, is declared, and on line 7 , the user is prompted to enter a phras The purpose of this program is to turn any exclamation marks (!) into dollar signs ( \((\) ) and to to remove any poun

The program loops as long as it is getting characters other than the end of file (Crrl+D under Linux, Crril+C 0 .
Crrl+D on other operating systems) (Remember that cine Ctrl+D on other operating systems). (Remember that cin. gete returns 0 for the end of the file). If the current
character is an exclamation point, tit sthrown away and the \(\$\) symbol is put back into the input buffer, it will character is an exclamation point, it is thrown away and the symbol is put back into the input buffer, it wiil
reat the next time. If the current ite is is on an exclamation point it is printed. The next character is "peeked
at, and when pound symbols are found, they are removed.

This example is not the most efficient way to do either of these things (and it will not find a pound symbol if is the first character), but it does illustrate how these methods work. They are relatively obscure, so do not spend a lo of time worryin
come in handy eventually.

\section*{Tip: peekfo and putuackto are typically used for parsing strings and other data, such as when writing a compiler.}

\section*{Output with cout}

You have used cout along with the overloaded insertion operator (<<) to write strings, integers, and other You have used cout along with the overioaded insertion operator (<<) Io write strings, integers, and other
numeric data to the screen. It is also possible to format the data, aligning columns and writing the numeric d Flushing the Output

You already saw that using end will flush the output buffer. end calls cou's member function fusho, which
writes all the data it is buffering. You can call the flusho method directly, either by calling the flusho memb writs ad one dataitis thutering. Fo ,
method or by writing the following:
```

cout << flus

```

This can be convenient when you need to ensure that the output buffer is emptied and that the contents are

\section*{Related Functions}

Just as the extraction operator can be supplemented with get/( and getineo, the insertion operator can b
supplemented with puw and wis The function put) is used to write a single character to the output device. Because put) returns an ostream reference and because cout is san ostream objee
operator. Listing 16.10 illustrates this idea.

InPut Listing 16.10 Using puto
(/ Listing \(16.10-\) Using put()
\#include <iostream.h>
int main()
cout.put
eturn 0 ;

Output
Hell
Note: Some compiers (other than GNU have
word Hello, you may wanto skip this isiting.
ANALYsIS \({ }^{\text {Line } 6}\) is evaluated like this. coutpu( obict. This leaves whites following:


The letter e is written, leaving cout.put( \((9)\). This process repeats, each letter being written and the cout objec returned until the final character ( (mn') is written and the function returns The function write0 works the same as the insertion operator (««), excepp that it takes
function the maximum number of characters to write. Listing 16.11 illustrates its use.

INPUT Listing 16.11 Using writed
// Listing 16.11 - Using write()
include <string.h>
int main()
char One[] = "one if by land";
nt fullLength \(=\operatorname{strlen}(0 n e)\)
Int tooshort \(=\) fulliLength -4 ;
int tooLong \(=\) fulliLength +6 ;



Outpu
One if by land
one if by
One if by land i?!
ote: The last line of output may look different on your computer-really it is just garbage

\section*{On line 15 , the complete plraa:
the correct phasas is printed.}

On line 16, the phrase is printed again, but it is four characters shorer than the full phrase, and that is reflecte in the output.
On line 17 , the phrase is printed again, but this time write is is intucted
phrase is writen, the next six byyes of contiguous memory are written.
Manipulators, Flags, and Formatting Instructions
The output stream maintains a number of tatet flags, deternining which base (decimal or hexadecimal) to use
how wide to make the fields, and what characeter to ouse to fillin in ields. Astate flag is a byte whose individual
 bits are each assigned aspecial meaning. Manipulating bits in this way is dis
Each of ostream's flags can be set using member functions and manipulators.

\section*{Using cout.width(}

The default width of your output will be just enough space to print the number, character, or string in the outp
buffer. You can change this by using widho. Because widtho is a member function, it must be invoked with a buffer. You can change this by using width). Because widho is a member function, it must be invoked with a
cout boject. It only changes the width of the very next output field and then immediately reverts to the default.
cot

Input
```

// Listing 16.12- Adjusting the width of output
nt main()
cout << "Start,
cout << "Start,
lol
cout<< \start
return 0;

```
Output
\begin{tabular}{ll} 
Start > & \begin{tabular}{l} 
123< End \\
Start \\
Start \(>123456<\) End
\end{tabular} \\
123< Next \(>456<\) End
\end{tabular}

ANalysIs, The first oupput, on lines \(6 .-8\) prints the number 123 within a field whose width is set to 25 on line The second line of output first prints the value 123 in the same field whose widh is set to 25 , and then prints ti
value 456 . Note that 456 is prined in a field whose width is reset to just large enough; as stated, the effect of ast olly as long as the very next output.

The final output reflects that setting a width that is smaller than the oupput is the same as seting a width thati
just large e enough.
 you may want to fill the area with other characteres, such as asterisks. To do this, you \(\mathrm{c}_{5}\)
parameere the character you want used as a fill character. Listing 16.13 illustrates this.
INPUT. Listing 16.13 Using fillo
/ Listing 16.13-fill(
include <iostream.h>
cout \(<\) "Start
cout width \((255) ;\)
cout \(\ll 1123 \ll\)
cout << "start
cout.width (25);
cout.finill \((* *)\);
```

cout<<12

```

Start >
Start \(>* * * * * * * * * * * * * * * * * * 123<~ E n d ~\)


\section*{Setting iostream Flag:}

The iostream object keep track of theis state by using flags
in one or another of the predefined enumerated constants.
Objects are said to have state when some or all of their data represents a condition that can change during the
course of the program.
For example, you can set whether to show trailing zeros (so that 20.00 does not become truncated to 20 ). To
turn trailing zeros sel
The enumerated constants are scoped to the iostram class (ios) and thus are called with the full qualification ios:flagname, such as ios:showpoin





INPUT. \(_{\text {Listing }} 16.14\) Using seff

\(\underset{\substack{\text { include <iostream.h> } \\ \text { \#include }<\text { iomanip. } h>}}{ }\)
int main()
const int number \(=185 ;\)
cout \(\ll\) "The number is" \(\ll\) number << endl;
cout << "The number is" << hex << number << endl;
cout. setf (ios: : showbase);
cout <<"The number is"
cout \(\ll\) "The number is ";
cout.width (10);
<< hex << number << en
cout. width (10) ;
cout. setf \(f\) (ios: : : left);
cout \(\ll\) hex \(\ll\) number
cout \(<\) "The num
cout width (10) ;
cout
cout. setf (ios:: internal),
out << hex \(\ll\) number \(\ll\) end1;
return 0; "The number is:" << setw (10) << hex << number << endl;

Ouvirut be displayed in hexadecimal form as b9. (The value \(b\) in hexadecimal form represents 11 . Eleven times 16 equals 176; add the 9 for a total of 185 .)

On line 12, the flag showbase is set. This causes the prefix \(0 x\) to be added to all hexadecimal numbers, as reflected in the output.

On line 16 , the width is set to 10 , and the value is pushed to the extreme right. On line 20 , the width is again set to 10 , but this time the alignment is set to the left, and the number is again printed flush left.

On line 25 , again the width is set to 10 , but this time the alignment is internal. Thus the 0 x is printed flush left, but the value b9 is printed flush right.

Finally, on line 29 , the concatenation operator setw() is used to set the width to 10 , and the value is printed again

\section*{Streams Versus the printf() Function}

Most C++ implementations also provide the standard C I/O libraries, including the printf() statement. Although printf() is in some ways easier to use than cout, it is much less desirable.
printf() does not provide type safety, so it is easy to inadvertently tell it to display an integer as if it was a character, and vice versa. printf() also does not support classes, and so it is not possible to teach it how to print your class data; you must feed each class member to printf() one by one.

On the other hand, printf() does make formatting much easier because you can put the formatting characters directly into the printf() statement. Because printf() has its uses and many programmers still make extensive use of it, this section briefly reviews its use.

To use printf(), be sure to include the stdio.h header file. In its simplest form, printf() takes a formatting string as its first parameter and then a series of values as its remaining parameters.

The formatting string is a quoted string of text and conversion specifiers. All conversion specifiers must begin with the percent symbol (\%). The common conversion specifiers are presented in Table 16.1.

Table 16.1The Common Conversion Specifiers
\begin{tabular}{ll}
\hline Specifier & Used For \\
\hline\(\hat{\text { A }}\) & Strings \\
\(\hat{\text { A }}\) & Integers \\
\(\hat{A}\) & Long integer \\
\(\hat{A}\) & Double \\
\(\hat{A}\) & Float \\
\hline
\end{tabular}

Each of the conversion specifiers can also provide a width statement and a precision statement, expressed as a float, where the digits to the left of the decimal are used for the total width, and the digits to the right of the decimal provide the precision for floats. Thus, \(\% 5 \mathrm{~d}\) is the specifier for a 5 -digit-wide integer, and \(\% 15.5 \mathrm{f}\) is the specifier for a 15 -digit-wide float, of which the final five digits are dedicated to the decimal portion. Listing 16.15 illustrates various uses of printf().

\section*{InPUT}

Listing 16.15 Printing with printf()
```

1: \#include <stdio.h>
int main()
{
printf("%s","hello world\n");
char *phrase = "Hello again!\n";
printf("%s",phrase)
int x = 5;
printf("%d\n",x);
char *phraseTwo = "Here's some values: ";
char *phraseThree = " and also these: ";
int y = 7, z = 35;
long longVar = 98456;
float floatVar = 8.8f;
printf("%s %d %d %s %ld %f\n",phraseTwo,y,z,
18:
f phraseThree,longVar,floatVar)
19:
20
21:
22
23:

```

\section*{Output}
```

hello world
Hello again!
Hel
Here's some values: 7 35 and also these: 98456 8.800000
Formatted: 7 35 8.800000

```

\section*{Analysis} The first printf) statement, on line 4 , uses the standard form: the term printf, followed by a quoted
strersion specifier (in this case \(\%\) s), followed by a value to insert into the conversion specifier The \%s indicates that this is a string, and the value for the string is, in this case, the quoted string "hello world.

The second printf() statement is the same as the first, but this time a char pointer is used, rather than quoting the string right in place in the printf() statement.

The third printf(), on line 10 , uses the integer conversion specifier, and for its value the integer variable x . The fourth printf() statement, on line 18 , is more complex. Here six values are concatenated. Each conversion specifier is supplied, and then the values are provided, separated by commas.

Finally, on line 21, format specifications are used to specify width and precision. As you can see, all this is somewhat easier than using manipulators.

\section*{Frequently Asked Questions}

\section*{FAQ: Can you summarize how I manipulate output?}

Answer (with special thanks to Robert Francis): To format output in C++, you use a combination of special characters, output manipulators, and flags.

The following special characters are included in an output string being sent to A \(\hat{A}\) using the insertion operator:

> In-Newline character
(r-Carriage return
tt-Tab
॥—Backslash
Iddd (octal number)—ASCII character
la—Alarm (ring bell)

\section*{Example:}
cout << "laAn error occuredlt"
Rings the bell, prints an error message, and moves to the next tab stop. Manipulators are used with the cout operator. Those manipulators that take arguments require that you include iomanip.h in your file.

The following is a list of manipulators that do not require imanip.h:
flush-Flushes the output buffer
endl-Inserts newline character and flushes the output buffer
oct-Sets output base to octal
dec-Sets output base to decimal
hex-Sets output base to hexadecimal
The following is a list of manipulators that do require imanip.h setbase(base) - Sets output base \((0=\) decimal, \(8=\) octal, \(10=\) decimal, \(16=\) hex)
setw (width)—Sets minimum output field width
setfill (ch)-Fills character to be used when width is defined
setprecision (p)-Sets precision for floating-point numbers
setiosflags (f)-Sets one or more ios flags
resetiosflags (f)-Resets one or more ios flags

\section*{Example}
cout \(\ll\) setw(12) \(\ll\) setfill('\#') \(\ll\) hex \(\ll\) x \(\ll\) endl;
sets the field width to 12 , sets the fill character to ' \(\#\) ', specifies hex output, prints the value of ' \(x\) ', puts a newline character in the buffer, and flushes the buffer. All the manipulators except flush, enal, and setw remain in effect until changed or until the end of the program. setw returns to the default after the current cout.

The following ios flags can be used with the setiosflags and resetiosflags manipulators:
ios::left-Left-justifies output in specified width
ios::right—Right-justifies output in specified width
ios::internal-Sign is left-justified, value is right-justified
ios::dec-Decimal output
ios::oct-Octal output
ios::hex-Hexadecimal output
ios::showbase-Adds \(0 x\) to hexadecimal numbers, 0 to octal numbers
ios::showpoint-Adds trailing zeros as required by precision
ios::uppercase-Hex and scientific notation numbers shown in uppercase
ios::showpos-The + sign shown for positive numbers
ios::scientific-Shows floating point in scientific notation
ios::fixed-Shows floating point in decimal notation
Additional information can be obtained from file ios.h and from the GNU library documentation.

\section*{File Input and Output}

Streams provide a uniform way of dealing with data coming from the keyboard or the hard disk and going out to the screen or hard disk. In either case, you can use the insertion and extraction operators or the other related functions and manipulators. To open and close files, you create ifstream and ofstream objects (for input and outpu
respectively) as described in the next few sections. Because ifstream objects are similar to ofsteam objects, the respectively) as described in the next few sections. Because ifstream objects are similar to ofsteam objects, the material covered is limited. Just apply ofstream techniques.

\section*{Using ofstream}

The particular objects used to read from or write to files are called ofstream objects. These are derived from the iostream objects you already used.

To get started with writing to a file, you must first create an ofstream object, and then associate that object with a particular file on your disk. To use ofstream objects, you must be sure to include fstream.h in your program.

Note: Because fstream.h includes iostream.h, you do not need to include iostream explicitly.

\section*{Condition States}

The iostream objects maintain flags that report on the state of your input and output. You can check each of these flags using the Boolean functions eof(), bad(), fail(), and good(). The function eof() returns true if the iostream object has encountered EOF, end of file. The function bad() returns true if you attempt an invalid operation. The function fail() returns true anytime bad() is true or an operation fails. Finally, the function good() returns true anytime all three of the other functions are false.

\section*{Opening Files for Input and Output}

To open the file myfile.cpp with an ofstream object, declare an instance of an ofstream object and pass in the filename as a parameter:
ofstream fout("myfile.cpp");
Opening this file for input works the same way, except that it uses an ifstream object:
ifstream fin("myfile.cpp");
Note that fout and fin are names you assign; here fout has been used to reflect its similarity to cout, and fin has been used to reflect its similarity to cin.

The input file is then closed, and the same file is reopened, this time in append mode, on line 22. After this ope (and every open), the file is tested to ensure that the file was opened properly. Note that if(ffout) is the same as testing if (fout.fail()). The user is then prompted to enter text, and the file is closed again on line 33.
inally, as in Listing 16.16, the file is reopened in read mode; however, this time fin does not need to b redeclared. It is just reassigned to the same filename. Again the open is tested, on line 36, and if all is well, the contents of the file are printed to the screen, and the file is closed for the final time

\section*{Do}

DO test each open of a file to ensure that it opened
DON'T try to close or reassign cin or cout.
successfully
DO reuse existing ifstream and ofstream objects
DO close all fstream objects when you are done using
them.

\section*{Binary Versus Text Files}

Some operating systems, such as DOS, distinguish between text files and binary files. Text files store everything as text (as you might have guessed), so large numbers such as 54,325 are stored as a string o everything as text (as you might have guessed), so large numbers such as 54,325 are stored as a string of
numerals ( \(5^{\prime},{ }^{\prime} 4\) ' ',', ' 3 ', \(2^{\prime},{ }^{\prime} 5\) '). This can be inefficient, but has the advantage that the text can be read and numerals ( \(5,4,,, 3,2,5)\) This can be inefficient, but has the advantage that the text can be read an
manipulated using the standard Linux commands like cat, more, vi, head, tail, grep, and so on, as well as simple programs such as the DOS type command.

To help the runtime library distinguish between text and binary files, C++ provides the ios::binary flag. On man systems, this flag is ignored because all data is stored in binary format. On some rather prudish systems, the ios::binary flag is illegal and will not compile-or even worse, is silently ignored.

Caution: Under Linux, the operating system views files as a stream of bytes-it is up to the programmers to
impose structure on that stream in their programs. The GNU compiler under Linux handles the ios:binary flag mpose structure on that stream in their programs. The GNU compiler under Linux handles the ios:binary flag properly. Under other operating systems, it may not.

Binary files can store not only integers and strings, but also entire data structures. You can write all the data at one time by using the write) method of fstream.
If you use write(), you can recover the data using read(). Each of these functions expects a pointer to character however, so you must cast the address of your class to be a pointer to charactel

The second argument to these functions is the number of characters to write, which you can determine using sizeof(). Note that what is being written is the data, not the methods. What is recovered is only data. Listing 16.18 illustrates writing the contents of a class to a file.
isting 16.18 Writing a Class to a File

\section*{include <fstream.h>}
class Animal
f
ublic:
Animal(int weight, long days):itsWeight(weight),
tsNumberDaysAlive(days) \{
\(\sim\) Animal () \{\}
int GetWeight()const \{ return itsWeight; \}
void SetWeight(int weight) \{ itsWeight = weight; \}
long GetDaysAlive()const \{ return itsNumberDaysAlive; \} void SetDaysAlive(long days) \{ itsNumberDaysAlive = days;
private:
int itsWeight;
long itsNumberDaysAlive;
\};
int main() // returns 1 on error
char fileName[80]
cout << "Please enter the file name: "
cin >> fileName
ofstream fout(fileName,ios::binary);
if (!fout)
\{
cout << "Unable to open " << fileName << " for writing. \(\backslash n\) "; return(1);
\}
Animal Bear \((50,100)\)
fout.write((char*) \&Bear,sizeof Bear);
fout.close();
ifstream fin(fileName,ios::binary)
if (!fin)
ई cout << "Unable to open" << fileName << " for reading. \(\backslash \mathrm{n}\) "; return (1) ;
\}
Animal BearTwo (1,1)
cout << "BearTwo weight: " << BearTwo.GetWeight() << endl;
cout << "BearTwo days: " << BearTwo.GetDaysAlive() << endl;
fin.read((char*) \&BearTwo, sizeof BearTwo);
cout << "BearTwo weight: " << BearTwo.GetWeight() << endl; cout << "BearTwo days: " << BearTwo.GetDaysAlive() << endl; fin.close();
return 0;

BearTwo days: 1
BearTwo weight: 50
BearTwo days: 1005

\section*{Analysi} for output in binary mode. An animal whose weight is 50 and who is 100 days old is created on line 34 , and its data is written to the file on line 35

The file is closed on line 37 and reopened for reading in binary mode on line 39 . A second animal is created
 on line 51 , wiping out the existing data and replacing it with the data from the file

\section*{Command-Line Processing}

Many operating systems, such as DOS and UNIX, enable the user to pass parameters to your program when th
program starts. These are called command-line options and are typically separated by spaces on the command program starts. p he:
line. For example:
SomeProgram Param1 Param2 Param3
These parameters are not passed to maino directly. Instead, every program's maino function is passed two parameters. The first is an integer count of the number of arguments on the command line. The program name itself is counted, so every program has at least one parameter. The example of a command line shown
previously has four. (The name somerosam pus the three paramerers make total of four previously has four. (The name SomeProgram plus the three parameters make a total of four command-line agnes.)
The second parameter passed to maino is an array of pointers to character strings. Because an array name is a
constant pointer to the first element of the array, you can declare this argument to be a pointer to a pointer to constant pointer to the first element of the array, you can dect,
char, a pointer to an array of char, or an array of arrays of char

Typically, the first argument is called argc (argument count), but you may callit anything you like. The seconc argument is often called argv (argument vector), but again this is just a convention.
It is common to test arge to ensure that you received the expected number of arguments and to use argy to acc the strings themselves. Note that argy[0] is the name of the program, and argv(I) is the first parameter to the the strings themselves. Note that argy(0) is the name of the program, and argvIII is the first parameter to the
program, represented as a string. .f your program takes two numbers as arguments, you will need to translate program, represented as a string. If your program takes two numbers as arguments, you will need to translate
these numbers to strings. On Day 21 you will see how to use the standard library conversions. Listing 16.19 illustrates how to use the command-line arguments
\begin{tabular}{|l|l}
\hline INPUT. & Listing 16.19 Using Command-line Argument \\
\hline
\end{tabular}
```

\#include <iostream.h>
int main(int argc, char **argv)
cout << "Received" << argc << " arguments.
for (int i=0; i<argc; i++)
return 0;

```

\section*{Outur}

Testrprogram Teach Yourself C++ In 21 Days
Received 7 arguments...
Received 7 arguments.
argument \(0:\) TestProgram.exe
argument 1:
1:
Teach
argument 2: Yourself
argument \(3: \mathrm{C+}\)
argument \(4:\) In
argument 4: In
argument 5: 21
argument 6: Days
Note: This works best when run from a command line. If you are using a GUI interface like \(K\) Deskop
Enviroment (KDE), you need to open a terminal session If you are using an operating system other than


ANALYSIS The function maino declares two arguments: argc is an integer that contains the count of commanc line arguments, and argy is a pointer to the array of strings. Each string in the array pointed to by argv is a command-line argument. Note that argy could have easily been declared as char *yrgyl or char argylIII. It is a
a matter of programming stye how you declare args; even though this program declared it as a pointer to a
lass Animal
public:
fitsNumberDaysAl, long days):itsWeight (weight),
fitsNumberDaysAlive (days) \{)
int GetWeight () const f return itsWeight,
void SetWeight (int weight) it itsWeight
long GetDaysAlive ()const \{return itsNumberDaysAlive; \(\}\)
void SetDaysAlive(long days) fitsNumberDaysAlive = days;
private
int itsWeight;
long itsNumberDaysAlive;
int main(int argc, char *argv[1) // returns 1 on error
if (argc != 2)
out << "Usage: " << argv[0] <<" <filename>" << endl; return (1);

\section*{[11, ios: binary)}
if (!fout)
cout <<" "Unable to open" << argv[1] <<" for writing. \(\backslash \mathrm{n}^{\prime \prime}\)

Animal Bear \((50,100)\);

fout.close();
if (!fin)
cout << "Unable to open" << argv[1] <<" for reading. \(\backslash \mathrm{n} "\); return (1);

Animal BearTwo (1,1)
cout << "BearTwo weight: " << BearTwo.GetWeight () << endl;
cout << "BearTwo days:" << BearTwo.GetDaysAlive() << endl;
fin.read((char*) \&BearTwo, sizeof BearTwo);
cout << "BearTwo weight:" << BearTwo.GetWeight() << end1;
cout << "BearTwo days:" << BearTwo.GetDaysAlive() << endl; return 0;

On lines 22-26, the program ensures that the expected number of arguments (exactly two) is received. If the use fails to supply a single filename, an error message is printed:

Usage TestProgram <filename>

Then the program exits. Note that by using argv[0] rather than hard-coding a program name, you can compile this program to have any name, and this usage statement will work automatically.

On line 28 , the program attempts to open the supplied filename for binary output. No reason exists to copy the filename into a local temporary buffer. It can be used directly by accessing argv[1].

This technique is repeated on line 40 when the same file is reopened for input, and it is used in the error condition statements when the files cannot be opened, on lines 25 and 31.

\section*{Summary}

Today streams were introduced, and the global objects cout and cin were described. The goal of the istream and ostream objects is to encapsulate the work of writing to device drivers and buffering input and output.

Four standard stream objects are created in every program: cout, cin, cerr, and clog. Each of these can be "redirected" by many operating systems.

The istream object cin is used for input, and its most common use is with the overloaded extraction operator ( \(\gg\) ) The ostream object cout is used for output, and its most common use is with the overloaded insertion operator (<<).

Each of these objects has a number of other member functions, such as get() and put(). Because the common forms of each of these methods returns a reference to a stream object, it is easy to concatenate each of these operators and functions.

The state of the stream objects can be changed by using manipulators. These can set the formatting and display characteristics and various other attributes of the stream objects.

File I/O can be accomplished by using the fstream classes, which derive from the stream classes. In addition to supporting the normal insertion and extraction operators, these objects also support read() and write() for storing and retrieving large binary objects.

\section*{Q\&A}

Q How do you know when to use the insertion and extraction operators and when to use the other member functions of the stream classes?
A In general, it is easier to use the insertion and extraction operators, and they are preferred when their behavior is what is needed. In those unusual circumstances when these operators don't do the job (such as reading in a string of words), the other functions can be used.

\section*{Q What is the difference between cerr and clog?}

A cerr is not buffered. Everything written to cerr is immediately written out. This is fine for errors that are written to the screen, but it may have a performance cost that is too high for writing logs to disk. clog buffers its output, and thus can be more efficient.
Q Why were streams created if printf() works well?
A printf() does not support the strong type system of \(\mathrm{C}++\), and it does not support user-defined classes.
Q When would you ever use putback()?
A When one read operation is used to determine whether a character is valid, but a different read operation (perhaps by a different object) needs the character to be in the buffer. This is most often used when parsing a file; for example, the \(\mathrm{C}++\) compiler might use putback \((\) (

\section*{Q When would you use ignore) ?}

A A common use of this is after using get(). Because get() leaves the terminating character in the buffer, it is not uncommon to immediately follow a call to get() with a call to ignore( \((1\), ' \(4 n\) '), . This is often used in parsing.
Q My friends use printf() in their \(\mathbf{C}_{++}\)programs. Can I?
A Sure. You'll gain some convenience, but you'll sacrifice type safety.

\section*{Workshop}

The Workshop contains quiz questions to help solidify your understanding of the material covered and exercise to provide you with experience in using what you learned. Try to answer the quiz and exercise questions before checking the answers in Appendix D, "Answers to Exercises and Quizzes," and make sure you understand the answers before going to the next day.

Quiz
1. What is the insertion operator, and what does it do?
2. What is the extraction operator, and what does it do?
3. What are the three forms of cin.get(), and what are their differences?
4. What is the difference between cin.read() and cin.getline()?
5. What is the default width for outputting a long integer using the insertion operator?
6. What is the return value of the insertion operator?
7. What parameter does the constructor to an ofstream object take?
8. What does the ios::ate argument do?

Exercises
1. Write a program that writes to the four standard iostream objects: cin, cout, cerr, and clog.
2. Write a program that prompts the user to enter her full name and then displays it on the screen.
3. Rewrite Listing 16.9 to do the same thing, but without using putback() or ignore().
4. Write a program that takes a filename as a parameter and opens the file for reading. Read every character of the file and display only the letters and punctuation to the screen. (Ignore all nonprinting characters.) Then close the file and exit.
5. Write a program that displays its command-line arguments in reverse order and does not display the program name.

\section*{Chapter 17}

\section*{Namespaces}

A new addition to ANSI C++ is the use of namespaces to help programmers avoid name clashes when using more than one library. Today you will learn:
- How functions and classes are resolved by name
- How to create a namespace
- How to use a namespace
- How to use the standard namespace std

\section*{Getting Started with Namespaces}

Name conflicts have been a source of aggravation to both C and \(\mathrm{C}++\) developers. The ANSI standardization offered a chance to resolve this problem through the use of namespaces, but be warned: Not all compilers support this feature yet. \(g_{++}+2.9 .5\) and newer do support this feature. \(g_{++}\)version 2.7 .2 is one of those that does not really support it and will give you the following message:
```

warning: namespaces are mostly broken in this version of g+-

```

A name clash happens when a duplicate name with matching scope is found in two parts of your program. The most common occurrence can be found in different library packages. For example, a container class library wil and Templates.")

It is not a surprise to find a List class also being used in a windowing library. Suppose that you want to maintair a list of windows for your application, or assume that you are using the List class found in the container class
library. You declare an instance of the window library's List to hold your windows, and you discover that the library. You declare an instance of the window library's List to hold your windows, and you discover that the
member functions you want to call are not available. The compiler has matched your List declaration to the List container in the standard library, but what you really want is the List found in the vendor-specific window library.

Namespaces are used to partition the global namespace and to eliminate, or at least reduce, name conflict. Namespaces are similar in some ways to classes, and the syntax is very similar.

Items declared within the namespace are owned by the namespace. All items within a namespace have public visibility. Namespaces can be nested within other namespaces. Functions can be defined within the body of th namespace or defined outside the body of the namespace. If a function is defined outside the body of the mespace, it must be qualified by the namespace's name.

\section*{Resolving Functions and Classes by Name}

As it parses source code and builds a list of function and variable names, the compiler checks for nam conflicts. Those conflicts that cannot be resolved by the compiler might be resolved by the linker.
The compiler cannot check for name clashes across translation units (for example, object files); that is the purpose of the linker. Thus, the compiler will not even offer a warning.
It is not uncommon for the linker (Id under Linux) to fail with the message Identifier multiply defined (identifier is so different translation units. You get a compiler error if you redefine a name within a single file having the same
```

// file first.cpp
int integerva,
int integervalue = 0;
return 0;
// file second.cpp
int integervalue = 0
// end of second.cpp
The GNU linker announces the following diagnostic messages:
second.cc: multiple definition of `integerValue'
If these names were in a different scope, the compiler and linker would not complain.
It is also possible to receive a warning from the compiler concerning identifier hiding. The compiler should
warn, in first.cpp, that integerValue in main() is hiding the global variable with the same name
To use the integerValue declared outside main(), you must explicitly scope the variable to global scope. Conside this example, which assigns the value 10 to the integerValue outside main() and not to the integerValue declared within main():

```
```

// file first.cpp

```
// file first.cpp
int integerValue =
int main()
int integerValue = 0 ;
    ::integerValue = 10 ; //assign to global "integerValue"
    return 0 ;
// file second.cpp
int integervalue = 0
// end of second.cpp
```

Note: Note the use of the scope resolution operator:., indicating that the integerValue being referred to is global,
not local.

The problem with the two global integers defined outside of any functions is that they have the same name and visibility, and thus will cause a linker error.

```
Note: The term visibility is used to designate the scope of a defined object, whether it is a variable, a class, of 
function. For example a variable declared and defined outside any function has file, or global, scope. The
Visibility of this variable is from the point of its definition through the end of the file. A variable having a block, or
loca,, scope is found within a block structure. The mos
```

int globalScopeInt $=5$
void $f()$
void $f()$
int localScopeInt $=10$
int main( )
int localScopeInt = 15
int anotherLocal $=20$;
int localScopeInt $=30$
return 0 ;

The first int definition, globalScopeInt, is visible within the functions fo) and main(). The next definition is found within the function f() and is named localScopeInt. This variable has local scope, meaning that it is visible only within the block defining it

Notice that this localScopentnt is hiding the localScopeInt defined just before the opening brace (the second Notice that this localScopeInt is hiding the localScopeInt defined just before the opening brace (the second defined resumes visibility. Any change made to the localScopeInt defined within the braces does not affect the contents of the outer localScopeInt.
$\begin{aligned} & \text { Note: Names can have internal and external linkage. These two terms refer to the use or availability of a name } \\ & \text { across multiple translation units or within a single translation unit. Any name having internal linkagac can only be }\end{aligned}$
across multiple translation units or within a single translation unit. Any name having internal linkage can only be
linkage can be shared by functions within the same translation unit. Names having external linkage are availation
to other translation units. The following example demonstrates internal and external linkage.
// file: first.cpp
const int $j=10$;
int main()
return 0 ;
// file: second.cp
extern int externalInt
const int $j=10$;
The externallnt variable defined in first.cpp has external linkage. Although it is defined in first.cpp, second.cpp can also access it. The two js found in both files are const and thus by default have internal linkage. You can overric the const default by providing an explicit declaration, as shown here:

```
// file: first.cpp 
// file: second.cpp
extern const int j
#include <iostream>
int main()
std::cout << "j is " << j << std::endl ;
return 0;
```

Note that this example calls cout with the namespace designation of std; this allows you to refer to all the "standard" objects in the ANSI standard library. When built, this example displays the following:

```
j is 10
```

The standards committee deprecates the following type of usage

```
static int staticInt = 10 ;
```

int main()
//...

The use of static to limit the scope of external variables is no longer recommended and may eventually become illegal. You should now use namespaces instead of static
Do Dor

DO use namespaces instead of the static keyword. DON’T apply the static keyword to a variable defined

## Creating a Namespace

The syntax for a namespace declaration is similar to the syntax for a struct or class declaration: First, apply the keyword namespace followed by an optional namespace name, and then an opening curly brace. The namespace is concluded with a closing brace but no terminating semicolon.

## Here is an example:

namespace Window
void move( int x, int y ) ;
The name Window uniquely identifies the namespace. You can have many occurrences of a named namespace. These multiple occurrences can occur within a single file or across multiple translation units. The $\mathrm{C}_{++}$standard library namespace, std, is a prime example of this feature. This makes sense because the standard library is a logical grouping of functionality.

The main concept behind namespaces is to group related items into a specified (named) area. The following is brief example of a namespace that spans multiple header files:

```
// header1.h
```

namespace Window
void move( int x , int y ) ;
// header2.h
namespace Window
void resize( int $x$, int $y$ ) ;

## Declaring and Defining Types

You can declare and define types and functions within namespaces. Of course, this is a design and maintenanc issue. Good design dictates that you should separate interface from implementation. You should follow this principle not only with classes but also with namespaces. The following example demonstrates a cluttered poorly defined namespace:

```
//. . . other declarations and variable definitions.
    void move( int x, int y) ; // declarations
    void resize( int x, int y)
    *)
    void move( int x, int y )
    if( x < MAX_SCREEN_X && x > 0)
            if( y < MAX_SCREEN_Y && y > 0) )
}
void resize( int x, int y 
    if( x < MAX_SIZE_x && x > 0)
            if(\begin{array}{l}{y < MAX_SIZE_Y && y > 0 )}\\{\mathrm{ platform.resize( x, y ) ; // specific routine}}\end{array}=\mp@code{l}
// . . . definitions continue
```

You can see how quickly the namespace becomes cluttered! The previous example is approximately 20 lines in length; imagine if this namespace were four times longer

## Defining Functions Outside a Namespact

 within h header file; the definitions can be placed into an implementation file.
## Here is an example:

## Adding New Members

New members can be addded to a namespace only within its body. You cannot define new members using
qualifier syntax. The most you can expect from this style of definition is a complaint from your compiler. Th qualifier syntax. The most you can expect $f$ f
following example demonstrates this error:

```
\amespace Window {
```

//... some code
The preceding line of code is illegal. Your (conforming) compiler will issue a diagnostic reflecting the error. 1
correct the error-or avoid italtogether-move the declaration within the namespace body.
All members within a namespace are public. The following code does not compile:
mespace Window
void move( int $x$, int $y$ );
Nesting Namespaces
A namespace can be nested within another namespace. The reason they can be nested is because the definition
of a namespace is also a declaration. As with any other namespace, you must qualify a name using the encl
namespace. If you have nested namespaces, you must qualify yach naumespace in turn . For example, the
following shows a named namespace nested within another named namespace:
following shows a named namespace nested within another named namesppace
and
namespace window $\mathfrak{t}$
mespace Pane
void size( int $x$, int $y)$ );
To access the function sizeo outside of Window, you must qualify the function with both enclosing namespac
int main()

Window::Pane::size( 10, 20)

## Using a Namespace

Now take a look at an example of using a namespace and the associated use of the scope resolution operaa
The code first declares all types and functions for use within the namespace window Atter defining everyth The code first declares all types and functions for use within the namespace Window. After defining everyhin,
required, the example then defines any member functions declared. These member functions are defined ous required, the example then defines any member functions declared. These member functions are defined oustial
of the namespace; the names are explicitly identified using the scope resolution operator. Listing 17.1 illustrate

Input Listing 17.1 Using Nand
: \#include <iostream>
: //Listing 17.1 Using a Namespace
: namespace Window

```
const int MAX_X = 30;
    class Pane
        Mane ();
        *)
            void show(),
        c
            int x
```

22: int Window::Pane:: :nt
23: Window::Pane:: Pane ():
24:





Window::Pane pane
pane. move ( 20,20
pane. show ();
return 0

## Introducing the using Keyword

The using keyword is used for both the using directive and the using declaration. The syntax of the using keyword deermines whether the context is a directive or a declaration.

## The using Directive

The using directive effectively exposes all names declared in a namespace to be in the current scope. You can refer to the names withou
shows the using directive:

```
Mamespace Window 1
    int value1 = 20;
```

Window::value1 = 10 ;
using namespace Window;
value2 $=30 ;$
The scope of the using directive begins at its declaration and continues on to the end of the current scope. I
that valuel must be qualified in order to reference it. The variable value2 does not
because the directive introduces all names in a namespace into the current scope.
The using directive can be used at any level of scope. This enables you to use the directive within block scope;
when that block goes out of scope, so do all the names within the namespace. The following sample shows thi
behavior
space Window
int value1 $=20 ;$
int value $=40 ;$
using namespace Window
value $2=30 ;$
value2 $=20$; //error!

The last line of code in f$)$, value $2=20$; is an error because value 2 is not defined. The name is accessible in the previous block because the directive pulls the name into that block. When that block goes out of scope, so dc the names in namespace Windo

Variable names declared within a local scope hide any namespace names introduced in that scope. This behavior is similar to how a local variable hides a global variable. Even if you introduce a namespace after a local variable, that local variable will hide the namespace name. The following example shows this:
nespace Window 1
int value1 $=20 ;$
int value2 $=40 ;$
void fí)
int value2 $=10$;
using namespace window

The output of this function is 10 , not 40 . This output confirms the fact that the value in namespace Window is hidden by the value in $f$ fo. If you need to use a name within a namespace, you must qualify the name with th namespace nam
An ambiguity can arise using a name that is both globally defined and defined within a namespace. The ambiguity surfaces only if the name is used, and not just when a namespace is introduced. This is demonstrate
in the following code fragment: in the following code fragment:
mespace Window 1
int value1 $=20 ;$
using namespace window
int value 1
void $f()$
value1 $=10$;
The ambiguity occurs within function $f($. The directive effectively brings Window:valuel into the globa namespace; because a valuel is already globally defined, the use of valuel in $f(0$ is an error. Note that if the lin
code in $f 0$ were removed, no error would exist. code in fo were removed, no error would exist.

## The using Declaration

The using declaration is similar to the using directive except that the declaration provides a finer level of control More specifically, the using declaration is used to declare a specific name (from a namespace) to be in the irrent scope. You can then refer to the specified object by its name only. The following example demonstra the use of the using declaration:
space Window
int value1 $=20$
int value $=40 ;$
int value $=60 ;$
using Window::value2 ; //bring value2 into current scope
Window::value1 = 10; //value1 must be qualified
Window::value $3=10$; //value3 must be qualified
Winder
The using declaration adds the specified name to the current scope. The declaration does not affect the other
names within the namespace. In the previous example, value2 is referenced without qualification, but valuel a names within the namespace. In the previous example, value2 is referenced without qualification, but valuel an value3 require qualification. The using declaration provides more control over namespace names that you bring ames in a namespace into scope.
After a name is brought into a scope, it it visible until the end of that scope. This behavior is the same as any
other declaration. A using declaration may be used in the global namespace or within any local scope. It is an error to introduce a name into a local scope where a namespace name has been declared. The reverse is also true. The following example shows this:
mespace Window 1
int value1 $=20 ;$
int
int value2 $=40$;
void $\ddagger()$
int value2 $=10 ;$
using Window
std::cout << value2 << std: multiple declaration
ste

The second line in function fo will produce a compiler error because the name value is already defined. The same error occurs if the using declaration is introduced before the definition of the local value2.
Any name introduced at local scope with a using declaration hides any name outside that scope. The following code snippet demonstrates this behavior
mespace Window
int valuel $=2$
int value $=20 ;$
int value2 $=40 ;$
t value2 = 10 ;
id f()
using Window:: value2
std:: :cout $\ll$ value2 $\ll$

## Using Namespace Alias

A namespace alias is designed to provide another name for a named namespace. An alias provides a shorthand term for you to use to refer to a namespace. This is especially true if a namespace name is very long; creating a alias can help cut down on lengthy, repetitive typing. Look at this example:
namespace the software_company
int value ;

$$
\begin{aligned}
& \text { int } \\
& \text { // . }
\end{aligned}
$$

$$
\begin{aligned}
& \text { the }
\end{aligned}
$$

the_software_company::value = 10
namespace TSC = the_software_company
TSC::value = 20
A drawback, of course, is that your alias may collide with an existing name. If this is the case, the compiler wil catch the conflict and you can resolve it by renaming the alias

## Using the Unnamed Namespace

An unnamed namespace is simply that-a namespace that does not have a name. A common use of unnamed spaces is to shield global data from potential name clashes between translation units. Every translation unit has its own unique, unnamed namespace. All names defined within the unnamed namespace (within each translation unit) can be referred to without explicit qualification. The following is an example of two unnamed namespaces found in two separate files:

```
file: one.cpp
```

namespace
int value ;
char p(char *p) ;
//.
// file: two.cpp
// file: tv
namespace
int value ;
char p( char *p ) ;
//. .
int main( )
char $c=p(p t r)$;

Each of the names, value and function $\mathrm{p}($, is distinct to its respective file. To refer to a (unnamed namespace) name within a translation unit, use the name without qualification. This usage is demonstrated in the previous example with the call to function p 0 . This use implies a using directive for objects referred to from an unnamed namespace. Because of this, you cannot access members of an unnamed namespace in another translation unit The behavior of an unnamed namespace is the same as a static object having external linkage. Consider this example:

```
static int value = 10 ;
```

Remember that this use of the static keyword is deprecated by the standards committee. Namespaces now exist to replace code as previously shown. Another way to think of unnamed namespaces is that they are global variables with internal linkage.

## Using the Standard Namespace std

The best example of namespaces is found in the C++ standard library. The standard library is completely encased within the namespace named std. All functions, classes, objects, and templates are declared within the namespace std.

You will, no doubt, see code such as the following
\#include <iostream>
using namespace std ;
Remember that the using directive pulls everything in from the named namespace. It is bad form to employ the using directive when using the standard library. Why? Because doing so defeats the purpose of using a namespace; the global namespace will be polluted with all the names found in the header. Keep in mind that al header files use the namespace feature, so if you include multiple standard header files and specify the using directive, everything declared in the headers will be in the global namespace. Please note that most of the examples in this book violate this rule; this action is not an intent to advocate violating the rule, but it is used to ensure brevity of the examples. You should use the using declaration instead, as in the following example:
\#include <iostream>
using std::cin ;
using std::cout
using std::endl ;
int main( )
int value $=0$;
cout << "So, how many eggs did you say you wanted?" << endl ;
cin >> value
cout << value << " eggs, sunny-side up!" << endl ;
return ( 0 ) ;
The following shows a sample run of the program
So, how many eggs did you say you wanted?
4
4 eggs, sunny-side up
As an alternative, you could fully qualify the names that you use, as in the following code sample:
\#include <iostream>
int main( )
int value $=0$;
std::cout << "How many eggs did you want?" << std::endl ;
std::cin >> value ;
std::cout << value << " eggs, sunny-side up!" << std::endl ;
return ( 0 ) ;
Sample output from this program is shown here
How many eggs did you want?


This might be appropriate for shorter programs but can become quite cumbersome for any significant amount code. Imagine having to prefix std:: for every name you use that is found in the standard library!

## Summary

 semicolon does not follow a namespace's closing brace. Second, a namespace is open, whereas a class is closed This means that you can continue to define the namespace in other files or in separate sections of a single file.Anything that can be declared can be inserted into a namespace. If you are designing classes for a reusable library, you should be using the namespace feature. Functions declared within a namespace should be defined outside that namespace's body. This promotes a separation of interface from implementation and also keeps the namespace from becoming cluttered.

Namespaces can be nested. A namespace is a declaration; this fact enables you to nest namespaces. Do not forget that you must fully qualify names that are nested.

The using directive is used to expose all names in a namespace to the current scope. This effectively pollutes the global namespace with all names found in the named namespace. It is generally bad practice to use the using directive, especially with respect to the standard library. Use using declarations instead.

The using declaration is used to expose a specific namespace name into the current scope. This allows you to refer to the object by its name only.

A namespace alias is similar in nature to a typedef. A namespace alias enables you to create another name for a named namespace. This can be quite useful when you are using a namespace with a long name.

Every file can contain an unnamed namespace. An unnamed namespace, as its name implies, is a namespace without a name. An unnamed namespace allows you to use the names within the namespace without qualification. It keeps the namespace names local to the translation unit. Unnamed namespaces are the same as declaring a global variable with the static keyword.

The C++ standard library is enclosed in a namespace named std. Avoid using the using directive when using the standard library; instead, use the using declaration.

## Q\&A

## Q Do I have to use namespaces?

A No, you can write simple programs and ignore namespaces altogether. Be sure to use the old standard libraries (for example, \#include <string.h>) rather than the new libraries (for example, \#include <cstring>.
Q What are the two forms of statements with the using keyword? What are the differences between

## those two forms?

A The using keyword can be used for the using directives and the using declarations. A using directive allows all names in a namespace to be used as if they are normal names. A using declaration, on the other hand, enables the program to use an individual name from a namespace without qualifying it with the namespace qualifier.

## Q What are the unnamed namespaces? Why do we need unnamed namespaces?

A Unnamed namespaces are namespaces without names. They are used to wrap a collection of declarations against possible name clashes. Names in an unnamed namespace cannot be used outside of the translation unit where the namespace is declared.

## Workshop

The Workshop contains quiz questions to help solidify your understanding of the material covered and exercises to provide you with experience in using what you learned. Try to answer the quiz and exercise questions before checking the answers in Appendix D, "Answers to Exercises and Quizzes," and make sure you understand the answers before going on to the next day.

## Quiz

1. Can I use names defined in a namespace without using the using keyword?
2. What are the major differences between normal and unnamed namespaces?
3. What is the standard namespace?

## Exercises

1. BUG BUSTER: What is wrong in this program?
```
#include <iostream>
int main()
{
    cout << "Hello world!" << end;
    return 0;
}
```

2. List three ways of fixing the problem found in Exercise 1.

## Chapter 18

## Object-Oriented Analysis and Design

 It is easy to become focused on the sybuild programs. Today you will learn

- How to use object-oriented analysis to understand the problem you are trying to solv
- How to use object-oriented design to create a robust, extensible, and reliable solution
- How to use the Unified Modeling Language (UML) to document your analysis and design


## Is $\mathrm{C}_{++}$Object-Oriented?

## C++ was created as a bridge between object-oriented programming and C , the world's most popular

 programming language for commercial software development. The goal was to provide object-oriented desigr to a fast, commercial software development platform.C was developed as a middle ground between high-level business application languages such as COBOL and C was developed as a midale ground between high-level business application languages such as Co "structured" programming, in which problems were "decomposed" into smaller units of repeatable activities called procedures.

The programs we are writing at the beginning of the new century are much more complex than those written the beginning of the last decade. Programs created in procedural languages tend to be difficult to manage, har to maintain, and impossible to extend. Graphical user interfaces, the Internet, digital telephony, and a host of new technologies have dramatically increased the complexity of programming projects at the same time that consumer expectations for the quality of the user interface are rising

In the face of this increasing complexity, developers took a long, hard look at the state of the industry. Wha they found was disheartening, at best. Software was late, broken, defective, bug-ridden, unreliable, and expensive. Projects routinely ran over budget and were delivered late to market. The cost of maintaining and building on these projects was prohibitive, and a tremendous amount of money was being wasted

Object-oriented software development offers a path out of the abyss. Object-oriented programming languages build a strong link between the data structures and the methods that manipulate that data. More important, in build a strong link between the data structures and the methods that manipulate that data. More important, in object-oriented programming,
instead about objects-things.

The world is populated by things: cars, dogs, trees, clouds, flowers. Each thing has characteristics (fast, friend brown, puffy, pretty). Most things have behavior (move, bark, grow, rain, wilt). We do not think about a dog' data and how we might manipulate it-we think about a dog as a thing in the world, what it is like, and what i does.

What Models Are
If we are to manage complexity, we must create a model of the universe. The goal of the model is to create a meaningful abstraction of the real world. Such an abstraction should be simpler than the real world but should also accurately reflect the real world so that we can use the model to predict the behavior of things in the real world.
A child's globe is a classic model. The model is not the thing itself; we would never confuse a child's glob the Earth but one maps the other well enough that we can learn about the Earth by studying the globe

There are, of course, significant simplifications. A child's globe never has rain, floods, globe-quakes and so forth, but I can use that globe to predict how long it will take me to fly from my home to Indianapolis should need to come in explain myself to the Sams senior masement when they ask me why my manuscin was late ("you see, I was doing great, but then I got lost in a metaphor and it took me hours to get out"),

A model that is not simpler than the thing being modeled is not much use. Steven Wright jokes about just such thing: "I have a map on which one inch equals one inch. I live at E5."

Object-oriented software design is about building good models. It consists of two significant pieces: a modelin language and a process.

Software Design: The Modeling Language
The modeling language is the least important aspect of object-oriented analysis and design; unfortunately, i tends to get the most attention. A modeling language is nothing more than a convention for how to draw a model on paper. We can easily decide to draw our classes as triangles and the inheritance relationship as a dotted line. If so, we might model a geranium as shown in Figure 18.1


## Figure 18.1 Generalization/specialization.

the figure, you see that a Geranium is a special kind of flower. If you and I agree to draw our inheritance (generalization/specialization) diagrams like this, we will understand each other perfectly. Over time, we n probably want to model lots of complex relationships to develop our own complicated set of diagramming
conventions and rules. conventions and rules

Of course, we need to explain our conventions to everyone else with whom we work, and each new employe collaborator will have to learn our conventions. We may interact with other companies that have their own conventions, and we need to allow time to negotiate a common convention and to compensate for the inevitab misunderstandings.

It would be more convenient if everyone in the industry agreed on a common modeling language. (For that matter, it would be convenient if everyone in the world agreed on a spoken language, but one thing at a time. The lingua franca of software development is UML-the Unified Modeling Language. The job of the UML i answer questions such as, "How do we draw an inheritance relationship?" The geranium drawing shown in Figure 18.1 would be drawn in UML as shown in Figure 18.2.


Figure 18.2 UML drawing of specialization
UML, classes are drawn as rectangles, and inheritance is drawn as a line with an arrowhead. Interestingly, arrowhead points from the more specialized class to the more general class. The direction of the arrow is解位-intuitive for most folks, but it does not matter much; when we all agree, the system works just fine.

The details of the UML are rather straightforward. The diagrams are not hard to use or to understand, and the are explained as we go along in this lesson, rather than trying to teach the UML out of context. Although it is possible to write a whole book on the UML, the truth is that 90 percent of the time, you use only a small subse of the UML notation, and that subset is easily learned

## Software Design: The Process

The process of object-oriented analysis and design is much more complex and important than the modelin language. So of course, it is what you hear much less about. That is because the debate about modeling

A methodologist is someone who develops or studies one or more methods. Typically, methodologists develo and publish their own methods. A method is a modeling language and a process. Three of the leading and publish their own methods. A method is a modeling language and a process. Three of the leading
methodologists and their methods are Grady Booch, who developed the Booch method; Ivar Jacobson, who developed object-oriented software engineering; and James Rumbaugh, who developed Object Modeling Technology (OMT). Together, these three men have created Objectory, a method and a commercial product from Rational Software, Inc. All three men are employed at Rational Software, where they are affectionately known as the Three Amigos.

This lesson loosely follows Objectory. I will not follow it rigidly because I do not believe in slavish adherenc to academic theory-I am much more interested in shipping product than in adhering to a method. Other methods have something to offer, and I tend to be eclectic; picking up bits and pieces as I go along and stitchi
them together into a workable framework.

The process of software design is iterative. That means that as we develop software, we go through the entire process repeatedly as we strive for enhanced understanding of the requirements. The design directs the implementation, but the details uncovered during implementation feed back into the design. Most important, w do not try to develop any sizable project in a single, orderly, straight line; rather, we iterate over pieces of the project, constantly improving our design and refining our implementation.

Iterative development can be distinguished from waterfall development. In waterfall development, the output from one stage becomes the input to the next, and there is no going back (see Figure 18.3). In a waterfall from one stage becomes the input to the next, and there is no going back (see Figure 18.3). In a waterfall
development process, the requirements are detailed, and the clients sign off ("Yes, this is what I want"); the development process, the requirements are detailed, and the chents sign off ("Yes, this is what I want"); the
requirements are then passed on to the designer, set in stone. The designer creates the design (and a wonder to requirements are then passed on to the designer, set in stone. The designer creates the design (and arn hands the
behold it is) and passes it off to the programmer who implements the design. The programmer in turn her conold to a QA person who tests the code and then releases it to the customer. Great in theory, disaster in practice
coder


Figure 18.3 The waterfall method.
In iterative design, the visionary comes up with a concept, and then we begin to work on fleshing out the requirements. As we examine the details, the vision may grow and evolve. When we have a good start on the requirements, we begin the design, knowing full well that the questions that arise during design may cause modifications back in the requirements. As we work on design, we begin prototyping and then implementing the product. The issues that arise in development feed back into design and may even influence our understanding of the requirements. Most important, we design and implement only pieces of the full product iterating over the design and implementation phases repeatedly.

Although the steps of the process are repeated iteratively, it is nearly impossible to describe them in such a cyclical manner. Therefore, I will describe them in sequence: conceptualization, analysis, design, implementation, testing, and rollout. Do not misunderstand me-in reality, we run through each of these steps many times during the course of the development of a single product. The iterative design process is just hard tt present and understand if we cycle through each step, so I describe them one after the other.

The following are the steps of the iterative design process:

1. Conceptualization-when the vision and overall purpose of the project is developed.
2. Analysis-determining the needs of the organization (figuring out what the software should do), this is where you model your classes.
3. Design-creating the blueprint for a solution to the problem set.
. Implementation-building the system from the design blueprint, this is where the code is developed (in C++ for example).
4. Testing-making sure the system does what it is supposed to do.
5. Rollout-getting the system out for the users.

These steps are a piece of cake. All the rest is details

Note: This lesson only covers the first three steps: conceptualization, analysis, and design. The other steps are
well beyond the scope of this introductory lesson. You can find plenty of books that will cover all of these topic well beyond the scope of this introductory lesson. You can find plenty of books that will cover all of these topics
in excruciating detail.

## Controversies

Endless controversies exist about what happens in each stage of the iterative design process, and even abou what you name those stages. Here is the secret: it does not matter. The essential steps are the same in just about every process: Find out what you need to build, design a solution, and implement that design.

Although the newsgroups and object-technology mailing lists thrive on splitting hairs, the essentials of object oriented analysis and design are fairly straightforward. In this lesson, I lay out a practical approach to the process as the bedrock on which you can build the architecture of your application.

The goal of all this work is to produce code that meets the stated requirements and that is reliable, extensible and maintainable. Most important, the goal is to produce high-quality code on time and on budget.

## Conceptualization: The Vision

All great software starts with a vision. One individual has an insight into a product he thinks would be good to build. Rarely do committees create compelling visions. The very first phase of object-oriented analysis and design is to capture this vision in a single sentence (or at most, a short paragraph). The vision becomes the guiding principle of development, and the team that comes together to implement the vision ought to refer back to it-and update it if necessary-as it goes forward.

Even if the vision statement comes out of a committee in the marketing department, one person should be designated as the "visionary." It is that person's job to be the keeper of the sacred light. As you progress, the requirements will evolve. Schedulingand time-to-market demands may modify what you try to accomplish in the first iteration of the program, but the visionary must keep an eye on the essential idea, to ensure that whatever is produced reflects the core vision with high fidelity. It is this ruthless dedication-this passionate commitment-that sees the project through to completion. If you lose sight of the vision, your product is doomed.

## Requirements Analysis

The conceptualization phase, in which the vision is articulated, is very brief. It may be no longer than a flash of insight followed by the time it takes to write down what the visionary has in mind. Often, as the object-oriented expert, you join the project after the vision is already articulated.

Some companies confuse the vision statement with the requirements. A strong vision is necessary, but it is not sufficient. To move on to analysis, you must understand how the product will be used and how it must perform. The goal of the analysis phase is to articulate and capture these requirements. The outcome of the analysis pha is the production of a requirements document. The first section in the requirements document is the use-case analysis.

## Use Cases

The driving force in analysis, design, and implementation is the use cases. A use case is nothing more than a high-level description of how the product will be used. Use cases don't just drive the analysis-they drive the design, they help you find the classes, and they are especially important in testing the product.
Creating a robust and comprehensive set of use cases may be the single most important task in analysis. It is here that you depend most heavily on your domain experts; the domain experts have the most information abol the business requirements you are trying to capture.

Use cases pay little attention to user interface, and they pay no attention to the internals of the system you are building. Any system or person who interacts with the system is called an actor

To summarize, the following are some definitions:

- Use case-A description of how the software will be used
- Domain experts-People with expertise in the domain (area) of business for which you are creating the product.
- Actor-Any person or system that interacts with the system you are developing. «uses» stereotype indicates that one use case is a superset of another. For example, it is not possible to withdra «uses» stereotype indicates hat one use case is a superset of another. For example, it is not possible to
cash without first logging on. We can show this relationship with the diagram shown in Figure 18.10.
$\square$
Figure 18.9 Use-case diagram.


Figure 18.10 The «uses» stereotype
Figure 18.10 indicates that the Withdraw Cash use case "uses" the Log In use case, and thus fully implements Log In as part of Withdraw Cash

that many developers have simply set aside eextends», feeling that its meaning is not sufficiently well-understood.
Personally, Iuse essess> when I would otherwise copy and paste the entire use case in place, and I use eexends»
when I only use the use case under certain definable conditions.

## Interaction Diagram

Although the diagram of the use case itself may be of limited value, you can associate diagrams with the use case that can dramatically improve the documentation and understanding of the interactions. For example, we customer, checking account, and the user interface. We can document this interaction with an interactio customer, checking account, and
diagram, as shown in Figure 18.11.


Figure 18.11 UML interaction diagram.
The interaction diagram in Figure 18.11 captures details of the scenario that may not be evident by reading th text. The objects that are interacting are domain objects, and the entire ATM/UI is treated as a single object, with only the specific bank account called out in any detail

This rather simple ATM example shows only a fanciful set of interactions, but nailing down the specifics of these interactions can be a powerful tool in understanding both the problem domain and the requirements o these interactions
your new system.

## Create Packages

Because you generate many use cases for any problem of significant complexity, the UML enables you to grou your use cases in packages.

A package is like a directory or a folder-it is a collection of modeling objects (classes, actors, and so forth). manage the complexity of use cases, you can create packages aggregated by whatever characteristics make sense for your problem. Thus, you can aggregate your use cases by account type (everything affecting checkii or savings), by credit or debit, by customer type, or by whatever characteristics make sense to you. More
importantly, a single use case can importantly, a single use case can appear in different packages, allowing you great flexibility of design

## Application Analysis

In addition to creating use cases, the requirements document will capture your customer's assumptions, In addition to creating use cases, the requirements document will capture your castomer's assumptions, constraints, and requirements about hardware and operating systems. Application requirements are your
particular customer's prerequisites-those things that you would normally determine during design and implementation but that your client has decided for you.

The application requirements are often driven by the need to interface with existing (legacy) systems. In this case, understanding what the existing systems do and how they work is an essential component of your analy
Ideally, you analyze the problem, design the solution, and then decide which platform and operating system b fit your design. That scenario is as ideal as it is rare. More often, the client has a standing investment in a on the existing system, and you must capture these requirements early and design accordingly.

## Systems Analysis

Some software is written to stand alone, interacting only with the end user. Often, however, you will be called on to interface to an existing system. Systems analysis is the process of collecting all the details of the systems with which you will interact. Will your new system be a server, providing services to the existing system, or will it be a client? Will you be able to negotiate an interface between the systems, or must you adapt to an existing standard? Will the other system be stable, or must you continually hit a moving target?

These and related questions must be answered in the analysis phase, before you begin to design your new system. In addition, you will want to try to capture the constraints and limitations implicit in interacting with 1 system. In addition, you will want to try to capture the constraints and limitations impliciti in interacting with
other systems. Will they slow down the responsiveness of your system? Will they put high demands on your other systems. Will they slow down the responsiveness
new system, consuming resources and computing time?

## Planning Documents

After you understand what your system must do and how it must behave, it is time to take a first stab at creatin a time and budget document. Often, the timeline is dictated, top-down, by the client: "You have 18 months to get this done." Ideally, you examine the requirements and estimate the time it will take to design and impleme
the solution. That is the ideal. The practical reality is that most systems come with an imposed time limit and the solution. That is the ideal. The practical reality is that most systems come with an imposed time limit and
cost limit, and the real trick is to figure out how much of the required functionality you can build in the allote time-and at the alloted cost
tost

Here are two guidelines to keep in mind when you are creating a project budget and timeline:

- If you are given a range, the outer number is probably optimistic.
- Hofstadter's Law: It always takes longer than you expect, even when you take into account

Hofstadter's Law.
Given these realities, it is imperative that you prioritize your work. You will not finish-it is that simple. It is important that when you run out of time, what you have works and is adequate for a first release. If you are
building a bridge and run out of time, if you did not get a chance to put in the bicycle path, that is too badd bu you can still open the bridge and start collecting tolls. If you run out of time and you are only halfway across th river, that is not as good.

An essential thing to know about planning documents is that they are wrong. This early in the process, it is virtually impossible to offer a reliable estimate of the duration of the project. After you have the requirements you can get a good handle on how long the design will take, a fair estimate of how long the implementation wi take, and a reasonable guesstimate of the testing time. Then you must allow yourself at least 20 to 25 percent
"wiggle room" which you can tighten "wiggle room," which you can tighten as you move forward and learn more

## Note: The inclusion of "wiggle room" in your planning document is not an excuse to avoid planning documents It is merely a warning not to rely on them to ouwh early on. As the project goes forward, you strengthen your understanding of how the system works, and your estimates will become increasingly precise.

## Visualizations

The final piece of the requirements document is the visualization. The visualization is a fancy name for the diagrams, pictures, screen shots, prototypes, and any other visual representations created to help you think through and design the graphical user interface of your product.

At the end of each phase of analysis and design, you will create a series of documents or "artifacts." Table 18. shows some of the artifacts of the analysis phase. These documents are used by the customer to make sure that you understand what the customer needs, by end users to give feedback and guidance to the project, and by the project team to design and implement the code. Many of these documents also provide material crucial both to your documentation team and to Quality Assurance to tell them how the system ought to behave,

Table 18.1Artifacts Created During the Analysis Stage of Project Development

| Artifact | Description |
| :--- | :--- |
| Use case report | A document detailing the use cases, scenarios, stereotypes, <br> preconditions, postconditions, and visualizations |
| Domain analysis | Document and diagrams describing the relationships among the <br> domain objects |
| Analysis collaboration diagrams | Collaboration diagrams describing interactions among objects in <br> the problem domain |
| Analysis activity diagrams | Activity diagrams describing interactions among objects in the <br> problem domain <br> Report and diagrams describing low-level and hardware systems <br> on which the project will be built |
| Systems analysis | Report and diagrams describing the customer's requirements <br> specific to this particular project |
| Application analysis document | Report describing performance characteristics and constraints <br> imposed by this client |
| Operational constraints report | Report with Gantt and Pert charts indicating projected scheduling, <br> milestones, and costs |
| Cost and planning document |  |

## Design

Analysis focuses on understanding the problem domain, whereas design focuses on creating the solutio Design is the process of transforming your understanding of the requirements into a model that can be implemented in software. The result of this process is the production of a design document.

The design document is divided into two sections: Class Design and Architectural Mechanisms. The Class Design section, in turn, is divided into static design (which details the various classes and their relationships characteristics) and dynamic design (which details how the classes interact)

The Architectural Mechanisms section of the design document provides details about how you will implemen object persistence, concurrency, a distributed object system, and so forth. The rest of this lesson focuses on the object persistence, concurrency, a distributed object system, and so forth. The rest of this lesson focuses on th class design aspect of the design
various architecture mechanisms.

## What Are the Classes?

As a C++ programmer, you are used to creating classes. Formal design methodology requires you to separate th C++ class from the design class, although they will be intimately related. The C++ class you write in code is the implementation of the class you designed. These are isomorphic: Each class in your design will correspond to a your code, but do not confuse one for the other. It is certainly possible to implement your design classe in another language, and the syntax of the class definitions might be changed.

That said, most of the time we talk about these classes without distinguishing them because the differences ar highly abstract. When you say that in your model your Cat class will have a Meow() method, understand that thi: means that you will put a Meow() method into your $\mathrm{C}++$ class as well.

You capture the model's classes in UML diagrams, and you capture the C++ classes in code that can be compiled. The distinction is meaningful, yet subtle.

In any case, the biggest stumbling block for many novices is finding the initial set of classes and understanding what makes a well-designed class. One simplistic technique suggests writing out the use-case scenarios and the creating a class for every noun. Consider the following use-case scenario:

Customer chooses to withdraw cash from checking. Sufficient cash is in the account, sufficient cash and receipts are in the ATM, and the network is up and running. The ATM asks the customer to
indicate an amount for the withdrawal, and the customer asks for $\$ 300$, a legal amount to withdraw at
this time. The machine dispenses $\$ 300$ and prints a receipt, and the customer takes the money and the his time. The machine dispenses $\$ 300$ and prints a receipt, and the customer takes the money and the receipt.

You might pull out of this scenario the following classes:

- Custome
- Cash
- Checkin
- Receipts
- ATM
- Network
- Amount
- Withdrawal
- Machin
- Money

You might then aggregate the synonyms to create this list, and then create classes for each of these nouns:

- Castomer
- Checking
- Account
- ATM (machine)
- Network

This is not a bad way to start, as far as it goes. You might then go on to diagram the obvious relationships among some of these classes as shown in Figure 18.12

## Transformations

What you began to do in the preceding section was not so much extracting the nouns from the scenario as to begin transforming objects from the domain analysis into objects in the design. That is a fine first step. Often distin the objects in the domain will have surrogates in the design. An object is called a surrogate to anstinguish between the actual physical receipt dispensed by an ATM and the object in your design that is merely an intellectual abstraction implemented in code.

You will likely find that most of the domain objects have an isomorphic representation in the design-that is, one-to-one correspondence exists between the domain object and the design object. Other times, however, a single domain object is represented in the design by an entire series of design objects. And at times, a series o domain objects may be represented by a single design object.

In Figure 18.12, note that we have already captured the fact that CheckingAccount is a specialization of Accour We did not set out to find the generalization relationship, but this one was self-evident, so we captured it Similarly, we knew, from the domain analysis, that the ATM dispenses both Cash and Receipts, so we captured that information immediately in the design

The relationship between Customer and CheckingAccount is less obvious. We know that such a relationship exists but the details are not obvious, so we hold off

These interface classes offer encapsulation of the interface protocol and thus shield your code from changes ir the other system. Interface classes allow you to change your own design, or to accommodate changes in the design of other systems, without breaking the rest of the code. As long as the two systems continue to support the agreed-on interface, they can move independently of one another

## Data Manipulation

Similarly, you will create classes for data manipulation. If you have to transform data from one format into another format (for example, from Fahrenheit to Celsius or from inch/pound to metric), you may want to encapsulate these manipulations behind a data manipulation class. You can use this technique when messaging data into required formats for other systems or for transmission over the Internet-in short, any time you mus manipulate data into a specified format, you will encapsulate the protocol behind a data manipulation class.

## Views

Every "view" or "report" your system generates (or, if you generate many reports, every set of reports) is a candidate for a class. The rules behind the report-both how the information is gathered and how it is to be displayed-can be productively encapsulated inside a view class.

## Device

If your system interacts with or manipulates devices (such as printers, modems, scanners, and so forth), the specifics of the device protocol ought to be encapsulated in a class. Again, by creating classes for the interfac to the device, you can plug in new devices with new protocols and not break any of the rest of your code; just create a new interface class that supports the same interface (or a derived interface), and off you go

## Static Model

When you have established your preliminary set of classes, it is time to begin modeling their relationships and interactions. For purposes of clarity, this lesson first explains the static model and then explains the dynamic interactions. For purposes of clarity, this lesson first explains the static model and then explains the dynami
model. In the actual design process, you will move freely between the static and dynamic models, filling in model. In the actual design process, you will move freely between the static and
details of both and, in fact, adding new classes and sketching them in as you go.

The static model focuses on three areas of concern: responsibilities, attributes, and relationships. The most The static model focuses on three areas of concern: responsibilities, attributes, and relationships. The most
important of these-and the one you focus on first-is the set of responsibilities for each class. The most important of these-and the one you focus on first-is the set of responsibilities for each class. The most important guiding principle is this: Each class should be responsible for one thing.

That is not to say that each class has only one method. Far from it; many classes will have dozens of methods. But all these methods must be coherent and cohesive; that is, they must all relate to one another and contribute to the class's capability to accomplish a single area of responsibility.

In a well-designed system, each object is an instance of a well-defined and well-understood class that is responsible for one area of concern. Classes typically delegate extraneous responsibilities to other, relatec classes. By creating classes that have only a single area of concern, you promote the creation of highly maintainable code.

To get a handle on the responsibilities of your classes, you may find it beneficial to begin your design work wit the use of CRC cards.

## CRC Cards

CRC stands for Class, Responsibility, and Collaboration. A CRC card is nothing more than a $4 \times 6$ index card. This simple, low-tech device enables you to work with other people in understanding the primary responsibilities of your initial set of classes. You assemble a stack of blank $4 \times 6$ index cards and meet around conference table for a series of CRC card sessions.

## How to Conduct a CRC Session

Each CRC session should be attended, ideally, by a group of three to six people; any more becomes unwieldy You should have a facilitator, whose job it is to keep the session on track and to help the participants capture what they learn. At least one senior software architect should be present, ideally someone with significant experience in object-oriented analysis and design. In addition, you will want to include at least one or two "domain experts" who understand the system requirements and who can provide expert advice in how things ought to work.

The most essential ingredient in a CRC session is the conspicuous absence of managers. This is a creative, free wheeling session that must be unencumbered by the need to impress one's boss. The goal here is to explore, to take risks, to tease out the responsibilities of the classes, and to understand how they might interact with one take risks
another.

You begin the CRC session by assembling your group around a conference table, with a small stack of $4 \times 6$ index cards. At the top of each CRC card you will write the name of a single class. Draw a line down the cente of the card and write Responsibilities on the left and Collaborations on the right.

Begin by filling out cards for the most important classes you have identified. For each card, write a onesentence or two-sentence definition on the back. You may also capture what other class this class specializes i that is obvious at the time you are working with the CRC card. Just write Superclass: below the class name an
fill in the name of the class this class derives from. fill in the name of the class this class derives from.

## Focus on Responsibilities

The point of the CRC session is to identify the responsibilities of each class. Pay little attention to the attribute capturing only the most essential and obvious attributes as you go. The important work is to identify the responsibilities. If, in fulfilling a responsibility, the class must delegate work to another class, you capture that information under collaborations

As you progress, keep an eye on your list of responsibilities. If you run out of room on your $4 \times 6$ card, it may make sense to wonder whether you are asking this class to do too much. Remember, each class should be responsible for one general area of work, and the various responsibilities listed should be cohesive and coherent-that is, they should work together to accomplish the overall responsibility of the class.

At this point, you do not want to focus on relationships, nor do you want to worry about the class interface or which methods will be public and which will be private. The focus is only on understanding what each class does.

## Anthropomorphic and Use-Case Driven

The key feature of CRC cards is to make them anthropomorphic-that is, you attribute humanlike qualities to each class. It works this way: After you have a preliminary set of classes, return to your CRC scenarios. Divide the cards around the table arbitrarily, and walk through the scenario together. For example, return to the following banking scenario

Customer chooses to withdraw cash from checking. Sufficient cash is in the account, sufficient cash and receipts are in the ATM, and the network is up and running. The ATM asks the customer to indicate an amount for the withdrawal, and the customer asks for $\$ 300$, a legal amount to withdraw at this time. The machine dispenses $\$ 300$ and prints a receipt, and the customer takes the money and the receipt.
How are these modeles differentiated？As you saw，they are differentiated by the engine size，body type，and performance characterisitics．These various discriminating characterisitics can be mixed and matched to create
various models．We can model this in the UML with the discriminator stereotype，as shown in Figure 18.18 ．苜菌国国 Figure 18.17 Modeling subtypes
Figure 18.18 Modeling the discriminator
The diagram in Figure 18.18 indicates that classes can be derived from Car based on mixing and matching
discriminating attributes．The size of the engine dictates how powerful the car is，and the performance discriminating attributes．The sizz of the engine dictates how powerful the car is，and the performance
characteristics indicate how sporty the car is．Thus，you can have a powerful and sporty station wagon，al power family sedan，and so forth．
BodyType $=\{$ sedan，coupe，minivan，stationwason $)$ It may turn out，however，that a simple value is insufficient to model a particular discriminator．For example，
the eertormance charactersict may be rather complex．Int this case the discriminnator can be modeled as a a clas．
and the discrimination can be encapsulated in an instance of that type．
about where the ensinine shifts and how fast it can turn．The UML stereetype for a class that encapsulates a seciminator，and that can be used to create instances of a c class（Car）that are logicially of different types（for
example，SportsCav versus LuxuryCCar）is cpowertyeer．In this case，the ereformance class is pertye for car．WH you instantiate car，you also instantiate a Performanance object，and you associate a given performanace object with ：
： Figur 18.19
覧
Powertypes enable you to create a variety of logical types without using inheritance．You can thus manage a
large and complex set of types without the combinatorial explosion you might encounter with inheritance． Typically，you implement the powertype in $\mathrm{C}+$＋with pointers．In this case，the Car class holds a pointer to an
instance of the Performancecharacerersisis class see Figure 18.20 ．I Ieave it as an exercise to the ambitious reade instance of the Performancechanacerisisis class see figure 18.20 ．2
to convert the body and engine discriminators into powertypes
品变
品变
Figure 18.20 The relationship between a Car object and its powertype．

\section*{| $\quad \begin{array}{l}\text { Car（1）；} \\ \sim \operatorname{Car}()\end{array}$ |
| :--- |}

other public methods elided
As a final note，powertypes enable you to create new types（not just instances）at runtime．Because each logica type is siffererentiated only by the eatributes of the associated powertype，these attributes can be parameters to t
powertyees s oonstuctor．This means shat you can，at runtime，create new types of cars on－the－fly．That is，by passing different engine sizes and shift points to the powertye，you can effectively create new performance
characteristics．By assigging those characteristics to various cars，you can effectively enlarge the set of types Dynamic Model
In addition to modeling the relationships among the classes，it is critical to model how they interact．For
example，going back to the banking scenario，the CheckingAccount，ATM，and Receipt clasess may interact
 18.21
 that the e ATM Class will delegate to the checking Account class al respons
the CheckingAccount will call on the ATM to manage display to the user．
Interaction diagrams come in two flavors．The one in Figure 18.21 is called a sequence ciagram．Another vie
on the same information is provided by the collaboration diagram．The sequence diagram emphasizes the
 generate a collaboration diagram directly from a sequence dia
automate this task at the click of a buton（see Figure 18.22 ）．
$=$
Figure 18.21 Sequence diagrar

## Figure 18.22 Collaboration diagran

As we come to understand the interactions among the objects，we have to understand the various possible stat
of each individual obiect．We can model the transitions amon the various states in a statat diagram（or state ot each individual object．We can model the transitions among the various states in a state diagram（or state
transition diagram．）Figure 18.23 shows the various states of the Customereccount class as the customer logs in
the transition
the systen


Every state diagram begins with a single start state and ends with zero or more end states．The individual state Ever samed，and the e tansititons may be labeled．The guard indicates a condition that must be satisfied for an
obiect to move from one state 0 anolher． Super States
The customer can change his mind at any time and decide not to log in．He can do this after he swipes his carc
to idenify his account or after he enters his password．In eithe case，the system must accept his request to to idenity his account or after he enters his password．In either
cancel and return to the＂not logged in state＂（see Figure 18.24 ） particularly annoying because canceling is an exceptional condition that should not be giv
diagram．You can simplify this diagram by using a super state，as shown in

## Do Not Rush to Code

One of the biggest problems in software development organizations is their rush to begin coding．There are
always time pressures，and the codingtesting portions of the project tend to be the largest part．And usually a always time pressures，and the coding ftestii
number of people are assigned to the task．

Of course, this is not an excuse to indulge in "analysis paralysis" where the analysis goes on and on forever with no results and no progress. There is a thin line between analysis being complete and going into paralysis Unfortunately, experience is the only way to determine the difference.

There is a concept of "good enough." It is extremely difficult to make anything truly "complete." There is an axiom in this industry that $80 \%$ of the results take $20 \%$ of the time and the remaining $20 \%$ takes $80 \%$ of the time. At some point you make the decision to move on to the next step or you will begin to suffer from what is known as Analysis Paralysis-never getting to the final product because you are trying to get the perfect model

Unfortunately, the best time to determine if the analysis and design is truly "good enough" is once the project is done. If the code delivered matches the organizational requirements, then it was good enough. Missing pieces o features missing, tend to indicate that not enough effort or time was expended. If there are features that no one wants or can figure out a purpose for, then that is a good indication that too much time was spent on the design.

Only retrospect and learning from the results will help you learn when the effort for a design is good enough.
How good or complete does your analysis have to be? How good is good enough? Only you and your organization can decide this. But think about it! Too many organizations do not think about the process and continually fail to deliver software on-time, on-budget, or with the needed features.

## Summary

This lesson provided an introduction to the issues involved in object-oriented analysis and design. The essence of this approach is to analyze how your system will be used (use cases) and how it must perform, and then to design the classes and model their relationships and interactions.

In the old days, we sketched out a quick idea of what we wanted to accomplish and began writing code. The problem is that complex projects are never finished; and if they are, they are unreliable and brittle. By investing up front in understanding the requirements and modeling the design, we ensure a finished product that is correct (that is, it meets the design) and that is robust, reliable, and extensible.

Issues relating to testing and rollout are beyond the scope of this book, except to mention that you want to plan your unit testing as you implement, and that you will use your requirements document as the foundation of your test plan prior to rollout.

## Q\&A

## Q In what way is object-oriented analysis and design fundamentally different from other

 approaches?A Prior to the development of these object-oriented techniques, analysts and programmers tended to think of programs as groups of functions that acted on data. Object-oriented programming focuses on the integrated data and functionality as discrete units that have both knowledge (data) and capabilities (functions). Procedural programs, on the other hand, focus on functions and how they act on data. It has been said that Pascal and C programs are collections of procedures, and C++ programs are collections of classes.
Q Is object-oriented programming finally the silver bullet that will solve all programming problems?
A No, it was never intended to be. For large, complex problems, however, object-oriented analysis, design, and programming can provide the programmer with tools to manage enormous complexity in ways that were previously impossible.

## Q Is C++ the perfect object-oriented language?

A C++ has a number of advantages and disadvantages when compared with alternative object-oriented programming languages, but it has one killer advantage above and beyond all others: It is the single mos popular object-oriented programming language on the face of the Earth. Frankly, most programmers do not decide to program in $\mathrm{C}++$ after an exhaustive analysis of the alternative object-oriented programming languages; they go where the action is, and currently the action is with $\mathrm{C}++$. There are good reasons for that; $\mathrm{C}++$ has a lot to offer, but this book exists because $\mathrm{C}++$ is the development language of choice at so many corporations and Linux is such a growth area.

## Workshop

The Workshop provides quiz questions to help you solidify your understanding of the material covered and exercises to provide you with experience in using what you have learned. Try to answer the quiz and exercise questions before checking the answers in Appendix D, "Answers to Exercises and Quizzes," and make sure you understand the answers before continuing tomorrow.

## Quiz

1. What is the difference between object-oriented programming and procedural programming?
2. What are the phases of object-oriented analysis and design?
3. How are sequence diagrams and collaboration diagrams related?

## Exercises

1. Suppose you had to simulate the intersection of Massachusetts Avenue and Vassar Street-two
typical two-lane roads, with traffic lights and crosswalks. The purpose of the simulation is to determine whether the timing of the traffic signal allows for a smooth flow of traffic.

What kinds of objects should be modeled in the simulation? What would the classes be for the simulation?
2. Suppose the intersection from Exercise 1 was in a suburb of Boston, which has arguably the most unfriendly streets in the United States. At any time three kinds of Boston drivers exist:

Locals, who continue to drive through intersections after the light turns red; tourists, who drive slowly and cautiously (in a rental car, typically); and taxis, who have a wide variation of driving patterns, depending on the kinds of passengers in the cabs.

Also, Boston has two kinds of pedestrians: locals, who cross the street whenever they feel like it and seldom use the crosswalk buttons; and tourists, who always use the crosswalk buttons and cross only when the Walk/Don't Walk light permits.

Finally, Boston has bicyclists who never pay attention to stoplights.
How do these considerations change the model?
3. You are asked to design a group scheduler. The software enables you to arrange meetings among individuals or groups and to reserve a limited number of conference rooms. Identify the principal subsystems.
4. Design and show the interfaces to the classes in the room reservation portion of the program discussed in Exercise 3.

Chapter 1
Templates
A powerful new tool for C++ programmers is "parameterized types" or templates. Templates are so useful that
the Standard Template Library (STL) has been adopted into the definition of the C++ language.
Today you will learn:

- What templates are and how to use them
- How to create class templates
- Whaw to creat function templates
Whe tandard Template Library is and how to learn to use it


## What Are Templates?

At the end of Week 2 , you saw how to build a PartsList object and how to use it to create a PartsCatalog. If you
want to build on the Parts is object to make a list of cats, you have a problem. Parts ist only knows about part To solve this problem, you can create a List base class and derive from it the PartsList and CatsList classes. You
could then cut and paste much of the PartsList class into the new CatsList declaration. Next week, when you war to make a list of Car objects, you would then have to make a new class, and again you would cut and paste. Neecless to say, this is not a satisfactory solution. Over time, the List class and its derived classes will hav Templates solve this problem, and with the adoption of the ANSI standard, they are an integral part of the

## Parameterized Types

Templates enable you to teach the compiler how to make a list of any type of thing, rather than creating a set of Templatess enable you to teach the compiler how to make a list of any type of thing, rather than creaing a
type-specific lists-a Partsist is is list of parts, a CatsList is a list of cats. The only way in which they differ is
ats type-specific lists - a PartsList is a list of parts, a catsList is a list of cats. The only way in which hiney
type of the thing on the list. With templates, the type of the thing on the list becomes a parameter to the
definition of the class. type of the thing on th
definition of the class.
A common component of virtually all C++ libraries is an array class. As you saw with Lists, it is tedious and
inefficient to create one array class for integers, another for doubles, and yet another for an array of Animals Temp let you declare a parameterized array class and then specify what type of object each instance of array will hold. Note that the Standard Template Library provides a standardized set of container classe including arrays, lists, and so forth. This lesson explores what it takes to write your own so that you fully
understand how templates work; however, in a commercial program, you would almost certainly use the STI classes rather than creating your owr

## Instantiating an Instance of a Template

the template
Parameterized templates provide you with the capability to create a general class and pass types as parae

## Template Definitior

You declare a parameterized Array object (atemplate for an array) by writing

| 1: template <class T> // declare the template and the parameter |
| :--- |
| 2: class Array | /he class being parameterized

( public:
Array ();
full class declaration here

The keyword template is used at the beginning of every declaration and definition of a template class. The paraemeters of the template follow the keryord template. The parameters are the things that will change with
instance. In the array template shown previously for example, the type of the objects stored in the array will instance. In the array template shown previously, for example, the type of the objects stored in the at
change. One instance might store an array of integers, and another might store an array of Animals.

In this example, the keyword class is used, followed by the identifier T . The keyword class indicates that this parameter is a type. The identifier $T$ is used throughout the rest of the template definition to refer to the
parameterized type. One instance of this class will substitute int every where $T$ appears, and another will parameterized
substitute Cat.

```
Array<int> anIntArray
```

Array<Cat> aCatArray:

The object anInA Array is of the type array of integerss, the object aCaAtray is of the type array of cats. You car
now use the type Array <int anywhere you would normally use a type-as the return value from a function, a now parameter
template.
$\qquad$
Listing 19.1 is not a complete program and will not compile.

INPUT. Listing 19.1 A Template of an Array Clas

```
#isting 19.1 A template of an array class
    include <iostream.h>
    template <class T> // declare the template and the parameter
    public:
        / constructors
        rray(int itsSize = DefaultSize)
        Array(const Array &rhs),
        // operators
        M& (%)
        // accessors
    rivate:
        T *pType;
``` In this case, the paramet
the parameterized type. From line 6 until the end of the template on line 24 , the rest of the declaration is like any other class declara
The only difference is that wherever the type of the object would normally appear, the identifier \(T\) is used The only difference is that wherever the type of the object would normally appear, the identifier \(T\) is used
instead. For example, operatorl would be expected to return a reference to an object in the array, and in fac instead. .or example, operatorn woulc
declared to return a reference to a \(T\).
When an instance of an integer array is declared, the operator= that is provided to that array will return a reference to an integer. When an instance
array will return a reference to an Animal.

\section*{Using the Name}

Within the class declaration, the word Array may be used without further qualification. Elsewhere in the
program, this class will be efer progran, this class will be referr
class declaration, you must write

\section*{INPUT Liting 192 The Innd}
\#include <iostream.h>
Iisting 19.2 . Tmplementing the template array
const int Defaultsize \(=10\);
/1 declare a simple Animal class so that we can
class Animal
public:
Animal
Animal (int
Animal)
Anin
int Getweight () const \(\left\{\begin{array}{l}\text { return itsweight; ; } \\ \text { void Display () const }\end{array}\right.\) (cout << itsweight; ;
private:
int
itsweight
Animal: : Animal (int weight)
itsweight (weight)
Animal: : Animal)
it sweight (0)
template <class I> // declare the template and the parameter
class Array
// the class being parameterized fubic.

\(\underset{\sim}{\text { Array (const Array }{ }^{\text {arhs }} \text { ); }}\)
// operators
Array\& operator= (const Array \(\alpha)\),
T\& operator [] (int offset)
ret

int GetSize() const ( return itssize; )

implementations follow
1 implement the constructor



itssize \(=\) rhs. GetSize ()
pType \(=\) new \(T[i t s s i z e] ;\)

/ operator=
template <class
Ampay<T>\& Array<T>: : operator=(const Array \&rhs)
if (this \(=\begin{aligned} & \text { qrhs) } \\ & \text { return *this; }\end{aligned}\)



return *thi
// driver program
int main()
Array<int> theArray; \(\quad\) // an array of integers
Array<nimal> theZoo;
An an array of Animals
Animal *panimal
// fill the arrays
for (int \(i=0 ; i\)
thearray \([\mathrm{i}]=\mathrm{i} * 2 ;\)
pAnimal \(=\) new Animal \((\mathrm{i} * 3)\)

// print the contents of the arrays
for (int \(j=0 ; j<\) theArray. Getsize (); \(j++\)


return 0
thearray
theAr \(\begin{gathered}\text { anay } \\ \text { thearray }\end{gathered}\) theArray [2
theArray 13
thearray 14
 \begin{tabular}{l} 
Array \([6]\) \\
Array \([7]\) \\
\hline
\end{tabular} thezool8:

ANalysIs Lines 8 -26 provide a strippec-d 0 andible to add to the arra
Line 29 declares hat what follows is atemplate and that the perameter to the eemplate is a type, designated as
The Aray ylass has two constructors as s shown, he first of which takes a size and defaults to the constant intege
Defendsise The assignment and offset operators are declared, with the latere declaring
The only aceessor rovided is cesizieo, which returns the size of the array.
 so forth would be required. All this is supplied by the STL container classes, as discussed toward the end of tu

\section*{Template Function}

If you want to pass an array object to a function, you must pass a particular instit
Therefore, if somefinction ( akes an integer array as a a parameter, you may writ
void SomeFunction (Array<int>\&); // ok

\section*{but you may not wri}
?

Here the function MyTemplaefinection is declared to be a template function by the declaration on the top line.
Note that like othere functions, template functions can have any name. Template functions can also take instances of the template in addition to the parameterized form. The followi
is an example:



\section*{Templates and Friends}

\section*{Template classes can declare three types of friends:}

A non-template friend class or function
A
\(A\) eneraral evpplate firiend class or function

\section*{-Tempate Friend Classes and Function}

It is possible to declare any class or function to be a friend to your template class. Each instance of the class
traat he friend properly, as if the declaration of frienship had been made in that particulari instance. Lisising


Listing 19.3 - Type-specific friend functions in templates
const int Defaultsize \(=10\)
// declare a simple Animal class so that we cai
\({ }_{1}^{\text {class Animal }}\)


private:
int
itswei ght
Animal: : Animal (int weight):
itswe ight (weight)
Animal: : : Animal
itssue ight ( 0 )
template <class \(\mathrm{T}>\) // declare the template and the parameter
class Array

rray (const Array = defaultSize)
// operators
Array\& operat


int Gectsize () const ( return itssize; )
// friend function
friend void Intrude (Array<int>),

/ friend function. Not a template, can only be used
// with int arrays: Intrudes into private data.
左 <

follow.
/ implement the constructor
template
<class
T>





/ operator=
template <class
rator=(const Array \(\&\) rhs




// driver progran
int main()



int
for
fic
\(i\)
cout << "thezool"
thezool \([j]\).Display
cout \(\ll\) endl

intrude (tnearray);
cout <<" "n\nDone. \(\backslash \mathrm{n}^{\prime}\);
feturn \(0 ;\)

\section*{To make this work, you need to declalre operator< to be a template function.}

Class Animal
pubicic:
aninalintin


\(\underset{\text { private: }}{\text { int }}\) itsweight;

Animal: : Animall
itswei ght (0)
template <class I>
class Array \(\quad \begin{aligned} & \text { declare the template and the parameter } \\ & \text { c/ the class being parameterized }\end{aligned}\)

Array (int itssirse
Array \(=\) Defanultsize
Array




int Getsize() const (return itssize; )



 output;
now...



itssize \(=\) rhs. 6 etsize
prype \(=\) new \(T\) I \(i\) ssize


(const Array \&rhs)



```

M

```
\(t\) main!

while (!stop)

if \(\begin{gathered}\text { (offset } \\ \text { break; }\end{gathered}\)
<
if (offset > 9)
cout << "**
cont inue,
cout << "Mhere's the entiire array: \(: \mathrm{nn}^{\prime \prime}\)
cout <<
ceturn \(0 ;\)
\[
\begin{gathered}
\text { private: } \\
\text { int itsweight; }
\end{gathered}
\]
// extraction operator for printing animals
thestream \(\ll\) theAnima1. Getweight (),
return theStreami
Animal: : Animal (int weight):
itsweight (weight)
|| cout << "Animal(int) \n"
Animali: : Animal
it swei ight ( 0 )
cout << "Animal() nn"; \(^{\text {n }}\)
Anima1: : ~Anima1)
template <class T> // declare the template and the paramete
class Array
// the class being parameterized public:



int Getsize() const ( return itsSize;
friend function
Criend ostreams ope

template <class T >
ostream\&
operatotorc

return output;
/ implement the constructor

itssize (size


itssize \(=\) rhs. 6 etsize (),
pType
new \(T\) Itssize


main!
Array<int> intarray;
Array \(\\) Animal
Int Finlifunct ion (intarray)


bool stop \(=\) false
int offtset, walue
while
whtop)

break;
use values between 0 and \(9 . * * *\) \begin{tabular}{c} 
cout \(\ll\) \\
continue \\
\hline
\end{tabular}
hearray loffset] = value;
\(\underset{\substack{\text { Animal } \\ \text { for } \\ \text { (int } \\ \text { panimal; }}}{ }\)
panimal \(=\) new Anima1;
panimal--setweimatin
theAray [i] \(= \pm\) PAnimali \(i\)
delete panimali; \(/ /\) a copy was put in the array
am

\#incluade <iostream.h>
isting \(19.6=\) serialiiting Template Implementations
const int pefaultsize \(=3\);
\[
\begin{aligned}
& \int_{\text {/ Atrivial class }}^{\text {class Animal }} \\
& \text { public }
\end{aligned}
\]
\[
\begin{aligned}
& \begin{array}{l}
\text { Animal (int) } \\
\text { Animal(i) } \\
\sim \text { Animal(); }
\end{array}
\end{aligned}
\]
 Ifriend operators
riend ostream\& ope
private
int \(i\) itsweight
// extraction operator for printing animals
streaema operatoprec
(ostream\& thestream, const Animal\& theanimal
thestream << theanimal. Getweight ();
return thestreamis
Animal: : Animal (int weight)
itsweight (weight)
cout << "animal (int) "
Animal: : Animall ()
itsweight (0)
cout << "animal()";
Animal: : ~Animal()
cout << "Destroyed an animal.
template <class T> // declare the template and the parameter
class Array
// the class being parameterized
public:
\(\begin{gathered}\text { Array } \\ \text { Array (const } \\ \text { (int } \\ \text { Array }\end{gathered}\) itssize \(=\) Deults )
Array() (detete (1) pyype; )


// accessors
int Getsize()
const ( return itssize;

\(\underset{\substack{\text { private: } \\ \text { I tpype; } \\ \text { int } \\ \text { int } \\ \text { itssize }}}{ }\)





for \((\) int \(i=0 ; i<i\) itssiz
pType \([i]=\) rns \([i] i\)
\(\underset{\text { return tinis; }}{\text { PTy }}\)
template <class T>
Array<T> : A Array (const Array \(\&\) rhs \()\)
itssize \(=\) rhs.Getsize();
pType \(=\) new \(T[i t s s i z i z e] ;\)




Array<Animal>>: Array (int Anima1ArraySize)
it sze (Animal Arraysize)
PType \(=\) new Animal [AnimalArraySize];

main()
Array<int> intarray;
Array<Anima1> animal



Void IntfillFunction (Array<int>\& thearray)
bool stop \(=\) false,
int offset,
in \(\begin{gathered}\text { value }\end{gathered}\)

if (offset \(<0\) )
braeaki,
if (offset \(>9\) )
cout <<"
continue;
\(\qquad\)








\section*{progam. andi is refectecd inthe output on lieses \(18-20.10\).}

\section*{remersed Tenpates}







\(\underset{\substack{\text { private } \\ \text { int itsee ight; }}}{\text { int }}\)



Hout «《 anmar (unt)

ani inal: :- - - nimalnal)




\({ }^{1 / 1}\) friend function

eity













 cout « \(n\) aninal arrays \(\mathrm{m}^{4}\)
 delete proterray
 \(=\)

ANALYSIS The declaration of the Animal class has been left out to save space. The Array class has added the static variable itsNumberArrays on line 74, and because this data is private, the static public accesso GetNumberArrays() was added on line 66.

Initialization of the static data is accomplished with a full template qualification, as shown on lines 77 and 78 The constructors of Array and the destructor are each modified to keep track of how many arrays exist at any moment.

Accessing the static members is the same as accessing the static members of any class: You can do so with an existing object, as shown on lines 132 and 133, or by using the full class specification, as shown on lines 126 and 127. Note that you must use a specific type of array when accessing the static data. One variable exists fo each type.

Note: Version 2.7.2 cannot compile this code-giving you the following messages:
1st19-07.cxx:78: sorry, not implemented: static data member templates
1st19-07.cxx:78: end of file read inside definition
Use version 2.9 .5 instead for all remaining examples in this lesson. It is the newer compiler and has the more
advanced features.

DO use statics with templates as needed
DO specialize template behavior by overridin
template functions by type

\section*{Ouse the parameters to template functions to narrow} their instances to be type-safe

\section*{The Standard Template Librar)}

A new development in C++ is the adoption of the Standard Template Library (STL). All the major compiler vendors now offer the STL as part of their compilers. Newer versions of the GNU compiler include the STL STL is a library of template-based container classes, including vectors, lists, queues, and stacks. The STL also includes a number of common algorithms, including sorting and searching

The goal of the STL is to give you an alternative to reinventing the wheel for these common requirements. Th STL is tested and debugged, offers high performance, and is free. Most importantly, the STL is reusable; after you understand how to use an STL container, you can use it in all your programs without reinventing it.

\section*{Containers}

A container is an object that holds other objects. The standard C++ library provides a series of container classe that are powerful tools that help C++ developers handle common programming tasks. Two types of Standard Template Library (STL) container classes are sequence and associative. Sequence containers are designed to provide sequential and random access to their members, or elements. Associative containers are optimized to access their elements by key values. As with other components of the standard C++ library, the STL is portabl between various operating systems. All STL container classes are defined in namespace std.

\section*{Understanding Sequence Containers}

The Standard Template Library sequence containers provide efficient sequential access to a list of objects. Th standard C++ library provides three sequence containers: vector, list, and deque

\section*{The Vector Container}

You often use arrays to store and access a number of elements. Elements in an array are of the same type and are accessed with an index. The STL provides a container class vector that behaves like an array, but that is mo powerful and safer to use than the standard C++ array

A vector is a container optimized to provide fast access to its elements by an index. The container class vector \(i\) defined in the header file <vector> in namespace std (see Day 17, "Namespaces," for more information on the use of namespaces). A vector can grow itself as necessary. Suppose that you have created a vector to contain 10 elements. After you have filled the vector with 10 objects, the vector is full. If you then add another object to the vector, the vector automatically increases its capacity so that it can accommodate the eleventh object. Her is how the vector class is defined
template <class T, class \(A=\) allocator<T>> class vector
// class members
\};
The first argument (class T) is the type of elements in the vector. The second argument (class A) is an allocator class. Allocators are memory managers responsible for memory allocation and deallocation of elements for the containers. The concept and implementation of allocators are advanced topics that are beyond the scope of this book

By default, elements are created using the operator new() and are freed using the operator delete(). That is, the default constructor of class T is called to create a new element. This provides another argument in favor of explicitly defining a default constructor for your own classes. If you do not define one explicitly, you cannot us the standard vector container to hold a set of instances of your class

You can define vectors that hold integers and floats as follows
\begin{tabular}{lll} 
vector<int> & VInts; & // vector holding int elements \\
vector<float> & vFloats; & // vector holding float elements
\end{tabular}

Usually, you would have some idea as to how many elements a vector will contain. For instance, suppose that your school, the maximum number of students is 50 . To create a vector of students in a class, you will want th vector to be large enough to contain 50 elements. The standard vector class provides a constructor that accepts the number of elements as its parameter. So you can define a vector of 50 students as follows
vector<Student> MathClass(50)
The compiler will allocate enough memory spaces for 50 students; each element is created using the defaul constructor Student::Student)

You can retrieve the number of elements in a vector using a member function, size(). In this example, vStudent.size() will return 50 .

Another member function, capacity(), tells you exactly how many elements a vector can accommodate before it size needs to be increased. You will see more on this later

A vector is said to be empty if no element is in a vector; that is, the vector's size is zero. To make it easier to test whether a vector is empty, the vector class provides a member function, empty), that evaluates to true if th vector is empty.

To assign a Student object Harry to the MathClass, you can use the subscripting operator []:
MathClass[5] = Harry; A deque is ike a double-ended vector-it it inherits the vector container class's efficiency in sequentiar read and
write operations. But, in addition, the deque containe class provides optimized forntend and back-end
operations. These operations are implemented similarly to the list container class, where memory allocations operations. These operations are implemented similarly to the list container class, where memory allocatio
engaged only for new elements. This feature of the deque class eliminates the need to reallocate the whole containe to a new memory location, as the vector class chas to do. Therefore, deques are ideally suited for
applications in which inesertions and deletions tare place at ither one or obth ends, and or which sequent applications in which insertions and deletions take place ate either one or both ends, and for which sequentiz
access of elements is important. An example of such an application is a train assembly simulator, in which access of elements is important. An examp
carriages can join the train a b both ends.

\section*{Understanding Associative Containers}

Although sequence containers are designed for sequential and random access of elements using the index or at
iterator, the associative containers are designed for fast random access of elements suing keys. The standard


\section*{The Map Container}

You have seen that a vector is like an enhanced version of an array. It has all the characteristics of an array an some additional features. Unfortunately, the vector also suffers from one of the signiticant weaknesses of
arrays: It has no provision for the random access of elements using key values other than the index or iterato arrays: It tas no provision for the random access of elements using key values other than
Associative containers, on the other hand, provide fast random access based on key value

The C++ Standard Library provides four associative containerss. map, multimap, set, and multiset. In Listins
19.10 , a map is used to implement the school class example shown in Listing 19.8.
INPUT Listing 19.10 A Map Container Class
```

\#include <iostrean>
\#incluae <lostresme

```
using namespace std;
Listing 19.10: Map container class
class student
public:
    Student ();
student (const string\& name, const int age),
    Student (const string\& name,
Student (const Student \(\AA\) rhs) ;
~Student ();


    Student \& operator=(const Student \& rhs);
private:
string itsName
;
Student: :Student ()
\(:\) itsName ("New Stu
Student : : Student (const string
: itsName (name), itsAge (age)
Student : : Student (const Student \& rhs
it
itsName (rhs GetName ()), itsAge (rhs
: itsName (rhs.GetName ()), itsAgge (rhs.GetAge ())
Student: : ~Student ()
void Student: : SetName (const string\& name)
        itsName \(=\) name \(;\)
string Student::GetName() const
        eturn itsName;
oid Student: : SetAge (const int age)
        itsAge \(=\) age;
    Student: :GetAge() const
        eturn itsAge;
    itsName \(=\) rhs.GetName ()
itsAge \(=\) rhs.GetAge()
    itsAge \(=\) rhs
return *this
os \(<\)
template<class T, class A>
void ShowMap (const
typedef map<string, student> SchoolClass;
int main()

Student Harry ("Harry", 18);
Student Sally ("Sally", 15); Student Sally "("Sally", 15);
Student Bill ("Bil"",17);
Student Peter ("Peter", 16);
SchoolClass Mathclass;
MathClass [Harry.Getname()]
MathClass [Harry.GetName ()] \(=\) Harry;
MathClass [Sally.GetName()] = Sally;
MathClass [Bill.GetName ()] \(=\) Bailli
Mathclass \([\) Peter.GetName ( \()]=\) Peter
cout << "Mathclass:
ShowMap (Mathclass);
 f years old \(\backslash \mathrm{n}^{\prime \prime}\);

// Display map properties
template<class T , class \(A>\)
void ShowMap (const map<T, \(A>\&\) V)
```

for (map<T, A>::const_iterator ci = v.begin();

```
    cout << endl;

Back on lines 81-84, four Student objects are created. The MathClass is defined as an instance of the SchoolClass on line 86 . On lines \(87-90\), we add the four students to the MathClass using the following syntax:
map_object[key_value] = object_value;
You could also use the push_back() or insert) functions to add a (key, value) pair to the map; you can check the GNU compiler's documentation for more details.

After all Student objects have been added to the map, we can access any of them using their key values. On line 95 and 96, we use MathClass["Bill"] to retrieve Bill's record.

\section*{The Other Associative Containers}

The multimap container class is a map class without the restriction of unique keys. More than one element can have the same key value.

The set container class is also similar to the map class. The only difference is that its elements are not (key, value) pairs. An element is only the key.

Finally, the multiset container class is a set class that allows duplex key values.

\section*{Stacks}

One of the most commonly used data structures in computer programming is the stack. The stack, however, is not implemented as an independent container class. Instead, it is implemented as a wrapper of a container. The template class stack is defined in the header file <stack> in the namespace std.
A stack is a continuously allocated block that can grow or shrink at the back end. Elements in a stack can only be accessed or removed from the back. You have seen similar characteristics in the sequence containers, notab vector and deque. In fact, any sequence container that supports the back(), push_back(), and pop_back() operations can be used to implement a stack. Most of the other container methods are not required for the stack and are, therefore, not exposed by the stack.
The STL stack template class is designed to contain any type of objects. The only restriction is that all elements must be of the same type.
A stack is a LIFO (last in, first out) structure. It is like an overcrowded elevator: The first person who walks in is pushed toward the wall, and the last person stands right next to the door. When the elevator reaches the destination floor, the last person in is the first to go out. If someone wants to leave the elevator earlier, all those who stand between her and the door must make way for her, probably by going out of the elevator and then coming back in.

By convention, the open end of a stack is often called the top of the stack, and operations carried out to a stacl are often called push and pop. The stack class inherits these conventional terms.

\section*{Note:}

The STL stack class is not the same as the stack mechanism used internally by compilers and operating systems, in which stacks can contain different types of objects. The underlying functnality, however, is very similar

\section*{Understanding Queues}

A queue is another commonly used data structure in computer programming. Elements are added to the queue one end and taken out at the other. The classic analogy is this: A stack is like a stack of dishes at a salad bar. You add to the stack by placing a dish on top (pushing the stack down), and you take from the stack by "popping" the top dish (the one most recently added to the stack) off the top.

A queue is like a line at the theater. You enter the queue at the back, and you leave the queue at the front. This is known as a FIFO (first in, first out) structure; a stack is a LIFO (last in, first out) structure. Of course, every once in a while, you are second-to-last in a long line at the supermarket when someone opens a new register ar grabs the last person in line-turning what should be a FIFO queue into a LIFO stack, and making you grind your teeth in frustration.

Like the stack, the queue is implemented as a wrapper class to a container. The container must support front), back(), push_back(), and pop_front() operations.

\section*{Algorithm Classes}

A container is a useful place to store a sequence of elements. All standard containers define operations that manipulate the containers and their elements. Implementing all these operations in your own sequences, however, can be laborious and prone to error. Because most of those operations are likely to be the same in
 container. The standard library provides approximately 60 standard algorithms that perform the most basic anc commonly used operations of container

Standard algorithms are defined in <algorithm> in namespace std
To understand how the standard algorithms work, you need to learn the concept of function objects. A functior object is an instance of a class that defines the overloaded operator 0 . Therefore, it can be called as a function. Listing 19.11 demonstrates a function object.
```

\#include <iostream>
using namespace std;
Listing 19.11: A function object
template<class T>
class Prin
public:
void operator()(const T\& t)
{
}
};
int main()
Print<int> DoPrint;
for (int i = 0; i < 5; ++i)
DoPrint(i);
return 0;
}

```

\section*{Outrut}

\section*{Non-Mutating Sequence Operation}

Non-mutating sequence operations perform operations that do not change the elements in a sequence. Thes .nd so forth 19.12 shows how to use a function object and the for_each algorithm to print elements in a vector

\section*{INPUT. Listing 19.12 Using the for each( Algorithm}
```

\#include <iostream>
\#include <vector>
\#include <algorithm>
using namespace std;
template<class I>
class Print
public
void operator()(const T\& t)
cout << t << " ";
;;
int main()
Print<int> DoPrint;
vector<int> vInt (5)
for (int i = 0; i < 5; ++i)
vInt[i] = i * 3;
cout << "for_each()\n";
for_each(vInt.begin(), vInt.end(), DoPrint)
return 0;

```

\section*{OUTrevt}

\section*{Mutating Sequence Algorithms}

Mutating sequence operations perform operations that change the elements in a sequence, including operations that fill or reorder collections. Listing 19.13 shows the fill) algorithm.

Listing 19.13 A Mutating Sequence Algorithm
```

\#include <iostream>
include <algorithm
using namespace std
// Listing 19.13: Mutating Sequence
emplate<class I
class Print
public:
void operator()(const T\& t
cout << t << " ";
}
int main()
Print<int> DoPrint;
fill(vInt.begin(), vInt.begin() + 5, 1);
(vint.begin() + 5, vInt.end(), 2);
for_each(vInt.begin(), vInt.end(), DoPrint)
cout << "\n\n";
return 0;

```

\section*{Output}

ANALYSIS The only new content in this listing is on lines 21 and 22 , where the fillo algorithm is used. The fill algorithm fills the elements in a sequence with a given value. On line 21 , it assigns an integer value 1 to the fi

\section*{Summary}

Today you learned how to create and use templates. Templates are a built-in facility of \(\mathrm{C}++\), used to create prameterize types - to create and use templates. Templates are a built-in facility of \(\mathrm{C}++\), used to create way to reuse code safely and effectively.

The definition of the template determines the parameterized type. Each instance of the template is an actual object, which can be used like any other object-as a parameter to a function, as a return value, and so fort
Template classes can declare three types of friend functions: non-template, general template, and type-specifi template. A template can declare static data members, in which case each instance of the template has its own set of static data.
ned ned speciaize behavior for some template functions based on the actual type, you can override a template function with a particular type. This works for member functions as well
Q\&A
Q Why use templates when macros will do?
A Templates are type-safe and built into the language.
A Templates are type-safe and built into the language.
Q What is the difference between the parameterized type of a template function and the
parameters to a normal function?
A A regular function (non-template) takes parameters on which it may take action. A template function enables you to parameterizze the type of a particular parameter to the function. That is, you can pass an
Array of Type to a function and then have the Type determined by the template instance.
Q When do you use templates and when do you use inheritance?
A Use templates when all the behavior, or virtually all the behavior, is unchanged, except in regard to the
type of the item on which your class acts. If you find yourself copying a class and changing only the type of one or more of its members, it may be time to consider using a template.
Q When do you use general template friend classes?
A When every instance, regardless of type, should be a friend to this class or function
Q When do you use type-specific template friend classes or functions?
A When you want to establish a one-to-one relation
should match iterator (int>, but not iterator<Animal>.
Q What are the two types of standard containers?
A Sequence containers and associative containers. Sequence containers provide optimized sequential an
random access to their elements. Associative contaners provides random access to their elements. Associative containers provides optimized element access using key values. A The class must define a default con operator.

\section*{Workshop}

The Workshop provides quiz questions to help you solidify your understanding of the material covered and exercises to provide you with experience in using what you learned today. Try to answer the quiz and exercise questions before checking the answers in Appendix D, "Answers to Exercises and Quizzes," and make sure you understand the answers before continuing to tomorrow.

\section*{Quiz}
1. What is the difference between a template and a macro?
2. What is the difference between the parameter in a template and the parameter in a function?
3. What is the difference between a type-specific template friend class and a general template friend class?
4. Is it possible to provide special behavior for one instance of a template but not for other instances?
5. How many static variables are created if you put one static member into a template class definition?
6. What are iterators in the C ++ Standard Library?
7. What is a function object?

\section*{Exercises}
1. Create a template based on this List class:
```

class List
{
private:
public:
List():head(0),tail(0),theCount(0) {}
virtual ~List();
void insert( int value );
void append( int value );
int is_present( int value ) const;
int is_empty() const { return head == 0; }
int count() const { return theCount; }
private:
class ListCell
{
public:
ListCell(int value, ListCell *cell =
):val(value), next (cell) {}
int val;
ListCell *next;
};
ListCell *head;
ListCell *tail;
int theCount;
;

```
2. Write the implementation for the List class (non-template) version.
3. Write the template version of the implementations.
4. Declare three list objects: a list of Strings, a list of Cats, and a list of ints.
5. BUG BUSTER: What is wrong with the following code? (Assume the List template is defined and Cat is the class defined earlier in the book.)
```

List<Cat> Cat_List;
Cat Felix;
CatList.append( Felix );
cout << "Felix is " <<
( Cat_List.is_present( Felix ) ) ? "" : "not " << "present\n";

```

HINT (this is tough): What makes Cat different from int?
6. Declare friend operator== for List.
7. Implement friend operator== for List.
8. Does operator== have the same problem as in Exercise 5?
9. Implement a template function for a swap that exchanges two variables.
10. Implement the SchoolClass in Listing 19.8 as a list. Use the push_back() function to add four students to the list. Then traverse the resulting list and increase each student's age by one.
11. Modify Exercise 10 to use a function object to display each student's record.

Chapter 20

\section*{Exceptions and Error Handling}

The code you have seen in this book has been created for illustration purposes. It has not dealt with errors so that you would not be distracted from the central issues being presented. Real-world programs must take error
conditions into consideration. conditions into consideratior

\section*{Today you will learn:}
- What exceptions are
- How exceptions are used, and what issues they raise
- How to build exception hierarchies
- How exceptions fit into an overall error-handling approach

\section*{Bugs, Errors, Mistakes, and Code Ro}

All programs have bugs. The bigger the program, the more bugs, and many of those bugs actually "get out th door" and into final, released software. That this is true does not make it okay, and making robust, bug-free programs is the number-one priority of anyone serious about programming.

The single biggest problem in the software industry is buggy, unstable code. The biggest expense in many programming efforts is testing and fixing. The person who solves the problem of producing good, soli bulletproof programs at low cost and on time will revolutionize the software industry

A number of discrete kinds of bugs can trouble a program. The first is poor logic: The program does just what you asked, but you have not thought through the algorithms properly. The second is syntactic: You used the wrong idiom, function, or structure. These two are the most common, and they are the ones most programme are on the lookout for.

Research and real-world experience have shown beyond a doubt that the later in the development process you find a problem, the more it costs to fix it. The least expensive problems or bugs to fix are the ones you manag to avoid creating. The next cheapest is the of energy into making more and more bugs show up at compile tim

Bugs that get compiled but are caught at the first test-those that crash every time-are less expensive to finc and fix than those that are flaky and only crash once in a while. A bigger problem than logic or syntactic bugs is unnecessary fragility: Your program works just fine if the use
enters a number when you ask for one, but it crashes if the user enters letters. Other programs crash if they run out of memory, if the floppy disk is left out of the drive, or if the modem drops the line

To combat this kind of fragility, programmers strive to make their programs bulletproof. A bulletproof progran is one that can handle anything that comes up at runtime, from bizarre user input to running out of memory, Another term for this process is defensive programming, which, much like defensive driving, is preparation fc the unexpected.

It is important to distinguish between bugs, which arise because the programmer made a mistake in syntax; logic errors, which arise because the programmer misunderstood the problem or how to solve it; and excep which arise because of unusual but predictable problems such as running out of resources (memory or dis space).

\section*{Exceptions}

Programmers use powerful compilers and sprinkle their code with asserts, which are discussed in Day 21 What s Next," to catch programming errors. They use design reviews and exhaustive testing to find logic errors.
Exceptions are different, however. You cannot eliminate exceptional circumstances; you can only prepare foi Exceptions are different, however. You cannot eliminate exceptional circumstances; you can only prepare for
them. Your users will run out of memory from time to time, and the only question is what you will do. Your choices are limited to the following
- Crash the program.
- Inform the user and exit gracefully
- Inform the user and allow the user to try to recover and continue.

Although it is not necessary or even desirable for every program you write to automatically and silently recov from all exceptional circumstances, it is clear that you must do better than crashing.
exception handling provides a type-safe, integrated method for coping with the predictable but unusual .

\section*{A Word About Code Rot}

Code rot is a well-proven phenomenon in which software deteriorates due to being neglected. A perfectly wel witen, fully debugged program will turn bad on your customer's shelf just weeks after you deliver it. After a begun to flake apart.

Besides shipping your source code in air-tight containers, your only protection is to write your programs so th when you go back to fix the spoilage, you can quickly and easily identify where the problems are

\section*{Note: Code rot is somewhat of a programmer's joke used to explain how supposedly bug-free code suddenly
becomes unreliable. It dees, however teach an important lesson. Programm are enomoustly complex, and buss,
errors, and mistakes can hide for a long time before turning up. Protect yourself by writing easy-to-maintain co}

A similar term has been applied to printed books. No matter how careful the author, tech editors, editor typesetters, and reviewers are, misteakes appear. But it seems the
market), the more frequently these errors appear. Think about it!

This means that your code must be commented even if you do not expect anyone else to ever look at it. Si ld

\section*{Exceptions}

In C++, an exception is an object that is passed from the area of code where a problem occurs to the part of th code that is going to handle the problem. The type of the exception determines which area of code will handle the problem, and the contents of the object thrown, if any, can be used to provide feedback to the user.

\section*{behind exceptons is fairly straightorwart}
- The actual allocation of resources (for example, the allocation of memory or the locking of a file) is usually done at a very low level in the program.
- The logic of what to do when an operation fails, memory cannot be allocated, or a file cannot be
- Exceptions provide an express path from the code that allocates resources to the code that can handle the error condition. If intervening layers of functions exist, they are given an opportunity to clean up memory allocations but are not required to include code whose only purpose is to pass along the error
condition. condition.

\section*{How Exceptions Are Used}

\section*{Catch Block}

A cach block is aseries of statemens, each of which begins with hhe keyword cach, followed by an exceppio
type in parentheses, followed by an opening brace, and ending with a closing brace ei is an example:
\(\stackrel{\text { try }}{\text { f }}\)
Function ();
(1)
catch (OutofMemo
// take action

\section*{Using try and catch Blocks}

\section*{More Than One catch Specification}

\section*{}
\#include <iostream. h>
Listing 20.2 . nult inle except ions
const int Defaultsize \(=10\);

\section*{\({ }^{c}\) class Array \\  \\  \\ array() ( deletete [] prype;}
\({ }_{\text {1/ operators }}^{\text {A.rray\& operatot }}\)

// accessors
int Getitssize() const ( return itssize; )
// friend function
friend ostreamk oper
define the exception classes
class xBoundary (1);



int \& Array: :operator [] (int offset)


\(\begin{aligned} & \text { throw xboundary (); } \\ & \text { return PType [0] ; }\end{aligned} /\) appease MF

return prype lof fset ;
throw xBundary \()\);
return prype [0]; // appease MFC

Array: : Array (int size
itssize (size)
if (size \(==0)\)
throw xzerol \()\)
if (size



\({ }_{1}^{\text {int main() }}\)

Array intarray (0)
for \((\) int \(j=0 ; j<100 ; ~\)
\(j++\)

Catch (Axray: :xBoundary)
tch (Array::xToosig)
cout << "This array is too big...
eatch (Array: :xToosmall)
cout << "This arra
cout <<"You asked for an array",
cout <<" of zero objects \(!\mathrm{n}^{\prime \prime} ;\)
eatch (...)
cout << "Something went wrong! !n"
cout
ceturn \(0 ;\)

\section*{Exception Hierarchies}

Exceptions are classes, and as such, they can be derived from. It may be advantageous to create a class \(x\) Size, and to derive from it xZero, xToosmall, xTooBig, and x Negative. Thus, some functions might just cal
and other functions might catch the specific type of \(x\) Size error. Listing 20.3 illustrates this idea.

INPUT. Listing 20.3 Class Hierarchies and Exception:
\[
\begin{aligned}
& \text { \#include <iostream.h> } \\
& \text { isting } 20.3 \text { Class Hie }
\end{aligned}
\]
cting 20.3: Class Hierarchies and Exceptions
const int DefaultSize \(=10\);
class Array
public: // constructors
Array (int itssi Array (const Array \&rhs) ~Array() ( delete [] pType;
// operators
Array幺 operator=(const Array \()\)
int \(\&\) operator \([1\) (int offset)
int\& operator [] (int offset) ;
const int\& operator [] (int of
// accessors int GetitsSize() const \{ return itsSize;
// friend function
friend ostream \(\AA\) ope
/ define the exception classes
define the except
class xBoundary ;
class xSize \(\} ;\)
class \(\times\) TToobig :
class xToobig: \(:\) public xSize
class xTooSmall : public xsize
class xTooSmall : public xSize \(\}\),
class xZero : public xToosmall \(\{\)

int \({ }^{\text {*pType; }}\)
int
itssize

Array::Array(int size):
-
if (size \(==0\) )
throw xZero()
if (size > 30000 )
throw \(\times\) Toobig ()
if (size \(<1\) )
throw xNegative();
\(\begin{gathered}\text { if } \\ \text { (size }<10 \text { ) } \\ \text { throw }\end{gathered}\) (Tol
pType = new int [size]);
for (int \(i=0 ; i<s i z e\) for (int \(i=0 ;\)
pType \([i]=0 ;\)
int\& Array::operator[](int offset)
int size \(=\) GetitsSize();
if (offset \(>=0 \& \dot{\&}\) offset \(<\) Getitssize())
return pType [offset];
throw xBoundary (); \(\qquad\)
const int\& Array::operator [1 (int offset) const
int mysize \(=\) GetitsSize()
if (offset \(>=0 \& \&\) offset
return pType [offset];
hrow xBoundary
return pType [0]; // appease MFC
int main()
\({ }_{i}^{\operatorname{try}}\)
Array intarray (0);
for (int \(j=0\) ) \(j<\)
intArray \([j]=j ;\)
cout \(\ll\) "intArray
catch (Array: :xBoundary)
cout << "Unable to process your input!\n"
Catch (Array::xTooBig)
cout << "This array is too big...\n";
catch (Array::xTooSmall)
cout << "This array is too small
catch (Array:: :xzero) cout << "You asked for an array
cout << " of zero objects! \({ }^{n}\) n";
catch (...)
cout << "Something went wrong! !n";
cout << "Done. nn" \(^{\prime}\);
cout \(\ll\)
return 0;
```

finclude <iostream.h>
*)
class xBoundary {},
lemplate <cla
/// constructors
Array(int itsSize = DefaultSize)
Array(const Array \&rhs);
// operators
T\& (aym\& operator=(const Array<T>\&);
const T\& operator[f(int offSet) const,

```
    // accessors
int Getitssize() const ( return itssize; \}
    // friend function
friend ostream\& oper
    / define the exception classes
    class xSize \{\};
    private:
        解
    template <class T>
Array<T>: \(:\) Array (int size)
itssize
    Array<T>: : :Arr
itssize (size)
        if (size \(<10 \|\) size \(>30000\)
throw \(\mathbf{x S i z e}\) () :
        throw XSizen ;
Type \(=\) new TIsize
        or (int \(i=0 ; i<\) size; \(i++\)

        if (this \(==\&\) rhs)
        delete [] prype;
        thsize \(=\) rhs. Getitssize
pType \(=\) new TIitssizel
        for (int \(i=0 ; i<i\) itss
pType \([i]=\) rhs \([i] ;\)
    template <class T>
        itsSize = rhs.GetitsSizel
        prype = new T[itssize];
            (1nt \(i=0 ; i<i t s s\)
pType[i] \(=\) rhs \([i]\) :
    template <class \(T>\)
T\& Array<T>: :operator [] (int offset)
        Int size \(=\) Getitssize();
if (offset \(>=0 \otimes \&\) offset
        return pType [offset];
        throw xBoundary ()
    template <class T>
        int mysize \(=\) GetitsSize()
if (offset \(>=0 \& \&\) fafse
            return pType loffset
row xBoundary ();
    template <class \(\mathrm{T}>\)
ostream\& operator

        output \(\lll " ;\)
eturn output;
    \({ }_{1}^{\text {int }}\) main()
        \({ }_{i}^{\operatorname{try}}\)
            Array<int> intarray (9);
for (int \(j=0 ; j<100 ;\)
            intArray \([j]=j ;\)
cout \(\ll\) "intanray \([" \ll j \ll "]\) okay \(\ldots "><\) endl
        catch (xBoundary)
            cout << "Unable to process your input! \n"
            atch (Array<int>::xSize)
            cout << "Bad Size! \({ }^{n}\) "
        cout <<
feturn 0;

This frequently becomes something of a religious argument, but a reasonable way to decide the question is to ask the following: Does use of exceptions in this way make the code easier or harder to understand? Are there fewer risks of errors and memory leaks, or more? Will it be harder or easier to maintain this code? These decisions, like so many others, will require an analysis of the trade-offs; no single, obvious right answer exists.

\section*{Dealing with Bugs and Debugging}

Nearly all modern development environments include one or more high-powered debuggers. The essential idea of using a debugger is this: You run the debugger, which loads your source code, and then you run your program from within the debugger. This enables you to see each instruction in your program as it executes and to examine your variables as they change during the life of your program.

All compilers will let you compile with or without symbols. Compiling with symbols tells the compiler to creat the necessary mapping between your source code and the generated program; the debugger uses this to point to the line of source code that corresponds to the next action in the program.

Symbolic debuggers make this chore easier. The GNU debugger, gdb, is a symbolic debugger running in text mode; some other symbolic debuggers support a windowed, full-screen mode. When you load your debugger, will read through all your source code. You can step over function calls or direct the debugger to step into the function, line by line.

With most debuggers, you can switch between the source code and the output to see the results of each executec statement. More powerfully, you can examine the current state of each variable, look at complex data structure examine the value of member data within classes, and look at the actual values in memory of various pointers and other memory locations. You can execute several types of control within a debugger that include setting breakpoints, setting watch points, examining memory, and looking at the assembler code.

\section*{Using gdb-GNU Debugger}

The first step in using the GNU debugger (or just about any other debugger) is to make sure that the compiler has included the required code in the executable file. With this information, the debugger knows where variables are stored, the names of variables and functions, and how to relate the source and binary code.

The first question that might come to mind is: "Why not always include the special debugging code in the executable?" The answer is simple; that code slows the execution of the program and causes it to occup additional memory. We all want our programs to run quickly and use the fewest system resources possible. So, unless we need it, we do not tell the compiler to include this information.

To tell the GNU compiler to include debugging information in your binary code, you use the -g option
g++ -g yourprog.c++ -o yourprog
By default, the debugging information is not included.
After you have compiled all your modules with the -g option and linked them into an executable, you are ready to use gdb to debug your program.

Table 20.1 shows some of the common commands for use within gdb.

Table 20.1Common gdbCommands
\begin{tabular}{ll}
\hline Command & Purpose \\
\hline & \\
\hline
\end{tabular}

Set a breakpoint when function in file is entered. file is optional.

\section*{break file:function}

\section*{list file:line}
next
print exp
quit
run arglist
set variable \(=\exp\)

\section*{undisplay}
watch
whatis exp

Short form of backtrace; display the program stack. Continue running program after breakpoint.
Show binary code as assembler code instead of source language (that is, \(\mathrm{C}++\) ).
Display exp whenever program stops.
Get general help or help on categories of commands, specific commands, and options for specific commands. Almost the most useful command of all.
Displays \(\pm 5\) lines of source code. file and line specify the source to display and are optional.
Run next program step without stepping into called functions.

Print exp; exp can be a variable, function name, or complex expression like the beginning of an array (array name) plus a value to show a specified element. This allows you to examine memory.
Exit from gdb. If you have a running program, you wil be prompted to confirm.
Run program from the beginning with optional command-line argument list arglist.
Set source code variable to expression exp. As with print, exp can be a variable, function name, or complex expression. The variable follows the syntax of the source language and recognizes things like array notation. This allows you to alter memory.
Alter environment of gdb.
Run next program step-stepping into any called functions.
Cancel display.
Create a watchpoint.
Display data type of exp.

Tip: Two quick tip for working with gdb:


The other tip involves a confusing error message:
est1.cc:6: No such file or directory (ENoent).



\section*{For more informatio
man page or info file}

The next sections describe what the terms breakpoint, watch point, examining and altering memory, an
disassembly mean.

\section*{Using Breakpoints}

Breakpoints are instructions to the debugger that when a particular line of code is ready to be executed, the program should stop. This allows you to ury your program unimpeded untid the ine in inuestion is reached.
Breakpoints help you analyze the current condition of variables just before and after a critical line of oode Using Watch Points
It is possible to tell the debugger to show you the value of a particular variable or to break when a particular
variable is read or written to. Watch points enable you to set these conditions, and at times even to modify variable is read or written to. Watch points enable
value of a variable while the program is running.

\section*{Examining and Altering Memory}

At imes it is inportant to see the actual values held in memory. Modern debuggers can show values in the fo
of the actual variable; that is, strings can be shown as characters, longs sas numbers rather than as 4 bytes. and
 forth. The sophisticated gab debugger can even show complete classes and provide the current value of all tha
member variabses, including the this pointer. Just as important, gda allows you to change variables as your
program is executing program is executing.
 Otten, the same e error can be generated in different parts of the code. Exceptions let you centralizz the
andling of errors. Additionally, the part of the code that generates the error may not be the best place to tetermine how to hande the error.
Why enenate an object? Why not just pass an error code?
Objects are more flexile and powerfut than error ocdes. They can convey more information, and the
 Qack to previous areas of the code, even when noneexceptional conditions exist? Yes, some C++ programmerrs use exceptions for just that purpose. The danger is thate exceptions mis
ceate memory leaks as the stack is unwound and some objects are inadvertenty left in the free store.


Does an exception have to be caugh in the same place where the try block created the exceptio
No. it is possible to catch an exceppion annywhere in the call stack. As the stack is unwound, the
exception is passed up the stack until it is handled. Q Why use a debugger when you can use cout with conditional (fifidef debug) compiling? The debugger provides a much more powerful mechanism for stepping through your code and

Workshop
The Workshop contains quiz questions to help solidify your understanding of the material covered and exerci to provide you with experience in using what you have learned. Try to answer the quiz and exercisis quee
beforo checking the answers in Appendix D, "Answers to Exercises and Quizzes," and make sure you understand the answers before continuing tomorrow
und

\section*{Quiz}
4. What is a cotch statatement? can exception contain?
5. When are exception objects created?
6. When are exception objects created? fould youss exceptions by value or by reference
7. Will a catch statement catch a derived exception if if it looking for the base class? . What does catch(...) mean?

Exercises
1. Create a try lock, a catch statement, and a simple exception.
2. Moodify the answer from Exercise 1 , put data into the exceptit
ise it in the catch boo
. derived objects and the base objects.
4. Modify the program from Exercise 3 to have three levels of function calls
. BUG BUSTER: What is wrong with the following code?


This exercise shows the intended bug. It is allocating memory to present the e error message, but it doess so
precisely when no memory is available to allocate after all, that is what this excention mane precisely when no memory is avaliable to allocate (after all, that is what this exception manages!,
this program by changing the line if (var \(=0\) ) to if \((1)\), which will force the exception to be thrown

Chapter 21

\section*{What's Nex}

Congratulations! You are nearly done with a full three-week intensive introduction to \(\mathrm{C}++\). By now you should have a solid understanding of \(\mathrm{C}_{++}\), but in modern programming, there is always more to learn. This day fills in some missing details and then sets the course for continued stud)
Most of what you write in your source code files is \(\mathrm{C}++\). This is interpreted by the compiler and turned into yo Before the compiler runs, however, the Crocessor runs, and this provides an opportunity fo conditional compilation. Today you will learn:
- What conditional compilation is and how to manage it
- How to write macros using the preprocessor
- How to use the preprocessor in finding bugs
- How to manipulate individual bits and use them as flags

\section*{The Preprocessor and the Compile}

Every time you run your compiler, your preprocessor runs first. The preprocessor looks for preprocessor instructions, each of which begins with a pound symbol (\#). The effect of each of these instructions is a chang to the text of the source code. The result is a new source code file-a temporary file that you normally do not see, but that you can instruct the compiler to save so that you can examine it if you want to

The compiler does not read your original source code file; it reads the output of the preprocessor and compile. that file. You have already seen the effect of this with the \#include directive. This instructs the preprocessor to find the file whose name follows the \#include directive and to write it into the intermediate file at that location. as if you had typed that entire file right into your source code, and by the time the compiler sees the source ode, the included file is there

\section*{Seeing the Intermediate Form}

Nearly every compiler has a switch that you can set either in the integrated development environment (IDE) o Neary every compier has a switch that you can set either in the integrated de
For the GNU compiler, that option is - E. For other compilers, check your documentation for the right switches to set if you would like to examine this file

\section*{The \#define Preprocessor Directive}

\author{
The \#define command defines a string substitution. If you write
}

\section*{define BIG 512"}
you have instructed the precompiler to substitute the string 512 wherever it sees the string BIG. This is not a string in the \(\mathrm{C}++\) sense. The characters 512 are substituted in your source code wherever the token BIG is seen. token is a string of characters that can be used wherever a string or constant or other set of letters might be use token is a string of
Thus, if you write
\#define BIG 512
int myArray [BIG]
The intermediate file produced by the precompiler will look like this
int myArray [512];
Note that the \#define statement is gone. Prec解

\section*{Using \#define for Constants}

One way to use \#define is as a substitute for constants. This is almost never a good idea, however, because \#define merely makes a string substitution and does no type checking. As explained in the section on constant tremendous advantages exist in using the const keyword rather than \#define

\section*{Using \#define for Tests}
A second way to use \#defin
Therefore, you could write
\#define BIG

Later, you can test whether BIG has been defined and take action accordingly. The precompiler commands to test whether a string has been defined are \#ifdef (if defined) and \#ifndef (if not defined). Both must be follow by the command \#endif before the block ends (before the next closing brace).
\#ifdef evaluates to true if the string it tests has been defined already. So, you can write
```

\#ifdef DEbuG

```
cout << "Debug defined"
\#endif

When the precompiler reads the \#ifdef, it checks a table it has built to see if you already defined DEBUG. If you have, the \#ifdef evaluates to true, and everything to the next \#else or \#endif is written into the intermediate file fo betwen \#ifdef D教 th if it place.

\section*{that point in the file}

\section*{Using the \#else Precompiler Commanc}

As you might imagine, the term \#else can be inserted between either \#ifdef or \#ifndef and the closing \#endi Listing 21.1 illustrates how these terms are used.

\section*{INPUT}
```

\#define DemoVersion
\#define NT_VERSION 5
\#include <iostream.h>
int main()
cout << "Checking on the definitions of Demoversion, NT_VERSION
f and WTNDOWS VERSION...\n";
\#1fdef DemoVersion
\#else
cout << "Demoversion not defined.\n";
\#ifndef NT_VERSION
\#else cout << "NT_VERSION not defined!\n"
cout << "NT_VERSION defined as:" << NT_VERSION << endl;
ifdef wINDows_VERSION
cout << "WINDOWS_VERSION defined!\n";
\#else}\mathrm{ cout << "wINDOWS_VERSION was not defined.\n";
\#endif
cout << "Done.\n";
return 0;

```

\section*{Outrut}

Finally, on line 23 , the program tests for wINDOWS_VERSIIN. Because you did not define windows_VERSIO

\section*{Inclusion and Inclusion Guards}

You will create projects with many different files. You will probably organize your directories so that each cl. You wilc create projects with many different files. You will probably organize your directories so that eac
has its own header file (for example, hpp with the class declaration and its own implementation file (for
example., cppp with the source code for the class methods.

Your maino function will be in its own .cpp file, and all the .cpp files will be compiled into .obj files, which wi
then be linked into a single program by the linker. inked into a single program by the linker.
Because your programs will use methods from many classes, many header files will be included in each file Also, header files often need to include one another. For
 include the file ANIMALL.hpp in Doc.hpp, or Dog will not be able to derive from Animal. The Cat header also include the file ANIMAL..hpp in DOG.hpp, o
includes ANIMAL.hpp for the same reason.
If you create a method that uses both a Cat and a Dog, you will be in danger of including ANIMAL.hpp twice. TI will generate a compile-time e error because it is not legal to declare a class (Animal) twice, even though the
declarations are identical. You can solve this problem with inclusion guards. At the top of your ANIMAL head declarations are identical.
file, you write these lines:
\#i fndef ANIMAL_HPP
didefine ANMMAL_HP
\(\ldots\) \#endif
This says, if you have not defined the term ANIMAL_HPP, go ahead and define it now. Between the \#define statement and the closing \#endif are the entire contents of the file.
The first time your program includes this file, it reads the first line and the tes
not yet defined ANIMAL HPP. So, it defines it and then includes the entire file
The second time your program includes the ANIMAL.hpp file, it reads the first line and the test evaluates to falso NIMAL.hpp has been defined. It therefore skips to the next telse in this case there is not one) or the next tend and the class is not declared twie.

The actual name of the defined symbol (ANIMAL_HPP) is not important, although it is customary to use the
filename in all uppercase with the dot (.) changed to an underscore. This is purely convention, however.
Note:
It never huts to use inclusion guard. Offen they will save you hours of debugging time.

\section*{Macro Functions}

The \#define directive can also be used to create macro functions. A macro function is a symbol created using \#he \#define directive can also be usce the same way that a function does. The preprocessor will substitute the
\#define; it takes an argument, in much substitution string for whatever argument it is given. For example, you can define the macro TwICE as follows \#define TWICE (x) ( \((\mathrm{x}) * 2)\)

\section*{and then in your code you write}

\section*{TWICE (4)}

The entire string TWICE(4) will be removed, and the value 8 will be substituted! When the precompiler sees th
4, it will substitute \((4) * 2)\), which will then evaluate to \(4 * 2\), or 8 . A macro can have more than one parame
Two common macros are MAX and MIN:

Note that in a macro function definition, the opening parenthesis for the parameter list must immediately follc the macro name, with no spaces. The preprocessor is not as forgiving of whitespace as is the compiler.

\section*{If you were to writt}
\#define MAX (x,y) ( \((x)\) > (y) ? (x) : (y)
and then tried to use max like this:
int \(x=5, y\),
\(z=\operatorname{MAX}(x, y) ;\)
the intermediate code would be
int \(x=5, y=7, z ;\)
\(z=(x, y)(x)>(y) ?(x):(y))(x, y)\)
A simple text substitution would be done, rather than invoking the macro function. Thus, the token MAX wou have sub
MAX.
If you remove the space between mAX and \((x, y)\), however, the intermediate code becomes

\section*{Why All the Parentheses?}

You may be wondering why so many parentheses are in many of the macros presented so far. The preprocesso
does not demand that parentheses be placed around the arguments in the substitution string, but the parenthese does not demand that parentheses be placed around the arguments in the substitution string, but the parenthes
help you to avoid unwanted side effects when you pass complicated values to a macro. For example, if you help you to avo
define MAX as
- \((x, y)\)
```

|/Listing 21.2 Macro Expansion
\#define CUBE (a) ((a) * (a) * (a)
int main()
long x = 5;
cout << "y:" << y << end1;
Iong a = 5,b=
z=THREE (a+b);
cout << "y:" << y << endl;

```
 \begin{tabular}{l} 
expan \\
125. \\
\hline
\end{tabular}
In the second use, on lines \(16-18\), the parameter is \(5+7\). In this case, CUBE \((57)\) ) evaluates to
( \(15+7\) ) * \((5+7\)
\((12) *(12) *(12)\)

\section*{}
\(\qquad\)

\section*{\(+(35)+(35)+7\)}

Macros Versus Functions and Templates
Macros suffer from four problems in C+H.T.The firsti is that they can be confusing if they eet large, because all character ( (), which is known os as line spplicing

The preprocessor distributed with the GNU compilier does not hand
remain on one line. Large macros quickly become difficult to manag
The second problem is that macros are expanded inline each time they are used. This means that if a macro
used a dozen times, the sussestution will appear 12 itimes in your program, rather than appearing once as function call will. On the other hand, they are usually quicker than a function call because the overtead of教

The fact that they are expanded inline leads to the third problem, which is that the macro does not appear in imemendiat sumce code
debugging maros tricky.
 any argument may be used with a macro, this completely undermines the strong typing of \(C+\) and so os
anale
antemato \(\mathrm{C}+\mathrm{t}\) programmers. Of course, the ight way to solve this is with templates, as you saw on Day 19 Inline Functions
 INPUT Listing 21.3 Using Inline Rather Than a Mact
 inline unsigned long Cubuare (unsigned long a)
(return \(a * a * a ;\) and
\(\underset{\substack{\text { unsigned } \\ \text { for }(i ; i)}}{ }\) long \(x=1\)


turn 0 ;

OUtput


\section*{Analysis}

为

ctions, it it sexactily sasif tit lis line had bad been writun like this


\section*{String Manipulation}

The preprocessor provides two special pepatatrs for manipulating strings in macros. The stringizing operator \((\) c
substitutes a quoted sting for whatever follows the stringizing operator. The concatenation operator bonds strings into

The stringizing operator puts quoeses around any characters following the operator, up to the next whitespace
Thus, if you write
```

and then ca

```
WBTTESTRTNG (This is a string)
the precompiler will turn it int
cout << "This is a string",
Concatenation

The concatenation operator allows you to bond more than one term into a new word. The new word is actually
token that can be used as a class name, a variable name, an offset into an array, or anywhere else series of
Assume for a moment that you
fFiverinin. You can then declare
+


An alternative is to use macros and the concatenation operator. For example, in most compilers, you could
\#def
\#define
public:
de
public:
Type\#ti
private:
int \(i\) itsLength;
For GNU, you would define it as follow

Of course, the paga is not wide enough to st
enough either, but you should get the idea.
This example is overly sparse, concatenation operatorbut the idea would
and data. When you were readd to create an Animal isis, you would write

\section*{Predefined Macros} of these names is surrounded by two underscore characters to reduce the likelihood that the names will conflict with names you are using in your program.

When the precompiler sees one of these macros, it makes the appropriate substitutes. For _DATE _ the curren date is substituted. For _TIME _, the current time is substituted. _LINE _ and __FILE _ are replaced with the source code line number and filename, respectively. You should note that this substitution is made when the
 current date; instead, you will get the date the program was compiled. These defined macros are very useful in debugging

\section*{assert}

GNU and many other compilers offer an assert() macro. The assert() macro returns true if its parameter evaluate true and takes some kind of action if it evaluates false. GNU and many other compilers will abort the program o an assert) that fails; still others will throw an exception (see Day 20, "Exceptions and Error Handling")
powe feature of the assert) macro is that the preprocessor collapses it into no code at all if NDEBUG is not defined. It is a great help during development, and when the final product ships, there is no performance penalty nor increase in the size of the executable version of the program.

Rather than depending on the compiler-provided assert(), you are free to write your own assert() macro. Listing 21.4 provides a simple assert0 macro and shows its use
isting 21.4 A Simple assert() Macro
```

/ Listing 21.4 ASSERTS
define DEBUG
include <iostream.h>
\#ifndef DEBUG
\#define ASSERT(x)
\#else
\#define ASSERT(x) if (! (x)) { cout << "ERROR!! Assert
\#x << " failed\n"; cout << " on line " << LINE
\n"; cout << " in file" << FILE << "\n";
\#endif
nt main()
int x = 5;
cout << "First assert: \n"
ASSERT (x==5);
Mout << "\Second assert: \n"'
ASSERT (x != 5);
cout << "\nDone.\n";
return 0;

```

\section*{Outpul}
```

First assert:
Second assert:
ERROR!! Assert x !=5 failed
on line 24
in file lst21-04.cxx

```

Done.
ANALYSIS On line 2, the term DEBUG is defined. Typically, this would be done from the command line (with the -D command-line argument for \(g++\) ) at compile time, so you can turn this on and off at will. On line (with 8 . the ASSERT() macro is defined. Typically, this would be done in a header file, and that header (ASSERT.hpp) would be included in all your implementation files.

On line 5, the term DEBUG is tested. If it is not define

The ASSERT() itself is one long statement because the GNU preprocessor does not support line splitting
Many other preprocessors do support line splitting-where the ASSERT() is split across seven source code lines as far as the precompiler is concerned. An example of code like that would be replace lines 8-14 in listing 21.4
8: \#define ASSER
9:
10:
11: cout << "ERROR!! Assert" << \#x << " failed 12 ";
out << " on line
\(\qquad\)
4: \}
On line 9 , the value passed in as a parameter is tested; if it evaluates false, the statements on lines 11-13 ar invoked, printing an error message. If the value passed in evaluates true, no action is taker

\section*{Debugging with assert()}

When writing your program, you will often know deep down in your soul that something is true: a function has a certain value, a pointer is valid, and so forth. It is the nature of bugs that what you know to be true might not be so, under some conditions. For example, you know that a pointer is valid, yet the program crashes. assert) ci help you find this type of bug, but only if you make it a regular practice to use assert) liberally in your code. Every time you assign or are passed a pointer as a parameter or function return value, be sure to assert that the pointer is valid. Any time your code depends on a particular value being in a variable, assert) that that is true.

No penalty is assessed for frequent use of assert(); it is removed from the code when you undefine debugging. It also provides good internal documentation, reminding the reader of what you believe is true at any given moment in the flow of the code

\section*{assert() Versus Exceptions}

Yesterday, you saw how to work with exceptions to handle error conditions. It is important to note that assert) not intended to handle runtime error conditions such as bad data, out-of-memory conditions, unable to open a file, and so forth. assert) is created to catch programming errors only. That is, if an assert() "fires," you know have a bug in your code.

This is critical because when you ship your code to your customers, instances of asserto will be removed. You cannot depend on an assert) to handle a runtime problem because the assert) will not be there

It is a common mistake to use assert) to test the return value from a memory assignment:
Animal *pCat \(=\) new Cat

Hatimemimanam
natin
and
.ew





default constructor oreates string of 0 bytes
tring: string

















on on tant of fet operator










tass anm



Asserr ITvaria

Const string\& nam



Asser (Invari iants(1):


int main)



O

\section*{Output}

AGE: \(\quad 5\)
(String Invariants Checked)
(String Invariants Checked)
String Invariants Checked)
(String Invariants Checked)
(String Invariants Checked)
(String Invariants Checked)
(String Invariants Checked)
(String Invariants Checked)
String Invariants Checked)
(String Invariants Checked)
(String Invariants Checked)
(String Invariants Checked)
(String Invariants Checked)
(String Invariants Checked)
(String Invariants Checked)
(String Invariants Checked)
(String Invariants Checked)
(String Invariants Checked)
(String Invariants Checked)
(Animal Invariants Checked)
(String Invariants Checked)
Sparky is (Animal Invariants Checked)
years old. (Animal Invariants Checked
(Animal Invariants Checked)
Sparky is (Animal Invariants Checked)
8 years old. (String Invariants Checked)
years old. (String Invariants Checked)
again with DEBUG \(=\) MEDIUM
AGE:
Sparky is 5 years old.
Sparky is 8 years

ANALYSIS On lines 11-21, the assert) macro is defined to be stripped if DEBUGLEVEL is less than Low (that to be stripped if DEBUG is less than MEDIUM; if DEBUGLEVEL is NONE or LOW, EVAL is Stripped

Finally, on lines \(30-35\), the PRINT macro is declared to be stripped if DEBUGLEVEL is less than HIGH. PRINT i sed only when DEBUGLEVEL is HIGH; you can eliminate this macro by setting DEBUGLEVEL to MEDIUM an used ony when DEBGLEEEL is HGH, you
still maintain your use of EVAL and assert).

PRINT is used within the Invariants 0 methods to print an informative message. EVAL is used on line 221 to PRINT is used within the Invariants 0 methods to print a
evaluate the current value of the constantinteger AGE.

\section*{Bit Twiddling}

Often you will want to set flags in your objects to keep track of the state of your object. (Is it in AlarmState? Ha this been initialized yet? Are you coming or going?
You can do this with user-defined Booleans, but when you have many flags, and when storage size is an issue, is convenient to be able to use the individual bits as flags

Each byte has eight bits, so in a 4 -byte long you can hold 32 separate flags. A bit is said to be "se"" if its value andear if its value is 0 . When you set a bit, you make its value 1 , and when you clear it, you make its val (Set and clear are both adjectives and verbs.) You can set and clear bits by changing the value of the long. b that can be tedious and confusing

\section*{Note:}

Appendix C, "Binary, Octal, and Hexadecimal Numbers and an ASCII Chart," "rovides valuable additional
information about binary, ccala and hexadecinal maniputain
provides bitwise operators that act upon the individual bits. These look like, but are different from, th operators, so many novice programmers confuse them. The bitwise 0 ere 21.1.

Table 21.1The Bitwise Operators
\begin{tabular}{ll}
\hline Symbol & operator \\
\hline\(\&\) & AND \\
1 & OR \\
\(\wedge\) & Exclusive OR \\
\(\sim\) & Complement
\end{tabular}

\section*{Operator ANL}

The AND operator \((\ell)\) is a single ampersand, in contrast to the logical AND, which is two ampersands. Whe the following: The result is 1 if bit 1 is set and if bit 2 is set.

\section*{Operator OR}

The second bitwise operator is OR (I). Again, this is a single vertical bar, in contrast to the logical OR, which . When you OR two bits, the result is 1 if either bit is set or if both tore

\section*{Operator Exclusive OR}

The third bitwise operator is exclusive OR ( 1 ). When you exclusive OR two bits, the result is 1 if the two bits

\section*{The Complement Operato}

The complement operator \((-)\) clears every bit in a number that is set and sets ev
value of the number is 1010 o011, the complement of that number is 01011100

\section*{Setting Bits}

When you want to set or clear a particular bit, you use masking operations. If you have a 4 -byte flag and you want to set bit 8 true, you need to OR the flag with the value 128 . Why? The value 128 is 10000000 in binary want the sulu 8 true, you need to or the flag with the value 128 . Why? The value 128 is 10000000 in binary
thus the value of the eighth bit is 128 . Whatever the current value of that bit (set or clear), if you OR it with th value 128 , you will set that bit and not change any of the other bits. Assume that the current value of the eight
bits is 1010011000100110 . ORing 128 to it looks like this:

987654321
\(1010011000100110 \quad / /\) bit 8 is clear
\(\mid \quad 0000000010000000 / / 128\)
\(\overline{1010} \overline{0} 0 \overline{110} \overline{1010} \overline{0110} \overline{-} / /\) bit 8 is set
There are a few things to note. First, as usual, bits are counted from right to left. Second, the value 128 is all zeros except for bit 8 , the bit you want to set. Third, the starting number 1010011000100110 is left unchange
by the OR operation except that bit 8 was set. Had bit 8 already been set, it would have remained set, which is by the OR opera
what you want.

\section*{Clearing Bits}

If you want to clear bit 8 , you can AND the bit with the complement of 128 . The complement of 128 is th number you get when you take the bit pattern of 128 ( 10000000 ), set every bit that is clear, and clear every bit
that is set ( 011111111 ). When you AND these numbers, the original number is unchanged, except for the eighth that is set ( 011111111\()\). Wh
bit, which is forced to zero

1010011010100110 // bit 8 is se
111111101111111 // ~128

class student
\[
\begin{aligned}
& \text { pub1ic: } \\
& \text { student (): } \\
& \text { myStatu }
\end{aligned}
\]
\[
\begin{aligned}
& \text { (student }(1) \\
& \text { sTATMS } \\
& \text { STATstat }
\end{aligned}
\]
rivate: \(\qquad\) ansigned myGradievel:
unsigned myHousing
unsigned myFoodilan
sTatus student::GetStatus()
if (mystatus)
return Fullime
;
return PartTime;
statas (Staius thestatus)
mystatus \(=\) thestatus;
student Jim;
if (Jim. Getstatus() )== PartTime)
Cout << "Jim is part time" \(\ll\) endl
else
cout << "Jim is full time" << end1,
Jim. SetStatus (PartTime);
if (Jim. Getstatus(l)
cout << "Jim is part time" << end1,
cout << "Jim is full time" << endl
so
char Plan [80]);
switch (Jim. GetPlan())
Case OneMeal: strcpy (Plan, "one meal"); break;
case A11Meals: strcpy (Plan,"A11 meais") ; break
Case Weeknis: strcpy (Plan,"weekend meal is"); break
case NoMeals: strcpy (Plan,"No Meals"); break;
default : cout << "Something bad went wrong \(!\) n"; brea
cout \(\ll\) P1
return \(0 ;\)

OUtput
Jim is part time
Jim is full thime
Jim is on the No
ANalvsis \(\begin{aligned} & \text { On lines } 4.7 \text {, several enumerated types are defined. Theses serve to definine the possible values for } \\ & \text { thithin the sudent class. }\end{aligned}\)



\section*{The class methods are writen as for any oth
fields and not integers or enumerated types.}

The member function Gestausus reads the Boolean bit and return an enumerated type, but this is not necessary
It could dust as asily have been witten to return the value of the bit field directly. The compiler would have done the transtation.
\(\qquad\)
\({ }_{1}^{\text {if }}\) (condition==true)
\(j=k ;\)
SomeFunction();

\section*{Long Lines}

Keep lines to the width displayable on a single screen. Code that is off to the right is easily overlooked, and
scrolling horizontally is annoying. When a line is broken, indent the following lines. Try to braak the line at
 In \(\mathrm{C}+\), fuuctions tend to be much shorter than they were in C , but the old, sound advice still applies. Try to
keep your functions short enough to print the entire function on one paga..

\section*{switch Statements}
switch (variable


\section*{Program Text}
an use several tips to create code that is easy to read. Code that is easy to read is easy to maintain. - Use whitespace to help readabiity,
 and sizeof
Binary
Band so
an onot
- Do not


- Place the pointer or refererence indicator next to the type name, not the variable name:
char \(\star\) foo;
int \(\&\) theInt \(;\)
rather than
char \({ }_{\text {floor }}^{\text {int }}\)
int
\&theInt;
- Do not

\section*{Identifier Names}
\begin{tabular}{|c|c|c|}
\hline \multicolumn{3}{|l|}{\multirow[t]{2}{*}{Idenifier names should be long enough to be descripitive.
Avoid cryptic abreviaions.}} \\
\hline & & \\
\hline \multicolumn{3}{|l|}{ve the time and energ} \\
\hline \multicolumn{3}{|l|}{do not use Hungarian notation. C++ is strongly typed and there is no r able name. With user-defined types (classes), Hungarian notation quic} \\
\hline \multicolumn{3}{|l|}{\multirow[t]{2}{*}{exceptions to this may be to use a prefix for pointers (p) and references (r), as variables (ints). On the other hand, some people swear by Hungarian notation.}} \\
\hline & & \\
\hline \multicolumn{3}{|l|}{\multirow[t]{2}{*}{\begin{tabular}{l}
- Short names (i, p, x, and so on) should be used only where their brevity
and where the usage is so obvious that a descriptive name is not needed. \\
otation
\end{tabular}}} \\
\hline & & \\
\hline \multicolumn{3}{|l|}{\multirow[t]{2}{*}{- The leng}} \\
\hline & & \\
\hline \multicolumn{3}{|l|}{anction (or method) names are ustally verbs or verb-noun phraes: Search, Resel), Find} \\
\hline \multicolumn{3}{|l|}{\multirow[t]{2}{*}{}} \\
\hline & & \\
\hline
\end{tabular}

\section*{ng and Capitalization of Name}

\section*{Spelling and capitilizal
include the following:}
- Use all uppercase and undersccore to separate the logical wordd of names. such as
Source FIE TEMPLATE. Note, however, that these are rare in \(C+\). Consider usin

 should begin with
lowerase elter.
- Enumerated con
enum TextStyle
tsplain,
tsbold,
tsplain,
tsBold,
tsItalic,
tstralic,
tsUnderscore,

\section*{
}

Higher--velec comments are infinitely more important than process details. Add value; do not merely
restate the code.
// n is incremented by one


- Use complete English sentences with appropriate punctuation and capitalization. The extra typing is
worth it. Do not be overly cryptic and do not abbreviate. What seems exceedingly clear to you as you


Use const wherever appropriate: for parameters, variables, and methods. Often there is a need for both a cons and a non-const version of a method; do not use this as an excuse to leave one out. Be very careful when explicitly casting from const to non-const and vice versa (at times, this is the only way to do something), but be ertain that it makes sense, and include a commen

\section*{Next Steps}

You just spent three long, hard weeks working at \(\mathrm{C}++\), and you are now a competent \(\mathrm{C}++\) programmer, but you are by no means finished. There is much more to learn, and you can get valuable information from many more are by no means finished. here is much more to learn, and yc
sources as you move from novice C++ programmer to expert.

The following sections recommend a number of specific sources of information, and these recommendations reflect only my personal experience and opinions. However dozens of books are available on each of these topics, so be sure to get other opinions before purchasing

\section*{Where to Get Help and Advice}

The very first thing you will want to do as a \(\mathrm{C}++\) programmer will be to tap into one of the \(\mathrm{C}++\) conferences on online service. These groups supply immediate contact with hundreds or thousands of \(\mathrm{C}++\) programmers wh can answer your questions, offer advice, and provide a sounding board for your ideas.
programmers participate in the \(\mathrm{C}++\) Internet newsgroups (comp.lang.c++ and comp.lang.c++.moderated) and I recommend them as excellent sources of information and support.

Also, you may want to look for local user groups. Many cities have C++ interest groups where you can meet other programmers and exchange ideas.

\section*{Magazines}

You can also strengthen your skills by subscribing to a good magazine on C++ programming. Some of the best magazines on this topic are: \(C++\) Report from SIGS Publications and \(C / C++\) Users Journal from Miller Freeman. Every issue is packed with useful articles. Save them; what you do not care about today will become critically important tomorrow

\section*{Staying in Touch}

If you have comments, suggestions, or ideas about this book or other books, feel free to contact the authors. If you have comments, suggestions, or ideas about this book or other books, feel free to contact the authors.
Jesse Liberty can be reached at jliberty@libertyassociates.com (Web page http://www.libertyassociates.com). Davi Jesse Liberty can be reached at jliberty@libertyassociates.com (Web page http://www.libertyas
Horvath can be reached at cpplinux@cobs.com (Web page http://www.cobs.com/~dhorvath).
\begin{tabular}{|l|l|}
\hline Do & Don't \\
\hline \begin{tabular}{l} 
DO look at other books. There's plenty to learn and \\
no single book can teach you everything you need to \\
know.
\end{tabular} & \begin{tabular}{l} 
DON'T just read code! The best way to learn C++ is \\
to write C ++ programs.
\end{tabular} \\
\begin{tabular}{l} 
DO subscribe to a good C++ magazine and join a good \\
C++ user group.
\end{tabular} & \\
\hline
\end{tabular}

\section*{Your Very Next Step}

Your very next step, after you catch your breath, is to learn more of the gory details of \(C^{++}\)and Linux. The bonus week of this book takes you there

\section*{Summary}

Today you learned more details about working with the preprocessor. Each time you run the compiler, th preprocessor runs first and translates your preprocessor directives such as \#define and \#ifdef.

The preprocessor does text substitution, although with the use of macros these can be somewhat complex. By using \#ifdef, \#else, and \#ifndef, you can accomplish conditional compilation, compiling in some statements under one set of conditions and in another set of statements under other conditions. This can assist in writing progran for more than one platform and is often used to conditionally include debugging information.
Macro functions provide complex text substitution based on arguments passed at compile time to the macro. It is important to put parentheses around every argument in the macro to ensure that the correct substitution takes place

Macro functions, and the preprocessor in general, are less important in C++ than they were in C. C++ provides number of language features, such as const variables and templates, that offer superior alternatives to use of the preprocessor.

You also learned how to set and test individual bits and how to allocate a limited number of bits to class members.

Q\&A
Q If \(\mathbf{C}++\) offers better alternatives than the preprocessor, why is this option still available?
A First, \(\mathrm{C}++\) is backward-compatible with C , and all significant parts of C must be supported in \(\mathrm{C}+\) Second, some uses of the preprocessor are still used frequently in \(\mathrm{C}++\), such as inclusion guards.
\(\mathbf{Q}\) Why use macro functions when you can use a regular function?
A Macro functions are expanded inline and are used as a substitute for repeatedly typing the same
commands with minor variations. Again, however, templates offer a better alternative.
Q How do you know when to use a macro versus an inline function?
A Often, it doesn't matter much; use whichever is simpler. However, macros offer character substitution stringizing, and concatenation. None of these is available with functions.
Q What is the alternative to using the preprocessor to print interim values during debugging? The best alternative is to use watch statements within a debugger. For information on watch statemen compiler or debugger if you are using something besides the GNU tools).
Q How do you decide when to use an assert() and when to throw an exception?
A If the situation you're testing can be true without your having committed a programming error, use an exception. If the only reason for this situation to ever be true is a bug in your program, use an assert().
\(\mathbf{Q}\) When would you use bit structures rather than simply using integers?
A When the size of the object is crucial. If you are working with limited memory or with
communications software, you may find that the savings offered by these structures is essential to the success of your product.
Q Why do style wars generate so much emotion?
A Programmers become very attached to their habits. If you are accustomed to the following indentation
if (SomeCondition)
\} // statements
it is a difficult transition to give it up. New styles look wrong and create confusion. If you get bored, try logging on to a popular online service and asking which indentation style works best, which editor is be for \(\mathrm{C}++\), or which product is the best word processor. Then sit back and watch as ten thousand messages are generated, all contradicting one another.

\section*{Q What is the very next thing to read?}

A Here are some books I've written to provide a course of study, although there are many other books of great value. C++ Unleashed, Beginning Object-Oriented Analysis and Design, and Career Change C++ are good starting points.
\(\mathbf{Q}\) Is that it?
\(Q\) Is that it?
A Yes! You've learned C++, but...no, you are not done! Ten years ago it was possible for one person to learn all there was to know about computers and programming, or at least to feel pretty confident that sh was close. Today it is out of the question. You can't possibly catch up, and even as you try, the industry is changing. Be sure to keep reding and stay in touch with the resources-magazines and online services-that will keep you current with the latest changes.

\section*{Where You Are Going}

At last, you have reached the bonus week, which covers the more specific and advanced topics of developins C++ programs with Linux. On Day 22, "Linux Programming Environment," you will learn about some of the advanced programming tools under Linux, and on Day 23, "Shell Programming," you will learn how to use th command shells under Linux and create simple shell scripts. Day 24, "System Programming," focuses on operating system interface and C++ runtime library functions and objects that you can use. On Day 25 , hiterprocess Communication," you will learn how to use system functions to communicate between process create applications that take advantage of the Graphical User Interface (GUI) tools available under Linux

\section*{Chapter 22}

\section*{The Linux Programming Environment}

In the days so far, you have been learning about the \(\mathrm{C}++\) programming language. By now, you know about th syntax of the language and how to design and write object-oriented programs. Today's lesson discusses the syntax of the language and how to design and write object-oriented programs. Today's lesson discusses the
environment that Linux provides to write programs. What are the tools available? How do you create a prog compile it, debug it, and track its history as you modify it over time?

This lesson touches on the following basics:
- Editors
- Compilers
- Building executables with make
- Libraries and linking
- Debugging with gdb
- Documentation

\section*{Philosophy and History}

In its early days, Linux was popular almost exclusively with developers-people who wanted to study or writt pieces of the kernel, or people who wanted to write tools and other utilities. If you are going to develop software, the first thing you need is development tools. After you have an editor, you need a compiler. As result, among the first things ported to Linux were compilers, assemblers, and linkers.
compiler is a difficult thing to write, and you certainly would not want to create one from scratch if you did not have to. It was natural, therefore, for the Linux developers to look for a free compiler, with source code available so that it could be ported. The obvious choice (perhaps the only choice) at the time was the GNU C compiler. The GNU tools adhered to a similar philosophy of open source, and they were widely available and known high quality.
Unlike many commercial operating systems today, the tools available for Linux are not limited to those that th vendor provides or that you can buy. Linux is open source software, and open source tools available on the Internet generally work with it.

If there is an open source tool that your distribution does not include and you can find it on the Internet, you ca probably download a version already built for Linux. On the other hand, almost all Linux distributions come probably download a version already built for Linux. On the
with a very complete set of GNU program development tools.

\section*{POSIX}

In the 1980s the Institute of Electrical and Electronics Engineers (IEEE) began developing a standard programming environment to promote the portability of applications between different UNIX environments. programming environment to promote the portability of applications between different UNIX environments.
The name given to this effort and the resulting standard was the "Portable Operating System Interface," and i The name given to this effort and the result ting standard was the Portable Operating System Interface, and it
commonly known as POSIX. The POSIX standard does not dictate how an operating system should behave, \(b\) it does define the programming interface ("API") that should be provided by the operating system to the application program writer. This may seem a bit confusing. Suffice it to say that Linux is "POSIX-compliant. In other words, it provides a standard set of system calls and services as defined by the POSIX standard

\section*{The X Window System}

If you have used Macintosh computers or computers running any version of Microsoft Windows, then you are familiar with what is known as the "graphical user interface" (or GUI). In these systems the GUI is an integra part of the operating system. Not so with Linux. In Linux the GUI is layered on top of windowing system software, which is essentially an accessory to the operating system. There are many advantages to this approa in performance, reliability and flexibility. The name " X Window System" is loosely applied to all of the known, began at MIT in the mid-1980's. The X Window System provided with Linux comes from the XFree8 project.

\section*{Using the Linux Editors}

If you are going to write code, clearly you need some way of entering it into a file. One of the first results of the GNU project was the emacs editor, and editors were among the early ports to Linux

\section*{ed, ex, vi, and the Variants of vi}
ed and ex are very early line-based editors historically available with UNIX systems. ed is generally used as a filter to modify text that is passed through it. ex is a serviceable but primitive line-based text editor. One of th earliest screen based editors in UNIX was vi, which is basically a screen-based front end written around ex. On most Linux systems, the open-source reimplementation of vi named vim is available. Many people consid
vi (and vim too) to be cryptic and difficult to use. It does have many advantages over other editors, however:
- It is very powerful
- It is available on practically all UNIX and Linux systems.

Hequires fewer system resources than emacs, and so will run even if the system is not completely
ioning
- It is not as highly customizable as emacs, which means that all implementations will behave in pretty

位 may choose to use another editor for daily use, but you should master the basics of vi. UNIX systems always come with vi, and because vim is the implementation of vi provided with Linux, vim is usually linked vi so that you can invoke it with the name vi. For the rest of this lesson, we use the terms vi and vim interchangeably. We are really referring to vim.

\section*{Starting vi}

Before you run any of the full-screen editors (including vi), you must have the TERM environment variable se to indicate your terminal type. This variable is normally set automatically to xterm when you start an xterm window.
You may want to view the online manual page for vi with the command: man vi. Note that the command man v gives you exactly the same man page

Extensive help is available within the editor as well

\section*{vi Concepts}
vi is a screen-based "mode editor" with three modes: command, insert, and ex. By default you are in command mode. This means that keys that you type are commands to the editor rather than input into the file you are working on

Most vi commands are a single letter, some with a scope modifier. Most can be preceded with a number, whic is the "repeat factor" (causing the subsequent command to be repeated that many times,
In insert mode, keystrokes are captured and entered into the file being edited. In this mode, the word INSERT displayed at the bottom of the screen.

The original vi was really a screen-based front end to the ex line editor, and vim emulates this behavior. The third mode available in vim is ex mode. In this mode, a "e" prompt appears at the bottom of the screen. You enter any ex command at this prompt. The ex commands are useful for switching files without leaving the enter any ex command at tis prompt. The ex commands are useful for switching files without leaving the
editor, and offer filename completion to make this easier. Most any other task that can be done with an ex
command can also be done from the vi command mode.

To return to command mode from any other mode, press the Esc key. If you are already in command mode, thi key has no effec
command mode
To access the online help as shown in Figure 22.2 with vi, type :help from command mode.


Using vi
Now try a sample vi session. You will create a simple C++ program named hello.cc.
1. Type vi hello.cc to begin editing the new file.
2. Type ito enter insert mode
\#include <stream.h>
main( int argc, char * argv[] )
cout << "hello world"
cout \(\lll\) hello
<< endl;
4. Press Esc to end insert mode.

Now that you are back in command mode note that the \(h, j, k\) and 1 keys allow you to move the curso left, down, up, and right respectively.

At any time, you can re-enter insert mode. The \(i\) key begins inserting before the cursor and the a key begins inserting after that. Note that while you are in insert mode, the word INSERT is displayed at the bottom of the screen. Pressing the Esc key exits insert mode.
5. Type \(\mathbf{Z Z}\) to write the file and exit.

\section*{GNU emacs}

Religious wars abound. Many Linux users swear that vi is too old-fashioned and the only editor worth using emacs. You have to decide for yourself which editor is easier to use. Certainly emacs has more features, and there are several substantial books written on it. It is far too complex to discuss in detail here, so this section ju covers the basics

\section*{Starting GNU emacs}
emacs can run within a terminal window or in a window of its own under X Windows. If you are running X , th command emacs \& will open a new window with the editor. To run within the current window, use the comman emacs -nw
me, it will create a "scratch" buffer for you and will show you a brie help display as shown in Figure 22.3

If you are running the X version of emacs, you can use the mouse to navigate the pull-down menus at the top o the screen. As you become proficient, you will find that there are keyboard shortcuts for all these pull-down commands.
Table 22.1 shows emac conventional notations.

Note: It is common in emacs documentation to refer to key sequences in an abbreviated notation. Table X below
describes the conventional emacs notation. The Meta key may be assigned differenty on different computers. On describes the conventional emacs notation. The Meta key may be assigned differently on different computers. O
some it is a key labeled "Meta" If your computer has no such key (most typical PC's do not) then in place of the
Meta key press and release the Esc key. So if the emacs documentation says \({ }^{*} M-\mathrm{v}\) " then you would press and Meta key press and release the Esc key. So if the emacs documentation says "M-v" then you would press and
celease Esc, and


Figure 22.3 Introductory GNU emacs window
Table 22.1emacs Conventional Notation
\begin{tabular}{ll}
\hline Command & emacs notation \\
\hline Enter & RET \\
Backspace & DEL \\
Escape & ESC \\
Control & C- \\
Meta & M- \\
Space & SPC \\
Tab & TAB \\
\hline
\end{tabular}

The info system (described later) provides online help for all GNU software. To view info, press \(C\) - \(h i\)
Other kinds of help are available within emacs. You can press \(C-h\) ? to find what sorts of help there are. Fo example, try this

> 1. \(C-h\) (displays help options) 2. \(c\) (selects "describe key briefly") 3. \(C-x C-c\) (selects the key to describe)

This will reveal that the sequence \(C-x C-c\) is a shortcut to the "save-buffers-Kill-emacs" command.

\section*{GNU emacs Concepts}

Unlike vi, emacs is known as a "modeless" editor. Typing in the emacs window causes characters to be entere into the buffe
control keys.
emacs has many commands, most of which have very long destiptive names. Many of these command
 bindings for a control key or sequence

One of the nice things about emacs is the wealth of coma learning the control sequences that these keys are bound to.

\section*{Using emacs}
\begin{tabular}{|c|c|}
\hline \multicolumn{2}{|l|}{\multirow[t]{9}{*}{\begin{tabular}{l}
1. Type emacs to start an edit session with no file. \\
2. Press \(C-x C-f\) to get a prompt that will let you select 3. Type hello.ce to open the file you created earlier. A spacebar. If emacs can unambiguously complete the fil \\
4. Press Enter and the file hello.cc will be displayed. \\
5. Use the arrow keys to move around the file. You ca \\
up, \(C-f\) to move forward one character, and \(C-b\) to mov \\
6. Type again and observe that the text is inserted in p \\
7. Press \(C-x C-s\) to write the file.
\end{tabular}}} \\
\hline & \\
\hline & \\
\hline & \\
\hline & \\
\hline & \\
\hline & \\
\hline & \\
\hline & \\
\hline
\end{tabular}

Many people use emacs exclusively. You can use it not only to edit files, but also to run commands, compile programs, debug, and run them. If you have compile-time errors, emacs can automatically jump to the line
containng the error. It is definitely worth your time to investigate emacs further.

\section*{ctags and etags}

\section*{When you are writing a large C+ program, chances are that you wil
several files. Each file will define the methods for a class or classes.}

Later when you are debugging the file, it can become difficult to navigate those sources. You may be editing
file \(A\) and there is a method invoked that is defined elsewhere, perhaps in file \(B\). The ctags and etags program file A and there is a method invoked that is defined elsewhere, perhaps in file B. The ctags and etags progral
create index or tay" files, which vi and emass cen use to help you navige your sources ctags is the oldder program, which generates vit tass by default. You can to to
generates emacs tags by default, but can also be told to generate vi tags.

\section*{ctags Example with vi}
```

Enter the following code into a file named helloMain.S

```
\#include <stream.h>
void doHello(int i)
for (int i=0; \(i<5 ; i++\) )
dohello(i);

\section*{\#include <stream.h>
void doHello (int i}
```

cout << "["<< i i Worl

```
return;

The new file tags has been created.
Begin an edit session with the command \(v\) i helloMain..c. . Now, using the \(h, j, k\) and 1 keys, line the cursor
the word dofello. Press \(C-\), and you can see that the editor automatically opens the file that contains the the word dothlo. Press \(C-\overline{1}\), and you can see that the editor automatically opens the
definition of that function (doHello.cc) and sets the cursor at the start of the function.
```

There is similar functionality in emacs. Use the online help to help exercise it

```

\section*{Languages}

Linux provides all the languages available with traditional UNIX systems, and then some. Many nontraditional
languages are available on the Internet. Most distributions come with C and C++ and often a Java languages are availaben he Interne. Most distributions come with and C+t and offen a Java
implementation. Scripting languages including perl, sed, and awk also are part of most distributions

\section*{gcc and g+1}

The GNU C compiler is named ;gcc and can compile C, C++, and Objective-C. The C compiler is ANSI
 compiler on any other UNIX system, you will find that gc
many commercial sites use it as their C Compiler of choic

\section*{Compiling with gcc}

You invoke the GNU compiler with the command gec. By default this will preprocess, compile, and link aC program. There are many options to oce
of the preprocessscompilelink sequence
```

up.

```

INPUT. Listing 22.1 Main Program dice
\#include <stdio.h>
int doRoll (void);
int i;
int iIter,
int Diel \([6]\)

Iter \(=\) atoi (argv [1]);

Grintf("\%d rolls\n", iIter)
printf ("\traceltrolis


\section*{Compiling with g+-}

The command gec is really a compiler "front end." By looking at the files it is given, it knows whether linkin
or compiling is requid or compiling is required. It also knows whether the source is C or \(\mathrm{C}+\). If you
alternatively invoke the \(\mathrm{g}+\mathrm{t}\) compiler directly with the command \(\mathrm{g}+\mathrm{f}\) fiename.
Although gec can compile C++ programs, it does not automatically link with the class libraries required. You
will need to use g++ for this. As a result, it's generally simpler to compile and link \(\mathrm{C}++\) programs with g+t. For instance
command:
g++ -c hellomain.c
-c doHello.cc
-o hello hellomain.o doHello.o
As with gcc, you can accomplish the same thing with the single invocation
g++ -o hello hellomain.cc doHello.cc

\section*{Scripting Languages (perl, sed, awk)}

As a complete UNIX-like environment, Linux also provides the typical scripting languages. perl, sed, and aw are three important languages provided by default with most Linux distributions. Other scripting languages
(Tcl/Tk, Python, and Expect among others) are available for Linux, but they are beyond the scope of this

When a program is compiled, an object file is generated, and when the program is linked, an executable binary When a program is compiled, an object fire is generated, and The linker must understand the format of the object files because the operating system must loz is created the executable, it must understand the format as well.
and run the ent
You learned earlier that the default executable is named a.out. Until not too long ago, the format of object and executable files was known as the a.out format. The a.out format is fairly old and has several shortcomings. The Frmat). ELF is much more versatile than a out was, and lends itself well to creating shared libraries.

You can tell the format of a file by using the file command:
file roll /usr/bin/file doRoll.o

\section*{Output}

ELF 32-bit LSB executable, Intel 80386 ELF 32-bit LSB executable, Intel 80386,
ELF 32-bit LSB executable, Intel 80386 ,
EIF 32-bit ISB relocatable, Intel 80386 ,

\section*{Shared Libraries}

Often several programs need to do some of the same things, I/O for instance. A long time ago the concept o libraries was developed to accommodate this. The common functions could be put into a file and then whenever a program is built, it extracts those functions that it needs from the librar

This was a breakthrough at the time, but it has several drawbacks. Executables become larger because they eac embed code copied from the libraries. If a bug is found in the library or a feature is added, the executable does
not make use of it unless it is rebuilt.

The solution to this is the shared (or dynamic) library. The mechanism of how shared libraries work is beyond the scope of this discussion. We will just look at how to make and use them linked them to create an executable

There seems to be a market for dice-rolling programs, and we think we can use the doRollo function in of products our new company will create. It makes sense to put the function into a library so that all our

First, we need to create the shared library. Compile the module with the following command
-fPIC -c doRoll.c
Now turn it into a shared library named libroll so
-shared -Wl,-soname,libroll.so.1 -o libroll.so.1.0 doroll.。

Finally, create a link to lo librol.so, so that the running program does not need to keep track of the versio information in the name of the shared library
ln -s libroll.so.1.0 libroll.so
-s libroll.so. 1 libroll.so
Now that we have the library, we must rebuild the main program so that it links with that library at runtime rather than incorporating the code within the executable
-o roll diceMain.c--T. -1roll
The -L. option tells the linker to look in the current directory for libraries, and the - -roll option tells it to look fo a library named libroll.so.

When you run the program, the operating system will dynamically load the correct library, but it has to kno
 where to locate additional libraries
setenv LD_LIBRARY_PATH /home/myname/mylibs
Finally, to see what libraries are used by a program, use the Idd command:

\section*{Output}
iibroll.so. 1 => /home/Dice/libroll.so. 1 ( \(0 \times 40014000\) ) /lib/ld-linux.so. \(2=>/ l \mathrm{ib} / 1 \mathrm{~d}-1 \mathrm{inux}\). so. 2 ( \(0 \times 40000000\)

\section*{Building}

You certainly do not want to be typing in all these gcc commands when building your program. With small programs like those demonstrated in this lesson, it's just slighty annoying. With larger programs involving fite and pertally impossibs

Linux comes with the GNU make utility. make reads a file known as the makefile for the inforion it buld your program. The utility is so important and widespread that it has been specified as a POSIX standard. The GNU version of make is POSIX-compliant.

GNU make automatically looks for a makefile named GNUmakefile. If it does not find that, it looks for makefil and faling that it looks for Makefie. These are the default names Y and faliling that it looks for Makefile. These are the default names. You can name your makefile anything you
like and explicitly direct make to use it. A makefile is an ordinary text file with a very specific syntax that the make utility can understand
make has a great number of built-in rules. For instance, it knows that files ending with .c are C sources, and those into oo object files. You can override any of these rules if you choose. In the simplest case, all that you need to specify in your makefile is the name your executable is to have and the .c that are needed to construct it.
```

roll: diceMain.o doroll.o

```
You can now build the program with a single command
```

Output

```
gcc -c diceMain.c -o diceMain.
gcc -c doroll.c -o doroll.
diceMain.o doroll.
Note: the line in the makefile that begins with "roil is known as the target. It defines the "dependencies" of th
progra
program. The next line is the build rule. make requires that the first
spaces. With spaces instead of a tab, the error make generates is this
make
```

Here is a slightly more complete makefile

```
\# Makefile to build the dice rolling program
CFLAGS \(=-\mathrm{C}\)
OBJS = diceMain.o doroll.
all: roll
roll: \$ (OBJS
\$(CC) \$(CFLAGS) -○ \$@ (OBJS)
directory

\title{
Sibcol1.s.o.1.o \(\$ 8\)
In s libroil.so \\ 1n -s librolil.so.1.0 1ibro11.so.
In -s libroll.so. 1 ibroll.so
}
( (RM) roll *. 0 librol1*

\section*{make Command-Line Option}

make is a very sophisticated program. One of the things it does is understand dependencies. For example, it it
knows that . ofiles are built from cf files. Your program may consist of several. \(c\) source files. If you chang knows that. of files are buil from. .c files. Your program may consist of several. . source files. If y yu chang
one, it is not neessaray to recompile them all every time that you build. You only need to recompile the on
 Sometimes you may war
following command:

This tells make to analyze the makefile and report what commands it will issue to build your program. make will
stop short of executing any commands. this tells make to anayzze the makefie
stop short of executing any command
Another useful faciility in make is the use of variables. Notice that we defined a variable named CFLAGS. make automatically passes this variable to the gcc when it compiles your program. You may want to change the value
of this variable one time without changing the makefile. You can specify a new value on the command line tha will override the one in the file

\section*{Debugging}
 also include xxglb, a version of the same debugger with a graphical
do not have an executable for xxgblb, it is certainly worth building.

List of classes of commands,
aliases -- Aliases of other
 -- Speci fying and examining files
\[
\begin{aligned}
& \text { support -- support facilities } \\
& \text { tracepoint --Tracing of program exed } \\
& \text { user-defined -- User-def ined commands }
\end{aligned}
\]

Type "help" followed by a class name for a list of commands in that class
Type "help" followed by conmand name for full documentation. and name abbreviations are allowed if unambiguous

The help command by itself shows you the classes
commandswithin a particular class, use help again:
(gdb) help breakpoint:

\section*{OUtPut}

Kigure 22.4 xxsdb
 command window. This shows you progran
move your mouse into this pane and type.
 Sample gab Debugging Session
int dorol1 (void);
\({ }_{1}^{\text {main ( }}\) int argc, char * argv[] )

if \((\operatorname{argc}<2)\)
print ("Usage:
exit \((1)\);
Iiter \(=\operatorname{atoi}(\operatorname{argv}[1]) ;\)

\[
\begin{aligned}
& \text { b) print } \\
& =0^{5} \text { next } \\
& \text { nor }
\end{aligned}
\]


Version Control

\section*{Documentation}
great deal of documentation is available with Linux, and most of it is shipped with the distributions. You h cess to man, info, and the Linux HOWTOs and FAQs. Because the documentation can be extensive, many of the distributions make installing the full set optional.

\section*{man Page}
ike any UNIX-like operating system, Linux includes man and man pages for all the command
Sometimes you will encounter a command with a man page in more than one manual section. Normally when you use man, it displays the first entry that it finds. Using man with th
pages for the selected command in all manual sections that have one.
```

The command xman provides a graphical interface to the basic man command. Type

```

\section*{xman -notopbox -bothshown}

You can click the left mouse button on the Sections panel to select another manual section. Left-click any command in the top panel to view the corresponding manual page in the bottom panel. Left-click the Options
button and choose the last entry (quit) to exit xman. Figure 22.5 shows the graphical interface to the manual pa

info
Online documentation in the form of man pages has been in UNIX almost from the start. A more rece tribution has the info system. Navigating info will seem natural to you if you are familiar with emacs you are not, it can be a little tough to learn.

You can invoke info from the command line with the command info. If you are already in emacs, press Esc, an You can invoke info from the command
then type - x info. Figure 22.6 shows info.

\section*{Figure 22.6 inf}

The info command is organized in nodes, where each node represents a major topic. As you read the inf in info, you can follow links to other topics or to more information. There are several single-character
commands in info. The best way to start to learn it is to access the tutorial with the h command from the info directory window). Use the \(m\) command to follow some menu item. Type \(m\) gec to follow the link into information on gcc. You can also move the cursor to a menu item and press Enter. At any time you can type return to the main info directory.
\[
\text { To exit info, type q. If you entered info through emacs, use the normal emacs exit sequence } C \text { - }-\mathrm{h}
\]

\section*{HOWTOs and FAQs}

There is a wealth of documentation contributed to Linux by users who have figured out how to do things and then written up their experiences. These write-ups are known as HOWTOS. There is also o long history in Linu
of various user groups and organizations providing help to users online. Aften the questions and corresponding answers will be collected into documents known as FAQs (Frequently Asked Questions). Most Linu
distributions ship with both the HOWTOs and FAQs. They can typically be found in the directory Idoc/HOWTO and / /us/docFFAQ. Your installation may place them elsewhere

The most recent FAQs and HOWTOs are available online at the Web site of the Linux Documentation Projec http://www.linuxdoc.org

Q I compiled and built my C++ program with gcc. It compiles fine but I have lots of link errors.
A gcc by default links with the standard C libraries, and not the libraries required by \(\mathrm{C}++\). You can specify these explicitly; however, it is simpler to use \(g++\) for the link phase of your build.
Q I've created a program usigshared ibrais, and I've created the libraries. Now when I run
the program I get an error indicating: "cannot open shared object file: No such file or directory"
A You have not told the loader where to find the libraries. The environment variable LD_LIBRARY_PA
is a list of colon-separated paths to directories where the shared libraries can be found.
Q When I type make, it says "make: nothing to be done for ..."
A make attempts to look at all the objects and their corresponding source files. It examines the time
than on these files. II the object is newer than the source, make concludes that object does not need to
be recompiled. If your target exists and all the component objects are newer than their sources, make
believes that there is nothing to be done.
Q When I run gdb, I cannot see my sources. When I run xxgd, the top pane does not show any
Ane ter for your program with the -g flag.
gone.
A When you check A When you check in with the ci command, RCS updates a history file in the RCS directory. Your file probably are not gone; you simply h h
check out locked or unlocked copies.

\section*{Workshop}

The Workshop provides quiz questions to help solidify your understanding of the material covered and
exercises to provide you with experience in using what you have learned. Try to answer the exercises to provide you with experience in wing you have learned. Try to answer the quiz and exercise questions before checking the answers in Appendix \(L\) L
understand the answers before continuing tomorrow.

\section*{Quiz}
1. What is POSIX?
2. What is \(X\) Window
3. What are the two primary text editors available with Linux?
4. What is a primary distinction between vi and GNU emacs?
4. What is a primary distinction between vi and GNU emacs?
5. What is an advantage of shared libraries over static libraries? Static libraries over shared libraries?
6. What is the utility used to compile and build programs? What is it's defaut input file?

\section*{Exercises}
1. Create an additional function for the dice program presented in today's lesson. The function should take as input a pointer to the Die array. For each face of the die, this function should print out the

Roll. Modify the makefile to link in the new function.
2.
3. Step through the program with gdb.

\section*{Shell Programming}

Athough you have been studying primarily the \(\mathrm{C}++\) language, you have also learned about the specifics and implications of this language in the Linux environment. Continuing in this vein, today's lesson addresses the
Linux shell and shell programming Linux shell and shell programming

Even though the shell is independent of the programming language you will be using, you should understan interact with the shell to to o anything. If you understand the basic mechanisms of the shell, you will find it a powerful tool that can help you in your work

\section*{This lesson touches on the following basics}
- What the shell is

What shells are available
Shell principles (redirection and backgrounding)
Shell commands construction (substitution and aliasing)
Environment variables
- Environment

\section*{What Is a Shell?}

\section*{Many people
is not the case.} The operating system is the software that runs the computer, talks to the
so on. The operating system is something that, basically, you never see.
When you see the computer prompt and enter commands in response, what you are dealing with is the shell. In tites past, the shell was known as a command-Iine interpreter
offen known as the CLI, and was named COMMAND.COM.
An interesting feature of UNIX and its derivatives including Linux is that shells are completely independent. A user is free to choose any one of a number of widely available sh
can choose to interact with one shell but write scripts in another.

\section*{Shells Available with Linu}

You can think of the shell as a program that sits in between the user and the operating system. It implements the
language that the user utilizes to control the operating system. The shell is started automatically when you log \(i\) to the syster The original shell available in UNXX systems was the Bourne shell ("sh"). Two later shells tha
were the C-shell ("csh") and the Korn shell ("ksh"). Each had unique and very useful features.

The Bourne shell was rewritten, and in a slight play on words the new version became known as the "Bourne again shell", or "bash." The C-Shell was rewritten and named theT-shell ("tcsh"). All three shells are available with Linux: bash, tcsh, and ksh. bash is probably the most frequently used shell, and in many installations it is

\section*{Shell Operation and Syntax Basics}

\section*{The shell performs a number of impotant function}

First of all, it is the program that enables you to interact with the operating system. The "command line" is the are input to the shell. The shell examines the command line, deternes if wat you have yped is the program, and if so then the shell starts that program. If the command is the name of a shell script then the appropriate shell is started and the script is fed to it as input.

All shell commands use the following general format:
command option1 option2 option3 ... optionn
This single line is known as the command line. A command line consists of a command name and one or more
\[
\begin{aligned}
& \text { options (or arguments). Whitespace on the command line is generally ignored. Command names and arguments } \\
& \text { are generally case-sensitive. You end the command by pressing the Enter (or Return) key. You can continue a }
\end{aligned}
\]
\[
\begin{aligned}
& \text { are generally case-sensitive. You end the command } \\
& \text { command on a new line with the backslash (1) key: }
\end{aligned}
\]
\[
\begin{aligned}
& \text { very-long-command-name long_option_1 long_option_2 } \\
& \text { long_option_3 } . . . \text { optionN }
\end{aligned}
\]
\[
\begin{aligned}
& \text { Iong_option_3 } \cdots \text { optionN } \\
& \text { You can concatenate more than one command on a line with a semicolon (;), like this: }
\end{aligned}
\]
nmand_1 option; command_2 option; command_3 option

Tip: You might often find the following sequence useful
clear; pwd; 1s
This will llear the terminal window, and
and \(a\) listof ofll he files in that directory.

\section*{Shell Features}

Shells have far too many features to discuss in great detail here. This section touches on the fundamentals behind shells in general. These features are present in almost all shells. The exact syntax may differ from on shell to another, and so specific examples are limited to the Bourne
shells does add some feature that may not be available in the others.

All the shells have complete and concise man page entries. In addition, the info file for bash is over 3,000 line
long. For more specific or advanced information, you should refer to this online documentation.

\section*{I/O Redirection}

When Linux programs are run, they have three I/O
input, the standard output, and the standard error.
By default the standard input is usually connected to the keyboard, and the standard output and standard ert
are usually connected to the screen. Because the shell starts the program you specify, it can reassign these are usually connected to the screen. Because the shell starts the program you specify, it can reassign these
default assignments in a process known as input redirection, output redirection, or more generically as \(/ / 0\) default assignments in a process known as as
redirection, before the program is started.

Suppose that you want to create a list of files in the /ussrinclude directory that includes other files. One way to Suppose that you want
this might be as follows
grep -l \#include /usr/include/*.h > IncludeList
The "\#" is a "wildcard" character. Wildcards are discussed in more detail later in this lesson. For now, just assume that grep pill check every file in / /usrfinclude whose name ends with ". h ". It will print the name of ever.
file in which it finds the string \#include.

The "»" is the output redirect character. It causes the shell to redirect the output of the grep command to a fil named Includel
look like this:
/usr/include/FlexLexer.
/usr/include/Fnlib.h
\(/\) usr/include/Fnlib_t
/usr/include/Imlib.h
/usr/include/ \(/\) mlib
include/Imlib_types.h
/usr/include/QwCluster.
/usr/include/QWSpriteField.h
To replace the /usr/include/ at the beginning of each filename, you can use the sed command
sed \({ }^{\text {s\#^/usr/include/\#+ \#' < IncludeList > ListIncludes }}\)
The command sed operates on data from the standard input. Using the """ input redirection character, you
redirect the standard input. Instead of reading from the keyboard, sed will read from the file Includel ist. As output redirection, the sed command is unaware that its standard input has been redirected (unless it explicitly output redirection, the sed command is unaware that its standard input has been red
checks for this). The output of the sed command is redirected to the file Listlncludes.
Note that the "‘" character does not redirect the standard error output, but only the standard output. Use >\& to
redirect the standard error as follows:
You can use redirection as follows
grep -1 \#include /usr/include/*. h > IncludeList
sort IncludeList > SortedList
Certainly there must be a better way than using two commands and a temporary file. The following is better
grep -l \#include /usr/include/*.h | sort > SortedList
The pipe character (|) strings two commands together, and connects (redirects) the standard output of the firs
one to the standard input of the second. A single command line can contain any number of pipes:
grep -l \#include /usr/include/*.h | sort | sed \ `s\#^/usr/include/\#+
SortedModifiedList

\section*{Variables}

Very often a program needs some specific information in order to work properly. For instance, the screen editors vi and emacs both need to know what type of terminal you are using. This information could be provided as a command-line option; however, it would be unnecessarily tedious to
require that you add this information every time you start the editor. In addition, there is more information tha does not change from one invocation to another beyond just the terminal type the editor requires.

Shells address this problem with environment variables. An environment variable is nothing more than name/string pair. The shell keeps a list of these variables and makes them available to any program.

There are in fact two types of variables: normal shell variables and environment variables. The distinction between these is subtle. Environment variables are also called "global" variables, while shell variables are als known as "local" variables. Global variables are passed on at the start of a new command or shell. To set a she variable in bash, use this syntax:

NAME=value
To set an environment variable, the syntax
export \(\mathrm{NAME=value}\)
If the variable includes whitespace or other special characters, the value can be enclosed in single or double quotes. Often it is useful to append to an already existing variable. You can do this as follows
export NAME="\$NAME new_string"
This appends "new_string" to the environment variable NAME.

\section*{Variables Used by the Shell}

解 of the more notable variables.
- DISPLAY-This variable is read by X programs to know where to display the
set to " ".0.", meaning that output will be displayed on the host computer monitor - PATH-A colon-separated list of directory names where the shell should search
you type any command name, the shell will look in these directories for a program by that name
- TERM-The type of terminal or terminal emulation. Programs such as editors must know the termina type so that they can send the proper commands to manipulate the screen and cursor.
- HOME-The user's home directory

\section*{Variables Set by the Shel}

Several variables are set by the shell and may be referenced by programs started from the shell. A partial list o important variables set by the shell includes the following
- SHELL-This is set by the shell and contains the full pathname of the current shell (that is, /bin/bash). - PWD-The current directory as set by the most recent cd command.

\section*{Backgrounding, Suspension, and Job Control}

Normally you type a command name on the command line and then wait for that command to finish. Suppose that the command you run will take a long time. Rather than wait, you could certainly open another xterm window and continue working there. Instead, however, you can make use of the multitasking nature of Linu You can run the command in the background with the special character " \(\downarrow\) " as follows
bash\# find
[1] 142
bash\#
This command will create a list of all files whose names end with ".cc". It will sort this list and place the output into a file named SortedList. The shell prints "[1] 142" and returns immediately, ready for another command. 7 output indicates that there is one task running in the background, and that the process I.D. (PtD) of this job
142. You can run more than one job in the background. In this case, this job is number 1 in the background queue

If you run a program from the command line and do not use the \& character, the shell waits fo the job to fini before prompting you for a new command. This job is said to be running in the foreground

If you run a job in the background, and then decide that you want to wait for it, you can "bring it to the foreground" with the command fg

If a job is running in the foreground and you want to pause it without killing it permanently, press Crrl+Z (th Control key and " \(z\) " at the same time). The job is now suspended. You can resume the job in the foreground with fg , or in the background with beg
Finally, if you want to kill the job permanently, you can use the kill commed Either kill it by its PID numbe (kill 142), or else by its place in your background queue (kill \(\% 1\) ).

To find out what jobs you are running in the background, use the jobs command.
Figure 23.1 demonstrates job control. The first "find" command is run in the background, and the "jobs" command indicates that it is running there. The next "find" command is run in the foreground and then suspended with CtrI+Z. Again, jobs shows this. Next, the suspended find command is resumed in the background, and this is verified with jobs. The fg command is used to bring the first background job into the
=

Figure 23.1 Backgrounding, suspension, and jobcontrol

\section*{Command Completion}
bash includes many shortcuts to reduce the amount of typing that you need to do. This is accomplished throug command completion and command substitution. The next section covers substitution and the wildcard nechanism.
Command completion refers to the ability of bash to "guess" what command or filename you are typing. Type the first few characters of a command, and instead of finishing the command, press the Tab key. If bash can uniquely ydentify the command, it will complete it for ,
time, and the candidate completions will be displayed.
The same mechanism works if you are specifying a filename as an argument to a commanc
For instance, type mo, press Tab, and the shell should beep. Press Tab again and you should see a list of commands that start with "mo." Type \(r\) and press Tab again. The shell should complete the command name more. Now press the spacebar, and then type /etc/in and press Tab. The shell beeps. Press Tab again and
files in /etc that start with "in" is displayed. Type e and then Tab. The filename inetd.conf is completed.

Command Substitution

\section*{Wildcard Substitution}

In many games of cards, a wildcard can be used to replace any other card. Similarly, in th
be used as a placeholder that the shell can replace with any other character or characters.
There are two important wildcard characters: the asterisk *) is matched by the shell by any sequence of zero or
more characters in a filename. The question mark ( (?) matches any single character. What does this mean? Perhaps you want to run the 1 s command on the current directory to see a list of all
filenames ending with ".".". You can simply enter is \(s\).h. The shell expands the wildcard character to a list of
filenames ending with ".h" before it invokes the \(1 s\) command. The full list of filenames is passed to the
command filenames
command
 "a??def*.xyz" to a list of filenames that statrt with " "" followed by any two characters, followed by "def,"
followed by zero or more characters and ending with "xyz".

Note: It is inportant to remember that wildards (and al
are expanded by the shell before the command is statred

\section*{String Substitutior}

\section*{bash allows substitution of specific character sequences. There are two styles of substitution}

You can specify a comma-separated list of strings in curly braces, and each one will be used in turn. For
```

bash\# ls a{b,c,de,fgh}z

```

The leading "a" and trailing " z " are combined with the " b " and then the " " " and then the "de" and then the
"fgh," which
You can specify ranges of characters with square brace
bash\# 1s a \([\mathrm{b}-\mathrm{h}] \mathrm{z}\)
abz acz adz ae
bash\#
abz
bash\#
The square braces can also specify specific characters:
bash\# 1 s a [bcde]z
abz acz adz ae
abz
bash\#
Finally, you can mix wildcards, curly braces, and square braces

\section*{Command Output Substitutior}
```

A Another form of substitution is

```
bash\# 1s -1 ‘find/usr/src -name Makefile -print’

This example will run the find command to locate all Makefiles in the /ustrssc directory tree. The list of files will
be presented on the command line to 15 , which will display the directory entries of these files.

\section*{Variable Substitution}

\section*{A final useful
command line}
```

bash\# echo \${HomE

```
bash\# echo \(\$\) \{f
/home/username

The value of the environment variable named "HOME" replaces the string \(\$\{\) HOME \(\mid\) in the command lin argument

\section*{Command History and Editing}
bash maintains a list of the commands you have typed in what's snown as a history list. The length of the list
depends on the environment, but is generally no longer than 500 command depends on the environment, but is generally no longer than 500 commands. bash will preserve this list from onc
session to the next. If you type the command history, the list of commands you have typed will be displayed. So for example consider the following segued histor, the lis


To invoke any previous comman
the preceding sequence, type \(: 2\).
Y-
You can also edit a previous command line before repeating it. Suppose that you type the command \(1 \mathrm{~s}-1 / \mathrm{TM}\)
 preceding command and proceeds with the substitution if possible. Emacs-style editing is also supported. Press
Ctrl+P and the previous command will be displayed. Use the arrow keys or Ctrlob (to move back a character) Aliases
You probably use certain commands or command sequences quite often. Yo
these, known as aliases. An alias is replaced by the shell with its definition.
For instance, the Is command shows you the names of files. With the command-line option -F, the type of th
file is also indicated ("**" for executable, " \("\) " for directory, and so on). You can create an alias to substitute file is also indicated (
commands as follows

The Is command with and without the - F argument display their results differently. When Is is aliased, then When you enter the command ls , the shell automatically substitutes the alias: \(\mathrm{ls}-\mathrm{F}\).
The command rm removes a file. With the - - option, it will first ask for confirmation. Many people find it usefu
to substitute the rm command top prevent accidentally removing file: alias rm="tm -1 "." The quote marks in this to substitute the rm command to prevent accidentally removing a file: alias rm="rm -ir . The quote marks in this
alias command are significant, because they group the whole right side of the equal sign as a single aryument alias command are significant, because they group the whole right side of the equal sign as a single argu.
the alias command. Without the quote marks, the alias command will try to interpert he \(-i\) as an option.

\section*{Shell Scripts}

Complex sequences of shell commands can be put into a file so that they can be repeated at any time. This i, features of standard languages including variables and control statementsts (which we have not yet discussed) All shell scripts begin with \#I/binsh. The first two characters tell the system that this is a script file, and the
/binsh starts the bash shell. After the shell is started, the remaining lines in the file are fed to it one at a time. Shell scripts can be simple one-line programs or complex and hundreds of lines. They are not as fast and
```

Shell scripts must have their "execut"" permission bit set. You can set this with the command: chmod a+x}\mathrm{ (fiename.

```
filename.

Variable
We've already discussed variables in shells. When you are running a script, some useful variables are already
defined:

\section*{- \({ }^{\text {S }}\) : \(0-\) The process ID of the shell program}

SO-The name of the script
\(\$ 1\) throunh so-The first nine command-line arguments passed to the script
\(\$ \$\) The number of command-line parameters passed to the script

\section*{Control Structures}
bash supports control statements analogous to many of those in the C and
different. These include statements such as if itthenelse, for, case, and while.
f!/bin/bash
```

3: MAX=9

# illustrate the iff-endif' control

exit 1 : SMAX or less arguments required"

```

```

c:e.
Mustrate the 'for-done' control
cise:
\#22:,
in
cosmol

```

```

con
M,

```

```

*)
ame
cin
3:

```

```

Matrate the unti-done' control
c,si-2t, fl:do
ci=s{(s+1)
si:

```
OUrour \(_{\text {m }}\)





Hex mixe



cose
case
case
detatlot case the the

winile example

\section*{Shell Startup File(s)}
 the user, upon Iogin an interactive bass sheli is
started and the script is fed to to ita line at a t ime
Upon login, the shell looks for a script named dect(profile. If that script is found, it is executed. If the file
.bash profile exists in the user's home directory, it is executed. If not, bass_login is checced.d If that is not found .bash_profilie exists if
ppofile is executed.

\section*{These statup fi
each login.}

\section*{Summary}

\section*{Redirection
Backgrounding
Bacer \\ Hastory
Command-line editing, wildcards, and substitution}

 provide an excellent reference for the synuax and features of the shell.
Q\&A
QI type a common Linux command but I get the error message "command not found"
A This is a common problem. You may have typed the command wrong, or it may not exist.


 "command not found."
is in your PATH environment list. The "". means "the current directory"
Workshop
The Workshop provides quiz questions to help solidify your understanding of the material covered and
 Quiz


What are the 3 standard /// ports available to programs?
What are the 3 forms of //O redid
command ine?
5. Are nvironment variables "Iocal" "o "global"? What about shell variables?
6. What bashe nenironment variable sests the command search path? Which one sis
prompt?
7. Name 2 command ine wildcard subssitution characters.
8. What needs to o be in a shell scrip file so that the shell kno

\title{

}

\section*{The /proc File System} Linux provides a /proc file system, which allows users to query information about the system and individual
running processes. /proc is not a traditional file system but a memory structure maintained by the kernel that provides system statistics. A user or process may obtain system and process information by issuing read anc write calls to the /proc file system
/proc is represented as a hierarchical file directory. The /proc directory contains a file entry for each running process, and each process has its own subdirectory named by its process ID. In addition to individual prod /proc contains system parameters and resource information. A program or system com
directories and files corresponding to a process to query information from the kernel.

You can query /proc at any given time to get information about the systen
Tip: The format of the data presented in the proc file system may change from kernel release to kernel relea

\section*{ocess State and Priority}

Before a discussion of process scheduling, you need to have a basic understanding of process state and proces priority.

During the lifetime of a process, it goes through many state transitions. The process state is the state, or mode, of the process at a given time. Process states can be found in the linux kernel header file sched.h and are define following.
- TASK_RUNNING-Process is waiting to run.
- TASK_INTERRUPTIBLE - Task is running and interruptible.
- TASK UNINTERRUPTIBLE-Task is running and is not interruptible
- TASK_STOPPED-The process is stopped, typically from receiving a signal.
- TASK_SWAPPING-The task is swapping.

Process priority is an integer assigned to a process or thread in which value is in a given range defined by the
scheduling policy. When a process is created, a static priority value is assigned. The value of the process scheduling policy. When a process is created, a static priority value is assigned. The value of the process
priority is the amount of time (in jiffies that the scheduler gives this process to run when it is allowed to priority ine arne the quantum.

There are two types of Linux processes, normal and real-time. Normal processes have a priority of 0 . A realtime process has priorities in the range of \(1-99\). Real-time processes are scheduled to have a higher priority th all the other non-real-time processes in the system. The priority of a process can only be altered using system calls such as niceo. If there is a real-time process ready to run, it will always run first. The nice system cal changes a proces
process priority.

\section*{Scheduling Algorithms}

The most precious resource in an operating system is CPU time. To efficiently use CPU time, operating systen use a concept called multiprocessing. Multiprocessing gives the user the impression that many processes are is always busy. This guarantees that the CPU is being used in the most efficient mand

The implementation of multiprocessing in an operating system consists of dividing its CPU time among many processes. The scheduler is the kernel process that determines which process to run next. The scheduler decid which process, of all the processes in the TASK_RUNNING state, to run. The critera, or rules, that determine which process to run next is called the scheduling algorithm.

The Linux scheduler uses process priority and process policy to determine which process to run next
There are three types of process policies: other, first in first out (FIFO), and round robin (RR). A normal proce has one type of policy, other. Real-time processes have two types of policies, round robin and first in first out. Other is implemented as
given an equal quantum.

First in first out schedules each runable process in the order that it is in on the run queue, and that order never changed. A FIFO process will run until it blocks on I/O or is pre-empted by a higher priority process

Round robin scheduling runs real-time processes in turn. The difference between a FIFO process and a rounc robin process is that a round robin \(p\)
placed at the end of the priority list.

Tip: A user must have root privilege to run real-time processe.

DO
DO run as root when running real Dor DON'T forget to have an intimate understanding of kernel data structures when reading about the froc file system.

\section*{Threads}

Ahread is defined as a sequence of instructions executed within the context of a process. Threads allow a threads.

Threads reduce overhead by sharing fund sections of a process. Remember that the environment of process consists of stack, data and text

The POSIX thread library package, called LinuxThreads, is an implementation of the POSIX 1003.1c thread package for Linux. POSIX 1003.1c is an API for multi-threaded programming standardized by IEEE as part o
the POSIX stand commonly referred to as pthreads. Pthreads and LinuxThreads are terms used interchangeably throughout today's lesson. LinuxThreads run on any Linux system with kernel 2.0 .0 or more recent and the glibc 2 library
Note: To compile with pthreads, you must include the pthread header file, finclude, and must link to the phtread
library. For example:
cc hello_world.c -o hello_world -1pthreads.

The thread API is the same as the pthreads, but the implementation is different from other operating system Other multithreaded operating systems, such as Windows 2000 , define a thread to be separate from a proces: LinuxThreads defines a thread as a context of executio



Single-Threaded

fine c-tcreate_h
lass Thread

int ceater) \(1 /\), crates the twreen
privates
Threads operator=(const Threads); / // disa11 ow copy
\begin{tabular}{|c|}
\hline \multirow[t]{2}{*}{\({ }_{\substack{\text { pethread_t } \\ \text { phread_atr_t } \\ \text { voided }}}\)} \\
\hline \\
\hline
\end{tabular}
*endif



Invot

1/ Listing 24.5 Producer/Consumer example
\#inclucue \(<\) iostreama h h
*include "toreate.h"
int data \(=0\);
void read_thread (void* param)
while (1) cout << "read:" << data << end1;
void write_thread (void param)
wiile
cout \(\lll{ }^{\text {write: }}\) "< data++ << end1;
nt main (int argc, char** argv)

thread. Createl);
threadal create ();
For (int i=0; \(1<10000 ; i+4)\)


 and

class synciock
publice
Sunciok)
syncococok)
(
(1)
Virtual int Acguiree() \(=0\);
virtual int Release()
private:
Synctocke operator=(const synctocke \()\) // disal1ow copy
: \#endif



Listing 24.8 nutex Producer/Consumer example
 \#include




hutext phutex \(=\) static_castexitext> (paran)
phatex- A acqui



for (int i - 0; \(1<100000 ; 1+\) )
. Destroy ();
Iock. Destrovor
 This sis nexample of reevtrant code.

\section*{Semaphore}




\(\qquad\)
int sem_trywait tseñ.t + se

int sen_getvalue (sen_t * seen,
int sen_destroy (sent. \(\%\) sem),
Condition Variable


ent Ofject Definition

\({ }^{1 \text { lass }}\) Event
fintor wex
\(\left.\begin{array}{l}\text { virtual int mait) } \\ \text { virtual } \\ \text { int } \\ \text { signa1t }\end{array}\right)=0\) or

\title{

}

Finclude epthread.hy

void dreatel
void
oestroy

\begin{tabular}{c} 
private: \\
Condavaci \\
\hline
\end{tabular}


3i: \#endif


, ining
 firistand onot





\section*{Chapter 25}

Interprocess Communication
Interprocess communication (IPC) is the means by which two processes communicate with one another.
Today's lesson discususes some of the methods available for processes to communicate with one another. Linux offers numerous methods for processes on the same system to communicate with one another. The
methods include, but are not timited to pipes, named pipes, and System \(V\) inererpocess communication. Lit


\section*{Background}

Before discussing today's topic, we will begin with the definition of an interface that is used for all today?:
classes. The Object class defined in Listing 25.1 implements a common interface for object creation and deletion
In order to implement a common interface for object creation and destruction, a pure virtual interface is define
Object All IPC objects defined in today's lesson are derived from the object lass Each of our objects inherits
sample code in Listing 25.1
\(\begin{array}{ll}\text { INPUT: } & \text { Listing } 25.1 \text { object Interface }\end{array}\)

Listing 25.1 object classs
\#if!defined (c_object_h)
\#define c_object_h
class object
public:
virtual
virtual int Create( \()=0 ;\)
virtual int Destroy () \(=0 ;\)
```

13: \#endif

```

Additionally, all our IPC methods use kernel data structures, resources, and handles. As such, it it no
 define the copy constructor as a private member function. Additionally, the copy constructor is no define the copy constructor as a private member function. Additionally, the copy constructor is not
implemented. By not implementing the copy constructor, other member tunctions cannot unitentionally call
the copy constructor; any attempt to call the copy constructor will result in a linker error.

\section*{Pipes}

A pipe is a simple method that allows one process to tatach its standard input to the standard ouput of anothe
process. A pipe is a a hal-duple communcation channel between a parent process and child proceess. Linux provides the pipeo system call. pipe returns two file descriptors as an array of file descriptors, fil descriptor 0 is the read end of
prototype is shown following:

\section*{pipe (int fal21) )}

A file descriptor is an integer used to refer to an oper
from a file, he file is identified by the file descriptor

Typical usage of a pipe is for a parent process to create a pipe and call fork( to create a child process
chid proceses descriptors that ereferencee to one ano pother. To completere the sestup for t a a half-duplex pipe for the contain fi
 descriptor. What remins is the write file descriptor for the child and the read file descriptor for the parent. Thi
sequence of events areates ald a pipe between two processe

\section*{Named Pipes (FIFOs)}

A named pipe is a method that allows independent processes to communicate. Named pipes are traditionally
called \(\operatorname{FIFOS}\) (first in first out) in the UNIX world. A named pipe is similar o o a pipe it is a half-duplex pipe
 between two processes. The difference between a pipe
the named pipe, the named pipe remains in the system,


After the pipe is created, data is passed between the two processes using standard read and write system call After the pipe is created, dat
fopen, fclose, fread, and fwite.
Listing 25.3 shows our implementation of a named pipe object
INPUT Listing 25.3 NamedPipe Class Definition
/ Listing 25.3 Namedpipe class
if
if lefined (c ne np \(h\) )
\#if !defined (c_
\#define c_nph
\#include <sys/types.h>
\#include <sys/stat.h>
Finclude <staio.h>
tinclude <object,
lass NamedPipe : public object
public:
NamedPipe();
N NamedPipe(),
int Create ();
int Create (char* name
int Create (char
int Destroy ()
int open();
int Close ();
int Read (char* buf, int len);
int Write (char* buf, int len);
private:
Namedpipe
bool
FILE*
init_;
fp-i
\#endif

ANALYSIS The Create function is overloaded to either take the name of a named pipe or use a default. After th named pipe is created, it must be opened. After the Namedipipe object is opened, the object can be read from o


When you are done using the Namedfipe, you should call Close to close the pipe handle and the Destroy function \(\stackrel{D}{\mathrm{Do}}\) Do use a named pipe
two urrelated proces.
System V IPC
AT\&T UNIX System V presented three added forms of interprocess communication for processes on the sam

\section*{Each of these methods of interprocess communication-messages, se es
described in this section along with a class demonstrating their usage.}

\section*{Key Creatio}

Each of the mentioned IPC methods uses a key to identify itself. Before delving into the details of System V
IPC, we will briefly discuss keys and develop a class to manage keys for the IPC methods. Creation of a key is the first step in using System V interprocess communication. A key is a non-negative
integer. In order for processes to access an IPC method, the method must use a key. A key is of type key Processes using System V interprocess communication must agree on the key. There are multiple ways fo
processes to agree on the same key value. In this section, we will call our processes clien and server.
 y the client, a file for example. msgee, semget, and smget are system V IPC routines that create instancas - The clientands. server can agree on a key and share the key in a common header file.
- The client and server can agree on a a aathname and an ID and call the system functio

The client and server can agree on a pathname and an ID and call the system function
a key. fook is asystem call that creates a key based on the input filename and process ID .
key_t Get (void);
private:
Keya ope
bool \(\left.\begin{array}{l}\text { init } \\ \text { key_t } \\ \text { key_i }\end{array}\right]\)
tena
: \#endif

ANaľsIs The first tep in obtaining a key is creating a key object and calling the Create member function. \(_{\text {Crate }}\). Create is overloaded, so you can pass a predecined key or have a key created by passing in
After the key is created, you can obtain the key value by calling the Get member function.

\section*{IPC Permission Structur}

Each method for IPC contains an ipc_perm structure. The ipc_ perm str
for the IPC object. Listing 25.5 Shows the Linux ipc_perm structure.
INPUT. Listing 25.5 Linux ipc_perm Structure


\section*{ipcs and ipcrm Commands}

System V interprocess communication methods are implemented in the kernel．If a process exits without
cleaning up its IPC methods，the method will remain in the kernel．Theips and ipcrm commands are used to display the status of IPC methods and clean up any residual IPC resources
dite
The ipcs command is used to obtain the status of all IPC objects．Following is the result of using the ipc command－two shared memory segments on my system，IDs 78592 and 78593.


The iprrm command is used to remove an IPC object from the system．Using the example in the preceding iperm command is used to remove an IPC object from the system．Using the example in the preca demonstrates the removal of shared memory id 78593 ．
resource deleted

Now that we have removed shared memory id 78593 ，we will rerun the ipsc command to verify that the shared memory method has been removed．Following is the results of rerunning the ipcs command．
\begin{tabular}{|c|c|c|c|c|c|}
\hline \[
\begin{aligned}
& \text { key } \\
& 0 \times 00001
\end{aligned}
\] & \[
\begin{gathered}
\text { shmid } \\
162 \mathrm{e} 78593
\end{gathered}
\] & \begin{tabular}{l}
owner \\
pau．
\end{tabular} & \[
\begin{gathered}
\text { perms } \\
666
\end{gathered}
\] & bytes
\[
64
\] & \[
\begin{aligned}
& \text { nattch } \\
& 0
\end{aligned}
\] \\
\hline key & Semaphore semid & \begin{tabular}{l}
Arrays－－ \\
owner
\end{tabular} & perms & nsems & status \\
\hline key & Message \(Q\) msqid & \begin{tabular}{l}
ueues－－－ \\
owner
\end{tabular} & perms & used－by & mess \\
\hline
\end{tabular}

As you can see the shared memory segment is removed．You may also note that there are no message queues o semaphores on the system．

\section*{Message Queues}

Message quueues are a method for processes to send messages（data）to one another．A message queue is a link
list of messages list of messages mainained by he kernel．Messages are added and removed from a AIFO queue by processes． Each message consists o
shows a message queue．

For each message queue，the kernel maintains msqiddds structure defined in the standard kernel heade ＜linux／msg．h＞．Listing 25.6 shows the Linux msqiidds structure．The msqid＿ds structure maintains pointers to th messages，number of messages on the queue and other related information about the messa
process id of the last process that sent a message and last process that received a message．
\(\square\)

Input．Listing 25.6 Linux msqid＿ds Structure


\section*{Analysis} own in Listing 25．6．

Before a message queue can be used，it must be created．The msggeto system call creates a new message queue，
msggetor returns the message queue identifer．The message queue identifer is an int．This identifer is used for the msgget（）returns the message queue identifer．The message queue identifer is an int．This identifer is used for th remaining message functions．Mssget）returns an integer mssqid value，and this mssqid value is passed into th

＊Get messages queue．
The key parameter is the key discussed in the previous section．The mssflg parameter is access permissions fo the message queue bitwise ORed with the following flags：
－IPC CREAT Create the queue if it does not exist，
信
Messages are posted to the message queue using the msgsndo system call．
 msgsnd
int
msssndo takes as arguments a type struct msgbur．msgbuf is defined in 〈linux／msg．h》，size，and mseflag．The mssbu type is a model definition for the actual message passed to msgbuf．Listing 25.7 shows the Linux msgbuf INPUT．Listing 25.7 Linux msgbuf Structur
```

lunct msgbuf
long mtype; (* type of message */

```

An IPC programmer typically redefines the message to conform to the standard mssbuf definition in msh．h．For
example，a message could be defined as shown in Listing 25.8 ． INPUT
INPUT．Listing 25.8 Sample User ipc＿message msgbuf Structur
```

typedef struct ipc_message
long mtype;
payload[8]

```

This message is defined as a message header，mype．The message type must be a non－negative integer，a head may be used as a flag to denote how the payload data should be interpreted．
mes．

Messages are read from a message queue by calling the msgrev0 system call．The parameter msgytpe allows yo exercise some granularity when reading messages．Msstype values and definitions are as follows．
＊Receive message from message queue．＊／
extern int msgrcv＿－\({ }^{\text {（ }}\)（int＿msqid，void＊＿msgp，size＿t＿＿msgsz，

The values for mseflg are defined as follows:
- If msstype is zero, the first message on the queue is returne.If msgtype is greater than zero, the message with an equal message type is returnedvalue of message is returned
licatio is done using the message queve, thesystem call msgetl is used to remove a message queue from the system.
* Message queue control operation.
Listing 25.9 defines a Message object. The Message object encapsulates all the message system calls defined in thi
\(\begin{aligned} & \text { Listing } \\ & \text { section. }\end{aligned}\)
Input. Listing 25.9 Message Obje
```

\#if Listing 25.9 Message clas
\#define c_message_h
\#include <sys/types.h
\#include <sys/ipc.h>
\#include <sys/msg.h>
\#include <object.h
class Message : public Object
public:
Message (Key k);
int Create(void);
int Create(int flags)
int Destroy(void);
virtual int Read(char* buf, int len, long type);
private:

```

```

typedef struct
long type;
char *payload;
bool init_;
Key key_;
38: \#endi:

```
\begin{tabular}{l}
\(36:\) \\
37 \\
37 \\
\hline
\end{tabular}
37
\(38 ;\)

ANALYSIS The Message object defined in Listing 25.9 contains the standard Create and Destroy member function The Message object defined in Listing 25.9 contains the standard Create and Destroy member function.
defined for all our objects. Note that the constructor takes a Key object. The client and server processes create a object by either passing in the same key value or using the same filename and ID; this way they will create a
Message object that references the same message queue. After the Message object has been created, the client and server processes send and receive messages using the Read and Write member functions.

The Message object encapsulates the message system calls described in this section. Sample code for this section
contains a demonstration of the Message abject: the program mestest uses the Message object to send and receive messages.

\section*{Semaphores}

Semaphores are a method used to synchronize operations between two processes. Semaphores are not a method ass data between processes, like messages, but rather a means for two processes to synchronize access to some shared resource.

For a process to obtain access to a resource, the process will test the value of the semaphore. If the value is great than zero, the process decrements the value by one. If the value is zero, the process blocks until the value is grea than one. When the process is done accessing the resource, it releases it by incrementing the semaphore value by one.

Note: A semaphore whose values are limited to oro and one is called a binary semaphore. System V semaphores are
not timited to zero and one but can be any non-negative integer.
Note: A semaphore whose values are inimited to zero and one
not limited to zero and one but can be any non-negative intege
described in the preceding section, the process of acquiring a semaphore requires two steps; first test the value, and then decrement the value. For this method to work properly among multiple processes, the test and decremer operations must be atomic. In order to make the test and decrement operations atomic, the implementation of the eperations is done in the kernel

For each semaphore in the kernel, the kernel maintains a semid ds structure. semid dd is defined in \&linuvem h> Listing 25.10 shows the Linux semid_ds structure.

INPUT. Listing 25.10 Linux semid_ds Structure Definition
```

    /* One semid dat
    ```

```

    Struct sem *sem_base; /* ptr to first semaphore in array */
    *)
    struct sem_queue **sem_pending_last; /* last pending operation */
    * /* undo requests on this array */
    ```
ANalrsis Before a semaphore can be used, it must be created. The semgett) system call creates a new semaphore
semget( returns a semaphore identifier or -1 if a semaphore cannot be created. The semaphore identifier is an int.
This identifer is used for the remaining semaphore functions. The function prototype for semget is shown in the
following lines:
/* Get semaphore.
Note the value of nsems. semget can create multiple semaphores in a single call. If semget is used to create multiple
semaphores, the semaphores are numbered beginning with 0 .
After a semaphore is created with semget), a semaphore is accessed using the semop0 system call.
/* Operate on semaphore.

as an argument a sembuf structure, which is defined in Listing 25

INPUT. Listing 25.11 Linux sembuf Structure Definition
```

* semop system calls takes an array of these. */
ruct sembuf {
sem_num;
unsigned short sem_num; /* semaphore index in a
/* semaphore index in array */
short sem_flg; /* operation flags */

```

For example, a member function that acquires the semaphore is implemented this way using the sembuf and emop system calls. Listing 25.12 shows the implementation of the Acquire member function, which is semop system calls. Listing 25.12 shows 1 .
implemented using the semop system call.

ANalrsis Listing 25.12 Semaphore::Acquire Function Demonstrating semop System Call
```

1:int Semaphore::Acquire(
2:{
3:// test the semaphore and decremen
4:static struct sembuf lock[1]
5:{
6:{0, 1, IPC_NOWAIT
8:
8:
10:
11:ir
12:
14:}

```

Semaphores have one more system call, semct(). semctl() is used to remove a semaphore from the system
/* Semaphore control operation. */
extern int semctl __P ((int __semid, int __semnum, int __cmd,
ext
Listing 25.13 defines a semaphore class that illustrates the uses of the semaphore system calls defined in this section

\section*{InPUT Listing 25.13 Semaphore Class Definitio}

\begin{abstract}
// Listing 25-13 Semaphore clas
define
include <sys/sem h>
\#include <key.h>
\#include "object
11: public
12: Semaphore (Key
15: int Crea)
16: int Create (int flags
17: int Destroy (void).
18:
19: int Queryvalue();
20: 21: int Acquire(void);
22: int Release(void);
23:
24: private:
25: Semaphore\& operator=(const Semaphoré); // disallow copy
27: bool init_;
28: Key key
32:\#endif
\end{abstract}
9: c1
10:
\(13:\)
\(14:\)
\(14:\)
\(29:\)
\(30:\} ;\)
31:

ANALYSIS The Semaphore object defined in Listing 25.13 contains the standard Create and Destroy membe functions defined for all our objects. Once again, note that the constructor takes a Key object; the client and server processes create a Key object by either passing in the same key value or using the same filename and ID This way they will create a Semaphore object that references the same semaphore. After the semaphore has been created, the client and server processes access the semaphore by calling the Acquire and Release membe functions.

The Semaphore object encapsulates the semaphore system calls described in this section. The semtest program contained on the CD-ROM demonstrates the use of the Semaphore object.

\section*{Shared Memory}

Shared memory is a method that allows two processes to share data directly. Refer back to the earlier discussic of message queues, pipes, and named pipes. Each of these interprocess communication methods allows multip processes to pass data between each other. A performance issue with these methods of interprocess communication is that the data copied between the processes is copied from the source process into the kernel, and then from the kernel into the destination process. These copies are expensive.
Shared memory allows multiple processes to get a pointer to an area of memory that is shared between the processes. Having a pointer to a segment of memory eliminates the expensive copies in and out of the kernel

Shared memory is not without its own drawbacks. Even though each of the processes have access to a segment of memory, the accesses to that memory must be synchronized

Process synchronization to shared memory is typically implemented using semaphores, as described in the previous section.

Before shared memory can be used, it must be created, the shmget) creates a shared memory sectio
/* Get shared memory segment. */
extern int shmget __P ((key_t __key, int __size, int __shmflg));
Each shared memory segment is maintained in the kernel by a shmid_ds structure. shmid_ds is defined in <kernel/shm.h>. Listing 25.14 shows the Linux shmid_ds structure.

INPUT. Listing 25.14 Linux shmid_ds Structure Definitio


The shmget) system call opens or creates a shared memory segment but does not provide access to the shared memory. In order for a process to access shared memory, the process must attach to shared memory. The shmat system call returns a pointer to the shared memory segment:
/* Attach shared memory segment. */ extern void *shmat _P ((int _shmid, _const void *__shmaddr,
[ic:ccc] int shmflg)) : [ic:ccc] int __shmflg));
When a process is finished with shared memory, it must detach the shared memory segment from the process. This is done by calling shmdt)
/* Detach shared memory segment. */
* Detach shared memory segment. \({ }^{* /}\) *_shmaddr))
extern int shmdt __P ((__const void

After all processes detach the shared memory, the shared memory segment must still be deleted. This is accomplished by calling the shmct1) system call
 Listing 25.15 defines a SharedMemory object that encapsulates all the shared memory system calls defined in thi section.
isting 25.15 SharedMemory Class Definition
```

// Listing 25.15 Shared memory classs
\#if !defined(c_smem_h)
\#define c_smem_h
\#include <sys/shm.h>
\#include <string.h>
\#include <key.h>
hinclude <object.h>
class SharedMemory : public Object
public:
SharedMemory (Key k);
~SharedMemory (void);
int Create(void);
int Create(int flags, int size);
int Destroy(void);
Char* Attach(void);
oid Detach(void);
int Read(char* buf, int len, int offset);
int Write(char* buf, int len, int offset);
private:
SharedMemory\& operator=(const SharedMemory\&); // disallow copy
bool init_;
Key
key_;
char* shmaddr
39: \#endif

```
37: \}
38:

ANALYSIS The SharedMemory object defined in Listing 25.15 contains the standard Create and Destroy member functions defined for all our objects. Note that the constructor takes a Key object. The client and server functions defined for all our objects. Note that the constructor takes a Key object. The client and server
processes create a Key object by either passing in the same key value or using the same filename and ID. Th processes create a key object by either passing in the same key value or using the same filename and ID. This
way they will create a SharedMemory object that references the same shared memory segment. After the shared memory object has been created, the client and server processes can read and write to shared memory directly using the pointer returned by the shmat system call. It is important to note that access to shared memory should be synchronized using a semaphore.

Linux also supports the ipc() system call. ipc is a centralized system call that is used to implement the System V IPC calls under Linux. ipco is only implemented under Linux and is not portable across other UNIX systems. Any of the System V IPC system calls mentioned today can be called using the ipc system call. Following is th ipc function prototype:
int ipc(unsigned int call, int first, int second, int
third, void *ptr, long fifth);

\section*{Summary}

Today we have defined various methods available to the Linux programmer for interprocess communication. Each of the sections of this lesson discusses the use of the methods as well as the details for using the methods

In the first half of this lesson, we discussed some of the first methods available to UNIX programs for interprocess communication, pipes and named pipes.

In the next half of the lesson, we discussed System V interprocess communication, messages, semaphores and shared memory and developed some C ++ object to use these IPC methods.

Q\&A
Q What advantage is there for using message queues over a pipe? A named pipe?
A Message queues are a full-duplex channel. Pipes and named pipes are half-duplex
Q Why is it important to synchronize access to shared memory?
A Because processes are running independently of one another, each of the processes could write to shared memory and corrupt the data read, or written, by the other process.

\section*{Workshop}

The Workshop provides quiz questions to help you solidify your understanding of the material covered and exercises to provide you with experience in using what you've learned. Try to answer the quiz and exercise exercises to proce checking the you understand the answers before continuing to tomorrow's lesson.

\section*{Quiz}
1. List the three routines used to create System V interprocess communication methods.
2. What signal is thrown if the read and write ends of a pipe are set up?
3. Why is shared memory faster than messages?
4. What is a binary semaphore?

\section*{Exercises}
1. Implement a client/server program where the client and server share data using the SharedMemory clas and synchronize access to shared memory using the Semaphore class.
2. Using pipes, set up a full-duplex communication between a parent and child process.
3. Extend the NamedPipe class to be able to open non-blocking pipes and a Read member function that does not block if there is no data available.

\section*{Chapter 26 \\ GUI Programming}

In the previous lessons, you have seen how the C++ language is constructed, its syntax and data types, and how you can create objects in it that can interact to produce useful and working programs on the Linux Operating System.

In this lesson, you will take these concepts further and see how you can use classes and objects to represent rea objects graphically so you can interact with them and your program with a mouse or other graphical input device

Today you will learn:
- What a Graphical User Interface (GUI) is, and why it is so important for Linux
- How GUIs represent familiar objects as widgets with which you can interact
- The history of GUIs on UNIX
- Two new and free GUIs or desktops for Linux
- How to write basic Linux GUI applications using the development toolkits that come with the GUI distributions
- How to write simple cross-platform C++ GUI applications using the wxWindows toolkit

\section*{The Linux Desktop}

Until recently, one of the most often heard complaints against Linux (as well as all other UNIXes) as an operating system was that it had only a command-line interface in the form of a "shell" or simple text on the screen. This means that you interact with the computer by typing commands as text and you get the output in the same way: as text. Not even pretty text either: awful chunky single-spaced typewriter text with no formatting, as likely as not.

This is fine for computer freaks and developers, and it is also fine for server systems that connect to other more user-friendly client applications that the real users sit at, but it is a massive obstacle in getting Linux where it belongs: onto the desktop, literally. Your average computer user wants results from the thing, not seemingly endless argument and complaints about bad inputs.

Things began to pick up in the early to mid-1980s when many applications appeared with an intermediate step in user interfaces between the command-line interface and the Graphical User Interface (GUI, usually pronounced "gooey").sThis was the non-graphical menu-based interface, which let you interact by using a mouse rather than by having to type in keyboard commands. Even so, the interface was still chunky and ugly because pictures and boxes were drawn using the line-characters in the extended ASCII character set. Not as nice as the GUIs of today.

Nowadays, of course, most important operating systems, such as all flavors of Windows, the Macintosh, and OS/2, provide a true GUI where the objects with which the user interacts are drawn cleanly and precisely pixel-by-pixel.

Applications typically use the elements of the GUI that come with the operating system and add their own GUI elements and ideas. A GUI sometimes uses one or more metaphors for objects familiar in real life: desktops, windows, or the physical layout of a building.

Note: When this lesson uses the word "window" or "windows" with a lower case "w," we are talking about the generic bounded object on a screen in which we view data; if we are talking specifically about Microsoft's family of operating systems, we use the capitalized "Windows."

Elements of a GUI include such things as windows, menus, buttons, check boxes, sliders, gauges, scroll bars, icons, emotional-icons that change their nature in real time as you traverse the file system, wizards, mice, joysticks, and dozens of other things. Multimedia devices now form part of most GUIs, and sound, voice, motion video, and virtual reality interfaces seem likely to become a standard part of the GUI for many applications.

A system's graphical user interface along with its input devices is sometimes referred to as its "look and feel." Another thing that tends to characterize GUI application is that they are typically event driven rather than procedural. This means that a GUI application typically sits and waits for you to ask it to do something rather than starting at the beginning and following a logical path to some end point, and then terminating. Note that an event driven program can still have procedural things happening inside it. An example of this is a GUI Integrated Development Environment which will wait for you to ask it to build the program (an event) and will then run make for you to do the build (procedural).

The GUIs familiar to most people in these modern operating systems and their applications originated at the Xerox Palo Alto Research Laboratory in the late 1970s. Apple used it in their first Macintosh computers; later, Microsoft used many of the same ideas in their first version of the Windows operating system for IBMcompatible PCs. Yet this is not true of Linux: Linux does not have a native GUI, unlike operating systems like Windows NT and the Apple Macintosh where the GUI is the operating system, at least in part.

So the thing that really sets Linux apart from other modern operating systems is that even though you can create a GUI for it, using a graphics card, the capabilities of the monitor, and software to drive the whole thing, the GUI is not actually part of the operating system itself. Rather, it is layered on top of it and disguises much of th operating system's command-line complexity from the user. This is generally reckoned to be A Good Thing by most users.

Although this may seem like a trivial and even pedantic difference between it and the operating systems that do have native GUI functionality, it is a difference that provides Linux-and you-with enormous scope for flexibility.

Linux is new, but its antecedents, the multifarious flavors of UNIX too numerous to mention, have been around for decades, and people have indeed attempted in the past to use the standard graphics capabilities of UNIX in the form of the X Windowing system to provide GUI functionality. But X programming is complex and difficult (see the sidebar "The X Protocol and More Recent Developments"), and even toolkits like Motif and it free clone LessTif, which are designed to hide much of the awful detail of X programming, barely succeed in making a bad job any better.

\section*{The X Protocol and More Recent Developments}

You have probably noticed in this lesson that we often mention " X " in connection with graphics on UNIX.
The X Org, a non-profit consortium made up of members, developed the X Protocol in the mid-1980s to answer the need for a network-transparent graphical user interface primarily for the UNIX operating system.

X provides for the display and management of graphical information, in much in the same way as Microsoft' Windows Graphics Device Interface (GDI) and IBM's Presentation Manager.

The chief difference is in the design of the X Protocol itself. Microsoft Windows and IBM's Presentation Manager simply display graphical information on the PC on which they are running, whereas the X Protocol distributes the processing of applications by specifying a client-server relationship at the application level.

The what to do part of the application is called an X client and is logically and often physically separated from the how to do part, the display, called the X server. X clients typically run on a remote machine that has excess computing power and displays on an X server. This is a true client-server relationship between an application and its display, and has all the attendant advantages of this relationship.

To achieve this separation, the application ( X client) is divorced from the display ( X server), and the two communicate using a socket-based asynchronous protocol that operates transparently over a network.

X further provides a common windowing system by specifying both a device-dependent and an independent layer. Essentially, the X Protocol hides the quirks of the operating system and the underlying hardware from the applications that use them. In theory, different hardware sets will present a common software interface and still have very different internals.

This is all very well in theory, but the reality of it is that the Xlib API, written in C, incidentally, is extremely complicated, and X programming is considered to be something of a black art.

People have written various layers on top of Xlib in recent times to try to simplify X programming, but these attempts, most notably Motif and its free clone, LessTif, are only partly successful (the API is still hideous) and, in any case, the "look and feel" is now starting to look dated.

Much more recently we have seen developments like GTK++ (a C API layered on top of GLib and GDK, which itself is a layer on top of Xlib), wxWindows (a C++ layer on top of GTK++) and TrollTech's Qt library (a C++ layer on top of Xlib on UNIX and on top of the GDI on Windows).

All this layering might begin to give you a headache, but it really is no cause for concern: Even the more familiar C++ graphics libraries like Microsoft's MFC and Borland's C++ Builder classes are little more than C++ layers on top of the native Windows C API, and the Graphics Device Interface (GDI). Indeed, GDK++, wxWindows, and Qt in particular seem to have at last got it right. The two latter libraries provide for platform independent source trees; that is, you write it once, compile it for your platform using the native library toolkit, and run it. This is somewhat like Java, except that the compilation is done before you send the code out and not when the user runs it.

The Qt and wxWindows examples later in this lesson will compile and run on both Linux and Windows platforms.

But here is a dilemma: we have a first-class, powerful, robust, and free operating system that your average user won't want to use because its GUI looks like what it is: tired, dated, and thoroughly unpleasant.

Fortunately for Linux and for us, the programmers who will be called upon to write the new applications, two new projects have put Linux firmly on the desktop with their incredibly successful attempts to turn the Linux GUI into a superb and eminently usable interface.

These two projects are the GNOME project (GNU Network Object Model Environment) from GNU, and the K Desktop Environment (KDE) from the KDE Project. Both of these are free software and released under the GPL. KDE does have one or two licensing issues to do with its underlying TrollTech Qt graphics library, but these seem to have been resolved in the main; however, this does appear to deter hard-core GNU aficionados from embracing the KDE as genuinely free software.

This lesson examines the GNOME and the KDE and creates some basic applications to show how simple writing GUI-driven software for these environments can be.

\section*{What Is the GNOME?}

The GNOME is the GNU Network Object Model Environment, and, of course, the GUI desktop from the GNU Project.

To quote from the original announcement from usenet newsgroup comp.os.linux.announce, GNOME is intended to be
"a free and complete set of user-friendly applications and desktop tools, similar to CDE and KDE but based entirely on free software."

And indeed, the GNOME is everything you would expect from a modern programming environment.In this respect, it is roughly the same as Common Desktop Environment (CDE), Win32, NextStep, and KDE.

The big difference is that, in contrast to the examples just mentioned, every single component of GNOME is free and released under either the GPL or the LGPL. Not only that, but the GNOME is extremely flexible compared to most desktop environments. As an added bonus for the user, you can easily customize it to suit your particular needs.

The GNOME uses the Object Management Group's Common Object Request Broker Architecture (CORBA) to allow software components to interoperate seamlessly, regardless of the language in which they are implemented, or even on what machine they are running. Not only this, but the GNOME developers are working hard on developing an object model called Bonobo, based on CORBA, which is similar to Microsoft's Object Linking and Embedding, v2 (OLE2).

Bonobo will allow programmers to export and import resources as components and will, for example, allow users to use whatever editor they like in their development environment, provided that their editor is supported via a CORBA standardized editor interface.

The GNOME uses the Gimp Tool Kit (GTK+++) as the graphics toolkit for all graphical applications. GTK+ has a lot of excellent features and sprang from the development of the GNU Image Processing Program (GIMP), which deserves a book all by itself. Because GTK++ underpins GNOME, they both use Imlib, an image library for the X Window System, which supports multiple image formats, from XPM to PNG, and multiple bit-depths, from 24-bit True Color down to 1-bit Black and White, all transparently to the programmer.

GNOME applications are session-aware: If you shut down, say, the GNOME word processor and then start it back up again, it will open up the document that you had open before, and put your cursor back at the same place. This is made possible by the X Session Management system, as implemented in the GNOME Session Manager.

The GNOME also supports the Uniforum standard internationalization and localization methods, allowing support for new languages to be added without even requiring a recompile of the application.

In truth the GNOME can be a bit of a pain to install owing to its many dependencies and so on, but Linux vendors now put it in the standard distribution. Red Hat Linux 6.1, for example, comes with GNOME, and you can install it as your default desktop at installation time

\section*{Getting the GNOME and Other GNOME Resources}

The GNOME home page is at http://www.gnome.org/. If you are impatient about getting the software, you car go directly to the Linux download page at http://www.gnome.org/start. You will find the GNOME FAQ at http://www.gnome.org/gnomefaq/ and a whole load of other useful information at http://www.gnome.org/gdp and http://www.gnome.org/mailing-lists/index.shtml has a comprehensive list of GNOME-related mailing lists
you are interested in developing GNOME itself or applications that use GNOME, it will be worth your time to look at http://developer.gnome.org/ where you will find lots of extra information and resources

The current release of GNOME is 1.0 .53

\section*{Getting GTK ++ and Other GTK ++ Resources}

The GTK++ home page is at http://www.gtk.org/ and you can download the latest release from ftp://ftp.gtk.org/pub/gtk/v1.2/. The GTK++ home page has a lot of links to other useful sites and essential documentation.

You can get GTK-, a free C++ wrapper library around GTK++, from http://gtkmm.sourceforge.net/. The distribution also includes gnome - a C++ wrapper for the GNOME libraries. At the time of writing, Januar 2000, gnome- seemed to be undocumented, and GTK - somewhat unstable. It may be worth your while checking out these projects again.

For GTK++ 1.1.7 (current version) and higher, you will need LibSigc++ 0.8.5 or higher. You can get this fron http://www.ece.ucdavis.edu/~kenelson/libsigc++/.

\section*{What Is the KDE?}

The KDE used to be "the Kool Desktop Environment", but, thankfully they have dropped the "Kool" from th name.

The KDE is a network-transparent modern desktop environment for UNIX workstations. It admirably fulfills the need for an easy-to-use desktop for UNIX workstations, similar to the desktop environments found under he MacOS and Windows.

With the advent of KDE, an easy-to-use modern desktop environment is now available for UNIX that rival anything out there in the commercial market. Like the GNOME, KDE can be a real pain to install, but it to comes with several vendors' releases of Linux. Red Hat Linux 6.1 also comes with KDE 1.1.2, and you can install it as your default desktop instead of the GNOME. In fact, you can configure your machine to switch you into either a GNOME session or KDE session each time you \(\log\) or

Note: Although installing the standard distributions of GNOME and KDE from the CD is an easy and tempting ay to get it up and running, one drawback you may find is that some of the latest releases of useful software, lik he CD. This is inevitable, as the CD releases have to be frozen at some point

In combination with Linux, KDE provides a completely free and open computing platform available to anyon free of charge including its source code for anyone to modify to suit their own ends.

Although this is generally true, there are some issues surrounding the Qt libraries that interface the KDE to the depths of the Xlib interface, and for many purists this means that KDE is not free software. Certainly you do no have to pay to use it or modify it or distribute it, but you cannot sell software you write using the KDE librarie unless you buy a professional license from TrollTech

Whatever the purists' objections, and although there will always be room for improvement, the KDE developers, a loosely coupled group of programmers linked together by the Internet, have delivered a viable alternative to some of the more commonly found and commercial operating systems/desktop combinations yo
a can get.

To the user, KDE offers, among other thing
An integrated help system allowing for convenient, consistent access to help on the use of the KDE desktop and its applications
Consistent look and feel of all KDE applications.
Standardized menu and toolbars, key-bindings, color-schemes, and so on.
Internationalization: KDE is available in more than 25 languages
A vast number of useful KDE applications.
- A vast number of useful KDE applications

In truth the default look and feel of KDE is remarkably like the Window 95 UI, something I believe not to be entirely coincidental.

\section*{Getting KDE and Other KDE Resources}

The KDE home page is at http://www.kde.org/, and you can download it from any one of the mirror sites listed at http://www.kde.org/mirrors.html/. Links from this page cover topics ranging from the KDE mailing list archives, to KDE documentation, to KDE T-shirts and cuddly toys (I am not making this up)

If you are thinking of developing KDE applications, you will need the free release of TrollTech's Qt libraries; you can get this from their Web site at http'//www troll nol. The latest version of Qt is 2.02 Note that the fre version covers you only for X -applications. If you want to compile and link your code on Windows platforms. or you want to sell your programs, you need to buy their Professional package.

KDE is currently running at version 1.1.2, but version 2.0 is promised for spring 2000
You can read about and download KDevelop by visiting http://www.kdevelop.org/. It is currently at versior 1.1beta1, but this number seems to be changing almost daily. By the time you are reading this book, it will hav been upgraded several times, so it will be worth checking the home page for the latest release.

\section*{Programming in C++ on the Linux Desktop} the desktop. Applications conform to a set format, or application framework; this is a term that will be familiar to anyone who has programmed using Microsoft's MFC or Borland's C++ Builder.

We will see a third framework in the wxGTK library. This library is similar to the Qt framework, and, indeed, the Microsoft and Borland frameworks in that it implements a Document/View framework where the Documen conceptually represents data objects and the View conceptually represents the user's view of those data.

\section*{Widgets}

In computing terms, a widget is an element of a graphical user interface (GUI) that displays information or provides a specific way for a user to interact with the operating system and application programs. Widgets include icons, pull-down menus, buttons, selection boxes, progress indicators, on-off checkmarks, scroll bars, windows, window edges (that let you resize the window), toggle buttons, forms, and many other devices for displaying information and for inviting, accepting, and responding to user actions.

In programming, a widget also means the small program that is written in order to describe what a particular widget looks like, how it behaves, and how it interacts with the user. Most operating systems include a set of predefined widgets that a programmer can incorporate in an application, specifying how it is to behave. You can create new widgets by extending existing ones, or you can write your own from scratch.

The term was apparently applied first in UNIX-based operating systems and the X Window System. In object oriented programming (OOP), each type of widget is defined as a class (or a subclass under a broad generic widget class) and is always associated with a particular window. In the AIX Enhanced X-Window Toolkit, for instance, a widget is the fundamental data type.

\section*{Fundamentals of GNOME Programming}

The GNOME is the GNU desktop and unlike the KDE is written entirely in C.
It is based upon the excellent GTK++ library, which is itself a C wrapper library around the GDK, which is a wrapper around the native Xlib libraries (see the sidebar "The X Protocol and More Recent Developments").

Although this is a book on how to learn \(\mathrm{C}++\), it is worth taking a brief look at the GTK++ and GNOME C APIs to illustrate some useful programming concepts that follow the more relevant \(\mathrm{C}++\) wrapper libraries examined later in this lesson.

First, because the GNOME uses GTK++ as its graphics engine, we will take a brief look at GTK++. GTK++ is a library for creating graphical user interfaces. It is licensed using the LGPL license, so you can develop open software, free software, or even commercial software using GTK++ without having to spend anything for licenses or royalties.

It is called the GIMP Toolkit because it was originally written for developing the General Image Manipulation Program (GIMP, another truly excellent free-software application), but GTK++ has now been used in a large number of software projects, including the GNU Network Object Model Environment (GNOME) project.

GTK++ is built on top of GDK (GIMP Drawing Kit), which is basically a wrapper around the low-level functions for accessing the underlying windowing functions (Xlib in the case of the X windows system). GTK++ is essentially an object-oriented API. Although written completely in C, it is implemented using the idea of classes and callback functions (pointers to functions).

There is a thin-layered C++ binding to GTK++ called GTK-, which provides a more C++-like interface to GTK++. In the current release, GDK++ 1.2, there is also a library called gnome-, which is a C++ wrapper around the GNOME's C API. But gnome- seems to be immature at the moment, a fledgling thing at best. This may have changed, of course, by the time you read this.

If you are determined to use \(\mathrm{C}++\), you have three choices:
- First, you can always use one of the specialized wrapper libraries that wraps the GDK++ C API in C++ classes in much the same way as GTK - does. In today's lesson, this is what we actually do because it ties in neatly with the KDE programming environment and thus makes meaningful comparisons easier. - Secondly, if you don't want to trust your applications to gnome- or yet another C++ wrapper library for whatever reason, you can use only the C subset of C++ when interfacing with GTK++ and then use the C interface. Remember that C is a valid subset of \(\mathrm{C}++\) and a valid C program is a valid \(\mathrm{C}++\) program, broadly speaking.
- Finally, you can use GTK++ and C++ together by declaring all callbacks as static functions in C++ classes, and again calling GTK++ using its C interface. The buttons example shown in Listing 26.1 does this.

If you choose this third approach, you can include as the callback's data value a pointer to the object to be manipulated (the this pointer).

Selecting between these choices is largely a matter of preference, because in both approaches you get C++ and GTK++. Personally I find the second solution rather ugly and inelegant, and I prefer to use a C++ wrapper library, but your mileage may vary.

None of these approaches requires the use of a specialized preprocessor, so no matter what you choose, you can use standard C++ with GTK++.

Now let us look at the GNOME itself, where it sits above GTK.
A true GNOME program is a GTK+++ GUI application that also uses the GNOME libraries. The GNOME libraries make it possible to have a similar look and feel among applications, and to make simple things simple to program. Not only that, but the GNOME libraries add a lot of widgets that do not properly belong in GTK++

Unfortunately for budding C++ programmers, the GNOME, like its supporting graphics library, GTK++, has a C API. To be sure, the API is strongly object-oriented and uses opaque structures and accessor functions and sc on, but it is nevertheless a C API. You do, of course, have the same three options to choose from for GNOME programming as you do for GTK++ programming.
class callback
public:
static void clicked (Gtkwidget *button, gpointer data);
static gint quit (GtkWidget *widget, GdkEvent *event, gpointer data);
callback::clicked(GtkWidget *button, gpointer data)
char string \(=(\) char*) data;
g_print (string)
q_print (""using
callback::quit (GtkWidget *widget, GdkEvent *event, gpointer data)
gtk_main_quit()
int

GtkWidget *app;
GtkWidget *button;
Gt kWidget *hbox;
/* Initialize GNome, this is very similar to gtk_init */
gnome_init ("buttons-basic-example", "0.1", argc, argv);
app \(=\) gnome_app_new ("buttons-basic-example","
hbox = gtk_hbox_new (FALSE,5);
gnome_app_set_contents (GNOME_APP (app), hbox);
* bind "quit_event to gtk_main_quit */
gtk_signal_connect (GTK++_OBJECT (app)
\(\underset{\text { GTK++_SIGNAL_EUNC }}{ }\) (callback: :quit), NULI);
button = gtk_button_new_with_label ("Button 1"); ;
gtk_box_pack_start (GTK+__BOX (hbox), button, FALSE, FALSE, 0);
gtk_box_pack_start (GTK++_BOX (hbor), button, FALSE, FA
gtk_signal_connect (GTK+_-OBJECT (button), "clicked",
button = gtk_button_new_with_label ("Button 2");
gtk_box_pack_start (GTK+__BOX (hbox), button, FALSE, FALSE, 0);
gtk_box_pack__start (GTK++_BOX hoos), button, FALSE,
gtk_signal_connect ( (TK + _OBECT (button), "clicked",
GTK++ SIGNAL FUNC (callback: clicicked),"Button \(2 \backslash\) ""
gtk_widget_show_all(app)
-
```

return 0;

```

ANALYsIS The first portion of the code, lines 7-27, contain the declaration of the class callback and its static callba functions that we created to respond to events that the GUI detects when we click the displayed buttons:
As you can see, the functions are simple and straightforward, with callback:.clickedo merely printing some text to stdout at line 19 (compare to cout in \(\mathrm{C}++\) ) and callback:quit) exiting the main program at line 25 .
The first call we make in main) is to gnome_ init) a t line 37 .
This is very similar to a pure GTK + + application where we would call gtk_ _init); similarly we use a call to gnome app_newo at line 38 to give us an instance of a new GNOME application. In a pure GTK ++ application, the
corresponding call would be to gtk_window_new \()\)

Although gnome_app_newo returns a GikWindow, we later verify that to a GnomeApp with the GNOME_APP macro. Th GnomeApp is the principal widget behind each application. It is the main window of the application containing the
document being worked on, the applications menus, toolbars and status bars, and so on. It also remembers the document being worked on, the applications menus, toolbars and status bars, and so on. It also remembers the
docked positions of menu bars and toolbars and such, so that the user gets the window the way the application left when it last shut down.

Creating a new GnomeApp widget is as easy as calling gnome app_newo with the name of the application name and t t
title of the main window. You can then create the content of the main window and add it to the GnomeApp widget b Cite of the main window. You can then create the content of the m:
calling gnome_app_set_contents 0 with your contents as the argument.
In this case, we create a horizontal box to hold the buttons we will shortly be creating and add it to the application widget at lines 40 and 42 .

The GnomeApp application framework is part of the libgnomeui library, and it it the part that truly makes a GNOME
application a GNOME application. It also makes programming GNOME applications reasonably simple and application a GNOME application. It also makes programming GNOME applications reasonably simple and
straightforward and makes the anplicatons comprehensive and consistent straightforward, and makes the applications comprehensive and consistent across the desktop. With plain GTK++ you have to do a lot of things by yourself, but GnomeApp takes care of the standard GNOME UI set
still allows the user to configure the behavior and have it be consistent over different applications.
Lines \(45-56\) link the callback:quit) event handler to its callback function and then create two buttons and link them the callback: :lickedo event handler, before adding them to the horizontal box

This means that when we click one of the buttons, the event will be routed to and processed by the callback:clicked callback function
Finally, at lines 58 and 59 , we tell the main application widget to show itself and all its contents, and then enter the main event loop and await the clicking mouse.

To build this program, you can use the following command
gcc -g -Wall 'gnome-config --cflags gnome gnomeui `LDFLAGS=' gnome-config
\(--1 i b s\) gnome gnomeui` buttons.cc -o buttons
The Figure 26.1 shows you what your buttons program ought to look like

\section*{Duver}

\section*{}

\section*{Figure 26.1 The output of the buttons program.}

\section*{Wrapping GTK++ with wxWindows}

In the last section, you saw how you can use C++ with the GNOME libraries to create GNOME applications. As far back as 1992 a group of programmers at Edinburgh University in Scotland created version 1.0 of the
wxWindows toolkit. The wxWindows toolkit is a set of pure C++ libraries that make GUI development me wxWindows tookiki. The wxW indows toolkit is a set of pure C++ libraries that make GUI development much easier
It allows \(\mathrm{C}++\) applications to compile and run on several different types of computers, with few if any source code changes.
During 1997 an effort to produce a standard Linux desktop environment was underway-GNOME. Its widget set
was GTK Linux \(4++\), built on top of X , and it looked as though GTK ++ based apps were to become the standard in the
and



class Myspep: public wxapp
    class myprane: pub1ic wxfrane
    \({ }^{\text {class MyP }}\)
    \(=\)
    tremen__Rep (Myappp













-



\section*{,}





S wypp: punic wxapp
Cass wyprane: pubitic wefrane
class MyF
public:


    Deciare__event_Itabe, ()

    mout

    Eveni_Taizi()
    Myppe: :onTnit ()



    Const wxsi zess size)




close (true);
WMessagebox ("rris is a waxindows hello worta sample


The first argument is the widget that will be emitting the signal, and the second is the name of the signal you
want to catch. The third is the function you want to be called when it is caught, and the fouth the data you want to catch. The third is the function you want to be called when it is caught, and the fourth, the data you want to have passed to this function

The function specified in the third argument is called a "callback function," and should generally be of the form void callback_func( Gtkwidget *widget, gpointer callback_data )
where the first argument is a pointer to the widget that emitted the signal, and the second a pointer to the data given as the last argument to the gtk_signal_connect) function as shown earlier.
wxWindows extends this idea of signals and callbacks into its own namespace and uses event tables to map events to actions.
You place an event table in an implementation file to tell wxWindows how to map events to member functions These member functions are not virtual functions, but they are all similar in form: They take a single wxEventderived argument, and have a void return type

You may notice that wxWindows' event processing system implements something very close to virtual metho in normal \(\mathrm{C}++\); that is, you can alter the behavior of a class by overriding its event-handling functions

In many cases this works even for changing the behavior of native controls. You can, for example, also filter ou a number of key-press events sent by the system to a native text control by overriding wxTextCtrl and defining a handler for key events using EVT_KEY_DOWN. This would indeed prevent any key-press events from being seI to the native control-which might not be what you want. In this case the event handler function calls Skipo indicate that the search for the event handler should continue.

In summary, instead of explicitly calling the base class version as you would have done with \(\mathrm{C}++\) virtual functions (that is, wxTextCtrl::OnChar()), you should instead call Skip().
In practice, that would look like this if the derived text control only accepts "a" to " z " and " A " to " Z "
01: void MyTextCtrl::OnChar(wxKeyEvent\& event)
02:
if ( isalpha( event.KeyCode() )
// key code is within legal range. we call event.Skip() so the / event can be processed either in the base wxWindows class // or the native control.
```

event.Skip();

```
\}
else
// illegal key hit. we don't call event.Skip() so the // event is not processed anywhere else
```

wxBell();

```
:
16: \}

You will see that line 3 checks to see if the key code was representing a letter, and calls Skip0 at line 8 to continue the search for a handler.

The normal order of event table searching by ProcessEvent is as follows
1. If the object is disabled, usually with a call to wxEvtHander::SetEvtHandlerEnabled \()\), the function skips
to step (6) to step (6).
2. If the object is a wxWindow, call ProcessEvent recursively on the window's wxValidator. Exit the function if this returns true.
until you run out of tables for an event handler for this event. If this fails, try the base class, and so on until you run out of tables or find an appropriate function, in which case the function exits.
4. Apply the search down the entire chain of event handlers (usually the chain has a length of one). Exi the function if this step succeeds.
5. If the object is
5. If the object is \(a\) wxWindow and the event is a wxCommandEvent, recursively apply ProcessEvent to the parent window's event handler. Exit if this step returns true.
6. Call ProcessEvent on the wxApp object

\section*{Adding a Menu to Your Own wxWindows Window Class}
we will add a menu to the simple application we have just developed. We functionality so much as providing a second way of invoking the functionality we already have

The code change we have to make is very small, so we do not have to include the entire source file: All we nee
to do is add some extra code (see Listing 26.4 ) into the MyFrame constructor Input

Listing 26.4 The wxWindows GNOMEHelloWorld Program with a Menu
```

MyFrame::MyFrame (const wxString\& title
const wxPoint\& pos,
wxFrame((wxFrame *) NULL, -1, title, pos, size)
wxSize panelSize = GetClientSize();
int height = panelSize.GetHeight()
nt width = panelSize.GetWidth()
_panel = new wxPanel(this, -1, wxPoint (0, 0), panelSize);
m_btnGreet = new wxButton(m_panel, ID_Greet, -T ("Greet"),
wxPine(50,20))
_btnQuit = new wxButton(m_panel, ID_Quit, _T("Quit"),
wxPoint(width/2+20, height/2-10)
wxSize(50,20));
wxMenu *menuApp = new wxMenu;
nenuApp ->Append( ID_Greet, "\&Greet..." );
menuApp ->AppendSeparator();
menuApp ->Append( ID_Quit, "\&Quit");
xMenuBar *menuBar = new wxMenuBar
enuBar->Append (menuApp, "\&Application")
SetMenuBar( menuBar );

```

ANALYSIS We have exactly the same code to begin with, and then simply append the code to create the men and associate the menu events with the event handlers we already have

At line 18 we create a new menu header on the main menu bar, and lines \(20-22\) populate the menu header anc determine what you will see when you select the menu header and the menu items drop into view.

We then create on line 24 a menu bar, the actual object you see across the top of the client area, and attach th menu header and its appended items to it on line 25 .
Finally, we can set the application's main menu to be the menu bar we have just created on line 27 .

\section*{To compile the program, use the following command}

\section*{etting wxWindows and Other wxWindows Resource}

The wxWindows home page is at http://www.wxwindows.org/. wxWindows comes in several flavors, so to
speak, and provides the widestrange of supported platforms I have ever seen for a package of this type. The you are primarily interested in for this book is wxGTK, and the download page for that is at you are primarily interested in for this bc
htppl/www.wxwindows. orgdll gik.htm.
The latest version of wxWindows is 2.1 .11 and is a substantial upgrade from version 1.1 I believe version 1 is
now considered to be obsolescent if not obsolete. You will need GTK++1.2 for wxGTK 2.1 .11. There is also a set of wxWindows mailing lists at htpp://www.wxwindows.orq/maillst2.htm. Some of these ar

\section*{ndamentals of KDE Programming}

The Q t library is the graphics library that underpins the entire KDE. It is a C+t tookkit
in Norway, and offers graphical elements you can use for creating X GUI applications. Additionaly, the toolkit offerr a complete set of classes and methods ready to use even for non-graphical
programming code, a solid framework of user interaction using virtual methods and the unique Qt signal and srogram nac coted a libechanimm and a libray of predefined GUI elements, GUI widgets, that you can use to create the visibl lements. It also offers predefinined dialogs that are often used in applications, such as progesess indicators and
elt

\section*{Creating Your First KDE Application-"Hello World"}
26.5 The KDEHelloWorld Program
```

KDEHelloworld.cpp

```
finclude \(<\) Kapp.h
include \(<k\) _mainwindow.h>
nt main( int argc, char **argv
KApplication MyApp( argc, argy )
KTMainWindow \(*\) MyWindow \(=\) new
KTMai inwindow \(*\) MyWindow \(=\) new KTMainwi
MyWindow \(\rightarrow\) setGeometry \((50,50,200,100)\);
MyApp. setMainWidget
MyWindow \(\rightarrow\) show (1)
(1)
return MyApp.exec ()
fese files with the following commands at the prom
 KDEHelloWorld KDEHell oWorld

ANALYSIS In this application, as with all KDE applications, you first have to instantiate a KApplication objec We pass the progr
maino unaltered.
Following this, we instantiate a KTMainWindow object and call it MyWindow on line 11; as the name suggests,
this will be the window you will eventually see displayed on the screnn. We size and place the window on the screen with the call to selGeometryy on line 12. We move it to coordinates ( 50 , 50 ), and change the size to
200x 100 p pixels (widhxheight). ,



\section*{Events as QEvent Objects}
The Kapplication sends events
decide what to do with them.
A widget receives the QEvent and calls QWidget:even(OEVent*), which then decides which event has been
detected and how to react; event is therefore the main event handler.

\section*{Adding Buttons to Your Own KDE Window Class}

So far these KDE applications have been pretty dull, to say the least; we may have a GUI but we are not using
to much effect.
In this third example, we will expand upon the code we have and add a couple of buttons to the application. A
the same time we will also introduce an important concept that underpins the whole of the Qt (and hence KDE application framework: signals and slots.
and
kapp->quit();
No
KMsgBox::message (0," KDEHelloworld", "Hello World!")
\(\qquad\)
close();
```

Compie. nase fles with the followngcommands at the prompt

```

g++-LSKDEDIR/1iil -1kdecore
KDEHelloworld KDEHelloworld.

ANALYSIS Now you can see that things are really starting to move. We have several new elements in the co
First, and most simply, you will notice a call to KMssBox:messageo at line 50

This puts a simple predefined dialog box on the screen giving the user some information; in this case it doesn't
say much other than "Hello World," but you can set it og give any message you like. The firss string parameter, say much other than "Hello World"" but you can set it to give any message you like. The first
"KDEHelloWorl", sets the title for the box, and the second, "Hello World", is the message itself.

The second, more important, thing you should notice is the odd keyword combination public sloss: in
KDEHelloworld. This in in not a standard C+ declaration but is part of the signals and slots mechanis KDEEEeloWorld.h. This is not a standard C++ declaration but is part of the signals and slots mechanism of the \(Q\)
library; it essentially tells the tenta Oject Compiler that these functions can ec called in response to signals library; it essentially tells the Meta Object Compiler that these functions can be called in response to signe
being emitted. If you look at the implementation of the KDEHElloWorld constructor,
 What all this means, essentially, is that by clicking the m buthreet widget you will send a clickedo signal and the
Qt library will invoke the KDEFHelloworld:Sloticreet) member function for you and display a messaga. Similarly by clicking the mobnenuit widget, you will send a c clickedo s signal and the Q t library will invoke the
K.

\section*{Adding a Menu to Your own KDE Window Class}

Finally, we will give our simple KDE application a final gloss of respectability by adding a menu to th
window. First we need to modify the source code so it looks like Listing 26.8.

\section*{INPUT. Listing 26.8 The KDEHelloWorld Program with Buttons and Menu}
```

KDEHelloworld.h

```
\#include <kapp.h>
\#include <ktmainwi
incluade <ktmainwindow.h>
include <kmenubar.h
class KDehelloworld : public KTMainWindow
Q_OBJECT
public:
public:
KDEHelloWorld ();
bubic slots:
void sloteupit ();
Private: \(\quad\) Pushbutton \(*_{\text {m_btngreet }}\);
QPushButton \({ }^{\text {m m_btnGreet }}\)
QPushButton \({ }_{\text {m_btneuit }}\)
KMеnuBar *m_Menu; \(^{\text {K }}\)
PopupMenu \(\star_{\text {m }}\) MenuApp ,

\section*{KDEHelloworld.cpp}
finclude "KDEHelloworld.mo
KDEHelloworld: : KDEHelloworld () : KTMainWindow()
m_btnGreet \(=\) new QPushButton ("Greet", this)
- btnGreet-> setGeometry \((45,30,50,20) ;\)
btngreet->show();
connect (m_btnGreet, SIGNAL (clicked()), this, SLOT(SlotGreet ()))
m_btnQuit \(=\) new QPushButton ("Exit", this)
mbtnuit->setGeometry \((105,30,50,20) ;\)
mbtneuit->show \() ;\)


m_MenuApp->insertItem("‘\&uit", this, SLOT (SlotQuit()) )
m_Menu \(=\) new KMenuBar (this);
m_Menu-> insertItem ("\&Application", m_MenuApp) ;
Void KDEHelloworld::closeEvent (QCloseEvent
app->quit();
oid KDEHelloworld::SlotGreet()
KMsgBox: :message (0,"KDEHelloworld","Hello World!");
oid KDEHelloworld:SlotQuit ()
close()

Compile these files with the following commands at the prompl.
g++ -c -I SKDEDIR/include -ISQTDIR -fno-rtti main.cp

g++-LSKDEDIR/lib -1kdecore -1kdeui -1gt -o KDEHelloWorld main.
.
ANALYSIS \(\begin{aligned} & \text { The first things to notice are the extra member variab } \\ & \text { added to the KDEFeloworld class on lines } 23 \text { and } 24 \text {. }\end{aligned}\)
```

The class-nar
lines $47-51$.

```

Output When you run this program, you should see something like the screen in Figure 26.7

\section*{,}

\section*{Building More Complex KDE Applications-KDevelo}

The sample applications you have seen so far are simple and serve only to illustrate the essentials of the KDE application framework. But for more complicated applications, switching between many source files and application framework. But for more complicated applications, switching between many source files and
navigating the classes and so on is a chore; in any case, no self-respecting GUI desktop would be without a
Integrated Development Environment, and the KDE is no exception.

The KDevelop Integrated Development Environment provides many features that developers need. In addition it wraps the functionality of third-party projects such as make and the GNU C++ compilers and
invisible, integrated part of the development process. KDevelop offers the following features:

- KAppWizard, which generates complete sample applications
- The creation of user handbooks writen with SGML and the automatic generation of HTML output with the KDE look and fee
Automatic HTML-based API-documentation for your project's classes with cross-references to the
used libraies Program debugging using KDbs
- The eninlusion of any other program you need for development by adding it to the Tools men
according to your individual needs

KDevelop makes it easy to work with all programs in one place and saves time by automating standard development processes as well as giving you direct and transparent access to all information you need to
your project. The integrated browsing mechanisms are designed to support documentation requests that your project. The integrated browsing mechanism
developers have in conjunction with their project.
- The KDE Core library,
your application may use
- The KDE UI library, co
The KDE U I library, containing user interface elements like menu bas The KFile library, containing the file selection dialogs
- The KHTMLW library, offering a complete HTML-interpreting widget that is used by various programs like KDEHelp, KFM, and Kdevelop
- The KFM library, enabling you to use the KDE file manager from within your application
The KAb library the KAddressook Provides addess bok - The KAh library, the KAddressBook. Provides address-book access, for example, for email
applications
applications
- The KSpell
- The KSpell library, offering widgets and functionality to integrate
checker, in applications like editors; used for the KEdit application

\section*{Summary}

Linux has finally come of age. The two new desktops you have seen today, the GNOME and the KDE, as well as their supporting libraries and applications, provide a choice of modern, intuitive, robust, and free GUIs for the Linux operating system. They also provide a set of powerful development frameworks, environments, and tools that take the drudgery out of writing GUI applications in X and enable you to concentrate on what your application ought to be doing rather than dragging you into a fight you can only lose with an intractable API

With the new Linux GUI interfaces, you have a genuine choice, and that can only be good for everyone as it will see a marked improvement in quality and price (especially as the Linux desktops are free) in the whole operating system marketplace

All the packages and libraries I have described in this lesson are being extended and improved almost daily by volunteers, and new ones are appearing almost as frequently

Their efforts and skill have firmly placed Linux where it belongs: on the desktop.

\section*{Q Why do we need the KDE and the GNOME?}

A The reason is that Linux, unlike Windows and MacOS, doesn't have a native GUI: You interact with using textual commands. The computer itself has graphics capability (usually), but Linux alone doesn't know how to drive it. The GNOME and the KDE provide a framework for you to interact with the computer graphically and to have your commands routed transparently to the operating system. In truth we don't need them both: They are different enough for them to coexist, and time will tell which one, if either, becomes more widespread.

\section*{QWhy can't I just use the Common Desktop Environment or Motif as my GUI?}

A You can if you want to. The beauty of Linux is that it is free, and if you use GNU software or KDE software, that is also free. You can get versions of the CDE and Motif for Linux, but they are proprietary software and you have to pay for them. Not only that, but to many people, the GUIs these provide are now starting to look a little old-fashioned.
Q If I am writing a GUI Application for GNOME, will it still run under KDE, and vice versa? A Maybe. If you write applications for one desktop, they will still run under the other if and only if you have the correct runtime libraries present on your machine. Moreover, if they run they will not behave as you might expect and want them to behave because the persistent session information from a GNOME application will be ignored by the KDE desktop (and similarly the other way around). More important, on some installations, GNOME applications, like the GIMP, upset the KDE color scheme when they have the focus. On my machine, when running Red Hat Linux 6.1, if I run the GIMP under KDE, the entire screen turns blue when the GIMP has the focus. If you just want a generic application, you can always use wxWindows, which uses the GTK++ but doesn't use GNOME.
Q What is cross-platform source code? Why might it be important?
A This is source code that you write once and compile using the correct toolkit on your target machine. The Qt library and wxWindows are both "platform-independent," meaning that the source code for a given application is the same no matter what platform you compile it on; obviously, the toolkit libraries are significantly different, but that isn't your concern. This is the area where some people think KDE falls down: The Qt libraries have a free license only for Linux and other UNIX builds. If you want to create software for Windows with the Qt libraries, you have to buy a license.
Q I want to use C++ but I don't want to use wxWindows or GTK- or any other wrapper. Can I use C++ direct on the GNOME/GTK++ libraries?
A Yes. Although it will probably lead to messier code, you can declare your callback functions as static members of your classes. You can then implement the classes that have the static callback members however you please and get them to do what you want.

\section*{Workshop}

The Workshop provides quiz questions to help solidify your understanding of the material covered and exercises to provide you with experience in using what you have learned. Try to answer the quiz and exercise questions before checking the answers in Appendix D, "Answers to Exercises and Quizzes," and make sure you understand the answers before continuing tomorrow.

\section*{Quiz}
1. What is the difference between event-driven programs and procedural ones?
2. What is a widget, in computing terms?
3. What is a callback function, and why are these prone to errors?
4. What is a Qt slot, and how does it react to Qt signals?
5. Signals and slots are typesafe. What does this mean and why is it a good thing?
6. How would you tie the following wxWindows event handler handler to a EVT_MENU event from a menu item with an ID of ID_MY_HELP? Assume OnHelp() is a member of the class MyFrame, and MyFrame is derived from wxFrame.
void OnHelp(wxCommandEvent\& WXUNUSED (event));
7. What does the IMPLEMENT_APP() macro do?
8. What does the Q_ObJECT macro do?

\section*{Exercises}
1. Using the buttons.cc source file as a base, extend the program so that the callback class puts up a dialog box when you click the buttons rather than writing directly to stdout.
2. Taking either the final wxWindows or KDE example as a basis, extend the code to display a filebrowse button that will enable you to select a file using a standard file-selection dialog.
3. Extend the code from Exercise 2 so that when you click OK and select the file, the program displays it in a text control on the main window.
4. Look at your results from Exercise 1 and consider how you might wrap the entire application into one class. Consider why you might want to implement the callback functions as private static members of the class and expose only "wrappers" that call them when the user triggers events. See the function do_message() in file ex26-01.cxx for example. Think how you might extend the idea to create a generic wrapper for the application and use virtual functions for setting the main window and connecting signals.

\section*{Appendix A \\ Operator Precedence}

It is important to understand that operators have a precedence, but it is not essential to memorize the precedence.

Precedence is the order in which a program performs the operations in a formula. If one operator has precedence over another operator, it is evaluated first.

Higher precedence operators "bind tighter" than lower precedence operators; thus, higher precedence operators are evaluated first. The lower the rank in Table A.1, the lower the precedence.

\section*{Table A.1Operator Precedence}
\begin{tabular}{|c|c|c|}
\hline Rank & Name & Operator \\
\hline 1 & Scope resolution (binary, unary) & : \\
\hline 2 & Function calls, parentheses, subscripting, member selection, postfix increment and decrement. & \[
\begin{aligned}
& \text { () [] } \\
& .-> \\
& ++--
\end{aligned}
\] \\
\hline 3 & sizeof, \(\mathrm{C}++\) casting, prefix increment and decrement, unary plus and minus, negation, complement, C casting,sizeof(), address of, dereference new, new[], delete,delete[] & \[
\begin{aligned}
& ++-- \\
& +-!-(\text { cast }) \& ~ *
\end{aligned}
\] \\
\hline 4 & Member selection for pointer & . ->* \\
\hline 5 & Multiply, divide, modulus & */ \% \\
\hline 6 & Add, subtract & + \\
\hline 7 & Bitwise shift & <<>> \\
\hline 8 & Inequality relational & <<= >> \(=\) \\
\hline 9 & Equality, inequality & == != \\
\hline 10 & Bitwise AND & \& \\
\hline 11 & Bitwise exclusive OR & \(\wedge\) \\
\hline 12 & Bitwise OR & | \\
\hline 13 & Logical AND & \& \& \\
\hline 14 & Logical OR & || \\
\hline 15 & Conditional & ?: \\
\hline 16 & Assignment operators \(=*=1=\%=\) & \[
\begin{aligned}
& +=-=\ll= \\
& \gg= \\
& \&=1=\wedge=
\end{aligned}
\] \\
\hline 17 & Comma & \\
\hline
\end{tabular}

\section*{Appendix B}

\section*{C++ Keywords}

Keywords are reserved to the compiler for use by the language. You cannot define classes, variables, or functions that have these keywords as their names. The list is a bit arbitrary, as it contains the standard keywords as well as some of the keywords that are specific to \(g_{++}\). Some of the keywords are not available with all compilers or with older versions of the g++ compiler. Your mileage may vary slightly.
\begin{tabular}{lll} 
asm & float & static \\
auto & for & static_cast \\
bool & friend & struct \\
break & goto & switch \\
case & if & template \\
catch & inline & this \\
char & int & throw \\
class & long & true \\
const & mutable & try \\
const_cast & namespace & typedef \\
continue & new & typeof \\
default & operator & typeid \\
delete & private & typename \\
do & protected & union \\
double & public & unsigned \\
dynamic_cast & register & using \\
else & reinterpret_cast & virtual \\
enum & return & void \\
explicit & short & volatile \\
extern & signed & while \\
false & sizeof & \\
\hline
\end{tabular}






\(\qquad\)```


[^0]:    

[^1]:    ＜＂＂（0）Quit

