

NOTE: Art still being compiled; not included in this draft.

III. THE PDP-11 FAMILY

last edit 1/18/78

latest edit 2/16/78

PDP-11 GLUE

This part could stand alone as a book on the evolution of the PDP-11 computer structure, although it does rely on the conceptual framework of Chapter 1, and the ISPS language description given in appendix 00. The 11 has evolved quite differently from the other computers in this book, and as such provides an independent and interesting story. The factors that have created the machines are clearly market and technology based; they have generated a large number of implementations (ten) over a relatively short (eight-year) lifetime. Because there are multiple implementations spanning a performance range at the same points in time, the PDP-11 provides problems and insight which did not occur in the evolutions of the traditional mini (8 Family); the best cost/performance machines (18-bit), and the high performance machines (the DECsystem 10). The 11 designs cover a range of 500:1 in system price (\$500 to \$250,000) and 500:1 in memory size (4 Kwords to 2 Mwords).

The part is divided into six sections.

1. Introduction.

The first chapter (00), published when the 11 was announced, introduces the architecture, gives its goals, and predicts how it might evolve.

The family notion is quite strong, although not specific about members. Chapter 00, ~~What Have We Learned From PDP-11~~, ^{The after three design generations} might be read next in order to get the broadest overview, and best immediate critique of the 11 evolution.

2. Conceptual Basis.

This section contains two papers that form some of the conceptual basis for the models. Strecker, Chapter 00, describes the cache memory mechanism for building high performance models (specifically the 11/70). Levy describes the intercommunication problem among physical components and how busses carry out this task.

3. Implementations.

Of the four implementation chapters, three are on specific implementations (LSI-11, CMU-11 and 11/60) and the fourth (Chapter 00) is a study of all models. This latter chapter provides comparative data on the various implementations, ~~together with how the designs fit a conceptual model.~~

4. Evaluation.

Chapter 00 evaluates the 11 as a machine for executing FORTRAN. ~~What~~ ^{The} ~~We Have Learned From PDP-11~~, ^{after three design generations} Chapter 00, discusses aspects of the PDP-11 evolution in terms of the ^{views} ~~models~~ introduced in Chapter 1.

5. VAX-11.

This paper, by the architect of VAX-11, discusses the new architecture and its first implementation, the VAX-11/780.

6. The multiprocessor research computers of ~~Carnegie-Mellon University~~.

built at Carnegie-Mellon University are discussed

Three multiprocessors: C.mmp (Chapter 00), Cm* (Chapter 00) and C.vmp (Chapter 00) give examples of a 16-processor multiprocessor, a set of computer modules based on LSI-11, and a triplicated, voting multiprocessor computer for high reliability.

A fourth multiprocessor built at DEC is PULSAR, is also discussed.

INTRODUCTION

A NEW ARCHITECTURE FOR MINICOMPUTERS--THE DEC PDP-11

It is somewhat anticlimactic to discuss this original PDP-11 description here because Chapter 00 explicitly discusses ~~What We Have Learned~~ from the PDP-11. The purpose of the chapter was originally threefold: to give the PMS and ISP architecture of the PDP-11 as it was first proposed, to describe the first (11/20) implementation, and to show possible extensions. This was attempted at a time when the whole family architecture had not been ~~worked out or even~~ fully considered.

The computer class definitions (given in 1970) of micro, mini and midi have stood the rest of time for they correspond quite closely to those of *view 3 in*

Chapter 1.

The major reasons (elaborated upon in Chapter 00) for the disparity between the predicted and actual evolution are:

1. The notion of designing with improved technology, especially for a family, was not understood in 1970. This understanding came later and was put forth in a paper in 1972 (Bell, Chen, Rege).
2. The Unibus proved unacceptable for intercommunications at the very high and low end designs. Although this chapter posits a multiprocessor and multiple Unibusses for high end designs, ^{such a particular} ~~this exact~~ structure did not evolve as a standard. Levy's chapter elaborates on the bus evolution.
3. The address space for both physical and virtual memory was too small.
4. The particular data-type extensions were not predicted. While floating point arithmetic was discussed, the character string and decimal operations were not described. These data types evolved in response to market need and COBOL -- factors which did not exist in 1970.

We have made a major change in the chapter by removing the original ISP description and replacing it with a correct and complete (by adding memory management and floating point) ISPS description given in Appendix 0 of the book.

CACHE MEMORIES FOR PDP-11 FAMILY COMPUTERS

Chapter 00 by Strecker is included for four reasons: it is a clear exposition of the cache memory structure and its design parameters; the cache is the basis of a fast PDP-8 [Bell et al, 1974], the 11/60 (Chapter 00) the 11/70 (page 00), and the KL10 (Chapter 00); the design methodology is well done--it is good engineering; and finally, the paper is well-written--in fact, it received the award for the Best Paper at the Third Computer Architecture Conference. We also publish it simply to serve as an encouragement and an example for those who should understand and describe their work -- such well-written and relevant papers are only too rare.

The cache design process is implicit in the way in which the work is carried out to determine the structure parameters. The relevant sensitivity plots (runs) are made to determine the effect of each parameter on the design. In the 11/60, Mudge (Chapter 00) uses Strecker's program traces and methodology. ¹ Footnote 1 Note that it is easy to collect statistics about PDP-11 program behavior since the trace bit in the PS word, permits the 11 to interpret itself on a single-instruction basis. One of the important parameters to understand is the time between changes of context because all real time and multiprogrammed systems have many context switches. To the best of our knowledge this study is unique. ~~The PDP-8 cache design (Chapter 00) shows the effect of segmenting the cache for instructions and data.~~

A cache design for a PDP-8 [Bell et al, 1974] is summarized on page 00.
The 8 study shows the effect of separating instructions and data whereas
the 11 ^{studies} does not. Strecker gives the performance evaluation in terms of
cache miss ratios whereas the reader is probably interested in performance
or speed-up. These two measures, see Fig. Cachespeed, are related (Lee,
1969) in the following way (assuming an infinitely fast processor):

p = total no. of memory accesses by the processor, P_c
 m = no. of memory accesses that are missed by the
cache and how to be referred to M_p
 $t.c$ = cycle time of cache memory, M_c
 $t.p$ = cycle time of primary memory, M_p
 R = $t.p/t.c$ (ratio of memory speeds), where R is
typically 3 to 10

the relative execution speeds are:

$$t(\text{no cache}) = pR$$

$$t(\text{to cache}) = p + mR$$

$$\text{speedup} = pR/(p + mR) = R/(1 + (m/p) R) = a = \text{miss ratio} = m/p$$

therefore

$$\text{speedup} = R/(1 + aR) = 1/(a + 1/R)$$

note that if $a = 0$ (100% hit), the speedup is R ; while
if $a = 1$ (100% miss), the speedup is $R/(1 + R)$, i.e.,
the speedup is less than 1 (i.e., time to
reference both memories)

IMPLEMENTATIONS

A MINICOMPUTER-COMPATIBLE MICROCOMPUTER SYSTEM: THE DEC LSI-11

Although the paper is a descriptive narrative about the design at each of the chip, board, and backplane levels, it lacks insight that the designers at Western Digital or DEC (Duane Dickhut, Lloyd Dickman, Rich Olsen, or Mike Titelbaum) might have provided. It was written from the viewpoint of a knowledgeable user. An account of the chip-level design is available (Soha and Pohlman, 1974). The design was done at Western Digital by Roberts, Soha, and Pohlman as a relatively general purpose microprogrammed computer that could be used to emulate many computers. When DEC started working with the designers to effect interpretation of the 11 ISP, the design took on more of the 11 structure. Two design levels are described in the paper: the 3 chip microprogrammed computer as it is used to interpret the 11 ISP, and the particular PMS-level components as they are integrated into a backplane to form a hardware system. Sebern points out the microprogramming tradeoff that took place between the chip and module levels to carry out functions normally in hardware: the time clock, the console, refreshing dynamic random access MOS memory, and power fail control.

The subtleties and uniqueness of the module structure are not described, nor are the design alternatives. For example, a bounded design that is typical of one board microcomputers was considered, though not described. However, a lower-cost, one-board system, like the VT78, (page 00) has evolved and is shown in Fig. RXT-11. The RXT-11 is an integrated system containing an LSI-11 chip set, 32 Kwords of memory, connectors for six EIA interfaces, and a controller for two floppy disk drives. One-hundred and seventy-five i.c.'s were used -- to implement the same functionality using standard LSI-11 modules, 375 i.c.'s would have been used.

Fig. RXT-11

The initial module-level design for the LSI-11 was predicated on a quad-sized form factor with a conventional backplane. The modules are shown ^(Fig. LSI-11/2) on page 00. Since there were not special ICs beyond the 3 chip processor, options tended to be relatively large and often occupied a full quad module. For options that were greater than a quad size, an ingenious packaging scheme was devised. It provided interconnection points on the extra half of the module (a double sized module) which was not used as the LSI-11 Bus (requires a double sized module). This permitted multiple board, complex options, e.g., a disk controller, to be packaged as a single option with no interconnection between the boards except via the second half of the quad board. As DEC began to build special ICs to interface to the bus, the option sizes decreased to occupy a double module. This system is now known as the LSI-11/2. A backplane system with modules is shown in Fig. LSI-11/2. The options available for the two systems have evolved as shown in Table LSI-11 Options. The effect of a lower cost LSI-11 system is to provide additional applications as we discuss in Chapter 1.

Fig. LSI-11/2
Table LSI-11 Options

Table of LSI-11 Options and Extensions (1978) for LSI-11/2

	<u>LSI-11</u>	<u>LSI-11/2</u>	<u>Additions</u>
<u>Basic</u>			
Processor + 4 Kw + end ^{serial} line	quad		
Processor		double	
Real time clock	quad		
Power controller/ sequences	double		
<u>1/18/78 Memories</u>			
4 Kw RAM	double		
4 Kw core	quad		
16 Kw RAM	quad		
4 Kw PROM	double		
4 Kw PROM; 256 RAM	double		
32 Kw RAM	quad	double	
<u>Interfaces</u>			
Parallel (16 line)	double		
4 line asynchronous	quad	double	
Asynchronous (serial)	double		
Instrument (IEEE-488)	double		
1 line synchronous	quad		

USING LSI PROCESSOR BIT-SLICES TO BUILD A PDP-11--A CASE STUDY IN
MICROCOMPUTER DESIGN

This paper by the designers of CMU-11 appears both in the module part and in this part on PDP-11 implementations. The Intel 3000 bit slice, ~~herein called a Microcomputer (for Microprogrammed Processor)~~, is used to interpret ^{the} a PDP-11 ISP. The purpose of the design was to test the assertion that ^a the bit-slice based arithmetic unit with register memory and microprogrammed control would simplify the design and construction of processors. The 11 was selected as a target problem in order to avoid the temptation of changing the problem (a real ISP) to fit the building blocks (the Intel 3000 processor). Indeed, the authors observed awkwardnesses that ultimately resulted in lower (than desired) performance. In retrospect, the Intel 3000 has not become the standard bit-slice architecture that the AMD 2900 series has; perhaps it suffered from being one of the earliest. Detailed comparisons including a breakdown of the various parts of the processor design are given and compared with the LSI-11 and 11/40 designs in terms of performance and cost (IC count and number of control bits in the microprogrammed controller).

A key part of the investigation was the evaluation of the computer-aided-design tools. The Stanford University Drawing System (SUDS) and the SAGE logic simulator were the major components. SAGE was predicated on being able to detect all the design and timing flaws prior to construction of a design. The authors claim that 95% of the errors were detected in this manner. The simulated-time/real-time ratio (~~10⁶/1~~) did

10⁶/1

constrain the design process, e.g., the number of different runs used to find worst-case conditions.

DESIGN DECISIONS FOR THE PDP-11/60 MID-RANGE MINICOMPUTER

Unlike the reports from an architect's or reporter's viewpoint this chapter is a direct account of the design from the close proximity of the project. A mid-range machine is an inherently difficult design for the reasons of the designer characteristics we presented in Chapter 00, page 00. ~~[CM: unclear]~~. As neither the lowest cost nor highest performance PDP-11, it has to be the right balance of features, price, and performance against criteria that are usually extremely vague.

Four interesting aspects of computer engineering are shown in the 11/60: the cache to reduce Unibus traffic; trace-driven design of floating-point arithmetic processors; providing writeable control store; and increasing the reliability, availability and maintainability.

Whereas the Unibus was previously thought to be inadequate for high performance systems, by using a cache most processor references do not use the Unibus and so leave it free for i/o traffic. This gives high performance without the attendant cost*. The cache work is based on Strecker's (Chapter 00). The study leading to determining the block size is given. The block size can only conveniently be one since fetching multiple words would tie up the Unibus with additional traffic.

*Using Amdahl's constants, page 00, the reader might compute the bus bandwidth (for i/o traffic) and the address space needs for this speed processor given the cache and compare these needs with the Unibus. [CM: We should do this for all models in "What We've Learned" Chapter]

The use of trace data to design the floating point arithmetic is described together with the resulting design. Note that the 11/60 performs roughly at 11/70 speeds but at lower cost. The implementation of the two can be compared in the following table.

Table - Implementation of 11/60 and 11/70 (count of printed circuit boards).

	11/60	11/70
Base Pc	54	8 [CM: check 11/70 sys. manual]
Floating point	4	4
Cache	1	4
Memory management	1	2
<hr/>		
Total	10	18

Microprogramming is used to provide both increased user-level capability and increased reliability, availability and maintainability. The ~~large~~ writeable control store option is described together with its ^{novel} use for data storage ~~and various applications~~. This option has been recently used for emulating the PDP-8 at the OS/8 operating system level.

A general discussion of microprogramming is also given, especially with respect to memory technology advances (see also Chapter 2, page 00). Other

*Using Amdahl's constants, page 00, the reader might compute the bus bandwidth (for i/o traffic) and the address space needs for this speed processor given the cache and compare these needs with the Unibus. [CM: We should do this for all models in "What We've Learned" Chapter]

Does this belong here too?

semiconductor technology improvements are described together with how they affect price and performance. It is interesting to note that the simple concept of tri-state logic* had such a great effect on the design.

all CAPS { Impact of Implementation Design Tradeoffs on Performance: The PDP-11, A Case Study

This chapter presents a ~~most~~ comprehensive comparison of the eight processor implementations used in the ten PDP-11 models. The work was carried out to investigate various design styles for a given problem, ^{*namely*} ~~the~~ interpretation of the PDP-11 ISP. The tables alone give more insight into processor implementations than is available from any single source we know. The usefulness of the data also comes from having an outside observer examine the machines and share his insight.

The tables include:

1. a set of instruction frequencies, by Strecker, for a set of ten different applications. The reader should note the frequencies do not reflect all uses, e.g., there are no floating point instructions, nor has operating system code been analyzed.
2. implementation cost (modules, ICs, control store widths) and performance (micro- and macro-instruction times) for each model; and
3. a canonical data path for all 11 implementations, against which each

*Ability to interconnect a number of subsystems together through a wired-or connection.

OR
CAPS

processor is compared.

With this background data, a "top-down" model is built which explains the performance (macro-instruction time) of the various implementations in terms of the micro-instruction execution, and primary memory cycle time.

Since these two parameters do not fully explain (model) performance, a ~~set~~ ^{approach is also used} of bottom-up ~~factors must be introduced~~. These factors include various design techniques and the degree of processor overlap. We believe that this analysis of a constrained problem should provide useful insight to both computer and general-digital-systems designers.

EVALUATION

TURNING COUSINS INTO SISTERS: THE ROLE OF SOFTWARE IN SMOOTHING HARDWARE DIFFERENCES

Since FORTRAN is ^{probably} ~~quite possibly~~ the most often executed language for the PDP-11 and the one we intended that it execute, it is important to observe the 11 architecture as seen by the language processor--its user. The first FORTRAN compiler and object (run) time system are described together with the evolutionary extensions to improve performance. The FORTRAN IV-PLUS compiler is only briefly discussed since its improvements, largely due to compiler optimization technology, are less relevant to the 11 architecture.

The chapter title overstates the compatibility problem since the five

variations of the 11 ISP for floating point arithmetic are made to be compatible by essentially providing five separate object (run) time systems and a single compiler. ~~This~~ ^{discussed in the chapter.} This transparency is provided quite easily using a concept called threaded code. This concept appears to be a very simple interpreter for the PDP-11--and might not be called an interpreter by many. With threaded code, one 1-word instruction requiring two memory cycle times is executed each time a high level operation code is to be interpreted, otherwise the processor is carrying out the desired op code. When a simple integer expression like $I = I + 1$, which occupies 2 memory words and requires 3 memory cycles to execute, is transformed into a threaded code version the program still only occupies 2 words, but instead requires 5 memory cycles to execute (nearly a factor of 2). For more complex operations requiring longer execution times, like floating point arithmetic, the overhead turns out to be quite low and the space utilization is quite good. It is also possible to move efficiently between threaded and directly executed code, although this is not done. Jim Bell (Ref) discovered the technique; it has been used extensively for other compilers because of its low time and space overhead. The ability to carry out the interpretation so elegantly was not part of the original PDP-11 design, but rather was a consequence of the generality of the 11's addressing modes.

The first version of the FORTRAN machine constructed was a simple stack machine. As such, the execution times turned out to be quite long. In the second version, by recognizing the special high-frequency-of-use cases, e.g., $A = 0$, $A = A + 1$, and by having better conventions for three-address operations (to and from the stack), speedups of 1.3 and 2.0 for floating

point and integers, were obtained.

It is interesting to compare the virtual machine described with the FORTRAN IV-PLUS machine which uses the floating point processors (on the 11/34, 11/45, 11/55, 11/60, and 11/70). If the FORTRAN machine described in the paper is implemented in microcode and made to operate at FPP speeds, the resulting machines turn out to operate at roughly the same speed and programs occupy roughly the same program space.

THE PDP-11 AFTER THREE DESIGN GENERATIONS ~~WHAT HAVE WE LEARNED FROM THE PDP-11?~~

This chapter is a substantially revised version of a paper called "What Have We Learned From the PDP-11?" written for the CMU Computer Science 10th Anniversary, September 1975. This paper was written to critique the original expository paper on the PDP-11 (Chapter 00) and to compare the actual with the predicted evolution. The four critical issues of technology, bus bandwidth (and PMS structure), address space and data-type evolutions are examined.

The first part of the chapter discusses how the technology is used as a basis for the evolution (something we did not understand when the machines were originally planned). The role of semiconductor memories is especially critical. The next section describes the evolution from the point of view of the various development projects and people. Some early (historical) design documents are introduced to further aid in understanding the design process.

The Unibus evolution is given and the case is made for its optimality. The Unibus has had greater longevity and use than any of the other DEC busses and compares favorably with the IBM I/O Channel Bus as a universal standard *for* of interconnection. *XSP*

We try to provide a set of evolutionary cost-performance metrics so the 11 can be compared with the other machines (18-, 12- and 36-bit) in the book (Chapter 00, 01, and 02). Also, here we go into the unique (within DEC) problem of designing a range of machines.

Although an ISP evaluation is given, it is quite weak. By comparison, Chapter 00 by Brender, gives a more useful evaluation of the architecture for FORTRAN execution. A complete section is given on the addressing extension, beyond the 11/45 and 11/70 extensions, which required a major perturbation: VAX-11.

The final section, addressed to the research community, describes some general problems encountered in structure design and engineering, together with how solving them might be useful in subsequent designs.

VAX-11/780: A VIRTUAL ADDRESS EXTENSION TO THE DEC PDP-11 FAMILY

This chapter
~~Chapter 00, by Bill Strecker,~~ provides a clean, somewhat terse, yet comprehensive description of the VAX-11 architecture. It is among the best papers on a specific architecture that we know. Since the VAX part of the architecture is so complete in terms of data-types, operators, addressing

and memory management, it can also serve as a textbook model and base, case study for architecture in general. Goals, constraints and various design choices are given. The first model, VAX-11/780, is also briefly described.

VAX-11 is the extension to PDP-11 ~~ISP~~ to provide a large, 32-bit virtual address for each user process. ~~When operating with 32-bit addresses, call native mode, the VAX part of the architecture is in effect.~~ The architecture includes a PDP-11 mode (compatibility) for PDP-11 programs written for the RSX-11/M program environment to run unchanged. In this way, PDP-11 programs can be moved among VAX and PDP-11 computers, depending on the user's address size, computational and generality needs.

~~THE MULTIPROCESSOR RESEARCH COMPUTERS OF CARNEGIE-MELLON UNIVERSITY AND
DEC'S PULSAR~~

Three computers, which use the 11 as a basis, were built at Carnegie-Mellon University to carry out research in computer structures and operating systems for multiprocessors. A fourth multiprocessor based on the 16 LSI-11s, called the DEC PULSAR is also discussed.

The first computer, C.mmp, is a 16 processor (11/40's and 11/20's) system with 2.5 million words of shared primary memory. It was built to investigate the programming (and resulting performance) questions associated with having a large number of processors.

The chapter on the second computer, Cm*, is located physically in the

Modules Part, as the modules that form Cm* are LSI-11's. Cm* was based on the premise that ultimately, the smallest modular unit, i.e., integrated circuit, used to build digital systems would be a computer. With this premise, we need to understand how to interconnect computers physically, how to assign parts of the application program to the various computers, and how to program the complete structure.

The third computer, C.vmp, was designed to investigate how a production microcomputer, the LSI-11, could be used to build a triplicating, voting high availability computer.

The goals of the first two are performance while the third has reliability as its goal.

We believe that technology will force the evolution of computing structures to converge to three styles of multiprocessor computers: (1) C.mmp-style, for high performance, incremental performance and availability; (2) loosely coupled computers like Cm* to handle specialized processing, e.g., front end, file, and signal processing; and (3) C.vmp-style for high availability *motivated* *ing* ~~based on~~ increased maintenance costs. *by*

The technology-force argument is based on history, near term technology, and resulting price extrapolations:

1. MOS technology is increasing in both speed and density faster than the technology, e.g., ECL, from which high performance machines are built.

2. The price per chip of the single-MOS-chip processors decreases at a substantially greater rate than for the low-volume high-performance, special designs. Chips in both designs have high design costs, but the special design enjoys a much lower volume.
3. Relative to all other costs of a system, the processor cost in a low end system is essentially zero.
4. Standards in the semiconductor industry tend to form more quickly for high volume products. For example, in the 8-bit microcomputer market, one type supplies about 50% of the market and three types supply over 90%.
5. A 16-bit processor-on-a-chip, with both an address space matching its performance and appropriate data-types, has been announced. Such a commodity will form the basis for nearly all future computer designs.

DEC's PULSAR, described below, is a good example of one of the more straightforward applications of this technology. As a result of these factors, the two classes of machines (MOS-processor-on-a-chip-based and low-volume, high-performance-processor-based) have rapidly diverging costs per operation per chip. Furthermore, large scale applications have been slow to form since problem complexity increases more rapidly than program size. Therefore, most subsequent computers will be based on standard, high volume parts. For high performance machines, since processing power is available at essentially zero cost from processor-on-a-chip-based

processors, large scale computing will come from arrays of processors, just as we build arrays of 64 Kbit ICs to form memory subsystems.

Insert (2)

Caps → C.mmp--A Multi-Mini-Processor

C.mmp was motivated by the need for more computing power to solve speech recognition/signal processing problems and to understand the multiprocessor software problem. Until C.mmp, only one large, tightly coupled multiprocessor had been built--the Bell Laboratory's Safeguard Computer (BSTJ issue?).

The introductory section describes the economic and technical factors influencing multiprocessor feasibility and argues for the timeliness of the research. Various problems to be researched are given together with a discussion of particular design aspects. For example, since C.mmp is predicated on a common operating system there are two sources of degradation: memory contention and lock contention. The machine's theoretical performance as a function of memory-processor interference is based on Strecker's (1971) work. In practice, because the memory was not built with low-order address interleaving, memory interference was greater than expected. This problem was solved by moving program segments.

As the number of memory modules and processors becomes very large, the theoretical performance (as measured by the number of accesses to the memory by the processors) approaches the memory bandwidth
($m/\text{memory-cycle-time}$) $\times 1/e$. [ref.?] Thus, with infinite processors, there

(2)

The multiprocessor research projects at CMU have emphasised synthesis and measurement. Operating systems have been built for them and the executions of user programs have been carefully analyzed. All the multiprocessor interferences, overheads, and synchronization problems have been faced; the resultant performance helps to put their actual costs in perspective.

Figure HARPY looks at one of applications, the HARPY speech recognition program, in detail. ~~the perf~~ Here the performance of C.mmp and Cm* is compared with three uniprocessors (KA-10, KL-10, and H/40).

is not a maximum limit on performance, provided all processors are not contending for the same memory.

Although there is a discussion outlining the design direction of the operating system, Hydra, later descriptions should be read [Wulf et al, 1975]. Since the ~~11's~~ ^{of the PDP-11} small address necessitated frequent map changes, 11/40s with writeable control stores were used to implement the operating systems ^{which} calls ~~to~~ change the segment base registers.

The C.mmp which was actually built is shown in Fig. PMSC.mmp.

There are three basic approaches to the effective application of multiprocessors:

1. System-level workload decomposition. If a workload contains a lot of inherently independent activities, e.g., compilation, editing, file processing, and numerical computation, it will naturally decompose.
2. Program decomposition by a programmer. Intimate knowledge of the application is required for this time-consuming approach.
3. Program decomposition by the compiler. This is the ideal approach. However, results to date have been disappointing [Ref. Illiac IV FORTRAN compiler projects].

C.mmp was predicated on the first two approaches. Since the original

paper, ALGOL 68, a language with facilities for expressing parallelism in programs, has been implemented. It has assisted greatly with program decomposition. See Figure x.

As can be seen in the paper, a model of the lock problem influenced the number of critical sections in the scheduler. Since the paper, the operating system Hydra has been designed and implemented. An extensive description is given in [Wulf et al, 1975].

Two experiments analyzing the performance of C.mmp and Hydra have been reported. The first, by ~~Ren~~^{Mar}athe [1977], used a hardware monitor to measure the degradation due to the locking mechanism which is invoked when shared data is accessed. The second, by Oleinick [1977], analyzed user-program-level synchronization. We summarize the results of both below.

The contention for shared resources in a multiprocessor system occurs at several levels. At the lowest level, processors contend at the cross-point switch level for memory. This memory interference is discussed in the chapter. On a higher level there is contention for shared data in the operating system kernel. At higher levels, processes contend for i/o devices and for software processes, e.g., for memory management. The following table points to models on experimental data at these different levels in C.mmp.

Contention Level

Reference

Table?

user-program

[Oleinick, 1977]

[Fuller and Oleinick, 1976]

Marathe and Fuller, 1977]

Hydra Kernel objects

[*Marathe*

cross-point switch

Chapter XX

[Fuller, 1976]

Mar
~~Mar~~athe's data show that the shared data of Hydra is organized into enough separate objects that a very small degradation (less than 1%) results from contention for these objects. Table LOCK is reproduced from his paper. He also built a queueing model which projected that the contention level would be about 5% in a 48 processor system.

*Table
LOCK*

Replace with section A

~~Oleinick uses a root finding algorithm to study various implementations of a single problem.~~

~~[CM: elaborate after studying his thesis chapter.]~~

Caps → Multi-Microprocessors: An Overview and Working Example

The Cm* work, sponsored by NSF and ARPA, is an extension of earlier NSF-sponsored research [Bell, *et al* etc. 1973] on register-transfer-level modules. As LSI and VLSI enable construction of the processor-on-a-chip, it is apparent that low-level, register-transfer modules are passe' for the construction of all but low-volume computers. Although the research is

[Aleinick's data on C.mmp
- for 11 glue - follows Marath's
data in the section on C.mmp]

Section (A)

Aleinick [1978] has used C.mmp to conduct an experimental, as opposed to theoretical, study of implementation of parallel algorithms on a multiprocessor. He ~~uses a root-finding algorithm~~ ^{studies} an extension of the bisection method for finding the roots of an equation. ~~tree~~ A decomposition of the binary search for a root into n parallel processes is to evaluate the function simultaneously at n points.

Under ideal conditions, ^{all} ~~each~~ processes would finish ~~evaluate~~ the function evaluation (required at each step) at the same time, and then some brief bookkeeping would take place to determine the next subinterval ^{this case.} for the n processes to work on. ^{Figure 4.2 shows} ~~since~~ ^{However,} because the time to evaluate the function is data dependent, some processes complete before others. Moreover, if the bookkeeping task is time consuming relative to the time to evaluate the function, the speedup ratio will suffer.

Aleinick systematically studies each source of fluctuation in performance. The dominant one is the mechanism used for process synchronization. Fig 4.2 ~~shows~~ ^{compares} the speed up obtained with ^{the} four different synchronization primitives ~~below~~ used. These are as follows: PM ϕ

PM1

Kernel

Spin lock - the processor ^{continues to} ~~simply waits~~ executes a short sequence of instructions which repeatedly test the lock.

Clearly the impact of process synchronization is a function of the ratio of synchronization time to function evaluation time. Figure 4.5 indicates that ~~spin locks~~ gives the ranges over which ~~the~~ each lock should ^{can} be used without dominating execution time.

predicated on structures employing a hundred or so processors, this chapter describes the culmination of the first (ten-processor) phase.

The authors motivate their work by appealing to the diseconomy-of-scale arguments which we advanced at the beginning of this section (page 00). To ^{and elaborate upon in Part IV.} provide additional context for their research, computer modules (Cm*), multiprocessors (C.mmp) and computer networks are described in terms of performance and problem suitability. The chapter gives a description of the modules structure, together with their associated limitations and potential research problems.

Insert section (5)

The final, most important, part of the chapter evaluates the performance of Cm* for five different problems.

Insert Section (3)

Caps { C.vmp: The Architecture and Implementation of a Fault Tolerant Multiprocessor

C.vmp is a triplicated, voting multiprocessor designed to understand the difficulty (or ease) of using standard, off-the-shelf LSI-11's to provide greatly increased reliability. There is concern for increased reliability because systems are becoming more complex, are used for more critical applications, and because maintenance costs for all systems are increasing. Because the designers themselves carry out and analyze the work, this chapter provides first-hand insight into high reliability designs and the design process--especially its evaluation. The system has operated for several months and the first phase of work is complete.

Section ③

A 50-Computer-Module CM^* is now under construction and is planned to be operational by the end of 1978 for evaluation in 1979. The extension of CM^* is known as $CM^*/50$ and is shown in Fig $CM^*/50$. It will be used to test the CMV ideas on parallel programming methods, fault tolerance, modularity, and the extensibility of the CM^* structure.

Section ⑤

The grouping of processor and memory into modules and the hierarchy of bus structures - LSI-11 bus, Map bus, and Intercluster bus - are radical departures from more conventional computer systems.

Several design goals are initially predicated and the work is carried out against the goals.

The goal of software and hardware transparency turned out to be easier to attain than expected, because of an idiosyncrasy of the floppy disk controller. Because the controller effects a word-at-a-time bus transfer from a one-sector buffer, voting can be carried out at a very low level. It is unclear how the system would have been designed without this type of controller; ^{at} as a minimum, some ^{part} ~~form~~ ^{the} of software transparency goal would ^{met} ~~not~~ have been ~~violated~~ and a significant controller modification would have been necessary.

A number of models are given by which the design is evaluated. ^{From the discussion of} ~~Various~~ component reliabilities ~~are used and~~ the reader should get ~~a great deal of~~ ^{some} insight into the factors contributing to reliability. It should be noted that a ^{custom-LSI-chip designed} ~~special hardware~~ voter is needed to get a sufficiently low cost for a marketable C.vmp. While the intent of C.vmp is not a product, it does provide much of the insight for such a product.

[PULSAR here]

Insert
Add Section (A)

Section (A)

Insert for
11 glue - Don't Lose!

Jega Arulpragasam
2 February 1978

Pulsar: A Performance Range mP System

16 LSI-11 multiprocessor computer for investigating

PULSAR is a project, begun in Autumn of 1976, aimed at investigating cost-effective ways of interconnection to create a symmetrical multiprocessing system covering a specific performance range of ~~This range is defined in terms of the total combined processor memory request rate and extends up to 5 megawords/sec. This is equivalent to a PDP-11 mip rate of about 2.0.~~ *(approximately 1 LSI-11 to ~~both show~~ on 11/70)*

the cost-effectiveness of multiple microprocessors

The particular interconnection chosen for breadboarding (so as to be able to explore more fully the software issues) is shown in Fig 1.

for simple instructions

(Fig. 1), similar is based on

structure

The breadboard system has the general functionality of the 11/70, including multiple interrupt levels and 22-bit physical addressing. It does not, however implement I&D space or supervisor mode, nor does it make provision for attaching Floating Point processors.

The processors communicate with each other (P-Boards), the Unibus Interface (UBI) and a Common Control Section (CC) via a high-bandwidth, synchronous bus. This bus could achieve a bandwidth of 8.33 megawords per second by restricting its length to <6".

Common Control and Cache

2 or 4-way interleaved

The CC contains a large (8K word), direct mapping, shared cache with a 2-word block size, interfacing to the 11/70 memory bus with 2 or more controllers providing block interleaving. Analysis shows this to be necessary to prevent the memory subsystem from becoming a severely limiting bottleneck, in spite of the large reduction in bandwidth demand provided by the cache. In addition, the CC provides all the mapping functions for both Unibus and processor accesses to memory. The Unibus map registers (à la 11/70) and the KP registers for each processor are held in a single monolithic bipolar chip array.

the control process map

The UBI provides all the Unibus control functions normally exercised by the CPU in conventional PDP-11 uniprocessor systems. It interfaces the Unibus to the P-bus to communicate with the CC for data transfer and console functions, and to the P-Boards for forwarding interruptions. The Interrupt Director mechanism in the UBI is capable of overlapping the handling of different priority levels on the Unibus side, but the forwarding of vectors to processors over the P-bus is not overlapped. Interrupts will be fielded by the first enabled processor reaching the Service Branch on I-fetch, with preferential treatment for any processor in WAIT state.

of a

The P-Boards contain two independent microprocessor chip sets with modified microcode and some support logic to provide the extended functionality

independent

Insert A, P3

needed. ~~The two chip sets act as independent processors with functionally independent adapters to the P-bus. Internal contention for the adapter is eliminated by running the two processors out of phase to each other. Such contention as does exist is resolved by the mechanism for arbitration of the P-bus itself.~~

The PULSAR has an ASCII console interfacing via a KMC-11, with modified microcode, ~~to the Unibus. In addition, there exists a debug panel with displays for every stage of the pipeline.~~

communications controller

Finally, 11/70 style Massbus Adapters are also provided.

OPERATION

The heart of the PULSAR System is a resource pipeline, ^{oriented} in which a given transaction is allocated every resource in the pipeline at specific intervals bearing a fixed relation to each other. ~~This implies that if a particular transaction does not use a specific resource, there is an interval of time when that resource will be idle because no other transaction is allowed to use it.~~

P-bus and controller

with specific time slots for each processor

This permits a single, simple arbitration mechanism, rather than separate (complex) ones for each resource. ~~That single point of arbitration is for the P-bus Address/Function lines.~~

Once these ^{pipeline is} ~~are~~ assigned to a transaction, the successive intervals of time are assigned to the following resources in order:

- 1) The mapping array.
- 2) The address translation logic.
- 3) The cache.
- 4) The ^{address} validation logic.
- 5) The data lines of the P-bus.

~~Of course, the memory subsystem itself, which is not a part of this resource pipeline, has an independent arbitration mechanism. Interfacing between these independent mechanisms is by means of queues.~~

There are some operations which require more than one access to the same resource in the pipeline. These operations are effectively handled as two transactions. Examples of such operations are Memory Writes and Internal I/O Page (say KT register) accesses. A memory write may need a second access to the cache for update,

while the Internal I/O Page may need another access to the map array.

There are other operations in which the timing does not permit the use of a particular resource in the specific interval that is allocated to that transaction. This happens, for instance, when a Read operation results in a cache miss. The data is not available in time. In this case too a second transaction takes place, initiated when backing store data becomes available.

In the case of interruptions, the ^{central control} CC does not participate, and direct communication is established between the UBI and an appropriate Pc, with the currently chosen interrupt handling algorithm. Consideration will be given to a change, ~~after the breadboard is running, that would take into account software~~ priority levels.

In that case, it is anticipated that the information regarding the software priority levels will be centralized in the CC.

In such a system a FORCE INTERRUPT message (which costs only the recognition hardware) suffices for all interprocessor communication.

The common cache has been proved to be more effective in both cost and in performance (provided the microprocessors' access time requirements are met) than a set of individual processor caches would be. It also eliminates the common stale data problem.

However, since the bandwidth of the pipeline is constrained by its slowest resource, the cache, the cache cycle has been separated into read - compare and compare - update cycles which are mutually exclusive during a single pass of a transaction through the pipeline. Careful design was needed to eliminate the stale-data-like problems this could cause, without resorting to cache invalidation.

If, however, P-bus bandwidth ever does become a problem, it would be possible to alleviate this with small on-board caches for Pc's that cached Read Only data.

A- Console operations are effected by the UBI interrogating or changing a save area for each processor, physically held in the Mapping Array, in response to ASCII console messages over the unibus. Each processor places all appropriate status in the save area on every HALT, and restores from the save area prior to acting upon every CONTINUE or START.

SOFTWARE CONTENTION

Fig 2 shows the effective performance of the system, with and without the effects of software contention.

Preliminary measurements on the 11/70 Program Development System suggest that average executive occupancy is about 5%. Accordingly, we plan to run the system with a single global lock on the executive, as the degradation expected is small.

However, executive occupancy levels of 30% have been measured over periods of a few minutes. We have therefore designed a hierarchical multiple lock system with a lock of locks, which is set for very brief periods of time. But an implementation will await the collection of data on the average time spent under each of the several different potential "sublocks."

COST COMPARISON

Cost projections indicate that a multiprocessor will have an increase in parts count over an equivalent performance uniprocessor ~~at each in the range. The cost penalty for~~ ^{each possible} ~~a~~ ^{increase} ~~approaching~~

This will range from about 20% for a 2 Pc mP, down to close to 0% at the top of the range. It is to be noted that the 20% premium can be reduced to ~~less than 10%~~ if provision is not made for expansibility over the entire range. Further, the premium is based on parts count only and excludes considerations of cost benefits due to volume.

Clearly a separate 1 Pc structure can be cost-effective (since this is the LSI-11).

mass production learning,
common spares and
manuals, lower engineering costs, etc.

Marathe, M. and Fuller, S. H. A Study of Multiprocessor Contention for Shared Data in C.mmp, Proceedings 1977 ACM SIGMETRICS Conferences, pp. 255-262.

Wulf, W. [and others] The Hydra Operating System, Fifth ACM SIGOPS Symposium on Operating Systems Principles (Nov. 1975).

new page

11 Glue Figures and Tables

Fig. Cachespeed - Structure of Pc, Mcache and Mp of cached computer.

Fig. LSI-11/2 - Photograph of double height modules forming LSI-11/2.

Fig. RXT11 - Bounded LSI-11.

Fig. HARPY

23 Fig. PMSC.mmp - PMS diagram of C.mmp.

Table Lock - Measurement of the locking behavior in Hydra/C.mmp.

Fig Cm*/50

1 Table LSI-11 Options

Table Implementation of 11/60 and 11/70.

Figure x

p. 24 April 68

Table ?

(in text)

last edit 2/16/78

latest edit 3/2/78

PDP-11 GLUE

This part could stand alone as a book on the evolution of the PDP-11 computer structure, although it does rely on the conceptual framework of Chapter 1, and the ISPS language description given in appendix 00. The 11 has evolved quite differently from the other computers in this book, and as such provides an independent and interesting story. The factors that have created the machines are clearly market and technology based; they have generated a large number of implementations (ten) over a relatively short (eight-year) lifetime. Because there are multiple implementations spanning a performance range at the same points in time, the PDP-11 provides problems and insight which did not occur in the evolutions of the traditional mini (8 Family); the best cost/performance machines (18-bit), and the high performance machines (the DECsystem 10). The 11 designs cover a range of 500:1 in system price (\$500 to \$250,000) and 500:1 in memory size (4 Kwords to 2 Mwords).

The part is divided into six sections.

1. Introduction.

The first chapter (00), published when the 11 was announced, introduces the architecture, gives its goals, and predicts how it might evolve.

The family notion is quite strong, although not specific about members. Chapter 00, The PDP-11 After Three Design Generations, might be read next in order to get the broadest overview, and best immediate critique of the 11 evolution.

2. Conceptual Basis.

This section contains two papers that form some of the conceptual basis for the models. Strecker, Chapter 00, describes the cache memory mechanism for building high performance models (specifically the 11/70). Levy describes the intercommunication problem among physical components and how busses carry out this task.

3. Implementations.

Of the four implementation chapters, three are on specific implementations (LSI-11, CMU-11 and 11/60) and the fourth (Chapter 00) is a study of all models. This latter chapter provides comparative data on the various implementations.

4. Evaluation.

Chapter 00 evaluates the 11 as a machine for executing FORTRAN. The PDP-11 After Three Design Generations, Chapter 00, discusses aspects of the PDP-11 evolution in terms of the views introduced in Chapter 1.

5. VAX-11.

This paper, by the architect of VAX-11, discusses the new architecture and its first implementation, the VAX-11/780.

6. Multiprocessor research computers.

Three multiprocessors built at Carnegie-Mellon University are discussed: C.mmp (Chapter 00) a 16-processor multiprocessor, Cm* (Chapter 00) a set of computer modules based on LSI-11, and C.vmp (Chapter 00) a triplicated, voting multiprocessor computer for high reliability. A fourth multiprocessor, PULSAR, ^{built at DEC} is also discussed.

INTRODUCTION

A NEW ARCHITECTURE FOR MINICOMPUTERS--THE DEC PDP-11

It is somewhat anticlimactic to discuss this original PDP-11 description here because Chapter 00 explicitly discusses what we have learned from the PDP-11. The purpose of the chapter was originally threefold: to give the PMS and ISP architecture of the PDP-11 as it was first proposed, to describe the first (11/20) implementation, and to show possible extensions. This was attempted at a time when the whole family architecture had not been fully considered.

The computer class definitions (given in 1970) of micro, mini and midi have

stood the rest of time for they correspond quite closely to those of view 3 in Chapter 1.

The major reasons (elaborated upon in Chapter 00) for the disparity between the predicted and actual evolution are:

1. The notion of designing with improved technology, especially for a family, was not understood in 1970. This understanding came later and was put forth in a paper in 1972 (Bell, Chen, Rege).
2. The Unibus proved unacceptable for intercommunications at the very high and low end designs. Although this chapter posits a multiprocessor and multiple Unibusses for high end designs, such a structure did not evolve as a standard. Levy's chapter elaborates on the bus evolution.
3. The address space for both physical and virtual memory was too small.
4. The particular data-type extensions were not predicted. While floating point arithmetic was discussed, the character string and decimal operations were not described. These data types evolved in response to market need and COBOL -- factors which did not exist in 1970.

We have made a major change in the chapter by removing the original ISP description and replacing it with a correct and complete (by adding memory management and floating point) ISPS description given in Appendix 0 of the book.

CACHE MEMORIES FOR PDP-11 FAMILY COMPUTERS

Chapter 00 by Strecker is included for four reasons: it is a clear exposition of the cache memory structure and its design parameters; the cache is the basis of a fast PDP-8 [Bell et al, 1974], the 11/60 (Chapter 00) the 11/70 (page 00), and the KL10 (Chapter 00); the design methodology is well done--it is good engineering; and finally, the paper is well-written--in fact, it received the award for the Best Paper at the Third Computer Architecture Conference. We also publish it simply to serve as an encouragement and an example for those who should understand and describe their work -- such well-written and relevant papers are only too rare.

The cache design process is implicit in the way in which the work is carried out to determine the structure parameters. The relevant sensitivity plots (runs) are made to determine the effect of each parameter on the design. In the 11/60, Mudge (Chapter 00) uses Strecker's program traces and methodology¹. One of the important parameters to understand is the time between changes of context because all real time and multiprogrammed systems have many context switches. To the best of our knowledge this study is unique.

A cache design for a PDP-8 [Bell et al, 1974] is summarized on page 00. The 8 study shows the effect of separating instructions and data whereas the 11 study does not. Strecker gives the performance evaluation in terms of cache miss ratios whereas the reader is probably interested in

¹Note that it is easy to collect statistics about PDP-11 program behavior since the trace bit in the PS word, permits the 11 to interpret itself on a single-instruction basis.

performance or speed-up. These two measures, see Fig. Cachespeed, are related (Lee, 1969) in the following way (assuming an infinitely fast processor):

p = total no. of memory accesses by the processor, P_c
 m = no. of memory accesses that are missed by the cache and how to be referred to M_p
 $t.c$ = cycle time of cache memory, M_c
 $t.p$ = cycle time of primary memory, M_p
 R = $t.p/t.c$ (ratio of memory speeds), where R is typically 3 to 10

the relative execution speeds are:

$$t(\text{no cache}) = pR$$

$$t(\text{to cache}) = p + mR$$

$$\text{speedup} = pR/(p + mR) = R/(1 + (m/p) R) = a = \text{miss ratio} = m/p$$

therefore

$$\text{speedup} = R/(1 + aR) = 1/(a + 1/R)$$

note that if $a = 0$ (100% hit), the speedup is R ; while

if $a = 1$ (100% miss), the speedup is $R/(1 + R)$, i.e.,

the speedup is less than 1 (i.e., time to

reference both memories)

IMPLEMENTATIONS

A MINICOMPUTER-COMPATIBLE MICROCOMPUTER SYSTEM: THE DEC LSI-11

Although the paper is a descriptive narrative about the design at each of the chip, board, and backplane levels, it lacks insight that the designers at Western Digital or DEC (Duane Dickhut, Lloyd Dickman, Rich Olsen, or Mike Titelbaum) might have provided. It was written from the viewpoint of a knowledgeable user. An account of the chip-level design is available (Soha and Pohlman, 1974). The design was done at Western Digital by Roberts, Soha, and Pohlman as a relatively general purpose microprogrammed computer that could be used to emulate many computers. When DEC started working with the designers to effect interpretation of the 11 ISP, the design took on more of the 11 structure. Two design levels are described in the paper: the 3 chip microprogrammed computer as it is used to interpret the 11 ISP, and the particular PMS-level components as they are integrated into a backplane to form a hardware system. Sebern points out the microprogramming tradeoff that took place between the chip and module levels to carry out functions normally in hardware: the time clock, the console, refreshing dynamic random access MOS memory, and power fail control.

The subtleties and uniqueness of the module structure are not described, nor are the design alternatives. For example, a bounded design that is

typical of one board microcomputers was considered, though not described. However, a lower-cost, one-board system, like the VT78, (page 00) has evolved and is shown in Fig. RXT-11. The RXT-11 is an integrated system containing an LSI-11 chip set, 32 Kwords of memory, connectors for six EIA interfaces, and a controller for two floppy disk drives. One-hundred and seventy-five i.c.'s were used -- to implement the same functionality using standard LSI-11 modules, 375 i.c.'s would have been used.

The initial module-level design for the LSI-11 was predicated on a quad-sized form factor with a conventional backplane. The modules are shown on page 00. Since there were not special ICs beyond the 3 chip processor, options tended to be relatively large and often occupied a full quad module. For options that were greater than a quad size, an ingenious packaging scheme was devised. It provided interconnection points on the extra half of the module (a double sized module) which was not used as the LSI-11 Bus (requires a double sized module). This permitted multiple board, complex options, e.g., a disk controller, to be packaged as a single option with no interconnection between the boards except via the second half of the quad board. As DEC began to build special ICs to interface to the bus, the option sizes decreased to occupy a double module. This system is now known as the LSI-11/2. A backplane system with modules is shown in Fig. LSI-11/2. The options available for the two systems have evolved as shown in Table LSI-11 Options. The effect of a lower cost LSI-11 system is to provide additional applications as we discuss in Chapter 1.

Table of LSI-11 Options and Extensions (1978) for LSI-11/2

	<u>LSI-11</u>	<u>LSI-11/2</u>
<u>Basic</u>		
Processor + 4 Kw + serial line	quad	
Processor		double
Real time clock	quad	
Power controller/ sequences	double	
<u>Memories</u>		
4 Kw RAM	double	
4 Kw core	quad	
16 Kw RAM	quad	
4 Kw PROM	double	
4 Kw PROM; 256 RAM	double	
32 Kw RAM	quad	double
<u>Interfaces</u>		
Parallel (16 line)	double	
Asynchronous (serial)	double	
4 line asynchronous	quad	double
1 line synchronous	quad	

Instrument (IEEE-488)	double
Direct Memory Access	quad
Foundation (general purpose)	quad
Fancy RT clock	quad
Analog-to-digital	quad
Digital-to-analog	quad
Floppy interface	double
Hard disk interface (RK05)	4 quads
Hard disk interface (RL01)	2 quads
line printer	double

Backplanes

4 x 8	for quads
2 x 4	
2 x 8	
2 x 12	

USING LSI PROCESSOR BIT-SLICES TO BUILD A PDP-11--A CASE STUDY IN
MICROCOMPUTER DESIGN

This paper by the designers of CMU-11 appears both in the module part and in this part on PDP-11 implementations. The Intel 3000 bit slice, is used to interpret the PDP-11 ISP. The purpose of the design was to test the assertion that a bit-slice based arithmetic unit with register memory and microprogrammed control would simplify the design and construction of processors. The 11 was selected as a target problem in order to avoid the temptation of changing the problem (a real ISP) to fit the building blocks (the Intel 3000 processor). Indeed, the authors observed awkwardnesses that ultimately resulted in lower (than desired) performance. In retrospect, the Intel 3000 has not become the standard bit-slice architecture that the AMD 2900 series has; perhaps it suffered from being one of the earliest. Detailed comparisons including a breakdown of the various parts of the processor design are given and compared with the LSI-11 and 11/40 designs in terms of performance and cost (IC count and number of control bits in the microprogrammed controller).

A key part of the investigation was the evaluation of the computer-aided-design tools. The Stanford University Drawing System (SUDS) and the SAGE logic simulator were the major components. SAGE was predicated on being able to detect all the design and timing flaws prior to construction of a design. The authors claim that 95% of the errors were detected in this manner. The simulated-time/real-time ratio ($10^6/1$) did constrain the design process, e.g., the number of different runs used to

find worst-case conditions.

DESIGN DECISIONS FOR THE PDP-11/60 MID-RANGE MINICOMPUTER

Unlike the reports from an architect's or reporter's viewpoint this chapter is a direct account of the design from the close proximity of the project. A mid-range machine is an inherently difficult design for the reasons of the designer characteristics we presented in Chapter 00, page 00. As neither the lowest cost nor highest performance PDP-11, it has to be the right balance of features, price, and performance against criteria that are usually extremely vague.

Four interesting aspects of computer engineering are shown in the 11/60: the cache to reduce Unibus traffic; trace-driven design of floating-point arithmetic processors; providing writeable control store; and increasing the reliability, availability and maintainability.

Whereas the Unibus was previously thought to be inadequate for high performance systems, by using a cache most processor references do not use the Unibus and so leave it free for i/o traffic. This gives high performance without the attendant cost. The cache work is based on Strecker's (Chapter 00). The study leading to determining the block size is given. The block size can only conveniently be one since fetching multiple words would tie up the Unibus with additional traffic.

The use of trace data to design the floating point arithmetic is described

together with the resulting design. Note that the 11/60 performs roughly at 11/70 speeds but at lower cost. The implementation of the two can be compared in the following table.

Table - Implementation of 11/60 and 11/70 (count of printed circuit boards).

	<u>11/60</u>	<u>11/70</u>
Base Pc	4	8 [CM: check 11/70 sys. manual]
Floating point	4	4
Cache	1	4
Memory management	1	2
	<hr/>	<hr/>
Total	10	18

Microprogramming is used to provide both increased user-level capability and increased reliability, availability and maintainability. The writeable control store option is described together with its novel use for data storage. This option has been recently used for emulating the PDP-8 at the OS/8 operating system level.

A general discussion of microprogramming is also given, especially with respect to memory technology advances (see also Chapter 2, page 00). Other semiconductor technology improvements are described together with how they

affect price and performance. It is interesting to note that the simple concept of tri-state logic* had such a great effect on the design.

IMPACT OF IMPLEMENTATION DESIGN TRADEOFFS ON PERFORMANCE: THE PDP-11, A CASE STUDY

This chapter presents a comprehensive comparison of the eight processor implementations used in the ten PDP-11 models. The work was carried out to investigate various design styles for a given problem, namely the interpretation of the PDP-11 ISP. The tables alone give more insight into processor implementations than is available from any single source we know. The usefulness of the data also comes from having an outside observer examine the machines and share his insight.

The tables include:

1. a set of instruction frequencies, by Strecker, for a set of ten different applications. The reader should note the frequencies do not reflect all uses, e.g., there are no floating point instructions, nor has operating system code been analyzed.
2. implementation cost (modules, ICs, control store widths) and performance (micro- and macro-instruction times) for each model; and
3. a canonical data path for all 11 implementations, against which each processor is compared.

*Ability to interconnect a number of subsystems together through a wired-or connection.

With this background data, a top-down model is built which explains the performance (macro-instruction time) of the various implementations in terms of the micro-instruction execution, and primary memory cycle time.

Since these two parameters do not fully explain (model) performance, a bottom-up approach is also used. These factors include various design techniques and the degree of processor overlap. We believe that this analysis of a constrained problem should provide useful insight to both computer and general-digital-systems designers.

EVALUATION

TURNING COUSINS INTO SISTERS: THE ROLE OF SOFTWARE IN SMOOTHING HARDWARE DIFFERENCES

Since FORTRAN is probably the most often executed language for the PDP-11 and the one we intended that it execute, it is important to observe the 11 architecture as seen by the language processor--its user. The first FORTRAN compiler and object (run) time system are described together with the evolutionary extensions to improve performance. The FORTRAN IV-PLUS compiler is only briefly discussed since its improvements, largely due to compiler optimization technology, are less relevant to the 11 architecture.

The chapter title overstates the compatibility problem since the five variations of the 11 ISP for floating point arithmetic are made to be compatible by essentially providing five separate object (run) time systems

and a single compiler. This transparency is provided quite easily using a concept called threaded code discussed in the chapter.

The first version of the FORTRAN machine constructed was a simple stack machine. As such, the execution times turned out to be quite long. In the second version, by recognizing the special high-frequency-of-use cases, e.g., $A = 0$, $A = A + 1$, and by having better conventions for three-address operations (to and from the stack), speedups of 1.3 and 2.0 for floating point and integers, were obtained.

It is interesting to compare the virtual machine described with the FORTRAN IV-PLUS machine which uses the floating point processors (on the 11/34, 11/45, 11/55, 11/60, and 11/70). If the FORTRAN machine described in the paper is implemented in microcode and made to operate at FPP speeds, the resulting machines turn out to operate at roughly the same speed and programs occupy roughly the same program space.

THE PDP-11 AFTER THREE DESIGN GENERATIONS

This chapter is a substantially revised version of a paper called "What Have We Learned From the PDP-11?" written for the CMU Computer Science 10th Anniversary, September 1975. This paper was written to critique the original expository paper on the PDP-11 (Chapter 00) and to compare the actual with the predicted evolution. The four critical issues of technology, bus bandwidth (and PMS structure), address space and data-type evolutions are examined.

The first part of the chapter discusses how the technology is used as a basis for the evolution (something we did not understand when the machines were originally planned). The role of semiconductor memories is especially critical. The next section describes the evolution from the point of view of the various development projects and people. Some early (historical) design documents are introduced to further aid in understanding the design process.

The Unibus evolution is given and the case is made for its optimality. The Unibus has had greater longevity and use than any of the other DEC busses and compares favorably with the IBM I/O Channel Bus as a universal standard for interconnection.

We try to provide a set of evolutionary cost-performance metrics so the 11 can be compared with the other machines (18-, 12- and 36-bit) in the book (Chapter 00, 01, and 02). Also, here we go into the unique (within DEC) problem of designing a range of machines.

Although an ISP evaluation is given, it is quite weak. By comparison, Chapter 00 by Brender, gives a more useful evaluation of the architecture for FORTRAN execution. A complete section is given on the addressing extension, beyond the 11/45 and 11/70 extensions, which required a major perturbation: VAX-11.

The final section, addressed to the research community, describes some general problems encountered in structure design and engineering, together

with how solving them might be useful in subsequent designs.

VAX-11/780: A VIRTUAL ADDRESS EXTENSION TO THE DEC PDP-11 FAMILY

This chapter provides a clean, somewhat terse, yet comprehensive description of the VAX-11 architecture. It is among the best papers on a specific architecture that we know. Since the VAX part of the architecture is so complete in terms of data-types, operators, addressing and memory management, it can also serve as a textbook model and base, case study for architecture in general. Goals, constraints and various design choices are given. The first model, VAX-11/780, is also briefly described.

VAX-11 is the extension to the PDP-11 to provide a large, 32-bit virtual address for each user process. The architecture includes a compatibility mode ^{which allows} ~~for~~ PDP-11 programs written for the RSX-11/M program environment to run unchanged. In this way, PDP-11 programs can be moved among VAX and PDP-11 computers, depending on the user's address size, computational and generality needs.

MULTIPROCESSOR RESEARCH COMPUTERS

Three computers, which use the 11 as a basis, were built at Carnegie-Mellon University to carry out research in computer structures and operating systems for multiprocessors. A fourth multiprocessor based on the 16 LSI-11s, called the DEC PULSAR is also discussed.

The first computer, C.mmp, is a 16 processor (11/40's and 11/20's) system with 2.5 million words of shared primary memory. It was built to investigate the programming (and resulting performance) questions associated with having a large number of processors.

The chapter on the second computer, Cm*, is located physically in the Modules Part, as the modules that form Cm* are LSI-11's. Cm* was based on the premise that ultimately, the smallest modular unit, i.e., integrated circuit, used to build digital systems would be a computer. With this premise, we need to understand how to interconnect computers physically, how to assign parts of the application program to the various computers, and how to program the complete structure.

The third computer, C.vmp, was designed to investigate how a production microcomputer, the LSI-11, could be used to build a triplicating, voting high availability computer.

The goals of the first two are performance while the third has reliability as its goal.

We believe that technology will force the evolution of computing structures to converge to three styles of multiprocessor computers: (1) C.mmp-style, for high performance, incremental performance and availability; (2) loosely coupled computers like Cm* to handle specialized processing, e.g., front end, file, and signal processing; and (3) C.vmp-style for high availability motivated by increasing maintenance costs.

The technology-force argument is based on history, near term technology, and resulting price extrapolations:

1. MOS technology is increasing in both speed and density faster than the technology, e.g., ECL, from which high performance machines are built.
2. The price per chip of the single-MOS-chip processors decreases at a substantially greater rate than for the low-volume high-performance, special designs. Chips in both designs have high design costs, but the special design enjoys a much lower volume.
3. Relative to all other costs of a system, the processor cost in a low end system is essentially zero.
4. Standards in the semiconductor industry tend to form more quickly for high volume products. For example, in the 8-bit microcomputer market, one type supplies about 50% of the market and three types supply over 90%.
5. A 16-bit processor-on-a-chip, with both an address space matching its performance and appropriate data-types, has been announced. Such a commodity will form the basis for nearly all future computer designs.

DEC's PULSAR, described below, is a good example of one of the more straightforward applications of this technology. As a result of these factors, the two classes of machines (MOS-processor-on-a-chip-based and

low-volume, high-performance-processor-based) have rapidly diverging costs per operation per chip. Furthermore, large scale applications have been slow to form since problem complexity increases more rapidly than program size. Therefore, most subsequent computers will be based on standard, high volume parts. For high performance machines, since processing power is available at essentially zero cost from processor-on-a-chip-based processors, large scale computing will come from arrays of processors, just as we build arrays of 64 Kbit ICs to form memory subsystems.

The multiprocessor research projects at CMU have emphasized synthesis and measurement. Operating systems have been built for them and the executions of user programs have been carefully analyzed. All the multiprocessor interferences, overheads, and synchronization problems have been faced; the resultant performance helps to put their actual costs in perspective. Figure HARPY looks at one of the applications, the HARPY speech recognition program, in detail. Here the performance of C.mmp and Cm* is compared with three uniprocessors (KA-10, KL-10, and 11/40).

C.mmp--A MULTI-MINI-PROCESSOR

C.mmp was motivated by the need for more computing power to solve speech recognition/signal processing problems and to understand the multiprocessor software problem. Until C.mmp, only one large, tightly coupled multiprocessor had been built--the Bell Laboratory's Safeguard Computer (BSTJ issue?).

The introductory section describes the economic and technical factors influencing multiprocessor feasibility and argues for the timeliness of the research. Various problems to be researched are given together with a discussion of particular design aspects. For example, since C.mmp is predicated on a common operating system there are two sources of degradation: memory contention and lock contention. The machine's theoretical performance as a function of memory-processor interference is based on Strecker's (1971) work. In practice, because the memory was not built with low-order address interleaving, memory interference was greater than expected. This problem was solved by moving program segments.

As the number of memory modules and processors becomes very large, the theoretical performance (as measured by the number of accesses to the memory by the processors) approaches the memory bandwidth $(m/\text{memory-cycle-time}) \times 1/e$.² ~~[ref. ?]~~ *Baskett and Smith, 1976* Thus, with infinite processors, there is not a maximum limit on performance, provided all processors are not contending for the same memory.

Although there is a discussion outlining the design direction of the operating system, Hydra, later descriptions should be read [Wulf et al, 1975]. Since the small address of the PDP-11 necessitated frequent map changes, 11/40s with writeable control stores were used to implement the operating systems calls which change the segment base registers.

The C.mmp which was actually built is shown in Fig. PMSC.mmp.

There are three basic approaches to the effective application of multiprocessors:

1. System-level workload decomposition. If a workload contains a lot of inherently independent activities, e.g., compilation, editing, file processing, and numerical computation, it will naturally decompose.
2. Program decomposition by a programmer. Intimate knowledge of the application is required for this time-consuming approach.
3. Program decomposition by the compiler. This is the ideal approach. However, results to date have been disappointing [Ref. Illiac IV FORTRAN compiler projects].

C.mmp was predicated on the first two approaches. Since the original paper, ALGOL 68, a language with facilities for expressing parallelism in programs, has been implemented. It has assisted greatly with program decomposition. ~~See Figure x.~~

As can be seen in the paper, a model of the lock problem influenced the number of critical sections in the scheduler. Since the paper, the operating system Hydra has been designed and implemented. An extensive description is given in [Wulf et al, 1975].

Two experiments analyzing the performance of C.mmp and Hydra have been reported. The first, by Marathe ^{and Fuller} [1977], used a hardware monitor to measure

the degradation due to the locking mechanism which is invoked when shared data is accessed. The second, by Oleinick [1978], analyzed user-program-level synchronization. We summarize the results of both below.

The contention for shared resources in a multiprocessor system occurs at several levels. At the lowest level, processors contend at the cross-point switch level for memory. This memory interference is discussed in the chapter. On a higher level there is contention for shared data in the operating system kernel. At higher levels, processes contend for i/o devices and for software processes, e.g., for memory management. ~~The following~~ ^{Contrefs} table points to models on experimental data at these different levels in C.mmp.

Contention Level

Reference

user-program

[Oleinick, 1978]

[Fuller and Oleinick, 1976]

Hydra Kernel objects

[Marathe and Fuller, 1977]

cross-point switch

Chapter XX (C.mmp chapter)

[Fuller, 1976]

Table Contrefs

References for experimental data on contention at each of three levels in the C.mmp system.

Marathe's data show that the shared data of Hydra is organized into enough separate objects that a very small degradation (less than 1%) results from

the Marathe and Fuller

contention for these objects. Table LOCK is reproduced from ~~his~~ paper. He also built a queueing model which projected that the contention level would be about 5% in a 48 processor system.

Oleinick [1978] has used C.mmp to conduct an experimental, as opposed to theoretical, study of implementation of parallel algorithms on a multiprocessor. He studies ^{rootfinder,} an extension of the bisection method for finding the roots of an equation. ^{natural} A decomposition of the binary search for a root into n parallel processes is to evaluate the function simultaneously at n points.

Under ideal conditions, all processes would finish the function evaluation (required at each step) at the same time, and then some brief bookkeeping would take place to determine the next subinterval for the n processes to work on. Figure ^I ~~ideal~~ shows this case. However, because the time to evaluate the function is data dependent, some processes complete before others. Moreover, if the bookkeeping task is time consuming relative to the time to evaluate the function, the speedup ratio will suffer.

Oleinick systematically studies each source of fluctuation in performance.

The dominant one is the mechanism used for process synchronization. Figure

Lockcomp

~~4.2~~ compares the speedup obtained with the four different synchronization primitives. These are as follows:

PM0

PM1

Replace with (2)

(2)

Four different locks for process synchronization are available to the C.mmp user. The spin lock is the most rudimentary. It does not cause an entry to HYDRA; it is a short sequence of instructions which continually tests a semaphore until it can set it successfully. The P and V operations are in fact the following POP-11 code sequences.

```
P:  CMP SEMAPHORE, #1      ;SEMAPHORE=1?
    BNE P                  ;loop until it is 1
    DEC SEMAPHORE          ;Decrement SEMAPHORE
    BNE P                  ;If neq 0 go to P

V:  MOV #1, SEMAPHORE     ;Reset SEMAPHORE to 1
```

Altho this repeating polling is extremely fast it has two major drawbacks. The first is that the processor is not free to do useful work. The second is that the polling process consumes memory cycles of the memobank that contains the semaphore.

The kernel semaphore, is implemented in HYDRA, is the low level synch^{ronization} mechanism used to be intended for system processes. When a process blocks or is to wake up, a state change is made for that process is made inside the kernel of HYDRA. If a process blocks while trying to P a semaphore, ~~its~~ the kernel swaps the process from the processor. However the pages belonging to the process are kept in primary memory.

The policy module semaphore is intended to be the user's primary mechanism for performing synchronization. When a process is blocked, it is not only swapped from the processor, but its pages are written back onto secondary storage. Often a process is blocked for just a few milliseconds, and ~~if a~~ ~~smaller~~ ~~length~~ ~~of~~ ~~time~~ than amount compared with the time it takes to update the pages on secondary storage (at least 32 milliseconds per page). The original policy module semaphore (PMØ) was modified by introducing a delay before beginning the updating. The delay time ~~for~~ is a parameter, e . The modified semaphore is referred to as ~~PM1~~ ~~to~~ ~~be~~ ~~PM1(e)~~ where

e is the delay time in milliseconds.

Figure lockcomp compares the performance of rootfinder for each of the four methods of synchronization: spin lock, kernel semaphore PMP, and PM1 ($e=300$). The distribution of the $F(x)$ computation was approximately normal with a mean of 72 msec & sd 18. ~~The curve for PMO sem-~~
~~shows that as soon as // is increased, degrad occurs,~~
~~because the overhead of synch is gr than the average time~~
~~The spin lock imple has the best speed up max of~~
~~about 2.8 for eight processes.*~~

of 72 milliseconds and a standard deviation of 18 milliseconds. The curve for the PMO semaphore shows that as soon as parallelism is increased, degradation occurs because the overhead of synchronization is greater than the average compute time. The spin lock implementation has the best speed up maximum of about 2.8 for eight processes.*

* The speed up of 2.8 for 8 processes appears ~~unimpressive~~ ^{unimpressive} until one observes that the rootfinder ~~of~~ algorithm inherently provides less than linear speedup. The theoretical speedup for 8 processes is about 2.9. ~~f~~

Kernel

Spin-lock - the processor continues to execute a short sequence of instructions which repeatedly test the lock.

Clearly the impact of process synchronization is a function of the ratio of synchronization time to function evaluation time. Figure ~~4.5~~ ^{Lockrange} gives the ranges over which each lock should be used without dominating execution time.

MULTI-MICROPROCESSORS: AN OVERVIEW AND WORKING EXAMPLE

The Cm* work, sponsored by NSF and ARPA, is an extension of earlier NSF-sponsored research [Bell, et al. 1973] on register-transfer-level modules. As LSI and VLSI enable construction of the processor-on-a-chip, it is apparent that low-level, register-transfer modules are passe' for the construction of all but low-volume computers. Although the research is predicated on structures employing a hundred or so processors, this chapter describes the culmination of the first (ten-processor) phase.

The authors motivate their work by appealing to the diseconomy-of-scale arguments which we advanced at the beginning of this section (page 00) and elaborate upon in Part IV. To provide additional context for their research, computer modules (Cm*), multiprocessors (C.mmp) and computer networks are described in terms of performance and problem suitability. The chapter gives a description of the modules structure, together with their associated limitations and potential research problems.

The grouping of processor and memory into modules and the hierarchy of bus structures - LSI-11 bus, Map bus, and Intercluster bus - are radical departures from more conventional computer systems.

The final, most important, part of the chapter evaluates the performance of Cm* for five different problems.

A 50-Computer-Module Cm* is now under construction and is planned to be operational by the end of 1978 for evaluation in 1979. The extension of Cm* is known as Cm*/50 and is shown in Fig. Cm*/50. It will be used to test the CMU ideas on parallel programming methods, fault tolerance, modularity, and the extensibility of the Cm* structure.

**C.vmp: THE ARCHITECTURE AND IMPLEMENTATION OF A FAULT TOLERANT
MULTIPROCESSOR**

C.vmp is a triplicated, voting multiprocessor designed to understand the difficulty (or ease) of using standard, off-the-shelf LSI-11s to provide greatly increased reliability. There is concern for increased reliability because systems are becoming more complex, are used for more critical applications, and because maintenance costs for all systems are increasing. Because the designers themselves carry out and analyze the work, this chapter provides first-hand insight into high reliability designs and the design process--especially its evaluation. The system has operated for several months and the first phase of work is complete.

Several design goals are initially predicated and the work is carried out against the goals.

The goal of software and hardware transparency turned out to be easier to attain than expected, because of an idiosyncrasy of the floppy disk controller. Because the controller effects a word-at-a-time bus transfer from a one-sector buffer, voting can be carried out at a very low level. It is unclear how the system would have been designed without this type of controller; at a minimum, some part of the software transparency goal would not have been met and a significant controller modification would have been necessary.

A number of models are given by which the design is evaluated. From the discussion of component reliabilities the reader should get some insight into the factors contributing to reliability. It should be noted that a custom-LSI-designed voter is needed to get a sufficiently low cost for a marketable C.vmp. While the intent of C.vmp is not a product, it does provide much of the insight for such a product.

PULSAR: A PERFORMANCE RANGE mP SYSTEM

PULSAR is a 16 LSI-11 multiprocessor computer for investigating the cost-effectiveness of multiple microprocessors. It covers a performance range of approximately 1 LSI-11 to better than an 11/70 for simple instructions.

PULSAR

The breadboard system (Fig. 1), is based on the 11/70 structure, including multiple interrupt levels and 22-bit physical addressing. It does not implement I&D space or supervisor mode, nor does it have Floating Point processors.

The processors communicate with each other (P-Boards), the Unibus Interface (UBI) and a Common Cache and Control via a high-bandwidth, synchronous bus.

The Common Cache and Control contains a large (8K word), direct mapping, shared cache with a 2-word block size, interfacing to the 2- or 4-way interleaved 11/70 memory bus. This prevents the memory subsystem from becoming a bottleneck, in spite of the large reduction in bandwidth demand provided by the cache. The control provides all the mapping functions for both Unibus and processor accesses to memory. The Unibus map registers (a la 11/70) and the process map registers for each processor are held in a single bipolar memory.

The UBI provides the Unibus control functions of a conventional PDP-11. Interrupts will be fielded by the first enabled processor with preferential treatment for any processor in WAIT state.

Each P-Board contains two independent microprocessor chip sets with modified microcode. Internal contention for the adapter is eliminated by running the two processors out of phase to each other. Such contention as does exist is resolved by the mechanism for arbitration of the P-bus itself. The PULSAR has an ASCII console interfacing via a KMC-11

communications controller, with modified microcode. In addition, a debug panel has displays for every stage of the P-bus and controller pipeline.

Console operations are effected by the UBI interrogating or changing a save area for each processor, physically held in the Mapping Array, in response to ASCII console messages over the unibus. Each processor places all appropriate status in the save area on every HALT, and restores from the save area prior to acting upon every CONTINUE or START.

OPERATION

The PULSAR System is pipeline oriented with specific time slots for each processor. This permits a single, simple arbitration mechanism, rather than separate (complex) ones for each resource.

Once the pipeline is assigned to a transaction, the successive intervals of time are assigned to the following resources in order:

1. The mapping array.
2. The address translation logic.
3. The cache.
4. The address validation logic.

5. The data lines of the P-bus.

The memory subsystem, which is not a part of this resource pipeline, has an independent arbitration mechanism. Interfacing between these independent mechanisms is by means of queues.

There are some operations which require more than one access to the same resource in the pipeline. These operations are effectively handled as two transactions. Examples of such operations are Memory Writes and Internal I/O Page (say ~~KT~~ ^{memory-management} register) accesses. A memory write may need a second access to the cache for update, while the Internal I/O Page may need another access to the map array.

There are other operations in which the timing does not permit the use of a particular resource in the specific interval that is allocated to that transaction. This happens, for instance, when a Read operation results in a cache miss. The data is not available in time. In this case too a second transaction takes place, initiated when backing store data becomes available.

COST COMPARISON

Cost projections indicate that a multiprocessor will have an increase in parts count over each possible equivalent performance uniprocessor in the range.

This will range from a 20% increase for a 2 Pc mP, approaching 0% at the top of the range. It is to be noted that the 20% premium can be reduced if provision is not made for expansibility over the entire range. Clearly a separate 1 Pc structure can be cost-effective (since this is the LSI-11). The premium is based on parts count only and excludes considerations of cost benefits due to production learning, common spares and manuals, lower engineering costs, etc.

Marathe, M. and Fuller, S. H. A Study of Multiprocessor Contention for Shared Data in C.mmp, Proceedings 1977 ACM SIGMETRICS Conferences, pp. 255-262.

Wulf, W. [and others] The Hydra Operating System, Fifth ACM SIGOPS Symposium on Operating Systems Principles (Nov. 1975).

Fuller, S. H., and Oleinick, P. N. "Initial measurements of parallel programs on a multi-processor multi-mini processor," IEEE Comp Con, pp 358-363 (1976).

Baskett, F. and Smith, A. J. "Interference in Multiprocessor Computer Systems with interleaved memory", CACM 19, 6 pp ~~327~~ 327-334 (June, 1976)

~~Fuller~~

Fuller, S. H. "Price/Performance Comparison of C.mmp and the PDP-10" Proc. of the Third Annual Symposium on Computer Architecture, pp ~~19-24~~ (1968) pp 195-202 (Jan 1976).

Wulf et al ~~5th Syops Nov 75~~

Oleinick, P. N. The Implementation of Parallel Algorithms on a Multiprocessor. Ph. D. Thesis, Carnegie-Mellon University Computer Science Dept., (in preparation).

11 Glue Figures and Tables

Fig. Cachespeed - Structure of Pc, Mcache and Mp of cached computer.

Fig. LSI-11/2 - Photograph of double height modules forming LSI-11/2.

Fig. RXT11 - Bounded LSI-11.

Fig. HARPY

Fig. PMSC.mmp - PMS diagram of C.mmp.

Table Lock - Measurement of the locking behavior in Hydra/C.mmp.

Figure Ideal
Fig. Lockcomp
Fig. Lockrange
Fig. Cm*/50

Figure PULSAR

Table LSI-11 Options (in text)

Table Implementation of 11/60 and 11/70 (in text)

~~Figure x p.24 Algol 68~~

Table */ Contrefs* (in text)

Note There are corrections here to

enter into file.

Gordon

last edit 12/16/77

latest edit 1/18/78

the

1/28/78. SB.

also, note Steve Teicher

needs to add
look at a

figure.

PDP-11 GLUE

This part could stand alone as a book on the evolution of the PDP-11 computer structure, although it does rely on the conceptual framework of Chapter 1, and the ISPS language description given in appendix 00. The 11 has evolved quite differently from the other computers in this book, and as such provides an independent and interesting story. The factors that have created the machines are clearly market and technology based, they have generated a large number of implementations (ten) over a relatively short (eight-year) lifetime. Because there are multiple implementations spanning a performance range at the same points in time, the PDP-11 provides problems and insight which did not occur in the evolutions of the traditional mini (8 Family); the best cost/performance machines (18-bit), and the high performance machines (the DECsystem 10). The 11 designs cover a range of 500:1 in system price (\$500 to \$250,000) and 500:1 in memory size (4 Kwords to 2 Mwords).

The part is divided into six sections.

1. Introduction.

The first chapter (00), published when the 11 was announced, introduces the architecture, gives its goals, and predicts how it might evolve.

The family notion is quite strong, although not specific about members. Chapter 00, What Have We Learned From PDP-11, might be read next in order to get the broadest overview, and best immediate critique of the 11 evolution.

2. Conceptual Basis.

This section contains two papers that form some of the conceptual basis for the models. Strecker, Chapter 00, describes the cache memory mechanism for building high performance models (specifically the 11/70). Levy describes the intercommunication problem among physical components and how busses carry out this task.

3. Implementations.

Of the four implementation chapters, three are on specific implementations (LSI-11, CMU-11 and 11/60) and the fourth (Chapter 00) is a study of all models. This latter chapter provides comparative data on the various implementations together with how the designs fit a conceptual model.

4. Evaluation.

Chapter 00 evaluates the 11 as a machine for executing FORTRAN. What We Have Learned From PDP-11, Chapter 00 discusses aspects of the PDP-11 evolution in terms of the models introduced in Chapter 1.

5. VAX-11.

This paper, by the architect of VAX-11, discusses the new architecture and its first implementation, the VAX-11/780.

6. The multiprocessor research computers of Carnegie-Mellon University.

Three multiprocessors: C.mmp (Chapter 00), Cm* (Chapter 00) and C.vmp (Chapter 00) give examples of a 16-processor multiprocessor, a set of computer modules based on LSI-11 and a triplicated, voting multiprocessor computer for high reliability.

INTRODUCTION

A NEW ARCHITECTURE FOR MINICOMPUTERS--THE DEC PDP-11

It is somewhat anticlimactic to discuss this original PDP-11 description here because Chapter 00 explicitly discusses "What We Have Learned from the PDP-11". The purpose of the chapter was originally threefold: to give the PMS and ISP architecture of the PDP-11 as it was first proposed, to describe the first (11/20) implementation, and to show possible extensions. This was attempted at a time when the whole family architecture had not been worked out or even fully considered.

The computer class definitions (given in 1970) of micro, mini and midi have stood the rest of time for they correspond quite closely to those of Chapter 1.

The major reasons (elaborated upon in Chapter 00) for the disparity between the predicted and actual evolution are:

1. The notion of designing with improved technology, especially for a family, was not understood in 1970. This understanding came later and was put forth in a paper in 1972 (Bell, Chen, Rege).
2. The Unibus proved unacceptable for ^{inter}most communications at the very high and low end designs. Although this chapter posits a multiprocessor and multiple Unibusses for high end designs, this exact structure did not evolve as a standard. Levy's chapter elaborates on the bus evolution.
3. The ~~physical memory~~ address space ^{for both physical and virtual memory} was too small.
^
4. The particular data-type extensions were not predicted. While floating point arithmetic was discussed, the character string and decimal operations were not described. These data types evolved in response to market need and COBOL -- factors which did not exist in 1970.

We have made a major change in the chapter by removing the original ISP description and replacing it with a correct and complete (by adding memory management and floating point) ISPS description given in Appendix 0 of the book.

CACHE MEMORIES FOR PDP-11 FAMILY COMPUTERS

Chapter 00 by Strecker is included for four reasons: it is a clear exposition of the cache memory structure and its design parameters; the cache is the basis of a fast PDP-8 [Bell et al, 1974], the 11/60 (Chapter 00) the 11/70 (page 00), and the KL10 (Chapter 00); the design methodology is well done--it is good engineering; and finally, the paper is well-written--in fact, it received the award for the Best Paper at the Third Computer Architecture Conference. We also publish it simply to serve as an encouragement and an example for those who should understand and describe their work -- such well-written and relevant papers are only too rare.

The cache design process is implicit in the way in which the work is carried out to determine the structure parameters. The relevant sensitivity plots (runs) are made to determine the effect of each parameter on the design. In the 11/60, Mudge (Chapter 00) uses Strecker's program traces and methodology. Note that it is easy to collect statistics about PDP-11 program behavior since the trace bit in the PS word, permits the 11 to interpret itself on a single-instruction basis. One of the important parameters to understand is the time between changes of context because all real time and multiprogrammed systems have many context switches. To the best of our knowledge this study is unique. The PDP-8 cache design (Chapter 00) shows the effect of segmenting the cache for instructions and data.

A cache design for a PDP-8 [Bell et al, 1974] is summarized ^{on page 00,} ~~in the~~
~~introduction to Part II. Two differences from the 11 work are of interest.~~

The 8 study shows the effect of separating instructions and data whereas the 11 does not. *

~~The second difference is that Strecker gives the performance evaluation in terms of cache miss ratios whereas the 8 data are~~ *reader is probably interested in performance or speed-up.*

~~in terms of a speed up factor.~~ These two ~~performance~~ measures, see Fig. Cachespeed, are related (Lee, 1969) in the following way (assuming an infinitely fast processor):

- p = total no. of memory accesses by the processor, Pc
- m = no. of memory accesses that are missed by the cache and how to be referred to Mp
- t.c = cycle time of cache memory, Mc
- t.p = cycle time of primary memory, Mp
- R = t.p/t.c (ratio of memory speeds), where R is typically 3 to 10

the relative execution speeds are:

$$t(\text{no cache}) = pR$$

$$t(\text{to cache}) = p + mR$$

$$\text{speedup} = pR / (p + mR) = R / (1 + (m/p) R) = a = \text{miss ratio} = m/p$$

therefore

$$\text{speedup} = R / (1 + aR) = 1 / (a + 1/R)$$

note that if $a = 0$ (100% hit), the speedup is R ; while
if $a = 1$ (100% miss), the speedup is $R/(1 + R)$, i.e.,
the speedup is less than 1 (i.e., time to
reference both memories)

IMPLEMENTATIONS

A MINICOMPUTER-COMPATIBLE MICROCOMPUTER SYSTEM: THE DEC LSI-11

Although the paper is a descriptive narrative about the design at each of the chip, board, and backplane levels, it lacks insight that the designers at Western Digital or DEC (Duane Dickhut, Lloyd Dickman, Rich Olsen, or Mike Titelbaum) might have provided. It was written from the viewpoint of a knowledgeable user. An account of the chip-level design is available (Soha and Pohlman, 1974). The design was done at Western Digital by Roberts, Soha, and Pohlman as a relatively general purpose microprogrammed computer that could be used to emulate many computers. When DEC started working with the designers to effect interpretation of the 11 ISP, the design took on more of the 11 structure. Two design levels are described in the paper: the 3 chip microprogrammed computer as it is used to interpret the 11 ISP, and the particular PMS-level components as they are integrated into a backplane to form a hardware system. Sebern points out the microprogramming tradeoff that took place between the chip and module levels to carry out functions normally in hardware: the time clock, the console, refreshing dynamic random access MOS memory, and power fail control.

The ^{design alternatives,} subtleties, and uniqueness of the ^{board level} module structure are not described ~~here~~ but Chapter 00 ~~or does~~ hence we felt compelled to discuss them in Chapter 00.

[note I took out page 8?]

Inset

The subtleties and uniqueness of the module structure are not described, nor are the design alternatives. For example, a bounded design that is typical of one board microcomputers was considered, though not described. However, a lower-cost, one-board system, like the VT78, (page 00) has evolved and is shown in Fig. RXT-11. The RXT-11 is an integrated system containing an LSI-11 chip set, 32 Kwords of memory, connectors for six EIA interfaces, and a controller for two floppy disk drives. One-hundred and seventy-five i.c.'s were used -- to implement the same functionality using standard LSI-11 modules, 375 i.c.'s would have been used. ~~[GB: Not announced yet -- probably March].~~

The initial module-level design for the LSI-11 was predicated on a quad-sized form factor with a conventional backplane. The modules are shown on page 00. Since there were not special ICs beyond the 3 chip processor, options tended to be relatively large and often occupied a full quad module. For options that were greater than a quad size, an ingenious packaging scheme was devised. It provided interconnection points on the extra half of the module (a double sized module) which was not used as the LSI-11 Bus (requires a double sized module). This permitted multiple board, complex options, e.g., a disk controller, to be packaged as a single option with no interconnection between the boards except via the second half of the quad board. As DEC began to build special ICs to interface to the bus, the option sizes decreased to occupy a double module. This system is now known as the LSI-11/2. A backplane system with modules is shown in Fig. LSI-11/2. The options available for the two systems have evolved as shown in Table LSI-11 Options.

~~Table LSI-11 Options~~

~~[Teicher to supply]~~

~~See next page~~

no IP
(The effect of a lower cost LSI-11 system is to provide additional applications as we discuss in Chapter 1, ~~page 00.~~)

USING LSI PROCESSOR BIT-SLICES TO BUILD A PDP-11--A CASE STUDY IN MICROCOMPUTER DESIGN

This paper by the designers of CMU-11 appears both in the module part and in this part on PDP-11 implementations. The Intel 3000 bit slice, herein called a Microcomputer (for Microprogrammed Processor), is used to interpret a PDP-11 ISP. The purpose of the design was to test the assertion that the bit-slice based arithmetic unit with register memory and microprogrammed control would simplify the design and construction of processors. The 11 was selected as a target problem in order to avoid the temptation of changing the problem (a real ISP) to fit the building blocks (the Intel 3000 processor). Indeed, the authors observed awkwardnesses that ultimately resulted in lower (than desired) performance. In retrospect, the Intel 3000 has not become the standard bit-slice architecture that the AMD 2900 series has; perhaps it suffered from being one of the earliest. Detailed comparisons including a breakdown of the various parts of the processor design are given and compared with the LSI-11 and 11/40 designs in terms of performance and cost (IC count and number of control bits in the microprogrammed controller).

A key part of the investigation was the evaluation of the computer-aided-design tools. The Stanford University Drawing System (SUDS) and the SAGE logic simulator were the major components. SAGE was predicated on being able to detect all the design and timing flaws prior to construction of a design. The authors claim that 95% of the errors were detected in this manner. The simulated-time/real-time ratio ($10^6/1$) did constrain the design process, e.g., the number of different runs used to find worst-case conditions.

DESIGN DECISIONS FOR THE PDP-11/60 MID-RANGE MINICOMPUTER

Unlike the reports from an architect's or reporter's viewpoint this chapter is a direct account of the design from the close proximity of the project. A mid-range machine is an inherently difficult design for the reasons of the designer characteristics we presented in Chapter 00, page 00. [CM: unclear]. As neither the lowest cost nor highest performance PDP-11, it has to be the right balance of features, price, and performance against criteria that are usually extremely vague.

Four interesting aspects of computer engineering are shown in the 11/60: the cache to reduce Unibus traffic; trace-driven design of floating-point arithmetic processors; providing writeable control store; and increasing the reliability, availability and maintainability.

Whereas the Unibus was previously thought to be inadequate for high performance systems, by using a cache most processor references do not use

the Unibus and so leave it free for i/o traffic. This gives high performance without the attendant cost*. The cache work is based on Strecker's (Chapter 00). The study leading to determining the block size is given. The block size can only conveniently be one since fetching multiple words would tie up the Unibus with additional traffic.

The use of trace data to design the floating point arithmetic is described together with the resulting design. Note that the 11/60 performs roughly at 11/70 speeds but at lower cost. The implementation of the two can be compared in the following table.

Table - Implementation of 11/60 and 11/70 (count of printed circuit boards).

	<u>11/60</u>	<u>11/70</u>
Base Pc	5	8 [CM: check 11/70 sys. manual]
Floating point	4	4
Cache	1	4
Memory management	1	2
	<hr/>	<hr/>
Total	10	18

Microprogramming is used to provide both increased user-level capability and increased reliability, availability and maintainability. The large

*Using Amdahl's constants, page 00, the reader might compute the bus bandwidth (for i/o traffic) and the address space needs for this speed processor given the cache and compare these needs with the Unibus. [CM: We should do this for all models in "What We've Learned" Chapter]

writeable control store option is described together with its use for data storage and various applications. This option has been recently used for emulating the PDP-8 at the OS/8 operating system level.

A general discussion of microprogramming is also given, especially with respect to memory technology advances (see also Chapter 1, page 00). Other semiconductor technology improvements are described together with how they affect price and performance. It is interesting to note that the simple concept of tri-state logic* had such a great effect on the design.

Impact of Implementation Design Tradeoffs on Performance: The PDP-11, A Case Study

This chapter presents a most comprehensive comparison of the eight processor implementations used in the ten PDP-11 models. The work was carried out to investigate various design styles for a given problem--the interpretation of the PDP-11 ISP. The tables alone give more insight into processor implementations than is available from any single source we know. The usefulness of the data also comes from having an outside observer examine the machines and share his insight.

The tables include:

1. a set of instruction frequencies, by Strecker, for a set of ten different applications. The reader should note the frequencies do not reflect all uses, e.g., there are no floating point instructions, nor

*Ability to interconnect a number of subsystems together through a wired-or connection.

has operating system code been analyzed.

2. implementation cost (modules, ICs, control store widths) and performance (micro- and macro-instruction times) for each model; and
3. a canonical data path for all 11 implementations, against which each processor is compared.

With this background data, a "top-down" model is built which explains the performance (macro-instruction time) of the various implementations in terms of the micro-instruction execution, and primary memory cycle time.

Since these two parameters do not fully explain (model) performance, a set of bottom-up factors must be introduced. These factors include various design techniques and the degree of processor overlap. We believe that this analysis of a constrained problem should provide useful insight to both computer and general-digital-systems designers.

EVALUATION

TURNING COUSINS INTO SISTERS: THE ROLE OF SOFTWARE IN SMOOTHING HARDWARE DIFFERENCES

Since FORTRAN is quite possibly the most often executed language for the PDP-11 and the one we intended that it execute, it is important to observe the 11 architecture as seen by the language processor--its user. The first

FORTTRAN compiler and object (run) time system are described together with the evolutionary extensions to improve performance. The FORTTRAN IV-PLUS compiler is only briefly discussed since its improvements, largely due to compiler optimization technology, are less relevant to the 11 architecture.

The chapter title overstates the compatibility problem since the five variations of the 11 ISP for floating point arithmetic are made to be compatible by essentially providing five separate object (run) time systems and a single compiler. This transparency is provided quite easily using a concept called threaded code. This concept appears to be a very simple interpreter for the PDP-11--and might not be called an interpreter by many. With threaded code, one 1-word instruction requiring two memory cycle times is executed each time a high level operation code is to be interpreted, otherwise the processor is carrying out the desired op code. When a simple integer expression like $I = I + 1$, which occupies 2 memory words and requires 3 memory cycles to execute, is transformed into a threaded code version the program still only occupies 2 words, but instead requires 5 memory cycles to execute (nearly a factor of 2). For more complex operations requiring longer execution times, like floating point arithmetic, the overhead turns out to be quite low and the space utilization is quite good. It is also possible to move efficiently between threaded and directly executed code, although this is not done. Jim Bell discovered the technique; it has been used extensively for other compilers because of its low time and space overhead. The ability to carry out the interpretation so elegantly was not part of the original PDP-11 design, but rather was a consequence of the generality of the 11's addressing modes.

The first version of the FORTRAN machine constructed was a simple stack machine. As such, the execution times turned out to be quite long. In the second version, by recognizing the special high-frequency-of-use cases, e.g., $A = 0$, $A = A + 1$, and by having better conventions for three-address operations (to and from the stack), speedups of 1.3 and 2.0 for floating point and integers, were obtained.

It is interesting to compare the virtual machine described with the FORTRAN IV-PLUS machine which uses the floating point processors (on the 11/34, 11/45, 11/55, 11/60, and 11/70). If the FORTRAN machine described in the paper is implemented in microcode and made to operate at FPP speeds, the resulting machines turn out to operate at roughly the same speed and programs occupy roughly the same program space.

WHAT HAVE WE LEARNED FROM THE PDP-11?

This chapter is a substantially revised version* of a paper written for the CMU Computer Science 10th Anniversary, September 1975. This paper was written to critique the original expository paper on the PDP-11 (Chapter 00) and to compare the actual with the predicted evolution. The four critical issues of technology, bus bandwidth (and PMS structure), address space and data-type evolutions are examined.

called "What Have We Learned From the PDP-11?"

The first part of the chapter discusses how the technology is used as a basis for the evolution (something we did not understand when the machines were originally planned). The role of semiconductor memories is especially

~~*The paper is 50% longer. The introductory overview has been deleted (and is now placed in Chapter 1) and the sections on the 11/45 and 11/70 have been greatly expanded to include the perturbations due to the memory address and protection extensions. A detailed evolutionary model of the cost, and performance characteristics has been added. More historical facts are introduced, particularly as they effect the design of the extensions. Finally the basis (need) for the VAX/11 extension is discussed.~~

critical. The next section describes the evolution from the point of view of the various development projects and people. Some early (historical) design documents are introduced to further aid in understanding the design process.

The Unibus evolution is given and the case is made for its optimality. The Unibus has had greater longevity and use than any of the other DEC busses and compares favorably with the IBM I/O Channel Bus as a universal standard of interconnection.

We try to provide a set of evolutionary cost-performance metrics so the 11 can be compared with the other machines (18-, 12- and 36-bit) in the book (Chapter 00, 01, and 02). Also, here we go into the unique (within DEC) problem of designing a range of machines.

Although an ISP evaluation is given, it is quite weak. By comparison, Chapter 00 by Brender, gives a more useful evaluation of the architecture for FORTRAN execution. A complete section is given on the addressing extension, beyond the 11/45 and 11/70 extensions, which required a major perturbation: VAX-11.

The final section, addressed to the research community, describes some general problems encountered in structure design and engineering, together with how solving them might be useful in subsequent designs.

VAX-11/780: A VIRTUAL ADDRESS EXTENSION TO THE DEC PDP-11 FAMILY

[MJ

add this to the 11 glue

when operating with 32-bit addresses, call native mode, the VAX part of the architecture is in effect.

ISP

is also ^{briefly} described, ~~to~~ ~~for~~

~~and the reader can observe~~

in terms of data-types, operators, addressing and memory management

~~Chapter 00, by Bill Strecker, is the cleanest, most comprehensive description of and architecture that~~
provides a clean, ^{somewhat} terse, yet comprehensive ~~descriptive~~ description of the VAX-11 Architecture. It is ~~among~~ ^{a specific} the best papers on architecture that we know ~~also 11/780 imp.~~ The first model, VAX-11/780,

VAX-11 is ~~the~~ the extension to PDP-11 to provide ~~a~~ ^{each} larger ~~size~~ ^{architecture}, 32 bit virtual address for ~~a~~ ^{to} user process. The extension includes a PDP-11 ^(compatibility) written for the RSX-11/M program environment to run mode for running PDP-11 programs unchanged, but In this way, the PDP-11 programs can be moved to VAX and among VAX and PDP-11 computers in

~~a completely compatible fashion.~~

Since VAX the VAX part of the architecture is so complete ^{textbook} it can also serve as a model for other ~~arch~~ ^{computer} companies

depending on the ^{user's} address size, computational and generality needs.

and base, case study for architecture in general

~~The description~~

~~The description proceeds around gives includes to the data-types~~

~~There is a~~

The architectural description is ~~also~~ ^{design} includes ~~commentary~~ ^{are design given as the} and constraints ~~and~~ that ~~affected~~ lead to various choices, and the resultant architecture is so comprehensive ~~The steps~~

~~that it can almost be used as a base study for all other architectures in that it~~

~~includes a wide variety of data-types and operators, and . The memory addressing~~

~~and segmentation scheme also serves as can serve as a model of comparison.~~

There

The descript

The evolution of the 11.

[Section on Strecker's paper to go here]

THE MULTIPROCESSOR RESEARCH COMPUTERS OF CARNEGIE-MELLON UNIVERSITY AND DEC'S PULSAR

Three computers, which use the 11 as a basis, were built at Carnegie-Mellon University to carry out research in computer structures and operating systems for multiprocessors. *A fourth multiprocessor, based on the 16 LSI-11s, called the DEC PULSAR is given in also discussed*

The first computer, C.mmp, is a 16 processor (11/40's and 11/20's) system with 2.5 million words of shared primary memory. It was built to investigate the programming (and resulting performance) questions associated with having a large number of processors.

The chapter on the second computer, Cm*, is located physically in the Modules Part, as the modules that form Cm* are LSI-11's. Cm* was based on the premise that ultimately, the smallest modular unit, i.e., integrated circuit, used to build digital systems would be a computer. With this premise, we need to understand how to interconnect computers physically, how to assign parts of the application program to the various computers, and how to program the complete structure.

The third computer, C.vmp, was designed to investigate how a production microcomputer, the LSI-11, could be used to build a triplicating, voting high availability computer.

The goals of the first two are performance while the third has reliability as its goal.

We believe that technology will force the evolution of computing structures to converge to three styles of multiprocessor computers: (1) C.mmp-style, for high performance, incremental performance and availability; (2) loosely coupled computers like Cm* to handle specialized processing, e.g., front end, file, and signal processing; and (3) C.vmp-style for high availability based on increased maintenance costs.

The technology-force argument is based on history, near term technology, and resulting price extrapolations:

1. MOS technology is increasing in both speed and density faster than the technology, e.g., ECL, from which high performance machines are built.
2. The price per chip of the single-MOS-chip processors decreases at a substantially greater rate than for the low-volume high-performance, special designs. Chips in both designs have high design costs, but the special design enjoys a much lower volume.
3. Relative to all other costs of a system, the processor cost in a low end system is essentially zero.
4. Standards in the semiconductor industry tend to form more quickly for high volume products. For example, in the 8-bit microcomputer market,

one type supplies about 50% of the market and three types supply over 90%.

5. A 16-bit processor-on-a-chip, with both an address space matching its performance and appropriate data-types, has been announced. Such a commodity will form the basis for nearly all future computer designs.

Pulsar | DEC's PULSAR, see described below, is a good example
← of one of the more straightforward applications of this technology.

As a result of these factors, the two classes of machines

(MOS-processor-on-a-chip-based and low-volume, high-performance-processor-based) have rapidly diverging costs per operation per chip. Furthermore, large scale applications have been slow to form since problem complexity increases more rapidly than program size. Therefore, most subsequent computers will be based on standard, high volume parts. For high performance machines, since processing power is available at essentially zero cost from processor-on-a-chip-based processors, large scale computing will come from arrays of processors, just as we build arrays of 64 Kbit ICs to form memory subsystems. *This type*

C.mmp--A Multi-Mini-Processor

C.mmp was motivated by the need for more computing power to solve speech recognition/signal processing problems and to understand the multiprocessor software problem. Until C.mmp, only one large, tightly coupled multiprocessor had been built--the Bell Laboratory's Safeguard Computer (BSTJ issue?).

The introductory section describes the economic and technical factors influencing multiprocessor feasibility and argues for the timeliness of the research. Various problems to be researched are given together with a discussion of particular design aspects. For example, since C.mmp is predicated on a common operating system there are two sources of degradation: memory contention and lock contention. The machine's theoretical performance as a function of memory-processor interference is based on Strecker's (1971) work. In practice, because the memory was not built with low-order address interleaving, memory interference was greater than expected. This problem was solved by moving program segments.

As the number of memory modules and processors becomes very large, the theoretical performance (as measured by the number of accesses to the memory by the processors) approaches the memory bandwidth $(m/\text{memory-cycle-time}) \times 1/e$. [ref.?] Thus, with infinite processors, there is not a maximum limit on performance, provided all processors are not contending for the same memory.

Although there is a discussion outlining the design direction of the operating system, Hydra, later descriptions should be read [Wulf et al, 1975]. Since the 11's small address necessitated frequent map changes, 11/40s with writeable control stores were used to implement the operating systems calls to change the segment base registers.

The C.mmp which was actually built is shown in Fig. PMSC.mmp.

There are three basic approaches to the effective application of multiprocessors:

1. System-level workload decomposition. If a workload contains a lot of inherently independent activities, e.g., compilation, editing, file processing, and numerical computation, it will naturally decompose.
2. Program decomposition by a programmer. Intimate knowledge of the application is required for this time-consuming approach.
3. Program decomposition by the compiler. This is the ideal approach. However, results to date have been disappointing [Ref. Illiac IV FORTRAN compiler projects].

C.mmp was predicated on the first two approaches. Since the original paper, ALGOL 68, a language with facilities for expressing parallelism in programs, has been implemented. It has assisted greatly with program decomposition. See Figure x.

As can be seen in the paper, a model of the lock problem influenced the number of critical sections in the scheduler. Since the paper, the operating system Hydra has been designed and implemented. An extensive description is given in [Wulf et al, 1975].

Two experiments analyzing the performance of C.mmp and Hydra have been reported. The first, by Ranathe [1977], used a hardware monitor to measure

the degradation due to the locking mechanism which is invoked when shared data is accessed. The second, by Oleinick [1977], analyzed user-program-level synchronization. We summarize the results of both below.

The contention for shared resources in a multiprocessor system occurs at several levels. At the lowest level processors contend at the cross-point switch level for memory. This memory interference is discussed in the chapter. On a higher level there is contention for shared data in the operating system kernel. At higher levels, processes contend for i/o devices and for software processes, e.g., for memory management. The following table points to models on experimental data at these different levels in C.mmp.

Contention Level

Reference

user-program

[Oleinick, 1977]

[Fuller and Oleinick, 1976]

Hydra Kernel objects

[Marathe and Fuller, 1977]

cross-point switch

Chapter XX

[Fuller, 1976]

Ranathe's data show that the shared data of Hydra is organized into enough separate objects that a very small degradation (less than 1%) results from

contention for these objects. Table LOCK is reproduced from his paper. He also built a queueing model which projected that the contention level would be about 5% in a 48 processor system.

Oleinick uses a root finding algorithm to study various implementations of a single problem.

[CM: elaborate after studying his thesis chapter.]

Multi-Microprocessors: An Overview and Working Example

The Cm* work, sponsored by NSF and ARPA, is an extension of earlier NSF-sponsored research [Bell, etc. 1973] on register-transfer-level modules. As LSI and VLSI enable construction of the processor-on-a-chip, it is apparent that low-level, register-transfer modules are passe' for the construction of all but low-volume computers. Although the research is predicated on structures employing a hundred or so processors, this chapter describes the culmination of the first (ten-processor) phase.

The authors motivate their work by appealing to the diseconomy-of-scale arguments which we advanced at the beginning of this section (page 00). To provide additional context for their research, computer modules (Cm*), multiprocessors (C.mmp) and computer networks are described in terms of performance and problem suitability. The chapter gives a description of the modules structure, together with their associated limitations and potential research problems.

The final, most important, part of the chapter evaluates the performance of Cm* for five different problems.

C.vmp: The Architecture and Implementation of a Fault Tolerant Multiprocessor

C.vmp is a triplicated, voting multiprocessor designed to understand the difficulty (or ease) of using standard, off-the-shelf LSI-11s to provide greatly increased reliability. There is concern for increased reliability because systems are becoming more complex, are used for more critical applications, and because maintenance costs for all systems are increasing. Because the designers themselves carry out and analyze the work, this chapter provides first-hand insight into high reliability designs and the design process--especially its evaluation. The system has operated for several months and the first phase of work is complete.

Several design goals are initially predicated and the work is carried out against the goals.

The goal of software and hardware transparency turned out to be easier to attain than expected, because of an idiosyncrasy of the floppy disk controller. Because the controller effects a word-at-a-time bus transfer from a one-sector buffer, voting can be carried out at a very low level. It is unclear how the system would have been designed without this type of controller; as a minimum, some form of software transparency goal would have been violated and a significant controller modification would have

been necessary.

A number of models are given by which the design is evaluated. Various component reliabilities are used and the reader should get a great deal of insight into the factors contributing to reliability. It should be noted that a special hardware voter is needed to get a sufficiently low cost for a marketable C.vmp. While the intent of C.vmp is not a product, it does provide much of the insight for such a product.

PULSAR here

Marathe, M. and Fuller, S. H. A Study of Multiprocessor Contention for Shared Data in C.mmp, Proceedings 1977 ACM SIGMETRICS Conferences, pp. 255-262.

Wulf, W. [and others] The Hydra Operating System, Fifth ACM SIGOPS Symposium on Operating Systems Principles (Nov. 1975).

11 Glue Figures and Tables

Fig. Cachespeed - Structure of Pc, Mcache and Mp of cached computer.

Fig. LSI-11/2 - Photograph of double height modules forming LSI-11/2.

Fig. RXT11 - Bounded LSI-11.

Fig. PMSC.mmp - PMS diagram of C.mmp.

Table Lock - Measurement of the locking behavior in Hydra/C.mmp.

Table of LSI-11 Options and Extensions (1978) for LSI-11/2

	<u>LSI-11</u>	<u>LSI-11/2</u>	<u>Additions</u>
<u>Basic</u>			
Processor + 4 Kw + send line	quad	double	
Real time clock Power controller/ sequences	quad double		
<u>Memories</u>			
4 Kw RAM	double		
4 Kw core	quad		
32 Kw RAM	quad	double	
4 Kw PROM	double		
4 Kw PROM; 256 RAM	double		
<u>Interfaces</u>			
Parallel (16 line)	double		
Asynchronous (serial)	double		
Instrument (IEEE-488)	double		
Direct Memory Access	quad		
Foundation (general purpose)	quad		
Analog-to-digital	quad		
Digital-to-analog	quad		
Floppy interface		double	

MJ - please ask Steve & Feichen ^{set} to add to this and check both

Table of LSI-11 Options and Extensions (1978) for LSI-11/2

Basic

Processor + 4 Kw + serial line quad
Processor

double

Memories

4 Kw ram.

double.

~~32~~ 16 "

quad
double
double.

double

4 Kw prom

4 Kw prom; 256 ram

4 Kw core

quad

32 Kw RAM

quad

double

Interfaces

Parallel (16 line)

double

~~Serial~~ 4 line serial asynchronous

double

Asynchronous (serial)

double

Instrument (IEEE-488)

double

Direct Memory Access

quad.

Foundation (general purpose)

quad

~~Fancy RI~~ ~~direct time~~ & clock

quad

Analog-to-digital

quad

Digital-to-analog

quad

Floppy interface

← double

Real time clock

quad

Power controller / sequencer

double

Hard disk interface (R005) 4 quads

(R601) 2 quads

Back planes line printer

double

4x8

for quads

2x4

2x8

2x12

f-line synchronous ~~set~~

quad

f-line ^{asy} ~~mix~~

quad

PDP-11 glue all caps

edited 12/15/77

is quite capable of standing alone as
This part could have stood alone as a ~~single~~ book on the PDP-11

although it relies heavily on the conceptual framework of Chapter 1, and the ISPS language description given in the appendix 00.

computer evolution, For a number of reasons the 11 has evolved quite differently than the other computers in this book, and as such ~~a~~ has ~~quite~~ provides an interesting story. ~~a much different id different story.~~ For Per Two factors seem

The factors ~~in~~ that seem to have created the machines are clearly ~~and technology~~ ^{large number (} market based, generating a number ^{ten)} of different implementations ^{to relatively seven year} over a short ~~to~~ lifetime. Thus, the fact ~~is~~ that there is an expanding ^{as we move along in time} range of machines causes ~~provides~~ the engineer ~~with~~ i with a set ^{and insight} of problems that is not natural to some of the

~~other designs~~ of the traditional mini ~~8 Family~~, of the best cost/performance (18-bit) ^{and high} of the ~~highest~~ performance ~~that one can build with~~ that is less than

500K (the DECsystem 10). Here, we The 11 designs cover a range of a factor of 500 in system price and ~~500~~ ^{to} 250,000 and 500 in memory size ~~5~~ 4 Kwords to 2 Mwords). ~~The other~~ Because

Because the 11 is a relatively low cost, it is an excellent vehicle to use in Computer Systems and C Computer Structures research to test various ideas in order that structures may evolve.

In

This section is ~~part~~ is really divided into ~~5~~ five sections consists

- 1 Introduction, ~~containing~~ ^{the first} the original paper on the 11 architecture ^{chapter published when the 11 was announced} This paper introduces the architecture, and gives various goals for it, and predicts how it might evolve. ~~The family notion is quite strong~~, although not specific. ^{Alternatively, Chapter 00, What Have We Learned From PDP-11,}
2. Conceptual ~~The~~ ^{Conceptual} basis for the evolution of. This section ~~cont~~ ~~A~~ ~~has~~ ~~been~~ ~~carried~~ ~~out.~~ ~~on~~ the Cache Memory both describes the concept and shows how the implementation is carried out.

May preferably ^{be} read next ~~after chapter 00~~ in order to get ~~an~~ the broadest overview, ~~of~~ and best immediate critique of chapter 00.

5. the Multip multiprocessor research computer of
Carnegie-Mellon University. ^{Three comp}
multiprocessors: C.mmp, ^(chapter 00) Cm* ^(chapter 00) and C.vmp ^(chapter 00)
~~described~~

give examples of ~~not~~ actual multipro a
16 processor multiprocessor, a ~~set~~ set of compa
modules based on LSI-11 and a voting
triplicated, voting ~~multi~~ multiprocessor
computer for high reliability.

contains two papers, that combined with Chapter 00,
form ~~the concepts many of the~~ some of the
conceptual basis for basis for the models.
Strecher, Chapter 00, describes the cache memory
mechanism as a basis for building high
performance models. ^(specifically the 11/70) Levy describes ~~various~~

~~the~~ the intercommunication ^{among} problem of physical components
and how busses carry out ~~the~~ this task.

3. Implementations ~~Four chapters~~ ~~con-~~ give of the

Four ~~cha~~

of four chapters, three are on specific implementations
(LSI-11, ~~11/60~~ and CMU-11 and 11/60) and the
fourth describes a (Chapter 00) ~~gives~~ is on
model for ~~the~~ all models. ~~the~~ This latter
chapter provides ^{both} a conceptual model ~~and~~ and it details
of the various implementations together with how ^{the design fits a} ~~conceptual~~
model. This is especially ~~very~~ useful to understand ^{the performance}.

4. Evaluation, Chapter 00 evaluates the 11 ~~as~~ as a
machine for executing Fortran. What We Have Learned From
PDP-11, Chapter 00 ~~evalua~~ disusses all aspects of
the evolution in terms of ~~Chapter 00~~ the introductory chapter (00).
over

Jardow

Note: When single-spaced,
all footnotes will
appear only once
on the appropriate
page.

~~PDP-11 Glue 12/7/77~~
~~INTRODUCTION~~

A New Architecture for Minicomputers--The DEC PDP-11

allcaps

It is somewhat anticlimactic to discuss this original PDP-11 description here because Chapter 00 explicitly discusses "What We Have Learned from the PDP-11". The purpose of the chapter was originally threefold: to give the PMS and ISP architecture of the PDP-11 as it was first proposed to describe the first (11/20) implementation at a time when the whole architecture had not been worked out or even fully considered; and to show possible extensions.

The reader might note that the computer class definitions ^{given in (1970)} of micro, mini and midi correspond quite closely to those of Chapter 1.

Although we comment on the disparity between the predicted and actual evolution in Chapter 00, some of the important reasons are:

1. The notion of designing with improved technology especially for a family was not understood then. The understanding for one of us (Bell), was put forth in a paper in 1972 (Bell, Chen, Rege).
2. The Unibus bandwidth proved unacceptable for all communications at the very high and low end designs. Whereas, this chapter posits a multiprocessor, and multiple Unibusses, this precise structure did not evolve. The bandwidth has subsequently been shown to be adequate for

all but the largest configurations when a cache is attached to the processor as in the 11/60 (see Chapter 00). Note the effect of a 90% cache ^h bit rate ^{is} to reduce the number of access to ^{primary memory} ~~Mr~~ via the Unibus by a factor of ten!

3. The memory address space was too small.
4. The particular data-type extensions weren't predicted. While floating point data was discussed, the character string and decimal operations weren't described. These data types evolved in response to market need (and COBOL) which didn't exist ^{for DEC} at the time the machine was designed.

We have made a major change in the chapter by removing the original ISP description and replacing it with a correct and complete (memory management and floating point) ISPS description that was used when the 11 was evaluated as a Computer Family Architecture (CFA) standard by the U. S. Defense Department.

Caps
Cache Memories for PDP-11 Family Computers

Chapter 00 by Strecker is included for ^{four} ~~three~~ reasons: it is a clear exposition of the cache memory structure and its design parameters; ~~since~~ the cache is the basis of the fast PDP-8 (Chapter 00), the 11/60 (Chapter 00), ~~and~~ the 11/70 (page 00), and the KL10 (Chapter 00); the design methodology is well done--it is "good engineering"; and finally the paper is "well-written"--in fact, it received the award for the Best Paper at the

Third Computer Architecture Conference at which it was presented. Thus, since a relevant paper that provides a clear exposition, while illustrating good engineering and being well-written, is so rare we also publish it simply to serve as an encouragement and an example for those who should expose, understand and describe their work.

The design process is implicit in the way the work is carried out to determine the structure parameters. It should be noted that it was easy to collect data statistics about the program's behavior since the trace bit, T permits the 11 to interpret itself on a one at a time basis. The relevant parameters are varied and sensitivity plots (runs) are made to determine the effect of each parameter on the design. In the 11/60, Mudge, also uses

and Straker's traces.

the same methodology (Chapter 00). One of the important parameters, ^{to understand is} the

time between changes of context, ^{because all} is especially important and to the best of

our knowledge ^{this study} is both unique and necessary for any real time or ^{is} ~~and~~

~~multiprogrammed system.~~ ^{behaves, have many context switches} The PDP-8 cache design (Chapter 00) shows the

effect of segmenting the cache for instructions and data. The PDP-8

chapter discusses the performance for a particular design, and the

performance numbers are in terms of a speed-up factor. Stracker gives the

performance evaluation in terms of cache miss ratios. ^{the} Performance measures, ^{see}

are related (Lee, 1969) in the following way (assuming an infinitely fast processor):

using Fig. Cachespeed,

- ~~P~~ = total no. of memory accesses by the processor, Pc
- ~~M~~ = no. of memory accesses that are missed by the

cache and how to be referred to Mp

t.c = cycle time of cache memory, Mc
t.p ~~top~~ = cycle time of primary memory, Mp
R = t.p/t.c (ratio of memory speeds), where R is
typically 3 to 10

the relative execution speeds are:

$$t(\text{no cache}) = pR$$

$$t(\text{to cache}) = p + mR$$

$$\text{speedup} = pR/(p + mR) = R/(1 + (m/p) R) = a = \text{miss ratio} = m/p$$

therefore

$$\text{speedup} = R/(1 + aR) = 1/(a + 1/R)$$

note that if a = 0 (100% hit), the speedup is R; while

if a = 1 (100% miss), the speedup is R/(1 + R), i.e.,

the speedup is less than 1

IMPLEMENTATIONS

Copy
A Minicomputer-Compatible Microcomputer System: The DEC LSI-11

This first chapter of the implementation section was written from the knowledgeable user viewpoint. Although the paper is a descriptive narrative of the design from the chips, boards and backplane levels, it

lacks ~~some~~ ^{that} insight ~~from one of~~ the designers at Western Digital or DEC (Lloyd Dickman, Rich Olsen, or Mike Titelbaum). ^{might have provided} An account of the chip-level design is available ^{ZORSA} (Soha and Pohlman, 1974). The design was done at Western Digital ~~corporation~~ ^{ZORSA S} by Roberts, ~~Zoha~~, and Pohlman as a relatively general purpose microprogram^{med} computer that could be used to emulate many computers. When DEC started working with them for use in the interpretation of the 11 ISP, the design took on more of the 11 structure. Two design levels are described in the paper: the 3 chip microprogrammed computer as it is used to interpret the 11 ISP; and the particular PMS module level components as they are integrated into a backplane to form a hardware system. Sebern points out ^{the microprogramming} ~~an interesting~~ ^{that took place} tradeoff between the chip and module levels ~~can be observed in the use of microprogramming to~~ carry out [↓] functions; ~~that are normally done with hardware;~~ the time clock, the console, refreshing dynamic random access MOS memory, and power fail control. [↑] The subtleties and uniqueness of the module structure are not described, ^{not are the design alternatives} The initial module level design was predicated on a quad-sized form factor and plugged into a conventional backplane. ^{for example,} The alternative, ^{was considered, though} more bounded designs ^{one board} that ~~are~~ ^{is} typical of other microcomputers, ^{is not} described. ^{A lower} However it should be noted that in order to get ~~more reduced~~ cost, ^{one board} such a system, like the VT78, ^(page 00) has evolved and will eventually be marketed. [↑] This structure is shown in Fig. KXT11. The modules are shown on page 00. Since there were not special ICs beyond the 3 chip processor, options tended to be relatively large and often occupied a full quad module. For options that were greater than a quad size, an ingenious packaging scheme was used for providing interconnection points on the extra half of the module ^(a double sized module) which was not used as the LSI-11 Bus. ^(requires a double sized module) This permitted

multiple board, complex options (e.g., a disk controller) to be packaged as a single option with no interconnection between the boards except the second half of the ^{quad}board. As DEC began to build special ICs to interface to the bus, the option sizes began to decrease to occupy a double module. This system is now known as the LSI-11/2. A backplane system with modules is shown in Fig. LSI-11/2. The options available for the two systems have evolved as shown in Table LSI-11 Options.

Table LSI-11 Options

[Teicher to supply]

The effect of ^alower cost LSI-11 ^{system}modules and backplane availability is to provide additional applications as we discuss in Chapter 1, page 00.

Using LSI Processor Bit-Slices to Build a PDP-11--A Case Study in
Microcomputer Design

Copy

This paper by the designers of CMU-11 appears both in the module part and in this part on PDP-11 implementations. The Intel 3000 bit slice, herein called a Microcomputer (for Microprogrammed Processor) ~~IC~~, is used to interpret a PDP-11 ISP. The purpose of the design was to test the assertion that the bit-slice based arithmetic unit with register memory and microprogrammed control would simplify the design ~~of~~ and construction of processors. The 11 was selected as a target problem in order to avoid the temptation of changing the problem (11 processor) to fit the building blocks (the Intel 3000 processor). As such, the authors observed the

awkwardness that ultimately resulted in lower (than desired) performance.

In retrospect, the Intel 3000 has not become the standard (~~now~~ ^{that} the AMD 2900 ^{series has} series) bit-slice architecture, but perhaps suffers from being one of the earliest. Detailed comparisons including a breakdown of the various parts of the processor design are given and compared with the LSI-11 and 11/40 designs in terms of performance and cost (IC count and number of control bits in the microprogrammed controller).

A key part of the investigation was to evaluate the computer aided design tools. The Stanford University Drawing System (SUDS) and the SAGE logic simulator were the key components. SAGE was predicated on being able to detect all the design and timing flaws prior to construction: The authors claim that 95% of the errors were detected by the simulation--and all errors were ultimately detectable once the simulation data or machine description was changed.

Caps
Design decisions for the PDP-11/60 Mid-Range Minicomputer

Unlike the reports from an architect's or reporter's viewpoint this chapter is a direct account of the design from the close proximity of the project. Because it is a mid-range machine, the 11/60 is a difficult design for the reasons of the designer characteristics, Chapter 00, page 00. The design is neither the lowest cost nor performs the ^{highest of the 11s} ~~best~~, but has to be the right balance of features, price, and performance against criteria that are usually extremely vague.

Four interesting aspects of computer engineering are shown in the 11/60:
the cache to reduce Unibus bandwidth; trace^xdrive^x design of floating-point
arithmetic; providing writeable control store; and increasing the
reliability, availability and maintainability. [↑] Whereas the Unibus had
previously thought to be inadequate for high performance, the cache is used
to unload the i/o traffic and provide high performance without the
attendant cost[†]. The work is based on Strecker's (Chapter 00). The study
leading to determining the block size is given. The block size can only
conveniently be one since additional traffic is generated by fetching
multiple words ^{which would tie} and ~~thereby tying~~ up the Unibus ^{with} for additional traffic.

The use of trace data to design the floating point arithmetic is described
together with the resulting design. ~~It should be noted~~ that the 11/60
performs roughly at 11/70 speeds ^{with lower cost, so the implementation}

^{of the two can be compared in the following table.}

Table - Implementation of 11/60 and 11/70.

	11/60	11/70
Base Pc	x	x
Floating point	x	
Cache	x	x
Memory management	0	x
<hr/>		
Total	x	x

^{The} Microprogramming is used to provide both increased user-level capability
*Ability to interconnect a number of subsystems together through a wired-or
connection.

and increased reliability, availability and maintainability. The large writeable control store option is described together with its use for data storage and various applications. This option has been recently used for emulating the PDP-8 with OS/8 operating systems.

*Ability to interconnect a number of subsystems together through a wired-or connection.

A general discussion of microprogramming is also given, especially with respect to memory technology advances (see also Chapter 1, page 00). Other semiconductor technology improvements are described together with how they effect price and performance. It is interesting to note that the simple concept of tri-state logic* had such a great effect on the design.

See page 7
Insert Front
40A
10B
Cups

* 6
EVALUATION

Turning Cousins into Sisters: The Role of Software in Smoothing Hardware Differences

Since FORTRAN is quite possibly the most often executed language for the PDP-11 and the one we intended that it execute, it is important to observe the 11 architecture as seen by the language processor--its user. The first FORTRAN compiler and (run) object time system are described together with the evolutionary extensions (~~evolution~~) to improve performance. The FORTRAN IV+ compiler is only briefly discussed since its improvements are largely a matter of compiler optimization technology and are less relevant to the 11 architecture.

The chapter title overstates the problem since the five variations of the 11 ISP for floating point arithmetic are made to be compatible by providing what amounts to five separate object (run) time systems and a single compiler. This transparency ^{is provided quite easily using} comes about because a concept called threaded code ^{This concept appears to be} is used to provide what is a very simple interpreter for the PDP-11--and might not be called an interpreter by many. With threaded code, one 1-word instruction requiring two memory cycle times is executed ^{high level} per transfer of control each time the next operation code is to be ^(the more complex operation)

*Using Amdahl's constraints, page 00, the reader might compute the bus bandwidth (for i/o traffic) and the address space needs for this speed processor given the cache and compare these with the Unibus.

constants

10A

Impact of Implementation Design Tradeoffs on Performance: The PDP-11, A Case Study

the implem interpretation of the PDP-11 ISF.

We believe

This chapter presents a ~~comprehensive~~ ^{comprehensive} comparison of ~~the eight PDP-11 processor implementations.~~ ^{one of the most comprehensive} The work was carried out to investigate ^{various design styles for a given fix problem -} ~~chpater.~~ Aside from any conclusions they ~~are~~ ^{are} significant tables that give more insight into processor implementations than is available from any source we know. Unlike some of the other dataa this is more objective ~~(and examinations include)~~

single

The significant tables include:

1. A ^{instruction use} set of frequencies, ~~that~~ ^{from} ~~Sto~~ ^{Sto} ~~Streecker~~ ^{Streecker} determined from operating system ~~usage.~~ ^{usage.} (The reader should ~~be aware~~ ^{note} ~~that~~ ^{the frequencies} they do not ~~reflect~~ ^{reflect} a fully general purpose ~~mix,~~ ^{use} ~~for~~ ^{eg.} there are no ~~any~~ ^{any} arithmetic and ~~or fixed point~~ floating point data in the instructions used.

2. All A canonical data path ^{is posited} ~~is presented~~ which for all 11 implementations, and each ~~the~~ processor is compared ~~and~~ ^{presented and} with it. This is insight that ~~we has not come from~~

3. All the execution times for the various models are present in place.

4. Implementation ~~pack~~ ^{modules} cost ~~measures~~ ^{measures} (packages and ICs, control store widths) ~~of the data also~~

The usefulness comes from having an outside ~~or~~ ^{an} observer examine ~~the~~ ^{machines and} and try to provide insight into the machine.

performance (the micro and macro-instruction times)

With this ~~background~~ ^{background} ~~the main purpose of the study~~ ^{the main purpose of the study} ~~is to build a model which explains the performance of various implementations given the detrd various metrics.~~ ^{time)} ~~The two metrics~~ ^{in (macro-instruction rate)} (as we might expect, are instruction time ~~macro~~ ^{macro} in microc

The two design parameters, micro-instruction execution time, and ~~main memory cycle time~~ ^{primary} are used (called the top-down approach) to

explain (to within 85% the behavior (macro-instruction times.

10B

Unfortunately, The design ~~tricks~~ and overlap enter the picture and must be used to finally explain the re ~~t~~the discrepancies of the model. These design ~~trc~~

Since these two parameters don't fully explain (model) performance, ~~the~~ a set of bottom-up factors ~~are~~ must be introduced. ~~to more~~ these factors include various design techniques and the ~~less~~ degree of ^{processor} overlap.

~~We believe the reader should study~~
~~these implementations as they in~~
We believe the fact that the problem is constrained should provide useful insight to ~~not only act~~ both computer ~~equipment~~ designers and ^{general} digital systems design.

otherwise the processor is
interpreted, ^{when} For a simple integer expression like $A = A + 1$, which occupies 2 memory words and requires 3 memory cycles to execute, is transformed into a threaded code version; the program still only occupies 2 words, but instead requires 5 memory cycles to execute (nearly a factor of 2). For ^{more complex} longer operations requiring longer execution times, like floating point arithmetic, the overhead turns out to be quite low and the space utilization is quite good. Jim Bell discovered this technique and it has been used extensively for other compilers because the time and space overhead is so low. The ability to carry out the interpretation so elegantly was not part of the original design, but rather turned out to be possible based on the generality in the 11's indexing modes.

The first version of the FORTRAN machine constructed was a simple stack machine. As such, the execution times turned out to be quite long. By recognizing the special ^{high use frequency} cases (e.g., $A = 0$, $A = A + 1$, etc.) and ^{by} having better conventions for three-address operations to and from the stack, speedups of 1.3 and 2.0 for floating point and integers, respectively, were obtained when a more complex machine was constructed for the second version.

Since this paper is really on the construction of an extension to the 11 to execute FORTRAN, it is interesting to compare it with the FORTRAN IV+ machine which uses the floating point processor (11/45, 11/55, 11/60, 11/70). ^{machines} The two ^{operate at roughly} turn out to be ^{roughly} the same speed and ^{programs occupy roughly} have the same program space ^{if} efficiency assuming the FORTRAN machine described in the

paper is microprogrammed and ^{made to operate at FPP speeds,}
*Using Amdahl's constraints, page 00, the reader might compute the bus bandwidth (for i/o traffic) and the address space needs for this speed processor given the cache and compare the needs with the Unibus.

Copy

What Have We Learned From the PDP-11?

This chapter is a substantially revised version* of a paper written for the CMU Computer Science 10th Anniversary, September 1975.

This paper was written to critique the original expository paper on the PDP-11 (Chapter 00) and to compare the actual with the predicted evolution. The four critical issues of technology, bus bandwidth (and PMS structure), address space and data-type evolutions are examined. ^PThe first part of the chapter discusses how the technology is used as a basis for the evolution (something we did not understand when the machines were originally planned). The role of semiconductor memories is especially critical. The next section describes the evolution from the point of view of the various development projects and people. Some early (historical) ^{design} documents are introduced to further ^{aid in} ~~attempt an~~ understanding ~~of~~ the design process.

The main section consists of: evaluating the Unibus, examining the cost-performance evolution, the ISP and asking why multiprocessors haven't evolved. ^{discussing the organizational issues behind} As a companion, ^{Chapter 00 on a C.mmp describes the technical} ~~problems associated with multiprocessors.~~

The Unibus evolution is given and the case is made for its optimality. The Unibus has had greater longevity and use than any of the other DEC busses and compares favorably with the IBM I/O Channel Bus as a universal standard of interconnection.

We try to provide a set of evolutionary cost-performance metrics so the 11
*The paper is 50% longer. The introductory overview has been deleted (and is now placed in Chapter 1) and the sections on the 11/45 and 11/70 have been greatly expanded to include the perturbations due to the memory address and protection extensions. A detailed evolutionary model of the cost, and performance characteristics has been added. More historical facts are introduced, particularly as they effect the design of the extensions. Finally the basis (need) for the VAX/11 extension is discussed.

can be compared with the other machines (18-, 12- and 36-bit) in the book (Chapter 00, 01, and 02). Also, here we go into the unique problem the 11 has of designing a range of machines.

Although an ISP evaluation is given, it is quite weak. By comparison, Chapter 00 by Brender, gives a useful evaluation of the architecture for FORTRAN execution.

A complete section is given on the addressing extension beyond the 11/45 and 11/70 extensions which required a major perturbation in the form of the VAX/11 extensions.

The final section ^{to the research communities} describes some general problems encountered in structure design and engineering, together with how solving them might be useful in subsequent designs. (This part has been expanded too.)

*The paper is 50% longer. The introductory overview has been deleted (and is now placed in Chapter 1) and the sections on the 11/45 and 11/70 have been greatly expanded to include the perturbations due to the memory address and protection extensions. A detailed evolutionary model of the cost, and performance characteristics has been added. More historical facts are introduced, particularly as they effect the design of the extensions. Finally the basis (need) for the VAX/11 extension is discussed.

MULTI-PROCESSOR

Central Caps

COMPUTERS
The RESEARCH MACHINES OF Carnegie-Mellon University

Inevitably

computer

These three multiprocessor computers which use the 11 as a basis
This series of machines were built at Carnegie-Mellon University
to as were designed to carry out three different aspects of Computer engineering and

and application program

science research. The first-

The Carnegie-CMU's first computer, C.mmp, was built to for is a

multit-minicomputer, multi, minicomputer processor computer (or

minicomputer - is a conventional, hardware structure

(model 11/40's)

16 processor multi-processor system with 1 million words of comp shared primary memory, and It was

built with the expressed purpose of evaluating the performance quest

to investigate programming (and resulting performance) questions associated with having a large number

of multiprocessors.

The second computer, C.m*, is most also, properly p discussed in the

part on modu the mo Modules Part as the module that forms C.m* are LSI-11 computers.

This C.m* evolved from a design as evolution was based on the premise (i.e. all ICs)

that ultimately, modules such as the smallest modular unit, that would be

used to build digital systems

that would be built . With this premise, it was important (ot)

undertake the research which showed how to interconnect them physically,

and to assign parts of the problems to the various computers and to

us understand how t the probramm programming structure. them.

The C.vmp, for voting multiprocessor, was designed to investigate the red problems associated with taking a an

available mult microcomputer, the LSI-11, and can be used

(turned into) to build a triplicaty, voting

high availability computer.

The goals, therefore, are per of the three are performance, perf first two are performance while the third uses multiple computers for redundancy and reliability.

To this end, Fig. Perf shows the effect of using multi processors to solve execute various algorithms (More to Come Here!)

EKISTICS ΟΙΚΙΣΤΙΚΗ

REVIEWS ON THE PROBLEMS AND SCIENCE OF HUMAN SETTLEMENTS
EKISTICS is published by the Athens Center of Ekistics • 24, Strat. Syndesmou St. • Athens 136, Greece, P.O. Box 471
EDITORIAL OFFICE: Page Farm Road • Lincoln, Massachusetts 01773 • USA • 617-259-9144

Dear Reader,


Thank you for ordering EKISTICS. We hope that you will find it useful and become a regular reader.

EKISTICS is interdisciplinary and deals with questions of human settlements as a unified whole. This perspective is part of a broad consensus among planners, environmentalists, teachers, researchers, architects, and policymakers.

Each month we explore a different topic which is determined in consultation with our expert advisory board and through suggestions from our readers. If you would take a moment to indicate on the form below the urban problems of greatest concern to you, I would appreciate it. I also invite you to become a part of the EKISTICS network by subscribing to the journal.

I look forward to your reply.

Cordially,


Gwen Bell
Editor

To: Gwen Bell, Editor, EKISTICS
Page Farm Road
Lincoln, Mass. 01773 USA

From:

I consider the most pressing problem confronting human settlements is: _____

Please enter my subscription to EKISTICS:

_____ One year (12 issues) \$24.00

_____ Two years (24 issues) \$44.00

For all intents and purposes, the processor cost of the low end is 0 relative to all other costs of the system!

~~We must point out~~

We believe that

~~We believe we can guarantee that the most likely evolution~~ technology forces will force the evolution

of ~~most~~ computing structures will be along the lines of some form of ~~evolution~~ ^{both} multiprocessor computers, such as C.mmp or less tightly coupled computers, such as C.mmp based solely on the evolving directions of technology. ~~Altho~~ ^{for high performance} ~~This part is clear. The forces~~ ^{more ad hoc specialized structures to handle specialized processing (eg. front end, file, signal processing); and C.vmp for high availability, hand on increased & maintenance costs.} ~~are and therefore~~ ^{technology force} ~~argument is simply based on history, and near term~~ ~~the direction is clear~~ ^{forces} ~~the technology which forms the low cost processor processors~~ ~~currently MOS~~ ^{is evolving at more rapid rate} ~~than~~

to density and the speed if is improvise at a substantially greater rate than the ECL state array technology from which large machines are formed; the price per chip of the MOS ^{one chip processors decreases} is coming down at a substantially greater rate ^{than low volume special designs} due to the much higher volumes; and very high design costs (for both designs).

and the resulting performance (in operations) per chip and cost per operation per chip are rapidly ^{high performance semiconductor} ~~is~~ ^{is} ~~substantially~~ ^{diversifying} from the ~~the large~~ ^{technology form} which conventional high performance machines are formed.

These factors ^{produce a better price and price/performance} ~~cost design costs for large machines is diversifying~~ ^{at a} ~~rate compared with large scale computers. Furthermore, large scale applications have been slow to form since there is a problem complexity increases more rapidly than program size~~ ^{high compared with the} ~~design cost per element of at the high volume products. This will tend to create a small number of micro standard microprocessors (eg. there are essentially only 2 dominant type with which covers 1 type supplies about 50% of the market and 3 types supply 95% of the market.~~

At a time when there is sufficient technology to have a ~~cost~~ ^{adequate} in terms of memory address space ~~processor on a chip, a standard will be created by which all other other computer will be constructed.~~

Therefore, ^{most} ~~all~~ subsequent computers will be based on standard, high volume parts. For high performance machines, since processing power is available at 0 cost from

16-bit A few processor (-on-a-chip) is eminent with an ^{acceptable (for the performance)} address space ~~that~~ and ^{appropriate} data-types is eminent.

a processor-on-a-chip based processors, large scale computing will come from this source! arrays (just as of processors, just as we build arrays of 16K bit ICs to form memory.

and resulting price technology extrapolations:
1. the ~~best~~ MOS technology which ~~drives~~ ^{is} increasing in ^{both} speed and density faster than the technology from (eg. ECL) from which high performance machines is formed;

at the low cost market end

4. Standards ^(in the semiconductor industry) for high volume products tend to form more quickly. For example, in the 8-bit microcomputer market, 1-type supplies about 50% of the market and 3-types supply over 90% of the market

acceptably (for the performance)

C.mmp--A multi-mini-processor

~~C.mmp came from several~~ The new C.mmp was motivated by the need for more computing power to solve ~~a particular set of problems~~ ^{recognition /} ~~(Speech signal processing)~~ ^{problems} to get more computing power and to ~~as a machine to study~~ ^{understand the} multiprocessor software ~~problem.~~ Until C.mmp, there only one ^{large} tightly coupled multiprocessor had been built --

only one large

~~had been for, and it was for a different problem~~ The safeguard Bell Laboratory's Safeguard Computer (?)

^{introductory} The first section ~~describes the~~ ^{feasibility} ~~economic and technical factors~~ ^{and argues that it's important} ~~behind influencing~~ multiprocessors ~~on together~~ ^{with their} ~~and why it is a matter of timeliness~~ to embark on the research that ~~makes them more and more practical in this time period.~~ ^{because} ~~Various problems to be researched are given together with~~ ^{of the timeliness.} ~~A number of research areas are given and a discussion of particular~~ ^{they are ultimately} ~~aspects of the design.~~ ^{time}

~~some of the particularly worrisome aspects of the design are presented in detail~~

~~For example, since a multiprocessor the multiprocessor C.mmp is predicated~~ ^{both} ~~on a common operating system there are both is both~~ ^{desredation} ~~due to memory~~ contention when all the processors use ~~the same block of memory~~ ^{and} ~~also, since~~ there are ^{common} ~~software synchronizing locks~~ ^{for parts of} ~~around the software, there can~~ ^{may occur when a processor must idle and wait to} ~~be significant desredation as waiting for the software to enter common critical sections.~~

~~The machine's structure is presented in detail and the~~ ^{due to} ~~maximum theoretical performance is computed based on~~ ^{memory processor interference} ~~based on~~ ^{this is computed based on the work of Strecker's (1971) work.}

~~Here it is interesting to note that the assumption that~~

~~the problem of interference~~ ^{did was} ~~is ultimately become a problem because~~ ^{with low order address interleaving,} ~~programs were the memory was not built on an interleaved (low order address~~

~~was not built, hence there was some interference more interference than the interference~~

BSTJ issue

of the timeliness.
they are ultimately
time

resulting in poorer than
expected performance

This problem ^{solved} moving ^{program} segments of the ~~did materialize and had to be compensated for by movement of programs within different~~ ~~to different parts within~~ ^{memory.}

~~part of physical memory.~~ ^{On} One aspect that ~~is counter-intuitive should~~
It should ^{be noted} that when the number of memory modules and the number of ~~processors are equal (for example)~~ ^{and} become very large, the performance ^(and speed)

(as measured by the number of accesses to the memory by the processor) approaches
the number of memory bandwidth $(m/\text{cycle}) \times 1/e$. Thus, there is not a maximum
limit of the on performance ~~when~~ with infinite processors provided they are all
not contending to access the same memory.
Although there is ~~an~~ ^a extensive discussion of the

~~operat~~ outlining the ^{design} ~~oper~~ direction of the operating system, Hydra, ~~the reader is urged to~~ ~~the~~

later descriptions of Hydra should be read (list the ~~see~~ references ?? here). Since ~~the~~ the 11's small

address ~~got in the way of~~ ^{impaired} ~~prob providing~~ use,
11/40's with ^{microprogrammed} writable control stores were used to ~~carry~~
~~but the~~ ^{memory} ~~mapping~~ fast changes implement ~~of~~ ~~an~~ operating
systems ^{calls} involving changing the ~~base~~ ~~ad~~ ~~base~~ ~~ad~~
segment base registers.

One of the most pleasant surprises ~~that came out of~~ ^{because} of
C.mmp was ~~the work on~~ Algol 68 implementation of ~~for the two~~
it enables parallel programs to be specified. The results (see page 00)
^{approaches} Note there are three ^{basic} approaches to ~~an~~ effectively ^{applying} using multiprocessors:
1. ~~only~~

in terms of performance are quite encouraging!

~~1. Only as app~~

~~1. conventional ^{as} partitioned mul~~

~~1. finish some~~

~~1. having lots of independent work to do through multiprogrammed, ~~for timesharing~~ ~~uses~~ and or multiple independent computers;~~

~~2. taking ^{using} a conventional language (eg. Fortran) and having the compiler ^{for a conventional language (eg. Fortran)} detect ~~all actual~~~~

~~and compile statements or sets of statements that can be executed in parallel; ~~and the~~~~

3. introducing ^{new co} ~~co~~ primitives into the programming language ^(eg. Algol 68) so that algorithms for parallel execution are specified by the programmer.

Algol 68 takes this

← although C.mmp was only predicated on use of ~~an~~ use of type 1, ^{the} Algol 68

implementation on C.mmp

makes type 3 use possible } so far ~~the~~ (see ?)

~~The results of it~~

I. Isolated AC Input

This module provides eight transformer-isolated 120-volt AC inputs, with the following features.

1. 8 Inputs per module; 2 wires per input.
2. Isolation to 1500 volts DC or peak AC between inputs or from input to ground.
3. Jumper strip provided to allow commoning of one side of all inputs. Strip mounts on screw terminals.
4. Frequency range: 47-63 Hz
5. Nominal input voltage for logic "1": 120 volts
6. Individual input indicators on AC side of isolation.
7. Nominal input loading: 10 mA
8. Response time (turn-on or turn-off): less than 1/2 cycle
9. Inputs protected against overvoltage transients by MOV's.

II. Isolated DC Input

This module provides sixteen bits of optically-isolated DC input. There is an optional interrupt capability.

1. 16 Inputs per module; 2 wires per input.
2. Isolation to 1500 volts DC or peak AC between inputs or from input to ground.
3. Jumper strip provided to allow commoning of one side of all inputs. Strip mounts on screw terminals.
4. Response time (turn-on or turn-off): 2.5 msec nominal

NSF ^{ARPA} sponsored
1

Multi-Microprocessors: An Overview and Working Example

~~Cm* is the name~~ ~~which can be interconnected~~ ~~to form~~
a system of high level modules, constructed
of
from LSI-11 computers,

work sponsored by NSF and ARPA

Which can be interconnected to form ^{computer structures.} a large ~~computer~~ system. The ~~Cm*~~ ^{are an extension} ~~are an extension~~
work came out of ^(Bell etc. 1973) a research on register transfer modules, ~~because~~ as LSI and VLSI enable the construction of

~~in which it became apparent that~~ ~~the~~ low level register ~~and~~
transfer modules ^{are} were passed ^{for the construction of low volume computers. Everyone is using and contributing} in that LSI and VLSI technology, ~~are~~ ^{enabling the construction of} ~~to the standard.~~
~~to build complete~~
were driving the construction of modules made up of computers. ~~current~~ Although the

research is ^{predicated on large structures employing}
up to a ^{hundred} ~~two~~ or so processors, this first

chapter ~~can~~ describes the ~~first~~ ^{culmination of the first}
- The
the
10 processor phase,

he proposal for the first computer modules was funded by NS the
National Science Foundation and the proposed direction was set described in

1973 (Bell et al 1973) ~~sdal;l;l;l;laaslsfljlsldfljlljlljkkjkkkllkllkllkkjkkjjjj~~

^A ^{of Cm*}
~~The motivation for the work is described based on the diversins cost~~

a diseconomy of scale for
large computer introductions

~~factors for large computers (i.e there is diseconomy of scale) over hte last
during two years 1975-1977. The need for the Computer modules, versus a~~

^S
~~multiprocessor structure (C.mmp) and computer networks are described in terms of
to further provide additional context for the research.
whet e the type of problems each wish be appropriate for.~~

performance and problem suitability

~~Most of this chapter is~~ The chapter gives a ^{straight forward} complete description of
their associated and

the modules structure, together with the ~~design problems and the problems~~ potential research topics problems (research).
~~problems associated iwth building such a structure~~

~~In a Since this chapter is the final~~

~~This chapter represents the culmination of the v first phase of
research project involving the construction of al alarse mut multiprocessor
structure, and although only 10 processors are eval used, there is a
very good evaluation of the structure agains in terms of a number of different
program types.~~

^{final}
^{most}
The important part of the chapter ~~is~~ ~~co~~ gives
the ~~Cm*~~ the performance of
evaluates Cm* for ~~so~~ five different problems.

C.VMP: The Architecture and Implementation of a Fault Tolerant Multiprocessor

C.VMP is a voting triplicated, voting multiprocessor designed to ~~help~~ understand

the difficulty (or ease) of using a standard off-the-shelf

LSI-11s to ^{provide} obtain greatly increased reliability. Since there is ~~a~~ ^{an} increased

growing concern for increased ^sreliability because systems are ~~getting~~ ^{becoming}

more complex, ~~and~~ are used for more critical applications, but ~~and~~

~~more~~ moreover there is a ^{basic} higher ^{cost} cost of maintenance for all systems that ~~are~~ ^{are} increasing.

~~might be improved.~~

^{Several} A number of design goals are predicated ^{initially} and the work is carried

out against these ~~design~~ ^{Two} goals. Some of the more interesting ~~designs~~

goals include ^{using} off the shelf hardware and software with no modifications to the components.

^{Because}

~~Here, the designer and because the a designer is also the~~ ^{and analyze} ~~also carries out the work and writes the paper, this paper provides a~~ ^{chapter}

great deal of ^{high reliability} insight into the design and design process. ~~Since the~~ ^{also}

systems ~~has~~ ^{end} been operated for several months, there is ~~a great deal to be~~

~~the fact that the first phase of work is also complete is especially valuable.~~

-- especially its evaluation

~~a One interesting observation is not possible~~

The goal of software and hardware transparency turned out ~~to be~~ ^{of the} more easy because of an idiosyncrasy possible because ~~we even with floppy id disks drives because~~ the controller, this controller of

for floppy disks operates on a word at a time transfer from a one ^{sector} sector buffer after a sector is transferred - thus

(It is unclear how the system ~~or~~ would have been designed without

this type of disk, but ^{controller} at the minimum, some form of software

transparency goal would have to be ^{been violated together with a} relinquished, or some form of

significant controller modifications.

voting is carried out ~~at a~~ at a very low level (i.e. as bus transfers are made).

A ~~number~~ number of models are ~~given~~ ^{by the} which ~~evaluated~~ ^{is evaluated,} ~~inter design~~ together.

~~These are based on various component reliabilities which come from~~

Various component reliabilities are used and these

~~read~~

~~is a great deal~~

reader should get a great deal of insight ~~as into~~

~~to~~ the factors contributing to reliability. It ~~is especially~~ should be

noted that a special router (hardware) is ~~highly~~ needed to in

~~no~~

order to ~~no~~ build a ~~practical~~, ~~m~~ marketable C.vmp.

While the intent of C.vmp is not a product, ~~there~~ ~~there is~~ it

~~much~~

~~much of~~

~~the~~

does provide ~~much~~ insight ^{about what} into ~~the factors~~ such a product might

~~con~~ for a & 1

for such a product

Correct and make 2 copies for Hi

Gordon

PDP-11 GLUE 12/15/77

edited
12/16/77

This part ^{can stand on} is ~~quite capable of standing~~ alone as a book on the PDP-11 computer structure evolution, although it relies ~~heavily~~ on the conceptual framework of Chapter 1, and the ISPS language description given in appendix 00. The 11 has evolved quite differently than the other computers in this book, and as such provides an ^{independent and} interesting story. The factors that have created the machines are clearly market and technology based, generating a large number (ten) of different implementations over its relatively short seven year lifetime. ^{Because} Thus, ~~the fact that~~ there is an expanding range of machines ^{as various} ~~we move along~~ in time provides problems and insight in addition to the traditional mini (8 Family); the best ^t cos/performance (18-bit), and the high performance less than \$500K (the DECsystem 10). The 11 designs cover a range of a factor of 500 in system price (\$500 to \$250,000) and 500 in memory size (4 Kwords to 2 Mwords).

This part is divided into five sections consisting of:

1. Introduction.

This ^{first} chapter ⁽⁰⁰⁾ published when the 11 was announced introduces the architecture, gives ^{its} ~~various~~ goals, ~~for it~~, and predicts how it might evolve. The family notion is quite strong, although not specific.

~~Alternatively~~, Chapter 00, What Have We Learned From PDP-11, ^{may} ~~might~~ preferably be read next in order to get the broadest overview, and best immediate critique of Chapter ~~00~~. ^{The 11 evolution}

2. Conceptual Basis.

This section contains two papers, that combined with Chapter 00, form some of the conceptual basis ~~for basis~~ for the models. Strecker, Chapter 00, describes the cache memory mechanism ~~as a basis~~ for building high performance models (specifically the 11/70). Levy describes the intercommunication problem among physical components and how busses carry out this task.

3. Implementations.

Of the four chapters, three are on specific implementations (LSI-11, CMU-11 and 11/60) and the fourth (Chapter 00) is ^{a study of} ~~on~~ all models. This latter chapter provides details of the various implementations together with how the designs fit a conceptual model. ~~This is especially useful to understand the performance.~~

4. Evaluation. Chapter 00 evaluates the 11 as a machine for executing FORTRAN. What We Have Learned From PDP-11, Chapter 00 discusses all aspects of the evolution in terms of the introductory Chapter (00).

5. The multiprocessor research computers of Carnegie-Mellon University. Three multiprocessors: C.mmp (Chapter 00), Cm* (Chapter 00) and C.vmp (Chapter 00) give examples of a 16-processor multiprocessor, a set of modules based on LSI-11 and a triplicated, voting multiprocessor computer for high reliability.

A NEW ARCHITECTURE FOR MINICOMPUTERS--THE DEC PDP-11

It is somewhat anticlimactic to discuss this original PDP-11 description here because Chapter 00 explicitly discusses "What We Have Learned from the PDP-11". The purpose of the chapter was originally threefold: to give the PMS and ISP architecture of the PDP-11 as it was first proposed to describe the first (11/20) implementation at a time when the whole architecture had not been worked out or even fully considered; and to show possible extensions.

The reader might note that the computer class definitions (given in 1970) of micro, mini and midi correspond quite closely to those of Chapter 1.

Although we comment on the disparity between the predicted and actual evolution in Chapter 00, some of the important reasons are:

1. The notion of designing with improved technology especially for a family was not understood then. The ^{is} understanding ~~for one of us~~ (Bell), was put forth in a paper in 1972 (Bell, Chen, Rege).
2. The Unibus bandwidth proved unacceptable for all communications at the very high and low end designs. Whereas, this chapter posits a multiprocessor, and multiple Unibusses, this precise structure did not evolve. ^{as a standard,} The bandwidth has subsequently been shown to be adequate for all but the largest configurations when a cache is attached to the

processor as in the 11/60 (see Chapter 00). Note the effect of a 90% cache hit rate is to reduce the number of access to primary memory via the Unibus by a factor of ten!

3. The memory address space was too small.
4. The particular data-type extensions weren't predicted. While floating point data was discussed, the character string and decimal operations weren't described. These data types evolved in response to market need (and COBOL) which didn't exist for DEC at the time the machine was designed.

We have made a major change in the chapter by removing the original ISP description and replacing it with a correct and complete (memory management and floating point) ISPS description that was used when the 11 was evaluated as a Computer Family Architecture (CFA) standard by the U. S. Defense Department.

CACHE MEMORIES FOR PDP-11 FAMILY COMPUTERS

Chapter 00 by Strecker is included for four reasons: it is a clear exposition of the cache memory structure and its design parameters; the cache is the basis of the fast PDP-8 (Chapter 00), the 11/60 (Chapter 00) the 11/70 (page 00), and the KL10 (Chapter 00); the design methodology is well done--it is "good engineering"; and finally the paper is "well-written"--in fact, it received the award for the Best Paper at the

Third Computer Architecture Conference at which it was presented. Thus, since a relevant paper that provides a clear exposition, while illustrating good engineering and being well-written, is so rare we also publish it simply to serve as an encouragement and an example for those who should ~~expose~~, understand and describe their work.

The design process is implicit in the way the work is carried out to determine the structure parameters. It should be noted that it was easy to collect data statistics about the program's behavior since the trace bit, T, permits the 11 to interpret itself on a one at a time basis. The relevant ~~parameters are varied~~ and sensitivity plots (runs) are made to determine the effect of each parameter on the design. In the 11/60, Mudge (Chapter 00) ~~also uses the same methodology~~ and Strecker's ^{program and methodology} traces. One of the important parameters to understand is the time between changes of context. To the best of our knowledge this study is unique, ^{and imp. this understandy is important} because all real time and multiprogrammed systems have many context switches. The PDP-8 cache design (Chapter 00) shows the effect of segmenting the cache for instructions and data. The PDP-8 chapter discusses the performance for a particular design, and the performance numbers are in terms of a speed-up factor. Strecker gives the performance evaluation in terms of cache miss ratios. The performance measures, see Fig. Cachespeed, are related (Lee, 1969) in the following way (assuming an infinitely fast processor):

p = total no. of memory accesses by the processor, Pc
m = no. of memory accesses that are missed by the
cache and how to be referred to Mp

t.c = cycle time of cache memory, Mc
t.p = cycle time of primary memory, Mp
R = t.p/t.c (ratio of memory speeds), where R is
typically 3 to 10

the relative execution speeds are:

$$t(\text{no cache}) = pR$$

$$t(\text{to cache}) = p + mR$$

$$\text{speedup} = pR / (p + mR) = R / (1 + (m/p) R) = a = \text{miss ratio} = m/p$$

therefore

$$\text{speedup} = R / (1 + aR) = 1 / (a + 1/R)$$

note that if $a = 0$ (100% hit), the speedup is R ; while

if $a = 1$ (100% miss), the speedup is $R / (1 + R)$, i.e.,

the speedup is less than 1 (ie. time to reference both memories)

IMPLEMENTATIONS

A MINICOMPUTER-COMPATIBLE MICROCOMPUTER SYSTEM: THE DEC LSI-11

This first chapter of the implementation section was written from the knowledgeable user viewpoint. Although the paper is a descriptive

narrative ^{about} of the design ^{of} from the chips, boards and backplane levels, it lacks insight that the designers at Western Digital or DEC (Lloyd Dickman, Rich Olsen, or Mike Titelbaum) might have provided. An account of the chip-level design is available (Soha and Pohlman, 1974). The design was done at Western Digital by Roberts, Soha, and Pohlman as a relatively general purpose microprogrammed computer that could be used to emulate many computers. When DEC started working with them for use in the interpretation of the 11 ISP, the design took on more of the 11 structure. Two design levels are described in the paper: the 3 chip microprogrammed computer as it is used to interpret the 11 ISP; and the particular PMS module level components as they are integrated into a backplane to form a hardware system. Sebern points out the microprogramming tradeoff that took place between the chip and module levels to carry out normally hardware functions: the time clock, the console, refreshing dynamic random access MOS memory, and power fail control.

Duane Dickman

The subtleties and uniqueness of the module structure are not described, nor are the design alternatives. For example, the alternative, bounded design that is typical of one board microcomputers was considered, though not described. A lower cost one board system, like the VT78, (page 00) has evolved and ~~will eventually be marketed.~~ This structure is shown in Fig. KXT11. The initial module level design was predicated on a quad-sized form factor and plugged into a conventional backplane. The modules are shown on page 00. Since there were not special ICs beyond the 3 chip processor, options tended to be relatively large and often occupied a full quad module. For options that were greater than a quad size, an ingenious

packaging scheme was used for providing interconnection points on the extra half of the module (a double sized module) which was not used as the LSI-11 Bus (requires a double sized module). This permitted multiple board, complex options (e.g., a disk controller) to be packaged as a single option with no interconnection between the boards except the second half of the quad board. As DEC began to build special ICs to interface to the bus, the option sizes began to decrease to occupy a double module. This system is now known as the LSI-11/2. A backplane system with modules is shown in Fig. LSI-11/2. The options available for the two systems have evolved as shown in Table LSI-11 Options.

Table LSI-11 Options

[Teicher to supply]

The effect of a lower cost LSI-11 system is to provide additional applications as we discuss in Chapter 1, page 00.

USING LSI PROCESSOR BIT-SLICES TO BUILD A PDP-11--A CASE STUDY IN
MICROCOMPUTER DESIGN

This paper by the designers of CMU-11 appears both in the module part and in this part on PDP-11 implementations. The Intel 3000 bit slice, herein called a Microcomputer (for Microprogrammed Processor), is used to interpret a PDP-11 ISP. The purpose of the design was to test the assertion that the bit-slice based arithmetic unit with register memory and microprogrammed control would simplify the design and construction of

processors. The 11 was selected as a target problem in order to avoid the temptation of changing the problem (11 processor) to fit the building blocks (the Intel 3000 processor). As such, the authors observed the awkwardness that ultimately resulted in lower (than desired) performance. In retrospect, the Intel 3000 has not become the standard bit-slice architecture that the AMD 2900 series has, but perhaps suffers from being one of the earliest. Detailed comparisons including a breakdown of the various parts of the processor design are given and compared with the LSI-11 and 11/40 designs in terms of performance and cost (IC count and number of control bits in the microprogrammed controller).

A key part of the investigation was to evaluate the computer aided design tools. The Stanford University Drawing System (SUDS) and the SAGE logic simulator were the key components. SAGE was predicated on being able to detect all the design and timing flaws prior to construction. The authors claim that 95% of the errors were detected by the simulation--and all errors were ultimately detectable once the simulation data or machine description was changed. ✓

DESIGN DECISIONS FOR THE PDP-11/60 MID-RANGE MINICOMPUTER

Unlike the reports from an architect's or reporter's viewpoint this chapter is a direct account of the design from the close proximity of the project. Because it is a mid-range machine, the 11/60 is a difficult design for the reasons of the designer characteristics, Chapter 00, page 00. The design is neither the lowest cost nor performs the highest of the 11s, but has to

be the right balance of features, price, and performance against criteria that are usually extremely vague.

Four interesting aspects of computer engineering are shown in the 11/60: the cache to reduce Unibus bandwidth; trace driven design of floating-point arithmetic; providing writeable control store; and increasing the reliability, availability and maintainability.

Whereas the Unibus had previously thought to be inadequate for high performance, the cache is used to unload the i/o traffic and provide high performance without the attendant cost*. The work is based on Strecker's (Chapter 00). The study leading to determining the block size is given. The block size can only conveniently be one since additional traffic is generated by fetching multiple words which would tie up the Unibus with additional traffic.

The use of trace data to design the floating point arithmetic is described together with the resulting design. Note that the 11/60 performs roughly at 11/70 speeds with lower cost. The implementation of the two can be compared in the following table.

Table - Implementation of 11/60 and 11/70.

	<u>11/60</u>	<u>11/70</u>
--	--------------	--------------

Base Pc

x

x

*Using Amdahl's constants, page 00, the reader might compute the bus bandwidth (for i/o traffic) and the address space needs for this speed processor given the cache and compare these needs with the Unibus.

Floating point	x	x
Cache	x	x
Memory management	0	x
	-----	-----
Total	x	x

Microprogramming is used to provide both increased user-level capability and increased reliability, availability and maintainability. The large writeable control store option is described together with its use for data storage and various applications. This option has been recently used for emulating the PDP-8 with OS/8 operating systems.

A general discussion of microprogramming is also given, especially with respect to memory technology advances (see also Chapter 1, page 00). Other semiconductor technology improvements are described together with how they effect price and performance. It is interesting to note that the simple concept of tri-state logic* had such a great effect on the design.

Impact of Implementation Design Tradeoffs on Performance: The PDP-11, A Case Study

This chapter presents a most comprehensive comparison of the eight PDP-11 processor implementations. The work was carried out to investigate various design styles for a given problem--the interpretation of the PDP-11 ISP.

Aside from any conclusions the tables give more insight into processor
*Ability to interconnect a number of subsystems together through a wired-or connection.

implementations than is available from any single source we know. The usefulness of the data also comes from having an outside observer examine the machines and provide insight.

The tables include:

1. A set of instruction use frequencies, by Strecker, for an operating system. *The reader should note the frequencies do not reflect general purpose use, e.g., there are no arithmetic and floating point instructions;*
2. Implementation cost (modules, ICs, control store widths) and performance (micro- and macro-instruction times) *for each model; and*
3. A canonical data path is ~~posited~~ for all 11 implementations, and each processor is presented and compared with it.

With this background data, a "top-down" model is built which explains the performance (macro-instruction time) of the various implementations in terms of the micro-instruction execution, and primary memory cycle time.

Since these two parameters don't fully explain (model) performance, a set of bottom-up factors must be introduced. These factors include various design techniques and the degree of processor overlap. We believe the fact that the problem is constrained should provide useful insight to both computer and general digital systems design.

~~*Ability to interconnect a number of subsystems together through a wired-or-connection.~~

EVALUATION

TURNING COUSINS INTO SISTERS: THE ROLE OF SOFTWARE IN SMOOTHING HARDWARE DIFFERENCES

Since FORTRAN is quite possibly the most often executed language for the PDP-11 and the one we intended that it execute, it is important to observe the 11 architecture as seen by the language processor--its user. The first FORTRAN compiler and object (run) time system are described together with the evolutionary extensions to improve performance. The FORTRAN IV+ compiler is only briefly discussed since its improvements are largely a matter of compiler optimization technology and are less relevant to the 11 architecture.

The chapter title overstates the problem since the five variations of the 11 ISP for floating point arithmetic are made to be compatible by providing what amounts to five separate object (run) time systems and a single compiler. This transparency is provided quite easily using a concept called threaded code. This concept appears to be a very simple interpreter for the PDP-11--and might not be called an interpreter by many. With threaded code, one 1-word instruction requiring two memory cycle times is executed each time the next high level operation code is to be interpreted, otherwise the processor is carrying out the desired op code. When a simple integer expression like $A = A + 1$, which occupies 2 memory words and requires 3 memory cycles to execute, is transformed into a threaded code version the program still only occupies 2 words, but instead requires 5

memory cycles to execute (nearly a factor of 2). For more complex operations requiring longer execution times, like floating point arithmetic, the overhead turns out to be quite low and the space utilization is quite good. Jim Bell discovered this technique and it has been used extensively for other compilers because the time and space overhead is so low. The ability to carry out the interpretation so elegantly was not part of the original design, but rather turned out to be possible based on the generality in the 11's indexing modes.

The first version of the FORTRAN machine constructed was a simple stack machine. As such, the execution times turned out to be quite long. By recognizing the special high use frequency cases (e.g., $A = 0$, $A = A + 1$, etc.) and by having better conventions for three-address operations to and from the stack, speedups of 1.3 and 2.0 for floating point and integers, respectively, were obtained when a more complex machine was constructed for the second version.

Since this paper is really on the construction of an extension to the 11 to execute FORTRAN, it is interesting to compare it with the FORTRAN IV+ machine which uses the floating point processor (11/45, 11/55, 11/60, 11/70). If the FORTRAN machine described in the paper is microprogrammed and made to operate at FPP speeds, the two machines turn out to operate at roughly the same speed and programs occupy roughly the same program space.

WHAT HAVE WE LEARNED FROM THE PDP-11?

Although ~~it is~~ not used, it is possible to more easily efficiently between threaded and directly executed code.

This chapter is a substantially revised version* of a paper written for the CMU Computer Science 10th Anniversary, September 1975. This paper was written to critique the original expository paper on the PDP-11 (Chapter 00) and to compare the actual with the predicted evolution. The four critical issues of technology, bus bandwidth (and PMS structure), address space and data-type evolutions are examined.

The first part of the chapter discusses how the technology is used as a basis for the evolution (something we did not understand when the machines were originally planned). The role of semiconductor memories is especially critical. The next section describes the evolution from the point of view of the various development projects and people. Some early (historical) design documents are introduced to further aid in understanding the design process.

The main section consists of: evaluating the Unibus, examining the cost-performance evolution, the ISP and discussing the organizational issues behind why multiprocessors haven't evolved. As a comparison, Chapter 00 on C.mmp describes the technical problems associated with multiprocessors.

The Unibus evolution is given and the case is made for its optimality. The Unibus has had greater longevity and use than any of the other DEC busses and compares favorably with the IBM I/O Channel Bus as a universal standard of interconnection.

*The paper is 50% longer. The introductory overview has been deleted (and is now placed in Chapter 1) and the sections on the 11/45 and 11/70 have been greatly expanded to include the perturbations due to the memory address and protection extensions. A detailed evolutionary model of the cost, and performance characteristics has been added. More historical facts are introduced, particularly as they effect the design of the extensions. Finally the basis (need) for the VAX/11 extension is discussed.

We try to provide a set of evolutionary cost-performance metrics so the 11 can be compared with the other machines (18-, 12- and 36-bit) in the book (Chapter 00, 01, and 02). Also, here we go into the unique problem the 11 has of designing a range of machines.

Although an ISP evaluation is given, it is quite weak. By comparison, Chapter 00 by Brender, gives a ^{more} useful evaluation of the architecture for FORTRAN execution. A complete section is given on the addressing extension beyond the 11/45 and 11/70 extensions which required a major perturbation in the form of the VAX/11 extensions.

The final section to the research community describes some general problems encountered in structure design and engineering, together with how solving them might be useful in subsequent designs. (This part has been expanded too.)

MULTIPROCESSOR

THE RESEARCH COMPUTERS OF CARNEGIE-MELLON UNIVERSITY

These three multiprocessor computers which use the 11 as a basis were built at Carnegie-Mellon University to carry out various computer structures, operating system and application program computer engineering and computer science research.

The first computer, C.mmp, is a 16 processor (model 11/40's) system with 1

million words of shared primary memory. It was built to investigate the programming (and resulting performance) questions associated with having a large number of processors.

The second computer, Cm*, is also discussed in the Modules Part as the modules that form Cm* are LSI-11 computers. Cm* evolution was based on the premise that ultimately, the smallest modular unit (i.e., ultimately all ICs) that would be used to build digital systems is the computer. With this premise, it is important to understand how to interconnect *them computers* physically, how to assign parts of the problem to the various computers and how to program *them, complete structure.*

W The C.vmp, for voting multiprocessor, was designed to investigate how a production microcomputer, the LSI-11, can be used to build a triplicating, voting high availability computer.

The goals of the first two are performance while the third uses multiple computers for redundancy and reliability. To this end, Fig. Perf shows the effect of using multiprocessors to execute various algorithms (More to come here!).

We believe that technology will force the evolution of computing structures to be along all three lines of multiprocessor computers: C.mmp ~~for~~ for high performance, incremental performance and availability; less tightly coupled computers, *like* Cm* ~~for~~ for more ad hoc structures to handle specialized processing (e.g., front end, file, signal processing); and C.vmp--for high

availability based on increased maintenance costs.

The technology force argument is based on history, near term technology and resulting price extrapolations. It and follows:

1. The MOS technology is increasing in both speed and density faster than the technology (e.g., ECL) from which high performance machines is formed;
2. The price per chip of the MOS one chip processors decreases at a substantially greater rate than low volume special designs due to the much higher volumes and very high design costs (for both designs);
3. For all intents and purposes, the processor cost of the low end is 0 relative to all other costs of the system! The resulting performance (in operations) per chip and cost per operation per chip are rapidly diverging from the high performance semiconductor technology ^{from} which conventional high performance machines are formed.
4. Standards (in the semiconductor industry) for high volume products tend to form more quickly. For example, in the ⁸18-bit microcomputer market, 1-type~~s~~ supplies about 50% of the market and 3-types supply over 90% of the market.
5. A 16-bit processor (-on-a-chip) with an acceptable (for the performance) address space and appropriate data-types is eminent.
Such a commodity will form the basis for all future designs.

¹
nearly

These factors product better (in terms of price and price/performance) computers at the low cost market end at a diverging rate compared with large scale computers. Furthermore, large scale applications have been slow to form since problems complexity increases more rapidly than program size. Therefore, most subsequent computers will be based on standard, high volume parts. For high performance machines, since processing power is available at 0 cost from a processor-on-a-chip based processors, large scale computing will come from arrays of processors, just as we build arrays of 16K bit ICs to form memory.

C.mmp--A Multi-Mini-Processor

C.mmp was motivated by the need for more computing power to solve speech recognition/signal processing problems and to understand the multiprocessor software problem. Until C.mmp, only one large, tightly coupled multiprocessor had been built--the Bell Laboratory's Safeguard Computer (BSTJ issue?).

The introductory section describes the economic and technical factors influencing multiprocessor feasibility and argues that it's important to embark on the research because of the timeliness. Various problems to be researched are given together with a discussion of particular design aspects. For example, since C.mmp is predicated on a common operating system there is degradation both due to memory contention when all the processors use the same memory block and since there are common synchronizing locks for parts of the software, degradation may occur when a

processor must idle and wait to enter common critical sections. The machine's theoretical performance due to memory processor interference is computed based on Strecker's (1971) work. In practice, memory interference was a problem resulting in poorer than expected performance because the memory was not built with low order address interleaving. This problem had to be solved by moving program segments of the memory. It should be noted that when the number of memory modules and processors become very large, the performance (as measured by the number of accesses to the memory by the processors) approaches the memory bandwidth ($m/\text{memory-cycle-time}$) $\times \frac{1}{e}$. Thus, there is not a maximum limit on performance with infinite processors provided they are all not contending to access the same memory.

Although there is a discussion outlining the design direction of the operating system, Hydra, later descriptions should be read (list the references?? here). Since the 11's small address impaired use, 11/40s with microprogrammed writeable control stores were used to implement operating systems calls involving changing the segment base registers.

One of the most pleasant surprises of C.mmp was the ALGOL 68 implementation because it enables parallel programs to be specified. The result (see page 00) in terms of performance are quite encouraging! There are three basic approaches to effectively applying multiprocessors:

1. having lots of independent work to do through multiprogrammed, timeshared and multiple independent computers;

2. having the compiler for a conventional language (e.g., FORTRAN) detect and compile statements that can be executed in parallel;
3. introducing new primitives into the programming language (e.g., ALGOL 68) so that algorithms for parallel execution are specified by the programmer.

Although C.mmp was only predicated on type 1 use, the ALGOL 68

implementation on C.mmp ^{explores} ~~makes type 3 use possible (see ?)~~. ^{permits parallel program}
(type 3) to be written. The results of some of this programming is shown in Fig. ?.

Multi-Microprocessors: An Overview and Working Example

Cm* is a system of high level modules, constructed of LSI-11 computers, which can be interconnected to form large computer structures. The Cm* work sponsored by NSF and ARPA are an extension of NSF sponsored research (Bell, etc. 1973) on register transfer modules. As LSI and ULSI enable construction of the processor-on-a-chip, it is apparent that low level

register transfer modules are passe' for the construction of low volume, specialized computers. ^{A complete industry will ultimately design, contribute to and use a} ~~Everyone is using and contributing to the standard,~~ ^{all but} Although the research is predicated on structures employing a hundred or so ^{computer.} processors, this chapter describes the culmination of the first ten processor phase.

A motivation of Cm* is based on a diseconomy of scale for large computer introductions during 1975-1977. ^(see page 00) Computer modules (Cm*), multiprocessors (C.mmp) and computer networks are described in terms of performance and

problem suitability to provide additional context for the research. The chapter gives a description of the modules structure, together with their associated and potential problems (research).

The final, most important part of the chapter evaluates the performance of Cm* for five different problems.

C.vmp: The Architecture and Implementation of a Fault Tolerant Multiprocessor

C.vmp is a triplicated, voting multiprocessor designed to understand the difficulty (or ease) of using standard, off-the-shelf LSI-11s to provide greatly increased reliability. There is concern for increased reliability because systems are more complex, are used for more critical applications, and basic maintenance cost for all systems are increasing. Because the designers carry out and analyze the work, this chapter provides a great deal of insight into high reliability designs and design process--especially its evaluation. The system has operated for several months and the first phase of work is complete.

Several design goals are initially predicated and the work is carried out against the goals. Two of the more interesting goals include using off-the-shelf hardware and software with no modifications to the components.

The goal of software and hardware transparency turned out to be easier
than expected

because of an idiosyncrasy of the floppy controller. This controller operates on a word-at-a-time transfer from a one-sector buffer after a sector is transferred--thus voting is carried out at a very low level (i.e., as bus transfers are made). It is unclear how the system would have been designed without this type of controller, but as a minimum, some form of software transparency goal would have been violated together with a significant controller modification.

A number of models are given by which the design is evaluated. Various component reliabilities are used and the reader should get a great deal of insight into the factors contributing to reliability. It should be noted that a special hardware voter is needed in order to build a marketable C.vmp. While the intent of C.vmp is not a product, it does provide much of the insight for such a product.

~~Created~~ created
12/7/77
Chen (P.)
11 glue

A New Architecture for Minicomputers-The DEC PDP-11

original PDP-11 description

It is ^{somewhat} quite anticlimactic to discuss this ~~paper~~ here because ~~the~~

Chapter 00 explicitly discusses the ~~subject~~ ^{what} We Have Learned

~~About the PDP-11. Nevert A sisini~~

The reader ^{might} ~~should~~ note that the ~~defi~~ ^{that} various definitions of micro-
computer classes ^{definitions of} into micro, mini and midi and ~~thees~~ ^{se} correspond

quite closely to the ^{those of} classes given in Chapter 1, ~~ev~~ ^{despite} ~~despite~~
~~the fact the definitions were made in 1969 and we have not tried to makee the-~~
~~two definitosn coincide.~~

The purpose of the ^{chapter originally} ~~paper~~ was threefold: to give the architecture of the
PDP-11 as it was first

proposed ^(11/20) ~~to~~ ^{or even fully considered} ~~specifically~~ describe the first implementation of the
11/20 (in at a time when the whole architecture had not been worked out; and
to ~~settly~~ show what extensions were possible and ~~might like ly follow~~ (even though
~~the s extensions were not even i at the proposal stage).~~

Although we comment in ~~chra~~ on the ~~disparity~~ ^{some of the important reasons are:} between the predicted
and actual evoution of the 11 in Chapter 00, it is worth of note ~~simply noting~~
that ~~there wer~~ ~~why the two~~ ~~herer why the two~~

are different. ~~Simply:~~
1. ~~The~~ ^{bandwidth proved unacceptable for} notion of desisnins with imporved technology (and famil especially for a family
in view of the family range, was not understood. ~~The evolutionary notion is given~~
~~posited in a paper~~ It is

then. The understanding for one of us (Bill), was
put forth in a paper in P. 1972 (Bell, Chen, Rege, 1972)
2. The Unibus ^{save out as a single bus} for all communications at the high and
low
end designs. ~~Whereas, the chapter~~ ^{key} this ~~paper~~ ^{out} posits a multiprocessor, the-

and especially multiple Unibuses, this precise form did not materialize.

The cache structure did a developer and hence the Unibus is can now still be shown to be adequate for even substantially higher performance machines when a cache is employed.

- 3. The ^{memory} address space was too small.
- 4. The particular data-type extensions weren't predicted. While floating point data was discussed, the character strings and decimal operations weren't described. These later data types evolved in response to market need (and Cobol) which DEC wasn't in at the time the machine was designed.

evolve. The ϵ bandwidth has subsequently shown to be adequate in terms of ϵ/ω for all but the largest configurations when a cache is attached to the processor as in the 11/60 (see chapter 00).

Note the effect of a 90% cache hit rate is to reduce the number of access to Mp via the unibus by a factor of ten!

didn't exist at the time

correct and

One final We have made a major change in the ~~appendix~~ to the ~~chapter~~ by removing the ^{original} ISP description and replacing it with ~~and~~ a complete (memory management and floating point) ISPS ~~des~~ ~~ac~~ the ~~and~~ ISPS description that was used ~~in~~ when the 11 was evaluated ~~by group~~ for standardization in the CFA. This description also includes the floating point and memory management extensions.

as a Computer Family Architecture (CFA) standard by the U.S. Defense Department

It show

Cache Memories For PDP-11 Family Computers

Chapter 00 by

Strecker's chapter is ~~given for~~ included ^{three} for ~~two~~ reasons: it is

a clear ^{exposition} description of the cache memory ^{and its design parameters} structure ~~which is used in~~

~~about the clearest explanation of the cache memory that we would expect~~

~~believe exists; it describes the parameters from a designer computer designers~~

~~standpoint~~

^{since the cache is} ~~As such the cache is the basis of a number of products systems~~

~~described here the PDP-8 (chapter 00), the 11/60 (chapter 00)~~

and the 11/70 (page 00); and the KL10 (chapter 00); the ~~method~~ ^{design}

methodology of the ~~design for~~ carrying out the ^{is well done}

~~design~~ it is "good engineering" ^{practice}; and finally the paper

is "well-written"--in fact, it received the award for the Best Paper

at the Third Computer Architecture Conference at which it was presented.

~~In this latter regard, it is rare~~

~~In all these regards it is rare to find~~

~~A combination of these three attributes: relevancy, a good end, and example~~

~~of good~~ ^{shows} provides a clear exposition, while

~~A paper that combines relevancy with illustrating good engineering and~~

~~is well-written is~~ ^{so} really rare. ^{and} Therefore we also publish it

~~to serve as an encouragement for those who should write (explain)~~

~~and expose, understand and describe their works,~~

~~The explanation is quite clear (and together with, The subtleties~~

~~of a design for working with an operating system environment~~

The design process is implicit in the way the work is carried out to determine

the parameters. The relevant parameters are ^{varied} tried, and sensitivity plots

(runs) are made to determine what effect of each parameter on the design.

One of the important parameters, the time between interrupts (or changes

of context is especially important, ^{and} ~~To our knowledge, this is the only~~

~~work that deals specifically with this issue. There is a The~~

~~The paper on the PDP-8 cache should be reused too, because together~~

~~they have nearly all the design parameters. In the case of The PDP-8, the cache~~

~~effect of segmenting the cache for instructions and data are examined.~~

It should be noted that it was easy to collect data on the behavior since the programs bit, T, permits the trace on at a time basis.

In the PDP-8, Mudge also used the same methodology (Chapter 00).

is both unique and relevant for multi-
any real time or
multi-programmed system.

shows the

↑

I in the original w

Chapters 1) and this chapter give different measures for

Performance

While the PDP-8 chapter discusses the performance for a particular design, the performance numbers are in absolutes. Since this work is

more general the performance evaluation is in terms of the miss rate ratio, cache miss ratios.

The two ~~may be compared~~ performance measures are related (Lee, 1969) in the following way:

(assuming an infinitely fast processor)

Using Fig. cachespeed

p = total no. of memory accesses by the processor, P_c
 m = no. of memory accesses that are missed by
(i.e. not in the cache and have to be referred to M_p)

t_{fast} = cycle time

t_c = cycle time of cache ~~M_c~~ ^{memory, M_c}

t_p = cycle time of primary ~~M_p~~ ^{primary memory, M_p}

$R = \frac{M_p}{M_c} = t_p / t_c$ (ratio of memory speeds), where
 R is typically 3 to 10.

The relative execution speeds are:

$$t (no\ cache) = p R$$

$$t (to\ cache) = p + m R$$

Speedup

$$Speedup = \frac{p R}{p + m R} = R / (1 + (m/p) R) =$$

$$a = \text{miss ratio} = m/p$$

Therefore

$$Speedup = R / (1 + a R) = 1 / (a + 1/R)$$

Note that if $a = 0$ (100% hit), the speedup is R ; while
if $a = 1$ (100% miss), the speedup is $R / (1 + R)$, i.e.
the speed is less than 1.

A Minicomputer-Compatible Microcomputer System: ^{The} DEC LSI-11

The LSI-11 This chapter is the first of ~~the~~ the implementation section on the various implementations. ~~Altho~~ It was written from ~~the~~ a user view knowledgeable user viewpoint, as ~~S~~ Sev Sebern was not a Bill Roberts of then of Western Digital

The ~~paper has none~~. Although the paper is a descriptive narrative of the design from ~~the~~ chips, boards and backplane levels, it ~~has none of the insight~~ that lacks the insight some insight that ~~one might get from~~

one of the designers (~~Roberts, Zohar, or Dickman~~) An account of the ~~design~~ ^{chips-level} design is given ~~in by Roberts~~ (Sohar and Pohlman, 1974) (MJ please ask Dickman or Teicher for a copy of this paper or find out if Roberts wrote a paper for the 74 Nerem Record Part 2 in Oct. 74) Since it ~~The~~

design was done at ~~Western Digital Corporation~~, it ~~was~~ a relatively general ~~purpose~~ ^{purpose} design that could be used to emulate ~~any~~ ^{many} computers. When DEC ~~DEC~~ started working with them for use ~~in~~ the interpretation of the 11 ISP, the design began to take on more of the structure of an 11.

Two design levels are described: ^{in the paper} the ~~chips~~ ^{3 chip} ~~micro~~ microprogrammed computer as it is used to interpret ~~the 11~~ ^{the 11} ISP; and the

^{module board level} ~~the~~ set particular PMS components as they are packaged on a board and integrated into a backplane to form the basis of a hardware system.

~~It is interesting to note that the set the board subtleties~~

The subtleties of the ~~board~~ ^{module} structure are not described and is not described. The initial ~~designs~~ ^{module level} were predicated on a quad-sized form factor which ~~and plusses~~ into a conventional backplane. The

at Western Digital or DEC

(Dickman, Olsen, or Mike Teitelbaum)
Lloyd Rich

by Roberts, Zohar, and Pohlman

Sebern points out an ~~one~~ very interesting tradeoff ^{between the chip and module levels} observed. Can be observed ~~as~~ in the ~~use~~ use of microprogramming to carry out ~~these~~ functions that are normally done with hardware: ~~clock~~ ^{the} time clock, ~~the~~ console, and refreshing dynamic ~~random~~ random access ~~semiconductor~~ MOS memory; and power fail control. and uniqueness

The alternative, more bounded designs that are typical of other microcomputers is not described. However it should

modules are shown in on page 00. ~~Since~~ Since there was not special ICs beyond the 3 chips from which the processor, ~~was formed~~, options tended to be relatively large and often occupied a full quad. For options that were greater than a quad size, an ingenious packaging scheme was

used for providing interconnection points on the extra half of the module

which was not used as a the LSI-11 Bus, (also known as a Q Bus. This a permitted permitted

multiple board, complex options (e.g. a disk controller) to be packaged as a single

option with no interconnection between the boards. As the level of

~~int IC we~~ ^{DEC} began to build special ^{ICS} options ^{to} interface to the ^{bus} Q Bus, the

average ~~size~~ ^{to} of the option sizes began ~~to decrease~~ ^{to occupy a double module} so that ~~a single~~ ^{is now} double board is needed for the interface. This system ~~was~~ ^{is now} to be known as the

LSI-11/2. ~~A set~~ A backplane system with modules is shown in Fig. LSI-11/2.

The ~~options are~~ options available for the two systems ~~are shown in Table~~ have evolved as shown in Table

LSI-11 Options

Table LSI-11 Options

Teicher to supply

~~Notes~~

lower cost module and backplane ~~and~~ availability
The effect of LSI-11 is to provide additional applications as

~~because it is available at a new, lower level~~

we discuss in Chapter 1, page 00.

be noted that in order to get more reduced cost, such a system, ~~like the~~ like the VT78, has evolved and will eventually be marketed.

forming the Q Bus LSI-11 Bus is only a ~~require~~ requires 1/2 of the pins, ~~or a~~ i.e. a double size.

except the second half of the board.

This structure is shown in Fig. LSI-11

Using LSI Processor Bit-Slices to Build a PDP-11--A Case Study in Microcomputer Design

by the designers of CMU-11

This paper appears both in the module section and part and in this part on PDP-11's.

implementations

Two

Here, the Intel 3000 bit slice bit slice, IC is used, herein called a Microcomputer

to interpret a PDP-11 ISP. (for Microprogrammed Processor)

The purpose of the design was to test the assertion of semiconductor

manufacturers to that the micro bit-slice ^{based,} arithmetic unit / with

with register memory and microprogrammed control would simplify

the design of and construction of processors. It is ~~worthy~~ to note

that in retrospect, ~~this~~ the Intel 3000 has not become ^{the} standard *

bit-slice architecture, but ~~it~~ perhaps suffers from being ~~one~~ one of the

earliest, and as such the current de facto standard

architecture. A American Micro by AMD has become the standard.

Since the three authors are also the principle designers, the

there is a great deal of insight into the design. ~~It should~~ There are

detailed comparisons ^{including a breakdown of} about the various parts of the processor design and

these are compared with the LSI-11 and 11/40 designs in

terms of ^{performance and cost} (IC count, number of control bits and in the

the microprogrammed controller) and performance.

A key part of the initial design goal ~~investigation~~ of the

project originally was to evaluate the computer aided design

tools available to assist in the design. Here, the Stanford University

Drawings System (SUDS) and the SAGE Logic simulator were the key ~~part~~ components.

The purpose of the design was to evaluate whether SAGE could detect all the

design flaws prior to and timing flaws prior to construction and in fact,

The authors claim that 95% were detected by the simulation - and all errors were ultimately

* ~~to~~ AMD 2900 series

detected ^{by} once the simulation data or sim machine description was changed.

as a target problem
The 11 was selected in order to avoid the temptation of changing the problem (processor) to fit the building blocks (the Intel 3000 processor). As such, the authors observed the awkwardness that ultimately resulted in lower (than ~~less~~ desired) performance. now the AMD 2900 series)

are given.

was predicated on being able to

Design Decisions for the PDP-11/60 Mid-Range Minicomputer

This In this ~~is~~ ^{chapter} This paper is a ~~direct~~ ^{reporter's} account of the design from the close proximity of the project, unlike several ~~several~~ ^{the} chapters of the papers that reports from an architect's or ~~design~~ ^{reporter's} descriptive viewpoint

Four interesting aspects of computer engineering
It shows ~~Several good design principles~~ are shown in the 11/60. Whereas the

Unibus had previously thought to be inadequate for high performance, the cache is used

to unload the ~~the~~ i/o traffic and provide high performance ~~without~~ ^{without} the attendant cost. ^{The} use of the Strecker work ^{is based on Strecker's} on modeling

cache performance is built on. ^{Chapter 00.} Like the PDP-8, the ~~the~~ ^{The study leading to} block size

can only conveniently be ~~one~~ ^{one additional} since ~~undue~~ ^{is} traffic would be generated by

fetching ~~two~~ ^{multiple} words and thereby tying up the Unibus for ~~an~~ ^{additional} traffic.

(This study is shown, described,

Because it is a mid-range machine,

Like all mid-range designs, the 11/60 is a ~~of~~ ^{difficult} design as for the

reasons we discuss ^{de of with} when we talk about the designer characteristics,

in Chapter 00, page 00. The design is neither the ~~cheapest~~ ^{lowest cost} or performs

the most, but has to be ~~in~~ ^{the} right balance of features, price, and performance

against criteria that are usually extremely vague.

reliability, availability, and maintainability

The use of trace data to design the floating point arithmetic is defined,

There are various algorithms, but ^{with a sketchy view of the} Here, It should be noted that the

11/60 performs ~~roughly~~ ^{roughly} at 11/70 speeds, for a processor floating

point ~~cost~~ ^{of} of x boards versus y boards and a ~~an~~ ^{overall}

described, together with the resulting design

	11/60	11/70
Base Pc	x	x
Floating Point	x	x
Cache	x	x

Memory Manager 0 X

Total X X

~~The third innovation that is char~~

The third aspect of the design that is carried forward to the user level

is micro user microprogramming using a writeable control store

~~the micro~~ is used to
Microprogramming ~~is discussed as it~~ provides both
increased ~~user~~ user-level capability ~~with the~~ when the
writable control store option is used and increased
reliability, ~~and~~ availability and maintainability. ~~Since a large,~~
The writable control store option is ~~described~~. ~~and~~
~~available~~ ~~des~~ described as together with its use for
data storage. ^{and various applications} This option has been ^{recently} used for ~~carry~~
emulating ^{the} PDP-8 ~~RT-8~~ with ~~RT~~ OS/8
Operating Syst operating system.

P A general discussion
of microprogramming is
also given, especially
with respect to
memory
technology
advances
(see also chapter 1,
page 100).

~~As~~ The ^{semiconductor} other ~~aspects~~ of technology improvements
are described ~~together~~ together with their use how they effect
price and performance. It is interesting to note that the simple
concept of tri-state logic* had ~~not~~ such a great effect
on the design.

* Ability to ~~dot wire~~ interconnect a number of subsystems together
through a wired-or connection.

Turning ~~Sisters~~ Cousins into Sisters: The Role of Software in Smoothing Hardware Differences

Since Fortran is ~~the~~ quite possibly the most often executed language for the PDP-11, it is important to have this ~~a d disc~~ ^{observe} ~~discussion~~ ^{the} of the ~~archit~~ 11 architecture ~~as seen~~ by the language processor. ~~On the other hand,~~ ^{- its user,} since ~~the~~ we designed the 11 ~~f~~ to be able to execute Fortran ~~e~~ easily, ~~this~~ paper also provided a critique of the 11 for the ~~this ind~~ intended function. The ~~original Fortran~~ ~~com~~ first Fortran compiler ~~are~~ and (run) object time system are described, together with the ^{evolutionary} extensions (evolution ~~to~~) to improve performance. The Fortran ~~IV~~ ^{IV+} compiler is only briefly discussed, ~~since it~~ ~~the~~ ~~in~~ its improvements are largely a ~~matter~~ matter of compiler and optimization technology and ~~have~~ are less relevant to the ~~architecture~~ ^{architecture} 11.

and the one we intended that it execute

~~The One of the problems~~

~~The title is not completely accurate, because it~~

The chapter title overstates the problem, ~~as it turned~~ ^{since}

~~out quite easy~~ to have the variations of the five of the 11 ISP ~~are~~ ^{for}

~~that dealt with floating point arithmetic~~ ^{are made to} be compatible ~~by~~ providing

what amounts to ~~per~~ five separate object (run) time systems, ^B and a

single compiler, ~~because~~. This transparency comes about ^{simply} because a

concept called threaded code is used to provide what is a very simple

interpreter for the PDP-11, ~~and~~ ^{with threaded code, one word} might not be called a interpreter

by many, ~~since~~ ^{requires} ~~only~~ 1 instruction is executed per transfer of control

for each ~~one~~ two memory cycle instruction is executed each time the

next ~~subroutine~~ ^{is} ~~is~~ operation code to be interpreted is ~~given~~

~~Thus~~ ^F for a simple integer expression like ~~A+1~~ ^{requires} ~~A=A+1~~, which

~~can be exec~~ occupies ² two memory words and takes 3 memory cycles to

requiring two memory cycle times

execute is interpreted— transformed into ^a ~~its~~ threaded code version;

the program still only occupies 2 words, but instead requires

(nearly a factor of 2)

5 memory cycles to execute. ~~Since I.C. Since, in the early machine—~~

Jim Bell ~~discovered~~ discovered this technique and it has been used extensively for

time and space
because the overhead is so low

other compilers, but the ability to carry out the interpretation

so elegantly was not part of the ~~original~~ ^{based on} design, but rather turned

out to be possible from the ~~seriality~~ ^{seriality} provided in the indexing modes.

For longer ~~operational~~ operations such as ~~—~~ require longer execution times,

like floating point arithmetic, the overhead turns out to be quite trivial.

low and the space utilization is quite good.

The first version of the Fortran machine constructed ^{was} a simple stack

machine. As such, the execution ^{time} turned out to be quite long. By ~~recog~~ recognizing

observing the special cases (e.g. $A=0$, $A=A+1$, etc) and having better conventions

for three-address operations to ~~stack~~ and from the stack, speedups of

1.3 and 2.0 were ~~observed~~ ^{observed} when a more complex machine was constructed ~~in~~ for the

second version.

for floating point and integers, respectively,

~~The finale of the paper describes~~

~~Since~~ Since this paper is really on the construction of an alternative extension

to the ~~machine~~ machine to execute Fortran, it is interesting to note that compare it with

in the final comparison the Fortran IV machine which uses the Floating Point

(11/45, 11/55, 11/60, 11/70). The two ~~take~~ have the same

Processor ~~turn~~ out to be roughly the same speed and have the same

program ^{space} efficiency assuming the

occupies the same number of words as a machine ^{that} microprogrammed to interpret the evolved

Fortran machine described in the paper is microprogrammed.

§
10

* Using Amdahl's constants, page 00,
the reader might compute
for with this speed processor
on the

given the cache and compare these needs with the Unibus.
the bus bandwidth induced needed (for i/o traffic)
and the address space needed based

and the address space needs

What Have We Learned From the PDP-11?

This chapter was originally written for the CMU 10th anniversary Computer Science ~~45th~~ 10th anniversary, September 1975. in the summer of 1975. The ~~or~~ chapter This

††† This chapter is a substantially revised version of a paper written

* The paper is 50% longer.

The ~~ins~~ introductory overview has ~~been~~ ^{been} deleted (and is now placed in Chapter 1), ^{and} the section on the 11/45 and ~~12~~ 11/70 have been ~~greatly~~ greatly

expanded to include the ~~a major~~ perturbations of the memory address ~~and~~ ^{due to} and protection.

extensions, and there has been a section added which gives a ~~bit~~ more ~~extensive~~ ^{A detailed} evaluation of the performance evolution of the machine in

terms of cost, ~~and~~ performance ~~as~~ ^{model of the} characteristics ~~is given~~ ^{my model is given} of the

~~has~~ been added. More historical facts are introduced, particularly as they effect the design of the ~~11/45~~ extensions. Finally the basis (need) for ~~the~~ VAX/11 extension is discussed.

This ~~was~~ ^{was} written to critique the original paper ~~was~~ intended to be a critique of the original

expository paper on the PDP-11 (chapter 00) and to show how the actual

compare the actual evolution with the predicted evolution. It does this ~~generally~~ ^{with respect to} the four critical issues of ~~as~~ described ~~earlier~~ in ~~part~~ on page 00, and it gives an elaborate comparison

of the differences especially vis a vis the issue of multiprocessors.

The first part of the chapter

A section discusses the use of ~~te~~ how the technology is used ~~as~~ as a basis

for the evolution (somethings we did not understand when the ~~the~~

machines were originally designed. ^{planned} ^{IP} The evolution of the Unibus ^{evolution} is given

~~busses is given~~ and the case is made for the Unibus optimality. Certainly the ⁽²⁾

Unibus has had greater longevity and use than any of the other DEC busses

and compares favorably with the IBM ~~and~~ I/O Channel ~~Bus~~ Bus as a universal standard of interconnection.

The next

A section ~~is~~ describes the aspects of the evolution of the from the point of

view of the people who led the various ^{projects development} together with how these interacted and people. Some early historical documents are introduced ~~as~~ ^{to} attempt further an understanding of the design process.

a single footnote

The role of memories is especially critical.

technology evolution, bus bandwidth, address space (and RMS structure) and data-type evolutions ~~are~~ ^{are} described examined,

4) Although a section is given on the ISP ^{an evaluation} is given, it is quite weak and the Chapter 00 by Brender, is far more useful in terms of the evaluation of the architecture. ^{By comparison, 4} Some for Fortran execution

6) The ^{find} first section describes ~~some~~ ^{some} of the general problems ^{encountered in} in computer structures and design and engineering, ^{solving them} together with how these might be useful in ~~to~~ ^{research problems to solve} subsequent designs.

8) The main ^{section} part of the paper or chapter 1 consists of evaluating: the Unibus, examining cost, the cost-performance evolution, the ~~and~~ the ISP and asking why multiprocessors haven't evolved.

3) We try to provide ~~also~~ a set of ~~evolutionary~~ ^{evolutionary} cost performance ^(18-bit, 12- and 36-bit) metrics so the 11 can be compared with the other machines in the book (chapter 00, 01, ^{and} 02 ^{of} ~~the~~ ^{the} book). Also here we go into the ~~two~~ ^{two} unique problem the 11 has of designing a range of machines.

5) A complete section is given on the ~~problems~~ ^{problems} addressing ~~&~~ [&] extension beyond the 11/45 and 11/70 extensions ~~so~~ ^{which} ~~cannot~~ ^{cannot} required a major perturbation in the form of ^{the} VAX/11 extensions.

PDP-11 GLUE 12/16/77

This part ^{could} ~~can~~ stand alone as a book on the ^{evolution of the} PDP-11 computer structure ~~evolution~~, although it ^{does rely} ~~relies~~ on the conceptual framework of Chapter 1, and the ISPS language description given in appendix 00. The 11 has evolved quite differently ^{from} ~~than~~ the other computers in this book, and as such provides an independent and interesting story. The factors that have created the machines are clearly market and technology based; ^{they have generated} ~~generating~~ a large number ~~(ten)~~ ^(ten) of ~~different~~ implementations ^{are multiple implementations spanning a performance range} over its ^a relatively short ^(eight) ~~(seven-year)~~ lifetime. Because there ^{is an expanding range of machines} ~~is an expanding range of machines~~ ^{at the same points in time, the PDP-11 provides} ~~versus time~~ ^{which did not occur in the evolutions of} provides problems and insight ~~in addition to~~ ^{machines} the traditional mini (8 Family); the best cost/performance ^(18-bit) ~~(18-bit)~~, and the high performance ^{machines} ~~less than \$500K~~ (the DECsystem 10). The 11 designs cover a range of ^{500:1} ~~a factor of 500~~ in system price (\$500 to \$250,000) and 500:1 in memory size (4 Kwords to 2 Mwords).

^{The} This part is divided into ^{six} ~~five~~ sections, ~~consisting of~~:

1. Introduction.

^{The} This first chapter (00), published when the 11 was announced, introduces the architecture, gives its goals, and predicts how it might evolve. The family notion is quite strong, although not specific ^{about members}. Chapter 00, What Have We Learned From PDP-11, might be read next in order to get the broadest overview, and best immediate critique of the 11 evolution.

2. Conceptual Basis.

This section contains two papers, ^{which when} that ~~combined with Chapter 00,~~ form some of the conceptual basis for the models. Strecker, Chapter 00, describes the cache memory mechanism for building high performance models (specifically the 11/70). Levy describes the intercommunication problem among physical components and how busses carry out this task.

3. Implementations.

Of the four ^{implementation} chapters, three are on specific implementations (LSI-11, CMU-11 and 11/60) and the fourth (Chapter 00) is a study of all models. This latter chapter provides ^{comparative data on} ~~details~~ of the various implementations together with how the designs fit a conceptual model.

4. Evaluation. ^{new line} Chapter 00 evaluates the 11 as a machine for executing

FORTRAN. What We Have Learned From PDP-11, Chapter 00 discusses ~~all~~ aspects of the ^{PDP-11} evolution in terms of the ^{models introduced in Chapter 1.} ~~introductory Chapter (00).~~

5. VAX-11 ^{this paper, by the architect of VAX-11, discusses the new architecture and its first implementation, the VAX-11/780.}

6. The multiprocessor research computers of Carnegie-Mellon University.

^{insert blank line} Three multiprocessors: C.mmp (Chapter 00), Cm* (Chapter 00) and C.vmp (Chapter 00) give examples of a 16-processor multiprocessor, a set of ^{computer} modules based on LSI-11, and a triplicated, voting multiprocessor computer for high reliability.

It is somewhat anticlimactic to discuss this original PDP-11 description here because Chapter 00 explicitly discusses "What We Have Learned from the PDP-11". The purpose of the chapter was originally threefold: to give the PMS and ISP architecture of the PDP-11 as it was first proposed, to describe the first (11/20) implementation, ^{family} at a time when the whole architecture had not been worked out or even fully considered, and to show possible extensions. *This was attempted*

~~The reader might note that~~ The computer class definitions (given in 1970) of micro, mini and midi ^{have stood the test of time for they} correspond quite closely to those of Chapter 1.

The major reasons (elaborated upon in Chapter 00) for

~~Although we comment on the disparity between the predicted and actual evolution in Chapter 00, some of the important reasons are:~~

1. The notion of designing with improved technology, especially for a family, was not understood ^{in 1970} ~~then~~. ^{It came later} This understanding ^{came later and} was put forth in a paper in 1972 (Bell, Chen, Rege).

2. The Unibus ~~bandwidth~~ ^{most} proved unacceptable for all communications at the very high and low end designs. ^{Although} ~~Whereas~~, this chapter posits a multiprocessor and multiple Unibusses, ^{for high end designs, this exact} ~~this precise~~ structure did not evolve as a standard. ^{Levy's chapter elaborates on the bus evolution.}

The bandwidth has subsequently been shown to be adequate for all but the largest configurations when a cache is attached to the processor as in the 11/60 (see Chapter 00). Note the effect of a 90% cache hit rate is to reduce the number of access to primary memory via the Unibus by a factor of ten!

*delete,
but make
sure points
are made
in Levy*

3. The ^{physical} memory address space was too small.
4. The particular data-type extensions were ~~not~~ ^{not} predicted. While floating point ^{arithmetic} ~~data~~ was discussed, the character string and decimal operations were ~~not~~ ^{not} described. These data types evolved in response to market need ~~and COBOL~~ ^{- factors which did not exist in 1970.} which didn't exist for DEC at the time the machine was designed.

We have made a major change in the chapter by removing the original ISP description and replacing it with a correct and complete ^{by adding} (memory management and floating point) ISPS description ^{given in Appendix B of the book.} that was used when the 11 was evaluated as a Computer Family Architecture (CFA) standard by the U. S. ~~Defense Department.~~

CACHE MEMORIES FOR PDP-11 FAMILY COMPUTERS

Chapter 00 by Strecker is included for four reasons: it is a clear exposition of the cache memory structure and its design parameters; the cache is the basis of ^a ~~the~~ fast PDP-8 ^[Bell et al, 1974] (Chapter 00), the 11/60 (Chapter 00) the 11/70 (page 00), and the KL10 (Chapter 00); the design methodology is well done--it is ^a good engineering[?]; and finally, the paper is ^a well-written[?]--in fact, it received the award for the Best Paper at the Third Computer Architecture Conference, ~~at which it was presented.~~ Thus, ~~since a relevant paper that provides a clear exposition, while illustrating good engineering and being well-written, is so rare~~ We also publish it simply to serve as an encouragement and an example for those who should

understand and describe their work. ^{and relevant} such well-written papers are only too rare.

^{cache} The design process is implicit in the way ^{in which} the work is carried out to determine the structure parameters.

^{Note} It should be noted that it ^{is} was easy to

collect ~~data~~ ^{PDP-11} statistics about the program's behavior since the trace bit, ^{in the PS word}

~~It~~ ^{single-instruction} permits the 11 to interpret itself on a one at a time basis. The

relevant sensitivity plots (runs) are made to determine the effect of each parameter on the design. In the 11/60, Mudge (Chapter 00) uses Strecker's program traces and methodology. One of the important parameters to

understand is the time between changes of context. To the best of our

knowledge this study is unique. ~~This understanding is important~~ because

all real time and multiprogrammed systems have many context switches. The

PDP-8 cache design (Chapter 00) shows the effect of segmenting the cache

for instructions and data. ^P ~~The PDP-8 chapter discusses the performance for~~

~~a particular design, and the performance numbers are in terms of a speed-up~~

~~factor.~~ Strecker gives the performance evaluation in terms of cache miss

ratios. ^{three two} ~~The~~ ^{whereas the 8 data are in terms of a speed up factor.} performance measures, see Fig. Cachespeed, are related (Lee,

1969) in the following way (assuming an infinitely fast processor):

p = total no. of memory accesses by the processor, Pc

m = no. of memory accesses that are missed by the cache and how to be referred to Mp

t.c = cycle time of cache memory, Mc

t.p = cycle time of primary memory, Mp

R = t.p/t.c (ratio of memory speeds), where R is

typically 3 to 10

→ confirm KL10

→ differs between 11 and 8

→

^P A cache design for a PDP-8 [Bell et al, 1974] is summarized in the introduction to Part II. Two differences from the 11 work are of interest. The 8 study shows the effect of separating instructions and data, whereas the 11 does not. * The second difference is that

the relative execution speeds are:

$$t(\text{no cache}) = pR$$

$$t(\text{to cache}) = p + mR$$

$$\text{speedup} = pR / (p + mR) = R / (1 + (m/p)R) = a = \text{miss ratio} = m/p$$

therefore

$$\text{speedup} = R / (1 + aR) = 1 / (a + 1/R)$$

note that if $a = 0$ (100% hit), the speedup is R ; while

if $a = 1$ (100% miss), the speedup is $R / (1 + R)$, i.e.,

the speedup is less than 1 (i.e., time to

reference both memories)

IMPLEMENTATIONS

A MINICOMPUTER-COMPATIBLE MICROCOMPUTER SYSTEM: THE DEC LSI-11

~~This first chapter of the implementation section~~ was written from the *viewpoint*
of a knowledgeable user, ~~viewpoint~~. Although the paper is a descriptive
narrative about the design ^{at each of} ~~of~~ the chips, boards, and backplane levels, it
lacks insight that the designers at Western Digital or DEC (Duane Dickhut,
Lloyd Dickman, Rich Olsen, or Mike Titelbaum) might have provided. ^{It} An
account of the chip-level design is available (Soha and Pohlman, 1974).

The design was done at Western Digital by Roberts, Soha, and Pohlman as a relatively general purpose microprogrammed computer that could be used to emulate many computers. When DEC started working with ~~them for use in the~~ ^{the designers to effect} interpretation of the 11 ISP, the design took on more of the 11 structure. Two design levels are described in the paper: the 3 chip microprogrammed computer as it is used to interpret the 11 ISP; and the particular PMS-~~module~~ level components as they are integrated into a backplane to form a hardware system. Sebern points out the microprogramming tradeoff that took place between the chip and module levels to carry out ^{functions} ~~normally~~ ⁱⁿ hardware: ~~functions~~ the time clock, the console, refreshing dynamic random access MOS memory, and power fail control.

→ RXT-11. The RXT11 is an integrated system containing an LSI-11 chip set, 32K words of memory, connectors for six EIA interfaces, and a controller for two floppy disk drives. 175 i.c.'s were used - to implement the same functionality using standard LSI-11 modules, 375 i.c.'s would have been used. [G.B.: Not announced yet - probably March]

The subtleties and uniqueness of the module structure are not described, nor are the design alternatives. For example, ~~the alternative~~ ^a bounded design that is typical of one board microcomputers was considered, though not described. ^{However, a} ~~lower-cost~~ ^{lower-cost} one-board system, like the VT78, (page 00) has evolved and is shown in Fig. ~~RXT11~~ ^P. The initial module-level design ^{for the LSI-11} was predicated on a quad-sized form factor ~~and plugged into~~ ^{with} a conventional backplane. The modules are shown on page 00. Since there were not special ICs beyond the 3 chip processor, options tended to be relatively large and often occupied a full quad module. For options that were greater than a quad size, an ingenious packaging scheme was ~~used for providing~~ ^{designed. It provided} interconnection points on the extra half of the module (a double sized module) which was not used as the LSI-11 Bus (requires a double sized module). This permitted multiple board, complex options, {e.g., a disk controller}, to be packaged as a single option with no interconnection

Handwritten initials/signature

between the boards except ^{via} the second half of the quad board. As DEC began to build special ICs to interface to the bus, the option sizes ~~began to~~ decreased to occupy a double module. This system is now known as the LSI-11/2. A backplane system with modules is shown in Fig. LSI-11/2. The options available for the two systems have evolved as shown in Table LSI-11 Options.

Table LSI-11 Options

[Teicher to supply]

The effect of a lower cost LSI-11 system is to provide additional applications as we discuss in Chapter 1, page 00.

USING LSI PROCESSOR BIT-SLICES TO BUILD A PDP-11--A CASE STUDY IN MICROCOMPUTER DESIGN

This paper by the designers of CMU-11 appears both in the module part and in this part on PDP-11 implementations. The Intel 3000 bit slice, herein called a Microcomputer (for Microprogrammed Processor), is used to interpret a PDP-11 ISP. The purpose of the design was to test the assertion that the bit-slice based arithmetic unit with register memory and microprogrammed control would simplify the design and construction of processors. The 11 was selected as a target problem in order to avoid the temptation of changing the problem (~~11 processor~~ ^{a real ISP}) to fit the building blocks (the Intel 3000 processor). ~~As such~~ ^{Indeed}, the authors observed the awkwardness ^{as} that ultimately resulted in lower (than desired) performance.

In retrospect, the Intel 3000 has not become the standard bit-slice architecture that the AMD 2900 series has; but perhaps ^{it} ~~suffered~~ ^{ad} from being one of the earliest. Detailed comparisons including a breakdown of the various parts of the processor design are given and compared with the LSI-11 and 11/40 designs in terms of performance and cost (IC count and number of control bits in the microprogrammed controller).

A key part of the investigation was ^{the} ~~to~~ ^{ion of} evaluate ~~the~~ computer-aided-design tools. The Stanford University Drawing System (SUDS) and the SAGE logic simulator were the ^{major} ~~key~~ components. SAGE was predicated on being able to detect all the design and timing flaws prior to construction ^{of a design}. The authors claim that 95% of the errors were detected ^{in this manner.} ~~by the simulation--and all errors were ultimately detectable once the simulation data or machine description was changed.~~

^{The simulated / real simulated-time / real-time ratio (10⁶/1) did constrain ^{the design process, e.g.,} the number of different runs used to find worst-case conditions.}

DESIGN DECISIONS FOR THE PDP-11/60 MID-RANGE MINICOMPUTER

Unlike the reports from an architect's or reporter's viewpoint, this chapter is a direct account of the design from the close proximity of the project.

A mid-range machine is ^{an} inherently difficult design because for the ~~Because it is a mid-range machine, the 11/60 is a difficult design for the~~

~~reasons~~ ^{we presented in} of the designer characteristics & Chapter 00, page 00. ^[CM: unclear] ~~The design~~

~~As neither the lowest cost nor performs the highest of the 11s, but has to~~ ^{highest performance PDP-11, it} be the right balance of features, price, and performance against criteria that are usually extremely vague.

Four interesting aspects of computer engineering are shown in the 11/60:

the cache to reduce Unibus ~~bandwidth~~^{traffic}; trace-driven design of floating-point arithmetic; providing writeable control store; and increasing the reliability, availability and maintainability.

Whereas the Unibus ~~had~~^{was} previously thought to be inadequate for high performance, ~~the cache is used to unload the i/o traffic, and provide high~~^{systems, by using a cache most processors/references do not use the Unibus and so leave it free for}

^{This gives high} performance without the attendant cost*. The ~~work~~^{cache} is based on Strecker's (Chapter 00). The study leading to determining the block size is given. The block size can only conveniently be one since ~~additional traffic is generated by~~ fetching multiple words ~~which~~ would tie up the Unibus with additional traffic.

The use of trace data to design the floating point arithmetic is described together with the resulting design. Note that the 11/60 performs roughly at 11/70 speeds ^{but at} with lower cost. The implementation of the two can be compared in the following table.

Table - Implementation of 11/60 and 11/70 (count of printed circuit boards)

	11/60	11/70
Base Pc	x 5	x 8 [CM: check 11/70 Sys Manual]
Floating point	x 4	x 4
Cache	x 1	x 4
Memory management	x 0	x 2
	<u>10</u>	<u>18</u>

*Using Amdahl's constants, page 00, the reader might compute the bus bandwidth (for i/o traffic) and the address space needs for this speed processor given the cache and compare these needs with the Unibus. [CM: we should do this for all models in "what we've learned" chapter].

Total **10* **18*

Microprogramming is used to provide both increased user-level capability and increased reliability, availability and maintainability. The large writeable control store option is described together with its use for data storage and various applications. This option has been recently used for emulating the PDP-8 ^{at the} with OS/8 operating system *level.*

A general discussion of microprogramming is also given, especially with respect to memory technology advances (see also Chapter 1, page 00). Other semiconductor technology improvements are described together with how they affect price and performance. It is interesting to note that the simple concept of tri-state logic* had such a great effect on the design.

Impact of Implementation Design Tradeoffs on Performance: The PDP-11, A Case Study

This chapter presents a most comprehensive comparison of the eight ~~PDP-11~~ processor implementations. *used in the ten PDP-11 models* The work was carried out to investigate various design styles for a given problem--the interpretation of the PDP-11 ISP. ~~Aside from any conclusions~~ *alone* The tables give more insight into processor implementations than is available from any single source we know. The usefulness of the data also comes from having an outside observer examine the machines and *show his* ~~provide~~ insight.

The tables include:

*Ability to interconnect a number of subsystems together through a wired-or connection.

1. a set of instruction ~~use~~ frequencies, by Strecker, for ~~an operating~~ *a set of ten different applications* system. The reader should note the frequencies do not reflect ~~general all~~ *uses* purpose use, e.g., there are no ~~arithmetic~~ and floating point instructions; *, nor has operating system code been analyzed.*
2. implementation cost (modules, ICs, control store widths) and performance (micro- and macro-instruction times) for each model; and
3. a canonical data path for all 11 implementations, ~~and~~ *against which* each processor is ~~presented and compared with it.~~

With this background data, a "top-down" model is built which explains the performance (macro-instruction time) of the various implementations in terms of the micro-instruction execution, and primary memory cycle time.

Since these two parameters ~~don't~~ *do not* fully ~~explain~~ *(set)* model performance, a set of bottom-up factors must be introduced. These factors include various design techniques and the degree of processor overlap. We believe ~~the fact~~ *that this* ~~that the problem is constrained~~ *analysis of a constrained* should provide useful insight to both computer and general-digital-systems designers.

EVALUATION

TURNING COUSINS INTO SISTERS: THE ROLE OF SOFTWARE IN SMOOTHING HARDWARE

DIFFERENCES

~~*Ability to interconnect a number of subsystems together through a wired-or-connection.~~

Since FORTRAN is quite possibly the most often executed language for the PDP-11 and the one we intended that it execute, it is important to observe the 11 architecture as seen by the language processor--its user. The first FORTRAN compiler and object (run) time system are described together with the evolutionary extensions to improve performance. The FORTRAN IV~~+~~-PLUS compiler is only briefly discussed since its improvements, ~~are~~ largely ~~a~~ *due to* ~~matter of~~ compiler optimization technology ~~and~~ are less relevant to the 11 architecture.

The chapter title overstates the ^{compatibility} problem since the five variations of the 11 ISP for floating point arithmetic are made to be compatible by ^{essentially} providing ~~what amounts to~~ five separate object (run) time systems and a single compiler. This transparency is provided quite easily using a concept called threaded code. This concept appears to be a very simple interpreter for the PDP-11--and might not be called an interpreter by many. With threaded code, one 1-word instruction requiring two memory cycle times is executed each time ~~the next~~ ^a high level operation code is to be interpreted, otherwise the processor is carrying out the desired op code. When a simple integer expression like $\overset{I}{A} = \overset{I}{A} + 1$, which occupies 2 memory words and requires 3 memory cycles to execute, is transformed into a threaded code version the program still only occupies 2 words, but instead requires 5 memory cycles to execute (nearly a factor of 2). For more complex operations requiring longer execution times, like floating point arithmetic, the overhead turns out to be quite low and the space utilization is quite good. ~~Although not used,~~ ^{also} it is possible to move ~~efficiently~~ ^{Although ~~this~~ this is not done.} between threaded and directly executed code, ^h Jim Bell

~~*Ability to interconnect a number of subsystems together through a wired-or-connection.~~

discovered ~~this~~ ^{the} technique; and it has been used extensively for other compilers because ~~the~~ ^{of its low} time and space overhead ~~is so low~~. The ability to carry out the interpretation so elegantly was not part of the original PDP-11 design, but rather ~~turned out to be possible based on~~ ^{was a consequence of} the generality ~~in~~ ^{of} the 11's ~~indexing~~ ^{addressing} modes.

The first version of the FORTRAN machine constructed was a simple stack machine. As such, the execution times turned out to be quite long. ~~By~~ ^{In the second} recognizing the special high-~~use~~ ^{of use} frequency-cases, e.g., $A = 0$, $A = A + 1$, etc., and by having better conventions for three-address operations (to and from the stack), speedups of 1.3 and 2.0 for floating point and integers, ~~respectively~~, were obtained, ~~when a more complex machine was constructed for the second version.~~

~~Since this paper is really on the construction of an extension to the 11 to execute FORTRAN,~~ ~~It~~ ^{the virtual machine described} is interesting to compare ~~it~~ ^{(on the 11/34,} with the FORTRAN IV+PLUS machine which uses the floating point processors (11/45, 11/55, 11/60, and 11/70). If the FORTRAN machine described in the paper is ~~microprogrammed~~ ^{implemented in microcode} and made to operate at FPP speeds, the ~~two~~ ^{resulting} machines turn out to operate at roughly the same speed and programs occupy roughly the same program space.

WHAT HAVE WE LEARNED FROM THE PDP-11?

This chapter is a substantially revised version* of a paper written for the CMU Computer Science 10th Anniversary, September 1975. This paper was written to critique the original expository paper on the PDP-11 (Chapter

*The paper is 50% longer. The introductory overview has been deleted (and is now placed in Chapter 1) and the sections on the 11/45 and 11/70 have been greatly expanded to include the perturbations due to the memory address and protection extensions. A detailed evolutionary model of the cost, and performance characteristics has been added. More historical facts are introduced, particularly as they effect the design of the extensions. Finally the basis (need) for the VAX/11 extension is discussed.

00) and to compare the actual with the predicted evolution. The four critical issues of technology, bus bandwidth (and PMS structure), address space and data-type evolutions are examined.

The first part of the chapter discusses how the technology is used as a basis for the evolution (something we did not understand when the machines were originally planned). The role of semiconductor memories is especially critical. The next section describes the evolution from the point of view of the various development projects and people. Some early (historical) design documents are introduced to further aid in understanding the design process.

The main section consists of: evaluating the Unibus, examining the cost-performance evolution, the ISP and discussing the organizational issues behind why multiprocessors haven't evolved. As a comparison, Chapter 00 on C.mmp describes the technical problems associated with multiprocessors.

The Unibus evolution is given and the case is made for its optimality. The Unibus has had greater longevity and use than any of the other DEC busses and compares favorably with the IBM I/O Channel Bus as a universal standard of interconnection.

We try to provide a set of evolutionary cost-performance metrics so the 11 can be compared with the other machines (18-, 12- and 36-bit) in the book (Chapter 00, 01, and 02). Also, here we go into the unique ^(within DEC) problem ~~the 11~~

~~*The paper is 50% longer. The introductory overview has been deleted (and is now placed in Chapter 1) and the sections on the 11/45 and 11/70 have been greatly expanded to include the perturbations due to the memory address and protection extensions. A detailed evolutionary model of the cost, and performance characteristics has been added. More historical facts are introduced, particularly as they effect the design of the extensions. Finally the basis (need) for the VAX/11 extension is discussed.~~

~~has~~ of designing a range of machines.

Although an ISP evaluation is given, it is quite weak. By comparison, Chapter 00 by Brender, gives a more useful evaluation of the architecture for FORTRAN execution. A complete section is given on the addressing extension, beyond the 11/45 and 11/70 extensions, which required a major perturbation; ~~in the form of the VAX/11 extensions.~~ VAX-11.

The final section, ^{addressed} to the research community, describes some general problems encountered in structure design and engineering, together with how solving them might be useful in subsequent designs. ~~(This part has been expanded too.)~~

VAX-11/780: A VIRTUAL ADDRESS EXTENSION TO THE DEC PDP-11 FAMILY
[Section on Strecker's paper to go here]
MULTIPROCESSOR

MULTIPROCESSOR
THE RESEARCH COMPUTERS OF CARNEGIE-MELLON UNIVERSITY

~~These Three multiprocessor~~ computers, which use the 11 as a basis, were built at Carnegie-Mellon University to carry out ^{research in} various computer structures, and operating systems ~~and application program computer engineering and computer-science research.~~ ^{for multiprocessors}

The first computer, C.mmp, is a 16 processor (~~model 11/40's~~ ^{and 11/20's}) system with ~~2.5~~ million words of shared primary memory. It was built to investigate the programming (and resulting performance) questions associated with having a large number of processors.

~~Chapter on the~~ ^{is located physically}
The second computer, Cm*, ~~is also discussed~~ in the Modules Part, as the
modules that form ~~Cm*~~ ^{Cm*} are LSI-11's computers. Cm* ~~evolution~~ was based on the
premise that ultimately, the smallest modular unit (~~i.e., ultimately all~~ ^{, i.e., integrated}
~~ICs~~) ^{would be a} that would be used to build digital systems ~~is~~ the computer. With
~~this premise,~~ ^{we need} ~~it is important~~ to understand how to interconnect computers
physically, how to assign parts of the ~~problem~~ ^{application program} to the various computers, and
how to program the complete structure.

The third computer,

C.vmp, ~~for voting multiprocessor,~~ was designed to investigate how a
production microcomputer, the LSI-11, ~~can~~ ^{could} be used to build a triplicating,
voting high availability computer.

The goals of the first two are performance while the third ~~uses multiple~~ ^{has}
~~computers for redundancy and reliability.~~ ^{as its goal} To this end, Fig. Perf shows the
~~effect of using multiprocessors to execute various algorithms (More to come~~
~~here!).~~

We believe that technology will force the evolution of computing structures
to ~~be along all three lines~~ ^{converge to styles} of multiprocessor computers: (1) C.mmp, ^{-style} for high
performance, incremental performance and availability; (2) ~~less tightly~~ ^{loosely} coupled
computers like Cm* ~~for more ad hoc structures~~ to handle specialized
processing, ^{e.g.,} front end, file, ^{and} signal processing; and C.vmp ⁽³⁾ ^{-style} for high
availability based on increased maintenance costs.

The technology-force argument is based on history, near term technology, and
resulting price extrapolations, ~~and follows:~~

1. ~~The~~ MOS technology is increasing in both speed and density faster than the technology, e.g., ECL, from which high-performance machines ~~is~~ *are built*.
~~formed;~~

2. The price per chip of the ^{single-MOS-}MOS ~~one~~ chip processors decreases at a substantially greater rate than ^{for the} low-volume ^{high-performance, ~~tasks~~} special designs. ~~due to the~~ *Chips*
in both designs have high design costs, but the special design enjoys a
~~much higher volumes and very high design costs (for both designs),~~
much lower volume.

3. ~~For all intents and purposes,~~ the processor cost ^{in a} of the low end ^{system} ~~is~~ *essentially zero.*

3. *R* ~~Relative to all other costs of the system,~~ ~~The resulting performance~~
~~(in operations) per chip and cost per operation per chip are rapidly~~
~~diverging from the high performance semiconductor technology from which~~
~~conventional high performance machines are formed.~~

4. Standards *3* in the semiconductor industry *3* ~~for high volume products~~ tend to form more quickly. For example, in the 8-bit microcomputer market, *one*
~~type~~ supplies about 50% of the market and *3* ~~types~~ supply over 90% of
~~the market.~~

5. A 16-bit processor *4*-on-a-chip *4*, with ^{both an address space matching its} ~~an acceptable~~ (for the ^{has been announced.}
performance) ~~address space~~ and appropriate data-types, ~~is eminent.~~ Such
a commodity will form the basis for nearly all future ^{computer} designs.

~~These factors produce better (in terms of price and price/performance)~~
~~computers at the low cost market end at a diverging rate compared with~~
~~large scale computers.~~ Furthermore, large scale applications have been

*As a result of these factors, the two classes of machines (MOS-processor-
on-a-chip ^{based} and low-volume, high-performance-processor-based) have
rapidly diverging costs per operation per chip.*

slow to form since problem complexity increases more rapidly than program size. Therefore, most subsequent computers will be based on standard, high volume parts. For high performance machines, since processing power is available at ^{essentially zero} cost from ~~a~~ processor-on-a-chip-based processors, large scale computing will come from arrays of processors, just as we build arrays of ⁶⁴ 16K bit ICs to form memory ^{subsystems}.

C.mmp--A Multi-Mini-Processor

C.mmp was motivated by the need for more computing power to solve speech recognition/signal processing problems and to understand the multiprocessor software problem. Until C.mmp, only one large, tightly coupled multiprocessor had been built--the Bell Laboratory's Safeguard Computer (BSTJ issue?).

The introductory section describes the economic and technical factors influencing multiprocessor feasibility and argues ~~that it's important to~~ ^{for} ~~embark on the research because of~~ ^{of the research} the 'timeliness'. Various problems to be researched are given together with a discussion of particular design aspects. For example, since C.mmp is predicated on a common operating system there ~~is~~ ^{are two sources of} degradation; ~~both due to memory contention when all the~~ ^{and lock contention.} processors use the same memory block and since there are common synchronizing locks for parts of the software, ~~degradation may occur when a~~ processor must idle and wait to enter common critical sections. The machine's theoretical performance ~~due to~~ ^{as a function of} memory-processor interference is ~~computed~~ based on Strecker's (1971) work. In practice, ~~memory interference~~

~~was a problem resulting in poorer than expected performance~~ because the *memory interference was greater than expected.*
memory was not built with low-order address interleaving, ^h This problem ~~had~~ *was*
~~to be solved by moving program segments of the memory.~~ It should be noted

AS ~~that~~ *theoretical* when the number of memory modules and processors becomes very large,
the performance (as measured by the number of accesses to the memory by the
processors) approaches the memory bandwidth (m/memory-cycle-time) x 1/e. [Ref?].
Thus, there is not a maximum limit on performance with infinite processors,
provided they are all *processors are for* not contending to access the same memory.

Although there is a discussion outlining the design direction of the
operating system, Hydra, later descriptions should be read (~~list the~~ *[Wulf et al],*
necessitated frequent map changes
1975] ~~references?? here~~). Since the 11's small address ~~impaired use,~~ 11/40s with
~~microprogrammed~~ writeable control stores were used to implement ^{the} operating
systems calls ~~involving changing~~ *strategy to change* the segment base registers.

One of the most pleasant surprises of C.mmp was the ALGOL 68 implementation
because it enables parallel programs to be specified. The result (see page
00) in terms of performance are quite encouraging! There are three basic
approaches to effectively applying multiprocessors:

1. having lots of independent work to do through multiprogrammed,
timeshared and multiple independent computers;
2. having the compiler for a conventional language (e.g., FORTRAN) detect
and compile statements that can be executed in parallel;

AS *The actual C.mmp which was actually built is shown in Figure PMSC.mmp.*

Insert pages 20 a, b, c

There are three basic approaches to the effective application of multiprocessors:

1. system-level workload decomposition.

If a workload contains a lot of inherently independent activities, e.g., compilation, editing, file processing, and numerical computation, it will naturally decompose.

2. program decomposition by a programmer.

Intimate knowledge of the application is applied—required for this time-consuming approach.

3. program decomposition by the compiler.

This is the ideal approach. However, results so date have been ~~dece~~ disappointing [Ref Illiac IV FORTRAN compiler ~~the~~ projects].

C.mmp was predicated on the first two approaches. Since ~~ALGOL 68~~ the original paper, ALGOL 68, a language with facilities for expressing parallelism in programs, has been implemented. It has assisted greatly with program decomposition. See Figure X.

As can be seen in the paper, a model of the lock problem influenced the number of critical sections in the scheduler. Since the paper, the operating system and ~~of~~ Hydra has been designed and implemented. ~~[Wulf et al, 1975]~~
An extensive description is given in [Wulf et al, 1975].

Two experiments ^{analyzing} on the performance of C.mmp and Hydra have been reported. The first, by Karathe [1977], ~~measured~~ used a hardware monitor to measure the degradation due to the locking mechanism which is invoked when shared data is accessed. The second, by Oleinick [1977], analyzed user-program-level synchronization. We summarize the results of both below.

The contention for shared resources in a multiprocessor system occurs at several levels. At the lowest level processors contend ~~for~~ at the cross-point switch level for memory. This memory interference is discussed in the ~~paper~~ ^{chapter}. On ~~at~~ a higher level, there is contention for shared data ^{in the operating system kernel}. At higher levels, processes contend for i/o devices and for software processes, e.g.,

for memory management. The following table points to ~~the~~ models or experimental data at these different levels in c.mmp.

Contention Level

Reference

User-program ~~system~~

[Oleinick ~~§~~, 1977]
[Fuller and Oleinick, 1976]

~~Hydra~~ - kernel

Hydra Kernel objects

[Marathe and Fuller, 1977]

cross-point switch

Chapter XX
[Fuller, 1976]

Marathe's data show that the shared data of Hydra is organized into enough separate objects that a very ~~small~~ small degradation (less than 1%) ~~is~~ results from contention for these objects. Table ~~§~~^{LOCK} is reproduced from his paper. He also built a queuing model which projected that the contention level would be about 5% in a 48 processor system.

Oleinick uses a root finding algorithm to study various implementations of a single problem.

[CH: elaborate after studying his thesis chapter]

3. introducing new primitives into the programming language (e.g., ALGOL 68) so that algorithms for parallel execution are specified by the programmer.

Although C.mmp was only predicated on type 1 use, the ALGOL 68 implementation on C.mmp permits parallel programs (type 3) to be written. The results of some of this programming is shown in Fig. ?.

Multi-Microprocessors: An Overview and Working Example

~~Cm* is a system of high level modules, constructed of LSI-11 computers, which can be interconnected to form large computer structures.~~ The Cm* work, sponsored by NSF and ARPA, ^{is} an extension of ^{earlier} NSF-sponsored research [Bell, etc. 1973] on register-transfer ^{-level} modules. As LSI and VLSI enable construction of the processor-on-a-chip, it is apparent that low-level, register-transfer modules are passe' for the construction of all but low-volume computers. ~~A complete industry will ultimately design, contribute to and use a standard computer.~~ Although the research is predicated on structures employing a hundred or so processors, this chapter describes the culmination of the first (ten-processor) phase.

The authors motivate their work by appealing to the diseconomy-of-scale argument which we advanced at the beginning of this section (page 00). To provide additional context for their research,
~~A motivation of Cm* is based on a diseconomy of scale for large computer introductions during 1975-1977 (see page 00).~~ Computer modules (Cm*), multiprocessors (C.mmp) and computer networks are described in terms of performance and problem suitability. ~~to provide additional context for the research.~~ The chapter gives a description of the modules structure,

together with their associated ^{limitations} and potential ^{research} problems. (~~research~~).

The final, most important, part of the chapter evaluates the performance of Cm* for five different problems.

C.vmp: The Architecture and Implementation of a Fault Tolerant

Multiprocessor

C.vmp is a triplicated, voting multiprocessor designed to understand the difficulty (or ease) of using standard, off-the-shelf LSI-11s to provide greatly increased reliability. There is concern for increased reliability because systems are ^{becoming} more complex, are used for more critical applications, and ^{because} basic maintenance costs for all systems are increasing. Because the designers ^{themselves} carry out and analyze the work, this chapter provides ~~a great~~ ^{first-} ~~deal of~~ ^{hand} insight into high reliability designs and ^{the} design process--especially its evaluation. The system has operated for several months and the first phase of work is complete.

Several design goals are initially predicated and the work is carried out against the goals. ~~Two of the more interesting goals include using off-the-shelf hardware and software with no modifications to the components.~~

The goal of software and hardware transparency turned out to be easier ^{to attain} than expected because of an idiosyncrasy of the floppy ^{disk} controller. ~~This~~ ^{Because the} controller ^{effects} ~~operates~~ on a word-at-a-time ^{bus} transfer from a one-sector buffer,

can be
~~after a sector is transferred--thus~~ voting ~~is~~ carried out at a very low level ~~(i.e., as bus transfers are made)~~. It is unclear how the system would have been designed without this type of controller, ~~but~~ as a minimum, some form of software transparency goal would have been violated ~~together and~~ *with* a significant controller modification *would have been necessary.*

A number of models are given by which the design is evaluated. Various component reliabilities are used and the reader should get a great deal of insight into the factors contributing to reliability. It should be noted that a special hardware voter is needed *to get a sufficiently low cost for* ~~in order to build~~ a marketable C.vmp. While the intent of C.vmp is not a product, it does provide much of the insight for such a product.

Karathe, M. and Fuller, S.H. A Study of
Multiprocessor Contention for shared Data
in C.mmp. Proc 1977 ACM SIGMETRICS
conferences, pp 255 - 262.

Wulf, W. [and others] The Hydra Operating
system, Fifth ACM SIGOPS Symposium on
Operating Systems Principles (Nov 1975).

11 glue figures & tables

Fig. Cachespeed Structure of P_c , Mcache and M_p of Cashed computer

Fig. LSI-11/2 Photograph of double height modules forming LSI-11/2

Fig RXT11 For Bounded LSI-11

Fig PMSC.mmp PMS diagram of C.mmp

Table Lock Measurement of the locking behaviour in Hydra / C.mmp.

It is somewhat anticlimactic to discuss this original PDP-11 description here because Chapter 00 explicitly discusses "What We Have Learned from the PDP-11". The purpose of the chapter was originally threefold: to give the PMS and ISP architecture of the PDP-11 as it was first proposed, to describe the first (11/20) implementation, ^{family} at a time when the whole architecture had not been worked out or even fully considered, and to show possible extensions. *This was attempted*

~~The reader might note that~~ The computer class definitions (given in 1970) ^{have stood the test of time for they} of micro, mini and midi correspond quite closely to those of Chapter 1.

The major reasons ^(elaborated upon in Chapter 00) for ~~Although we comment on~~ the disparity between the predicted and actual evolution ~~in Chapter 00, some of the important reasons are:~~

1. The notion of designing with improved technology, especially for a family, was not understood ^{in 1970} ~~then~~. This understanding ^{came later and} was put forth in a paper in 1972 (Bell, Chen, Rege).

2. The Unibus ~~bandwidth~~ ^{most} proved unacceptable for ~~all~~ communications at the very high and low end designs. ^{Although} ~~Whereas~~ this chapter posits a multiprocessor and multiple Unibusses, ^{for high end designs, this exact} ~~this precise~~ structure did not evolve as a standard. ^{Levy's chapter elaborates on the bus evolution.} The bandwidth has subsequently been shown to be

adequate for all but the largest configurations when a cache is attached to the processor as in the 11/60 (see Chapter 00). Note the effect of a 90% cache hit rate is to reduce the number of access to primary memory via the Unibus by a factor of ten!

*delete,
but make
sure points
is made
in Levy*

*B
S*