Ed Anton

digital

INTEROFFICE MEMORANDUM

TO: LIST

DATE: 27 Dec. 1978
FROM: Steve Jenkins S Jenkins
DEPT: VAX 11 Engineering
EXT: 247-2395
LOC/MAIL STOP: TW/CØ4

SUBJ: VENUS PROJECT PROPOSAL

Attached is a proposal for the development of a VAX 1178Ø replacement system with equivalent cost and higher performance. This proposal needs funding action for FY79 Q3, Q4 to meet an FY82 Q1 FCS date.

A "product description" package including a more detailed view of functionality with subsystem block diagrams will be distributed before the end of January.

LIST (ACTION):

BRIAN CROXON

BILL DEMMER

CC (INFO):

GORDON BELL
ULF FAGERQUIST
GEORGE HOFF
BERNIE LACROUTE
PETER CONKLIN
AL AVERY
JOEL BERMAN
DON AMES
DON CHACE
FRED BOWERS
ED KENNEY
HENK SHALKE
JIM DREW
BOB HALL
JOE CARCHIDI
DICK HUSTVEDT
DAVE CUTLER
VENUS ENGINEERING GROUP
DAVE RODGERS
AL RYDER
PER HJERPPE
DAVE BEST
RON BINGHAM
MARION DANCY
TOM EGGERS
JOHN JORGENSEN
CHARLES DENNEY
ALAN KOTOK
VIC KU
KATHIE NORRIS
MIKE POWELL
FRANCO PREVID
JOHN SHEBELL
ED SLAUGHTER
BILL STRECKER
JIM MARSHALL
JOHN GROARK
JOHN MUCCI
LARRY WADE
BARBARA SLAUGHTER

VENUS PROJECT PROPOSAL

AUTHOR: Steve Jenkins
DATE: December 27, 1978

PREFACE

This document will be superseded by a preliminary VENUS PROJECT PLAN in February 1979. At that time more detailed information will be available about Project schedules and development expenses along with product cost and performance refinements.

The following people contributed to this document:

Sas Durvasula
Al Helenius
Jud Leonard
Dave Potter
Dave Rodgers
Bob Stewart

CONTENTS

1.0 INTRODUCTION

VENUS is a high performance general purpose computer system which implements the VAX-11 Architecture and runs the VAX/VMS Operating System and associated layered products.

The product is primarily targeted at the high end VAX-11 Business ($250K System Product Cost) and as a replacement system for the VAX 11/780 in the early 1980's to meet and overcome the competitive pressures expected from newer 32 bit systems offered by SEL,DG, and Interdata.

The primary goals of the product are:

* 11/780 System Cost
* FCS at Q1, FY'82
* 3.5x 11/780 Performance

2.0 COST AND PERFORMANCE GOALS

* Same transfer cost as 11/780
* 3.5x CPU Performance of 11/780
* 13.3 MB/SEC Bus Bandwidth coupled with Improved Memory Bandwidth and High Priority I/O Servicing.

2.1 COST

The goal here is to provide an entry level system with a selling price of under 100K, while allowing a large system price of up to 250K. Each of these systems would have higher performance tape drives and at least twice the main memory and disc capacity as the minimum configuration for the 11/780 Unibus and Massbus based systems.

Entry Level System-

The target FY'83 transfer cost will be 25K (less FA&T)

Configuration includes:

CPU
Power System & Cabinet
1 MB Memory
Console Terminal and load device
Unibus Adapter
8 Asynchrorous Communication lines
>100 MB Disc Capacity

Selling price goal appears to be achievable as a Dock Merge Product.

Large System-

The target FY'83 transfer cost will be 45K (less FA&T)

Configuration includes:

CPU
Power System & Cabinet
2MB Memory
Console terminal and load device
Unibus Adapter
8 Asynchronous Communication lines
2 Massbus Adapters
300 MB Disc capacity
125 ins./sec. Magtape

2.2 CPU PERFORMANCE

The CPU cycle time improvement using advanced technology ECL components can easily reach 4x 11/780 CPU cycle time because of a gate speed improvement of 5x that of Schottky TTL. Interconnection and clock skew delays become more important and prevent a complete correspondence between gate speed and cycle time.

A Main Memory organization similar to the 11/70 configuration will provide a 2.5x improvement relative to the 11/780 for CPU to Main Memory access times.

These speed improvements coupled with a high speed ECL memory cache will provide 3.5x better overall CPU performance than the 11/780.

2.3 I/O THROUGHPUT

The I/O to Memory connection uses the 11/780's SBI BUS, UNIBUS Adapter, and Massbus Adaptor.

The SBI will remain a 13.3 million byte/second bus. A practical improvement in I/O performance will follow from removing processor-to-memory traffic from the SBI and providing improved memory bandwidth for read-modify-write operations. A faster CPU will decrease interrupt service and process context switch times, thus creating another improvement in real I/O performance.

The Massbus adaptors provide an aggregate throughput rate of 2 million bytes/sec per adapter.

The Unibus adaptor provides an aggregate throughput rate of 1.5 million bytes/sec with by Buffered Data paths and a 500 thousand bytes/sec. through the Direct Data path.

2.4 VAX FAMILY PRICE/PERFORMANCE

All measurements made with respect to 1178Ø and includes CPU with Power and Package, Unibus Adapter, Terminal, Software Update Media, and 256 KB Memory!

Machine	Xfer Cost	Performance	P/P	FCS
VENUS	2Ø.5K	3.5	3.5	Q1 FY82
1178Ø	2Ø417	1	1	Q2 FY78
COMET	7269	.6	1.68	Q2 FY8Ø
NEBULAE	365Ø	.33	1.84	Q1 FY81

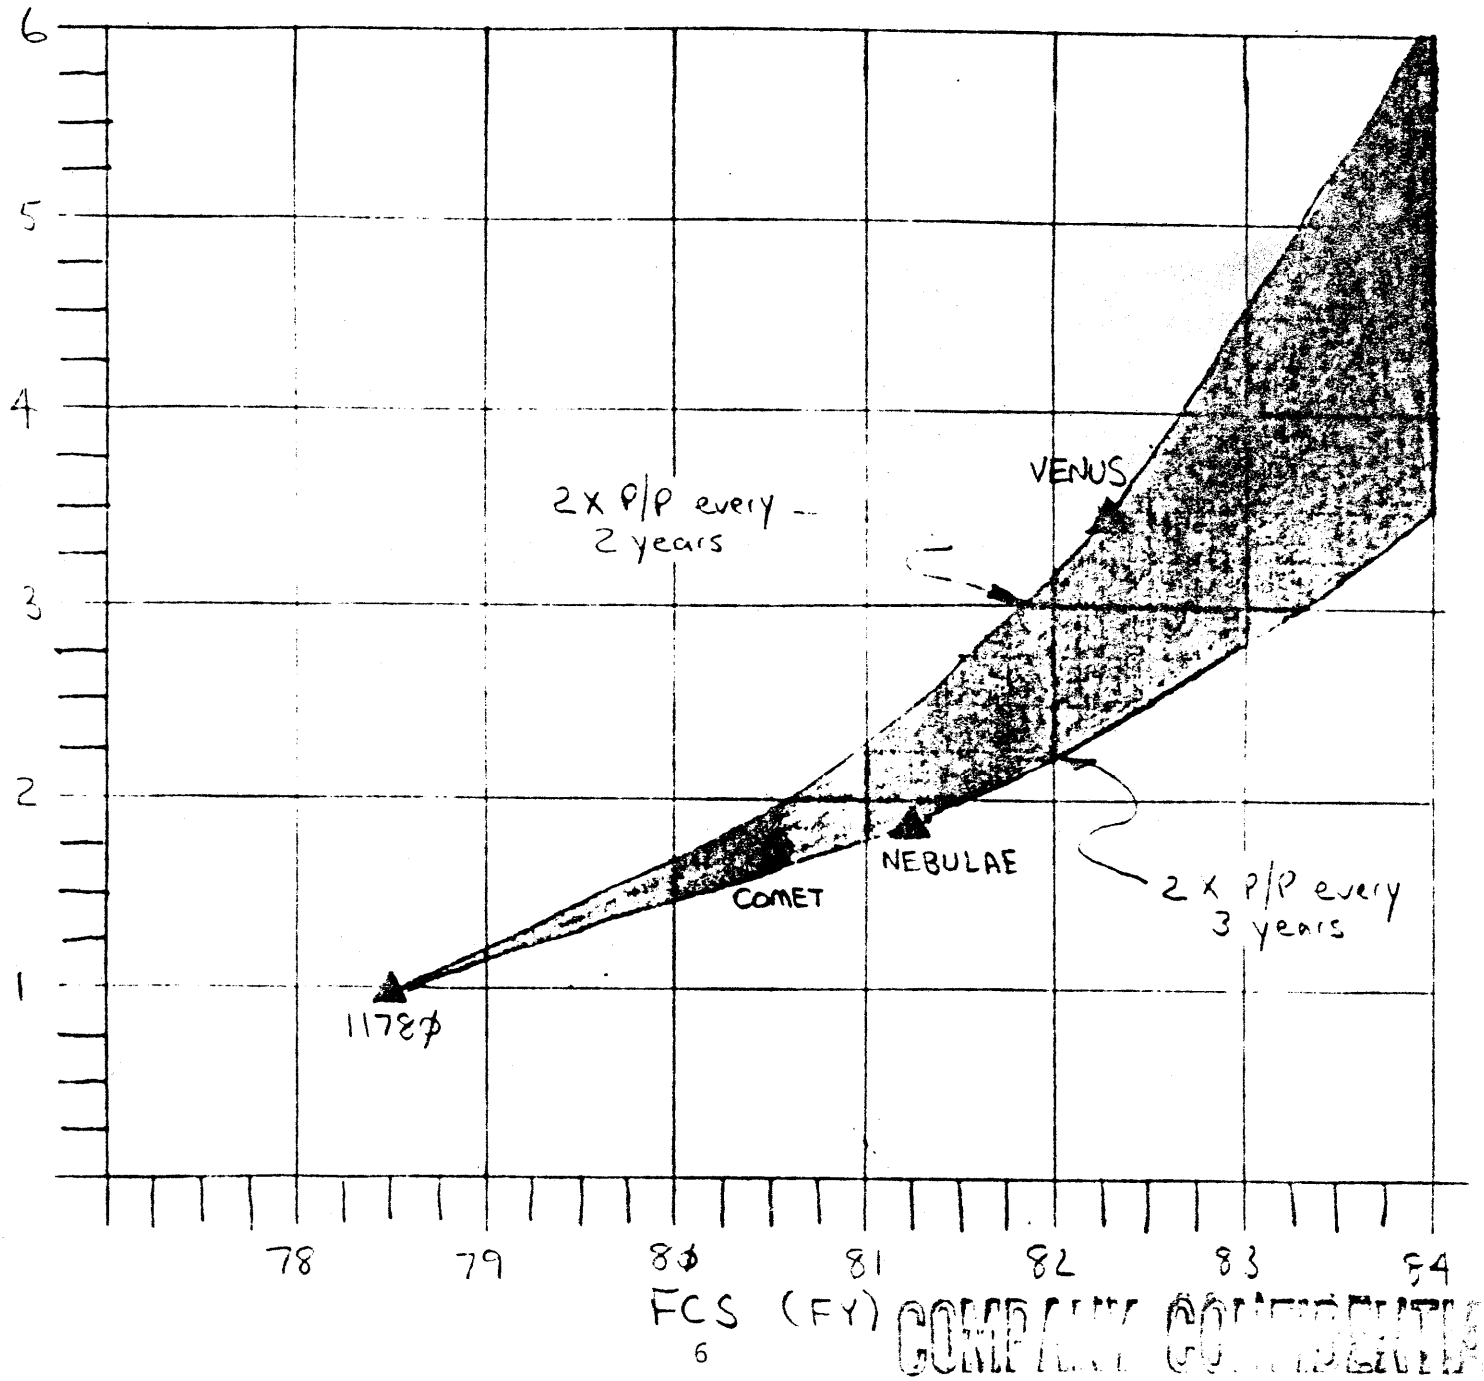

3.0 RAMP GOALS

* $\geq$ MTBF as 11/780
* $\leq$ MTTR as 11/780
* $\geq$ Availability as 11/780

3.1 Reliability

The MTBF of a system is determined by its component count and the quality index of those components. Our design goal is not to decrease the MTBF from that of the 11/780 (1300 Hrs.) for CPU,Memory,UBA,MBA,and console components.

3.2 Maintainability

The maintainability will be equivalent to that of the 11/780 with the same design concepts being applied for fast maintenance steps and low MTTR.

Modular system throughout
No cable connections on modules
Diagnostic processor
Parity on all RAMS and major busses
History recorder on the SBI
Microdiagnostics

3.3 Availability

System availability will be the same as the 11/780 with the use of Power Fail Restart, On-Line Diagnostics, Track Offset Retry, etc. Improved Availability can be achieved through use of redundant peripherals and peripheral controllers.

4.0 TECHNOLOGY

* ECL Macro Cell Arrays
* 4 Kbit ECL RAMS
* 64 Kbit MOS Memory
* Fine line/Standard Grid 4 Layer Extended Hex Modules
* Press Pin Backplanes
* Modular Power System

4.1 ECL MCA's

The Macro Cell Array (MCA) is an ECL array in which Macro logic functions (e.g. flip-flop,latches, multiplexers) may be interconnected on the chip through the use of 2 levels of metal. The basic chip diffusion (at the transistor and resister level) is fixed for all designs. The diffused silicon may be inventoried and then "programmed" to its final logical function through the addition of the final 2 levels of metal at a later time.

A joint effort between DEC-Marlboro and Motorola has resulted in an MCA design which includes 500 to 700 gate equivalents on a chip with built in I^2L diagnostic logic. Macro propagation delays are approximately 1.3 nanoseconds.

The design will be supplemented with 10K series ECL SSI and MSI parts.

4.2 ECL RAM's

4 Kbit ECL Rams with access times of 25 nanoseconds will be used for control storage and Data cache.

Less dense memory units such as scratchpads and translation Buffer will use 1K bit ECL Rams with access times of 10 nanoseconds.

4.3 64K MOS

The 64 Kbit dynamic RAM with an access time of 150 nanoseconds will be used as Main Memory. ECC will be applied to allow correction of all single bit errors and detection of all double bit errors.

4.4 Module Form Factor

The 4 layer Extended Hex module has proven to be the most cost effective format for logic modules and will be used in the design.

The density of interconnects on a module will dictate the use of fine line etching (8 mil line widths). Fine Line Etching coupled with standard grid layout will achieve the necessary produceability and manufacturing quality.

4.5 Backplane

Press Pin technology will be applied to facilitate the use of 1/10 inch center pins and reverse pin cable connections to the backplanes. This will improve the quality of the backpanel by eliminating solder shorts and warpage.

4.6 Power Systems

A Modular Power System will be used to provide a versatile configuration with improved manufacturability.

The power supply system will exhibit high immunity to power line disturbances, high efficiency, and low cost.

5.0 COMMONALITY OF COMPONENTS

* Reduced Development Cost
* High Volume Manufacturing Product Cost Benefits
* Reduced Field Service Training and Sparing expenses.

5.1 Reduced Development Cost

The FCS product will support all existing SBI options such as Massbus Adaptor, Unibus Adaptor, DMA Interface, and Multi-port Memory. There will be only a small development cost for testing these options with the VENUS System.

Any new SBI option can be added to both the 11/780 and VENUS as product enhancements for nearly half the development cost as would have been true for two systems with different busses.

The MCA technology will be shared with the Dolphin project thereby saving additional development and tooling funds.

The Dolphin Power and Packages system may be used in VENUS if consistent with product cost goals.

5.2 High Volume Manufacturing

The benefits of high volume buying and manufacturing will directly reduce product cost. Manufacturing overhead rate will be reduced as a result of a more favorable balance between Direct Labor and Production Control and Manufacturing Engineering personnel (overhead).

5.3 Field Service

The commonality of parts between VENUS and 11/780 Systems will allow more coverage of spares and thereby reduce Field Service inventory. Course Development costs for the FCS SBI options on VENUS already have occurred and new costs only have to be incurred once for future extensions.

6.0 SOFTWARE

* Microdiagnostics
* On-Line Diagnostics
* Repair Level Diagnostics
* Minimum changes to VMS

6.1 Diagnostics

The Microdiagnostic strategy incorporated in the 11/780 has proven to be very successful for quickly isolating faults. This same strategy will be used in VENUS with possibly better fault resolution due to the diagnostic logic within the MCA chips and a smaller number of modules.

6.2 VMS

System Software changes should be minimal due to the use of RH780 and DW780 adaptors, and instruction set compatability. Minimum changes to VMS is a goal.

7.0 PRODUCT DESCRIPTION

* I and E Boxes provide overlap of Instruction Decode and Processing.
* High Speed Data Cache and Address Translation Buffer.
* Internal Main Memory connection for faster CPU Access times and a more balanced arbitration scheme.
* Console Microprocessor.
* Optional Floating Point Accelerator.
* WCS Tools
* Multiprocessing
* High Speed DMA Customer Interface.

7.1 CPU Organization

Separate Instruction decode and execution processing allows operand evaluation in parallel with instruction execution. This coupled with a high speed Data Cache and Address Translation Buffer produces fast CPU operation.

A Console Microprocessor manages Control Store loading for instruction set processing or diagnostic evaluations. Operator commands and Bootstrap sequences are also performed by the Console Microprocessor. The console command language is compatible with all other VAX processors.

Physical Memory capabilities at FCS will be 8MB located in the main CPU cabinet. Expansion capabilities to 32MB will be provided at slightly increased access time through use of an expansion cabinet.

An optional Floating Point Processor speeds up instructions which perform Floating Point and integer multiplication and division arithmetic.

7.2 BUS OPTIONS

The use of the MA780 Multiport Memory will provide multiprocessing at FCS using Star, Comet and Venus processors.

The DR780 allows special high speed Customer DMA devices to be interfaced to the Venus system. This provides a migration path for customers with these devices on Star and Comet systems to convert to a higher performance processor with minimum difficulty.

Existing SBI options will be fully supported on the Venus system. The DW780 and RH780 Unibus and Massbus Adaptors will be available at FCS to interface peripheral devices to the system.

8.0 FUTURE EXTENSIONS

* Optional Vector Processor
* 32MB Physical Memory Capability
* IBM Channel Interface
* DEC I/O Architecture Interface
* Comet-Like Unibus Adaptor

9.0 PROJECT SCHEDULE

* Field Test - Q3 FY81
  FCS - Q1 FY82
  (with additional Q3, Q4 FY79 funding)
* Field Test - Q1 FY82
  FCS - Q3 FY82
  (without additional Q3, Q4 FY79 funding)

9.1 KEY MILESTONES

CONCEPTUAL DESIGN PHASE

-Subsystem Block Diagrams Done; mid Jan., 1979
-Specification reviews complete; April, 1979

IMPLEMENTATION PHASE

-Start Logic/Microcode Design; May, 1979
-Complete Logic Design; Dec., 1979
-Complete Microcode Design; Q4, FY80

BUILD & VERIFY PHASE

-Start MCA Layout; July, 1979
-Start Module Layout; Dec., 1979
-Power On Breadboard; Q4, FY80
-Power on Prototype; Q2, FY81

PRODUCT EVALUATION PHASE

-Start DMT/102 Evaluation; Q2, FY81
-Ship Field Test Unit(s); Q3, FY81

PRODUCTION PHASE

-Limited Release Sign-off; Q2, FY81
-FCS; Q1, FY82

9.2 CONCEPTUAL DESIGN PHASE

During this phase simulations of sample programs will be run to determine instruction and cache characteristics of VAX code. The overall system organization has been determined and this will lead to a more detailed view of each subsystem.

A software model will be written to study the impact of a higher CPU speed and I/O throughout with the memory configuration under consideration.

Detailed block diagrams and specifications will be generated for each of the subsystems and MCA chip types and partition will be understood.

9.3 IMPLEMENTATION PHASE

Complete system logic design will take place including MCA chip design and simulation. It is possible that SAGE simulations will be accomplished at the subsystem level. Module partition will be determined.

A MIMIC (TUMS?) simulator will be written based on functional specifications of the system.

Complete system microcode will be written with key areas debugged on the simulator.

9.4 BUILD AND VERIFY PHASE

Approximately 30 to 40 unique MCA chips and 12 to 15 extended hex modules will be processed through the Design Services group.

A single breadboard will be constructed using first pass etch boards and MCA fabrications but possibly using alternate power supply and cabinet designs.

Second pass edits to modules and MCA chips will be processed and parts used with the products power system and package to construct several prototypes.

9.5 PRODUCT EVALUATION PHASE

Prototype and Pilot Production units will be used to insure that the product meets all of it's specifications, is reliable, and can be manufactured by Digital Equipment Corporation.

Evaluation will continue to Production Release phase.

9.6 PRODUCTION PHASE

The manufacturing organization will start building units from released prints according to the established process. Product and Process maturity will evolve from a limited release to a production release atmosphere through the completion of DMT and PMT evaluations.

10.0 MANPOWER AND BUDGET

* Planned Team Build-up

10.1 MANPOWER REQUIREMENTS (Start-up)

It is important to recognize the need to plan for the hiring and training of the development team. Incremental increases of manpower will be planned during the quarter preceding their use.

The current stage of Conceptual Design requires the addition of four senior level engineers to develop the detailed conceptual designs of each of the subsystems. The collection of people involved in this design phase are expected to produce the necessary block diagrams and specifications to define the product.

During the specification writing phase additional engineers will be added and trained in order to perform the logic and microcode designs. This will be followed by the inclusion of Tools people for SUDS/SAGE/MIMIC support.

The early part of the Design Phase will be utilized to hire and train the necessary technicians and diagnostic engineers. The Power Supply and Packaging Support people will also be brought in during this time.

Manpower Summary:

	Current	Addition	Total
FY79, Q2	4.5E	4 E	8.5E
Q3		9.5E 1 P 3 U	18 E 1 P 3 U
Q4		2.25M 3 E 3 U 4 P 14 T	4.25M 19 E 6 U 5 P 14 T

E=Engineers; M=Managers/Supervisors; T=Technicians
U=Microcoders; P=Programmers

10.2 BUDGET

Budgeted funds reflect the anticipated movement of CC375 personnel to the Venus project. Other Tewksbury funds could be moved to decrease the variance indicated.

	Q3	Q4
E20-02214	42.5	42.5
E20-02216	45.0	67.5
E20-02201		2.5
E20-02215		5.0
Budget	87.5	117.5
Forecast	152.5	370.2
Variance	(65.0)	(252.7)

TOTAL DEVELOPMENT COST (5 Year Product Lifecycle)

FY79	FY80	FY81	FY82	
565.7K	3795.5K	3530.1K	1517.1K	
FY83	**FY84**	**FY85**	**FY86**	**TOTAL**
219.0K	128.0K	140.8K	154.4K	10.05M

NOTE: Cost of development is for FCS product and does not include Future Extensions.

11.0 PROJECT RISKS

* New Technology
* ECO Impact
* Development Resources
* Product Cost

11.1 NEW TECHNOLOGY

Any new technology has certain risks associated with it and the MCA development will have to be monitored closely to avoid problems in the manufacturers process and with Digital's design tools. The Dolphin project is timed in a way that will give us early warning of component or logistic problems.

11.2 ECO IMPACT

Changes to the MCA designs could result in significant schedule impact during the design phase of development or high engineering costs during manufacturing start-up.

The former risk will be addressed by comprehensive design simulations at the chip and subsystem level. A strategy will be developed to deal with MCA replacement at the debug stages.

Material during Manufacturing start-up will have to be strictly controlled to reduce the exposure to high engineering costs as a result of scrapping obsolete parts.

11.3 DEVELOPMENT RESOURCES

The major risk to the project schedule will be the availibility of sufficient resources for MCA production and Computer Aided Design.

The combined requirements of Dolphin and Venus for MCA layouts and manufacturing will surely strain the capacity of Motorola and DEC Worcester. Careful planning will be necessary to phase the requirements to meet the capacity limits.

Extensive SUDS, SAGE, and MIMIC activity could cause problems with shared PDP10 resources at the Tewksbury facility. Detailed analysis of Comet's PDP10 usage will help us determine our needs and set requirements on computer time or determine impact on schedule.

11.4 PRODUCT COST

Meeting the Entry-Level system price goal represents high risk at this time. Further study and analysis are necessary.

Large system price goal should be attainable but could be adversely affected by increased costs of peripherals which are currently in the development stages.

Meeting the cost objectives of this project will be top priority. If a choice must be made, functionality and performance will be traded off to the extent that the product remains competetive.

APPENDIX A

SBI BANDWIDTH CONSIDERATIONS

WHY 13MB IS PLENTY OF BUS BANDWIDTH

Jud Leonard

There is a good deal of misunderstanding surrounding the SBI and Memory Subsystem on the 11/780, and this has caused fear that the SBI may be inadequate as a DMA path for VENUS.

While it is true that the SBI is capable of sustaining 13.3 Mbytes/Second for arbitrary periods, this bandwidth is not available to any single nexus or from a single memory controller, and a variety of circumstances can reduce the available bandwidth far below that level. In particular, masked writes (as generated by the STAR CPU for all writes) not only use the SBI bandwidth relatively inefficiently, they invoke a logic sequence in the memory controller which results in memory throughput falling from 10.0 Mbyte/Sec to 3.3 Mbytes, or less.** Furthermore, this loss of memory throughput is compounded by the fact that the CPU can get into a lock-step sequence with the memory controller in which even higher priority devices (like the MBA) are excluded from getting memory cycles.

Obviously, having encountered these problems in the 11/780, we will design to avoid them in VENUS, and will conduct thorough simulation studies to test the design. In fact, many of the design ideas to solve these problems are already implemented in the MA780 and/or are going into the DOLPHIN bus design.

1. CPU traffic will have lowest priority for memory access. In 11/780, we mistakenly believed that this was a simple consequence of the fact that the CPU has lowest priority for the bus. In fact, one does not follow from the other.

2. CPU/memory traffic will not take place over the SBI. In 780, the CPU is a major consumer of SBI bandwidth, particularly for writes, and most especially during lengthy multiple-write sequences as in PUSHR, CALL, or MOVC. We could not reasonably achieve the projected performance levels for VENUS if all memory requests by the processor had to use the SBI.

3. VENUS raw memory bandwidth will be well in excess of that required to support all SBI requests. In 780, the memory controller design was optimized to match the memory throughput to the bus throughput. We did not understand at that time either the degradation that masked writes would cause, or the catastrophic avalanche effects that could be caused by a temporary reduction of throughput.

In combination with these improvements in the memory subsystem, we expect that future bandwidth requirements from disks will be lower, not higher, than today's. It has long been understood that the average level of I/O activity, in even our busiest systems, does not justify high bandwidth busses. It is only the peak loading due to multiple simultaneous disk and/or tape transfers that has forced us to build high bandwidth busses and to specify complex configuration rules.

Fortunately, intelligent disk controllers with full sector buffering have recently become a cost-effective alternative to traditional hardwired disk controllers, and they seem to be the only viable solution to many of the technological problems facing disk designers. A happy side effect of this buffering is that it is no longer necessary to make the controller-to-main memory transfer at the same rate as the data comes off the disk. On the other hand, it is feasible to transfer between controller and main memory faster than the disk data rate if the controller and the connection to it will support such a rate, and on the other hand, the only cost of being slower is a small apparent increase in disk latency. Thus, the advent of NDS with full sector buffering in the controller will allow us to plan for the average transfer rate of our largest systems, rather than the peak rate.

The average I/O rate of any computer system is a function of a number of variables, but an upper limit is imposed by the processor's ability to process or generate data. Amdahl has estimated this limit at 1 bit of I/O per instruction executed. Strecker believes this estimate to be correct only for batch systems, and that interactive systems may reach 1 Byte/instruction. In any case, generously rating VENUS as capable of 3.5 Megabytes/second, 13.3 megabytes/second is a factor of 4 greater than Strecker's estimate.

** DOING SUCCESSIVE OPERATIONS OF SAME FUNCTION

QUAD RDS	=	10MB/Sec	(800ns)
QUAD WRITES	=	8MB/Sec	(1000ns)
MASK READS	=	6.66MB/Sec	(600ns)
MASK WRITES (4 BYTES)	=	2.85MB/Sec	(1400ns)

DIAGRAM A

VENUS SYSTEM DIAGRAM

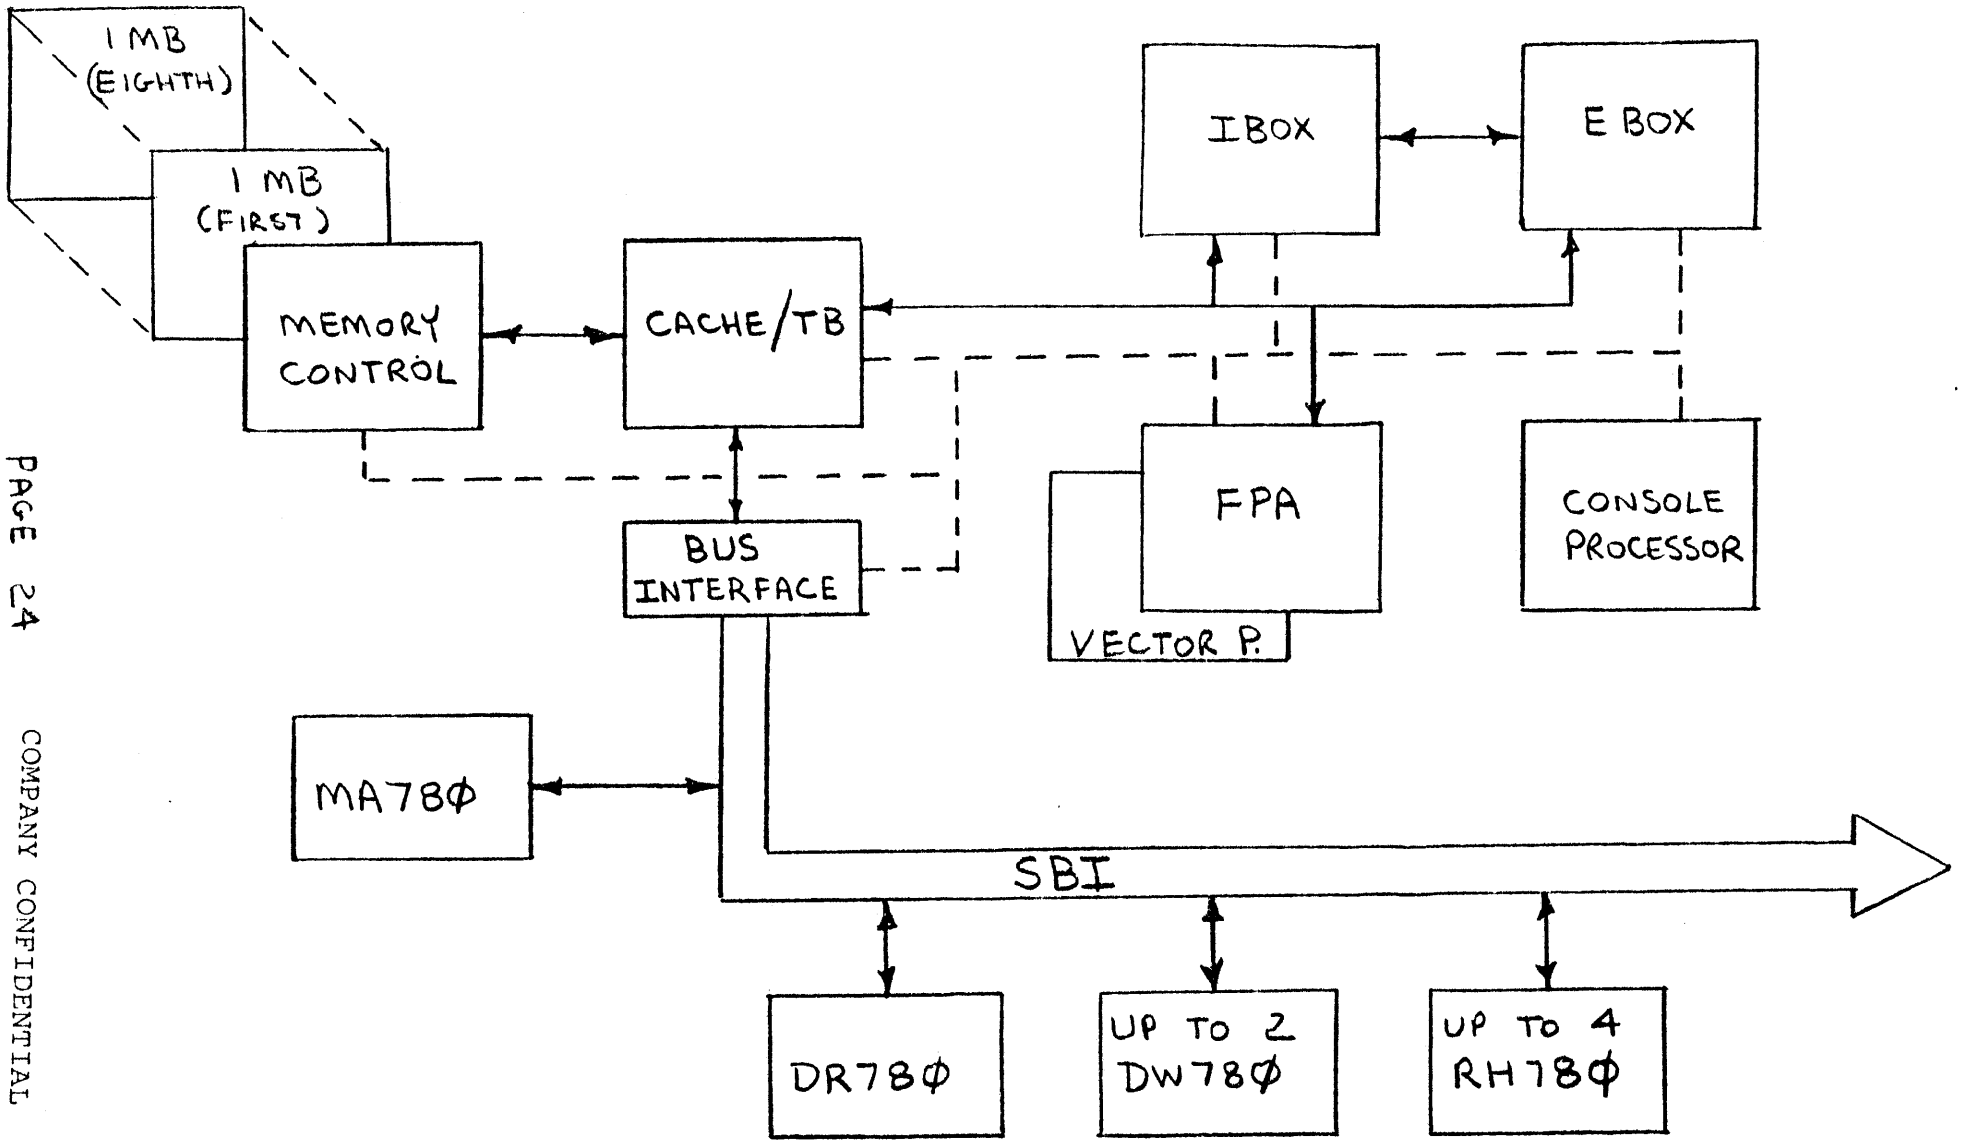

DIAGRAM B

VENUS ARRANGEMENT DRAWING