



VCU

Virginia Commonwealth University
VCU Scholars Compass

Theses and Dissertations

Graduate School

1983

A comparison of processor technologies

Eddie R. Wachter

Follow this and additional works at: <https://scholarscompass.vcu.edu/etd>



Part of the [Computer Sciences Commons](#), and the [Mathematics Commons](#)

© The Author

Downloaded from

<https://scholarscompass.vcu.edu/etd/5617>

This Thesis is brought to you for free and open access by the Graduate School at VCU Scholars Compass. It has been accepted for inclusion in Theses and Dissertations by an authorized administrator of VCU Scholars Compass. For more information, please contact libcompass@vcu.edu.

College of Humanities and Sciences

Virginia Commonwealth University

This is to certify that the thesis prepared by Eddie R. Wachter has been approved by his committee as satisfactory completion of the thesis requirement for the degree of Master of Science.




Dr. James E. Ames, IV
Director of Thesis




Dr. Richard Allan
Committee Member




Dr. Francis R. Kane
Committee Member



Dr. Donald Hartford
Committee Member



Dr. William E. Haver
Department Chairman



Dr. Elkse v. P. Smith
Dean, College of Humanities
and Sciences

May 12, 1983

Date

A Comparison of Processor Technologies

A thesis submitted in partial fulfillment of the requirements for the degree of Master of Science at Virginia Commonwealth University.

by

Eddie R. Wachter

Director: Dr. James E. Ames, IV

Affiliate Assistant Professor

Department of Mathematical Sciences

Virginia Commonwealth University

Richmond, Virginia

May, 1983

Acknowledgements

I wish to thank Dr. Ames for his support and assistance in preparing the study, Dr. Richard Allan for his experience and insights which guided me in the right directions, and above all, my family for understanding the time involved in pursuing a degree at night over the last four years.

Table of Contents

List of Figures	iv
Abstract	v
Introduction	1
Processor Technology	4
Central Processor Organization	17
Main Storage Systems	40
Input/Output Systems	48
Reliability, Availability, Serviceability	57
Applications	63
Futures	69

List of Figures

Figure 1 - Switching Speed vs. Power Consumption . . .	7
Figure 2 - Amdahl 580 Bus Organization	23
Figure 3 - CDC CPU Block Diagram	29
Figure 4 - IBM 308X Processor Units	32
Figure 5 - IBM Central Processor Flow Diagrams . . .	34
Figure 6 - Univac 1100/90 System Configuration . . .	38

Abstract

The purpose of this paper is to present a discussion of the technology implementation and design of four very high performance mainframe computer systems. The systems evaluated are:

Amdahl 580 Series

CDC 170 Series 800

IBM 308x Series

Univac 1100/90 Series

Included in this evaluation is a survey of the technology used, its characteristics, packaging and performance. Each system component is evaluated on the basis of design philosophy, technology, and the total system design with regards to reliability, availability, and performance.

Introduction

Basic to all computer systems is a pre-defined and established architecture and the physical implementation of this architecture. Often times these two areas are confused when a study is undertaken and therefore some background is presented here to distinguish between the two structures.

The architecture of a machine is not a blueprint for the design of a computer system, but a description of the logical appearance, the conceptual structure and functional behavior of the processor, as viewed by the programmer. This information about the architecture of the machine defines how the hardware will respond to the software. When programs running on different machine implementations produce identical results as defined by a single architecture, the implementations are considered compatible. Most major computer manufacturers, such as those studied here, have developed families of compatible computer systems in which the architecture of the machines within the family are the same, yet the implementation of that architecture may vary from one

machine model to another.

The reasons that manufacturers develop different implementations are manifold. One is to take advantage of technological changes that have been developed in the laboratory and are now commercially available. Another is to functionally enhance a product with specialized hardware, microcode and features which add to the versatility of the machine. But most of all, implementations are changed to provide varying levels of processing speed, also known as computer performance, within a product line. These different implementations provide different operating speeds, yet are functionally compatible.

This is especially true where plug compatible mainframes (PCM) compete, such as IBM and Amdahl. Both manufacturers conform to the IBM System/370 Architecture, yet the physical implementation by each vendor of that architecture is completely different. As long as all machines within a family of systems or all PCM machines meet the specific architectural requirements, the system evaluation criteria changes to other components, such as price, flexibility, reliability, and performance.

In order to understand these particular areas just mentioned, a closer look at the physical implementation is required. How the manufacturer has designed the functional components of the processor from the type of circuitry used, design of the functional parts, and how they all communicate to achieve the architectural standard has a direct bearing on performance, reliability, and flexibility.

It is this aspect of processor design that is presented here.

Processor Technology

The circuitry design and technology used in mainframe computers are the primary determinants of the processing power of the computer. The resulting performance of the Central Processing Unit (CPU) is determined by the cycle time and the number of instructions that can be performed in a given cycle. The number of instructions in a given cycle depends on the particular architecture of the machine, the number of logic levels between driving and receiving registers and the degree of parallelism in the design. Since the cycle time of the machine is very dependent on the technology chosen for the hardware implementation, namely the circuits, packaging and interconnections, substantial emphasis has been placed on research in this area. Improvements in processor throughput can then be made simply by changing to a faster logic circuit without changing the basic architecture of the computer.

Using the traditional engineering 1/3 rule for delay estimates, the cycle time of a processor is roughly

allocated as follows:

- 1/3 circuit switching
- 1/3 loading and unloading of power
- 1/3 interconnection transmission.

For example, if the cycle time is to be 24 ns (nano seconds, or billionths of a second), then each function should be completed in 8 ns. If they can be completed in less than 8 ns, the cycle time can be reduced accordingly.

The circuit switching time is the total circuit delay, not that of an individual circuit. Since there are typically 6 to 8 levels of logic per circuit, the net delay per circuit should be less than 1 ns each to stay within the total allocation of 8 ns.

In order to understand the performance associated with the CPU's studied, it is necessary to understand the types of circuit technology used by the manufacturers in each component of their CPU. Each technology has its own characteristics and applicability to different processes, and the industry has standardized on a few of them.

Those currently in manufacture, regarded as "state of the art", are;

- . TTL - Transistor-Transistor Logic
- . ECL - Emitter Coupled Logic
- . NMOS - Negative Channel Metal Oxide
Semiconductor
- . GaAs - Gallium Arsenide
- . Josephson Junctions

Figure 1 diagrams the switching speeds and power consumptions of each of these technologies.

Transistor-Transistor Logic, TTL, developed in the mid 1960's, has provided both the standard for interfaces from computers to peripheral equipment, and mainframe logic circuits. Since it is a more mature technology, TTL manufacturing costs are low and yields are high. Moderate power is required to power the TTL gate, and switching time is slow, on the order of 1/3 to 1/4 as fast as ECL. TTL circuits require special termination of each gate to eliminate transmission line like effects, and are generally mass produced in common form, using part of the circuits on each chip.

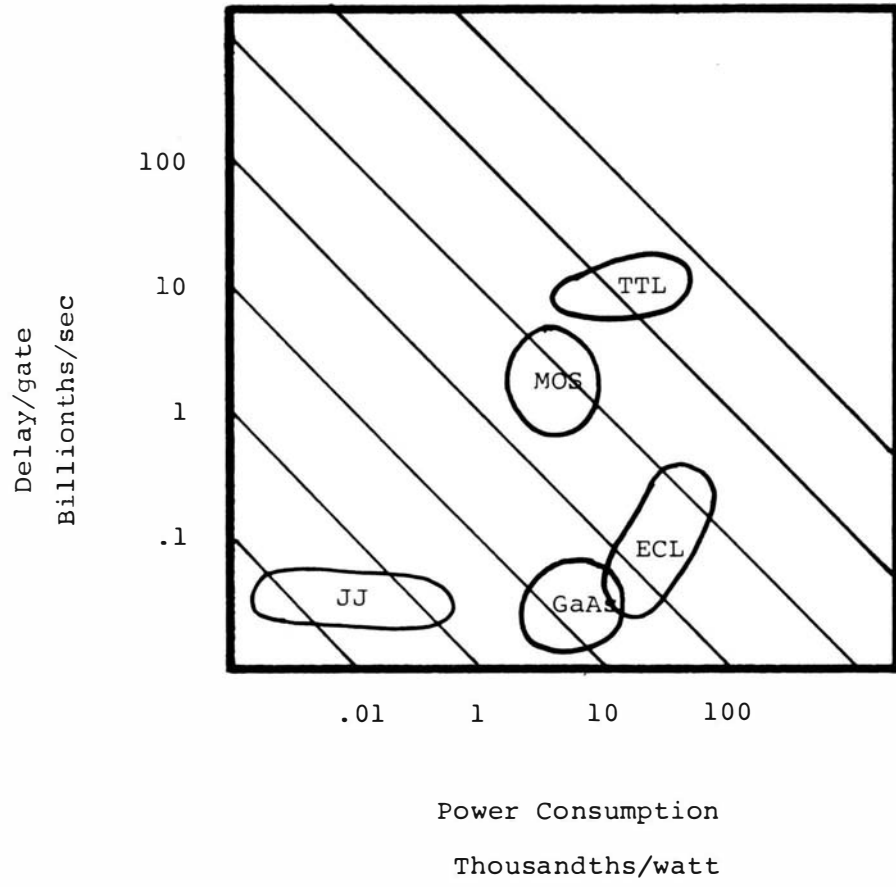


Figure 1

ECL, Emitter Coupled Logic, is the fastest commercially available technology and has been around for approximately 10 years. It operates 3 - 4 times faster than TTL, although its power consumption is higher, on the order of twice as much. Another attribute of ECL is its capability to drive more circuits (place signals to many other circuits) thus requiring less circuitry to perform a logic function. This leads to higher reliability due to fewer circuits being used. Considerations in using this type of circuitry is that ECL is very critical of circuit length, which is compensated with higher power requirements. ECL is used for main logic units and high speed memory.

Active element or volatile memories are the latest technology for fabricating computer memory. Early memory devices such as magnetic cores have been replaced with large scale integration (LSI) memories which are much faster, more compact, and require less power.

NMOS, Negative Channel Metal Oxide Semiconductor, is the dominant memory technology. It offers moderate speed in switching, very high density, low power consumption, and a reasonable manufacturing cost.

Using NMOS technology, 256k (256,000) chips are possible, whereas current processors are now using 16k chips.

The primary reason that 64k and above chips are not currently used is due to manufacturing costs associated with the fallout percentage of denser chips. Once this problem is solved, 64k to 256k chips will be standard.

GaAs, Gallium Arsenide (compound semiconductor), a new technology, offers three to four times the device speed potential of silicon-based circuits. This technology has a number of shortcomings to be dealt with, such as manufacturing difficulty and high power consumption, but offers the advantages of high speed, a high degree of fanout, and increased speed at 0° F. GaAs offers an alternative to Josephson Junctions with a reasonable environment, but as with any temperature sensitive device, control of the environment becomes the determining factor if a technology is to be used or not.

Josephson Junctions are transistors that have a very fast switching speed due to their superconductivity, a result of its operation near absolute zero temperature. Their development has been limited to laboratory experimentation due to the temperature requirements

and therefore are not currently applicable to commercial use.

Once the basic technology is established by a manufacturer, the chip and its associated packaging levels are designed according to the architectural features of the machine. The number of circuits per chip and the interconnections of the chips weigh heavily on the processing speed of the CPU. If pulses must travel great distances (greater than 10 cm) to the next logic circuit, the speed of the machine is reduced accordingly. It is therefore a prime design criteria to place as many circuits on a chip that are related to a single operation, and to place chips with related functions in proximity to one another to reduce the interconnection distance.

The following is a list of implied "rules" each manufacturer follows in designing the physical structure of their processor in order to achieve speed, reliability, and performance.

1. Use high speed switching circuitry as the basis for logic circuits.
2. Maximize the number of circuits per chip.
3. Reduce the number of interconnections and the distance between them.
4. Place needed microcode, where used, close to circuitry that requires it.
5. Cool the technology to provide the necessary speed and reliability of the circuitry.

How each manufacturer has implemented these "rules" depends on the technology they are familiar with and their design objectives for the processor, such as speed, flexibility, redundancy, cooling technology, and price. Let us now examine each manufacturer against the components (chips and technology) described above.

Amdahl Corporation's 580 series of processors uses ECL circuit technology, placing 400 circuits on a chip with a 95% use factor. Because of ECL circuitry, gate delays are in the 400 picosecond (trillionth of a second), power required per chip is 3.0 watts, and the fanout or ability to drive other circuits is 2 to 5 circuits. Due to the high power requirements, uniform cooling of the chip is necessary and this is accomplished by the use of

a cooling fin bonded directly to the chip. Air from the computer room is drawn over the chip across the cooling fins and out of the mainframe to be cooled by the computer room air conditioning. Amdahl processors therefore require no coolant refrigeration units or chilled water units to dissipate heat generated by the processor.

Each chip is soldered into an appropriate multiple chip carrier, (MCC), of which 8 are required to make up the logic circuitry of the processor. Each MCC contains 121 chips made up of logic and Random Access Memory (RAM) chips to hold the microcode. The RAM chips are placed next to logic chips requiring the microcode to reduce the interconnect time of the signal. Each MCC is a 14 layer printed circuit board and the 8 MCC's are housed in a MCC stack with side panels for interconnections, which are also multi-layered printed circuit boards, and the whole MCC stack occupies 5.6 cubic feet. The 8 MCC's are designated as follows:

- 5 - CPU
 - 1 - Instruction Unit
 - 1 - Execution Unit
 - 1 - Storage Unit
 - 2 - High Speed Buffer

- 1 - Input/Output Processor (2nd optional)
- 1 - Console Processor
- 1 - Memory Bus Controller

The functional replaceable unit therefore consists of the 121 chip MCC, each of which performs a particular function of the CPU.

Control Data Corporation's Cyber 170 Model 875 processor utilizes ECL circuitry on plug-in circuit boards that are not generally considered large scale integration, LSI. The boards are mounted in a logic chassis that is a functional unit, i.e. CPU, I/O control, and memory. Diagnostics are done to individual chips, and replaced at that level. Utilizing ECL technology, CDC has advantages of high fanout, low gate delays, but a somewhat high power requirement. Cooling is accomplished by a closed loop chilled water system.

The CPU consists of 3 main bays or frames, each housing a particular component; the central processor, the input/output unit, and the central memory. The central processor consists of 9 independent functional units and a controlling central processor.

IBM Corporation's 308x series of processors uses TTL technology for its primary CPU circuitry. The TTL chip has a theoretical 704 gates per chip with an effective utilization of approximately 60% or 400 circuits. Switching time per gate is 1200 picoseconds with a low power requirement of .4 to 2.7 watts per chip. Fanout ability is 1 to 3 circuits. Cooling of the chip is accomplished by use of a heat conduction mechanism of a cylinder touching a chip, transferring the heat of the chip to a helium chamber, through a lexan interposer to a water jacket where heat is carried out of the processor by a series of water hoses. The unit, called a Thermal Conduction Module, TCM, contains 133 chips per unit and is the field replaceable unit of the processor. Chips can either contain logic circuitry or hold microcode for use by neighboring chips. The central processor consists of, depending on CPU model, 19 to 54 TCM's. The basic

uniprocessor model 3083 is configured with TCM's as follows:

- 8 - CPU
 - 1 - Execution Element
 - 1½ - Instruction Element
 - 3 - High Speed Buffer
 - 2 - Control Store Element (microcode control)
 - ½ - Variable length instruction execution

- 6 - External Data Controller - channel controller

- 5 - System Controller

Each major function, CPU, EXDC, and SC, are housed in their own multilayered printed circuit board. Functions are placed in proximity to other frequently referenced functions, and microcode is distributed to memory chips on each TCM. Communication between CPU, EXDC, and SC, is accomplished via cable connections.

The Sperry Univac 1100/90 series of computers uses high speed ECL technology incorporating gate array implementation, with 168 gates per chip. Each chip dissipates approximately 5 watts of heat. Univac calls its new implementation High Performance Packaging which is 10 times as dense as previous models. A custom rectangular ceramic chip carrier is used with 54 connecting leads from the chip to the circuit board. The field replaceable unit is this chip carrier. The circuit boards are placed in pairs in a stack arrangement with a liquid cooling plate between each board of the pair. Board to Board interconnection is accomplished via two side panels and a backplane. Boards are mounted in 3 functional frames: CPU, I/O processing, and memory. Cooling of the units is accomplished by chilled water systems and a radiator effect of drawing air over a radiator type dissipator and between the cards.

Central Processor Organization

Once the basic circuit technology is established and the packaging has been designed, the next most critical factor in processor speed is the component construction of the central processor. It is the primary function of the central processor to take an instruction, decode it into its component parts, perform the requested operation, and either store the results or indicate the results by use of setting some type of flag or condition code. The components of the central processor, how they perform their tasks, and their interrelation all contribute to overall system performance. Instruction timing, micro-coding of operations and instruction control are the components of the central processor and the functions that must be managed.

The cycle time of a particular machine does not necessarily directly translate to processor speed. Moreover, it is the number of cycles that it takes to execute an instruction that determines the MIPs, or millions of instructions per second. Theoretically, if

all machines had similar cycle times and performed one instruction per cycle, then these machines would perform equally. Since technology has its limiting factors in number of circuits and switching times, other design concepts come into play. One of these is pipelining instructions where 2 or more instructions are in various phases of execution at the same time. Each instruction to be executed must be verified for a valid operation, have its operands fetched, set a condition code, be executed, and have its results stored. If each function can be performed on a different instruction in a different phase, theoretically this results in one instruction executed every machine cycle, provided a good algorithm for branch determination is used.

Another way that instruction execution time can be improved is to utilize multiple processors to perform instruction execution in parallel. This parallelism can be accomplished either by universal processors, one that executes all instructions, or specialized processors that execute particular types of instructions such as boolean, floating point, decimal, etc. The degree of overlap or parallelism controlled by the instruction processor has a direct bearing on the throughput of the processor.

Should these concepts be used together, such as pipelining and parallelism, quantum factors in throughput can be achieved.

In order to make the machine as flexible as possible, where new features can be added without significant changes to the processor, microcode is used to control the flow of logic. While somewhat slower than hardware, microcode is much more flexible. Not only is it used for logic control, but diagnostic functions utilize microcode to facilitate scanning of circuits to determine their status so that complete system monitoring can be accomplished.

There are two types of microcode used in high speed processors, horizontal and vertical. Horizontal microcode is a wide word or bit string which generally controls one machine cycle. Each bit within the word controls a data path of the operation, a 1 signaling to take the path, a 0 not to. Horizontal microcode is used where high speed and flexibility are needed, since the testing of a bit determines the data path. Vertical microcode, on the other hand, is characterized by being

a set of instructions, much like a mini program, that are tailored to the specific task. Vertical microcode is easily written and modified, and is generally not hardware dependent. As a result, simulation of routines is a prime candidate for vertical microcode.

Controlling the flow of instructions, data, and results between main memory, instruction units, and other components of the processor is the purpose of a system controller. Where there is no overlap and everything is done sequentially, (each task must be completed before the next begins), no system controller is needed. But in order to take advantage of overlap, pipelining and parallelism, control of the system must be undertaken. A prime implementation of using a controller is the use of a bus system architecture, where the controller provides paths and message traffic control between the major components, insuring that as many components of the processor are busy as possible. The efficient use of all system resources along with minimal serialization and efficient system control enhances the throughput of the machine.

Each manufacturer has designed his CPU to exploit the features of their product, based on the target market for their machine and the limits of the technology. Also key to the internal design of some machines is providing the ability to have different MIP rates of the same machine by varying internal communications. This provides the manufacturer to build one basic machine and upgrade a processor by enabling intra-system communication and overlap. This concept is most prevalent in the commercial processor environment.

Amdahl

The Amdahl 580 computer system is designed as a uniprocessor rated at 13 MIPs. This rated speed is achieved by utilizing a cycle time of 24 ns and a pipeline architecture whereby several instructions are in various stages of execution at the same time. The CPU is composed of 8 multiple chip carriers (MCC's), each of which has a particular function within the processor complex. They are housed within a stack 5.6 cubic feet in size and are mounted horizontally in the stack and connected by two uni-directional communication buses.

The buses are distributed by 2 multi-layered printed circuit boards that form the sides of the stack.

(Figure 2). Five of the stack implemented functional units compose the CPU:

- . Instruction Unit (I-Unit): Fetches, decodes, and controls instruction execution
- . Execution Unit (E-Unit): Provides computational facilities of the CPU
- . Storage Unit (S-Unit): Controls instruction operand storage and retrieval facilities
- . Instruction Buffer (I-Buffer): High speed buffer storage for instruction streams
- . Operand Buffer (O-Buffer): High speed buffer storage for operand data

Within the Amdahl 580, two sets of functions are performed simultaneously: instruction fetch, which provides a doubleword of instruction stream every cycle holding it in the I-Unit for execution, and instruction execution. Instruction fetch for each cycle looks at the instruction needed to fill the instruction buffer, be it a branch address or next logical instruction, and has it fetched and prepared for holding in the instruction buffer. Extensive buffering of target instruction streams allows for early decision making on branch

580 BUS ORGANIZATION

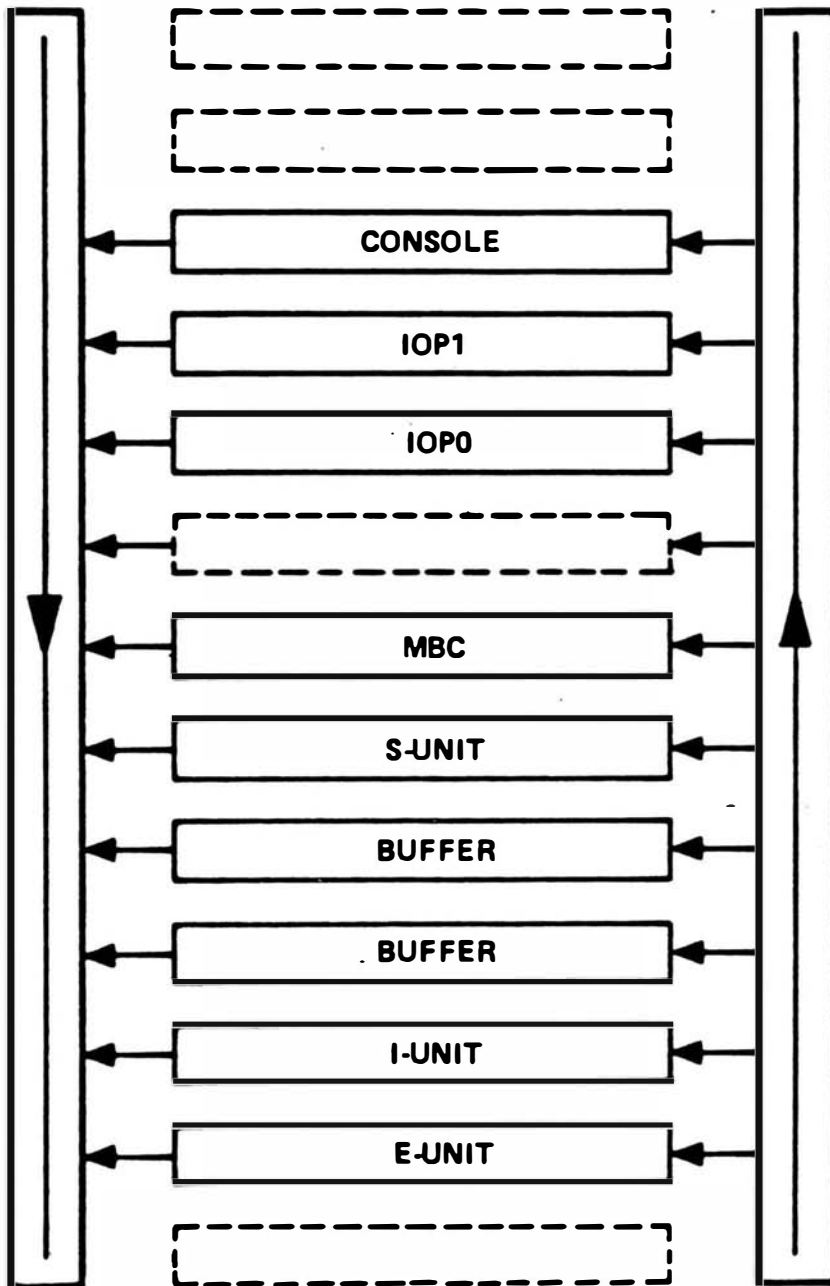


Figure 2

instructions and minimal loss of wasted instruction fetch.

Instruction execution is accomplished by presenting a new instruction to the pipeline (execution process).

The pipeline consists of the following 5 phases:

- . Generate - decode the instruction, validate the opcode, generate operand addresses, and send the opcode to the E-Unit
- . Buffer - access the operand buffer for operands
- . Luck - logical functions, comparisons, etc. condition codes are set
- . Execution - calculations are done
- . Write - results are stored in registers or O-Buffer

Extensive use of microcode is used in the Instruction Unit and Execution Unit for instruction control.

Horizontal microcode is used for fast access to control logic flow. The microcode is loaded on 4K RAM chips on the MCC associated with the function, and the RAM chips are located close to the logic chips that use it. Access time of the RAM chips is 7ns, within the 1/3 of the cycle time of the machine, 24ns.

The Storage Unit along with the High Speed Buffer provide storage for instructions and operands. The Storage Unit consists of one MCC, while the High Speed Buffer consists of two MCC's, one for instructions and one for operands. Both are 32K bytes in size and have an access time of one cycle, 24ns. The Storage Unit receives and processes all data traffic between the CPU data buffers and main memory, and all instruction unit data requests for instruction preparation, taking advantage of the bus architecture.

With each MCC being a functional unit and the side panels acting as buses, control of all processor functions is critical. Another MCC, the Memory Bus Controller, provides communication paths and message traffic control between the major functional parts. Each MCC can direct requests to another MCC via the bus system, and the MBC determines whether or not it needs to be involved to control access. A pass-thru circuit on the MBC allows for a direct connection between functional MCC's.

Amdahl has built in concepts for providing additional performance processors using the same design. A slower

speed model is achieved by changing the instruction pipeline or increasing the cycles per instruction. A larger throughput processor is achieved by attaching the buses of two processors together forming a multiprocessor configuration.

Current announcements by Amdahl are a range of 580 processors with a MIP range of 10 to 24 MIPs.

- . 5850 - uniprocessor - 64K HSB - 10 MIPs
- . 5860 - uniprocessor - 64K HSB - 13 MIPs
- . 5870 - attached processor - 128K HSB - 22 MIPs
- . 5880 - multiprocessor - 128K HSB - 24 MIPs

CDC

The CDC Cyber 170 875 is a uniprocessor design using a central processing unit, nine independent functional units, a storage move unit, and a central memory control. It is classified as a uniprocessor only for the reason that a second processor with nine additional functional units can be added to form an attached processor configuration.

Each of the nine functional units is a specialized arithmetic unit with an algorithm for performing a portion of the central processor instructions. The processor is designed with a 25ns cycle time, rated at 19 MIPS.

The central processing unit consists of operating registers and control logic to prefetch instructions and pass them to the appropriate functional unit, and store the results to central memory control. Each unit is independent of the others and multiple units can be in operation at any one time providing overlapped instruction execution. The functional units are:

- . Boolean unit: logical operations
- . Shift unit: shift left, right, pack, unpack
- . Normalize unit: floating point normalization
- . Floating add unit: add, subtract floating point
- . Long Add unit: add, subtract extended float pt
- . Multiply unit: multiply floating point
- . Divide unit: divide floating point
- . Population count unit: count 1 bits in operand
- . Increment unit: address generation

The Storage Move Unit performs all block copy transfers to and from extended memory. The Central Memory Control unit controls the flow of data between central memory and the system components. (Figure 3).

The central processor uses a 12 word instruction stack that performs a function similar to high speed or cache memory. The instruction stack receives instructions from main memory and can hold up to 48 instructions. Executed instructions are not discarded but are retained, providing a facility to loop within the stack or branch to a stacked instruction.

Instruction execution is accomplished by presenting operand data via registers to the appropriate functional unit by the central processor. The operation is performed within the number of clock cycles required and the results are placed in output registers. The central processor then reads these registers and places the results in memory. Quasi-pipelining is used on most functional units since data passes through a set of registers during each cycle within the functional unit. Therefore a new set of operands can be started into the

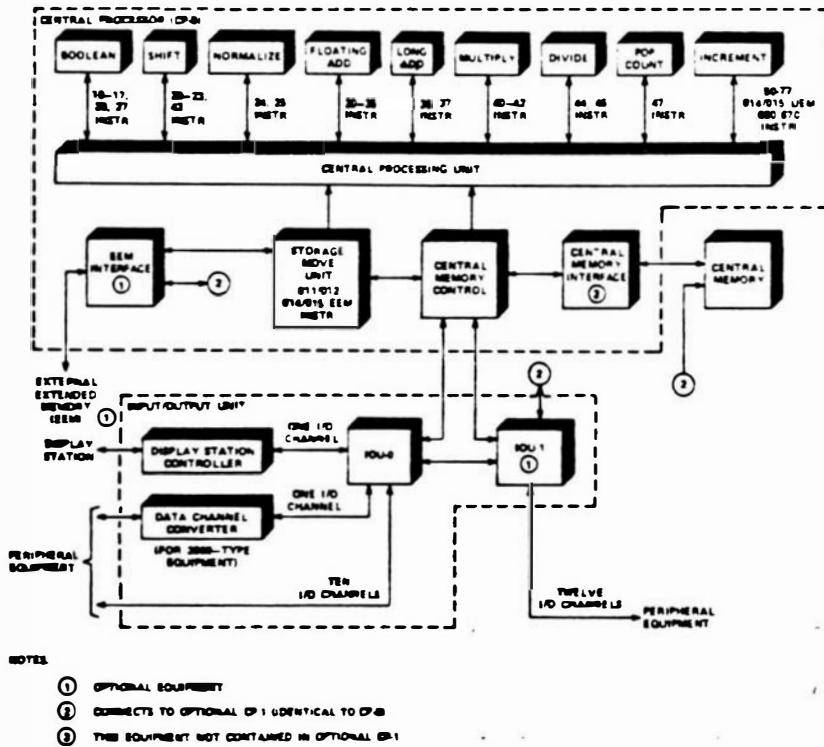


Figure 3

functional unit each clock cycle, even though it takes multiple clock cycles to complete an operation. This approach provides a theoretical maximum of one instruction executed per functional unit per cycle.

All central processor components are hard wired with a very high use of specialized instructions that would otherwise be delegated to microcode. Extensive use of systems registers provide for rapid task switching and concurrent execution of programs. This complement of registers also greatly reduces the number of main storage references required.

Additional processing capabilities rated at 32 MIPS can be achieved by the addition of a second central processor forming an attached processor configuration.

IBM


The IBM 308x series of processors was originally designed as a dyadic processor, the central processing unit consisting of two identical central processors running under the control of a system controller.

Variations of this configuration have been developed and will be discussed later. The 3083 uniprocessor, model J, will be used as the base machine for descriptive purposes. The 3083 is a 26ns cycle machine rated at 7.5 MIPs. This MIP value is lower than Amdahl or CDC, even though the cycle time is relatively the same, the reason being that IBM does not use bus or pipeline architecture.

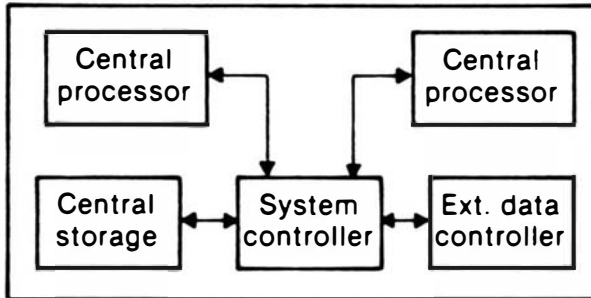
The basic 308x Processor Unit consists of four components: central storage; system controller (SC); external data controller (EXDC); and one or more central processors. (Figure 4). Each component, except central storage, is composed of multiple thermal conduction modules, (TCMs), sets of which are designed for a specific purpose. The TCMs are mounted vertically to large backplanes that are multi-layered circuit boards. These boards are connected together to form the processor unit.

The system controller provides all communication facilities between components. All storage requests, high speed cache requests, I/O requests and dispatching

3083



3081



3084

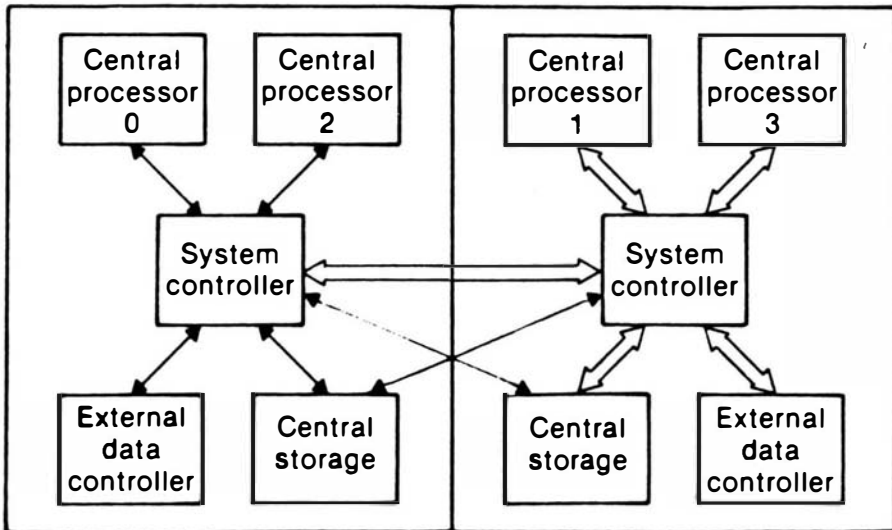


Figure 4

of work takes place within the system controller. Each component acts as an individual processor to the SC and is handled accordingly. Some models of the 308x perform overlap of instruction fetch and execution under the control of the SC.

The central processor consists of five functional elements each of which is located on one or more TCMs. (Figure 5). The components and their function are:

- . Instruction Element - controls instruction sequencing, initiates instruction requests, decodes instructions, generates operand address and executes most arithmetic and logical operations. Controlled by horizontal microcode
- . Variable Field Element - executes all variable field length storage to storage instructions. Controlled by horizontal microcode.
- . Execution Element - executes fixed point multiply and divide instructions and all floating point operations. Unit is hardwired, no microcode used.
- . Buffer Control Element - 32K cache memory using a store-in algorithm.
- . Control Store Element - controls sequencing of microcode throughout the central processor.

Within each TCM are both logic and microcode chips. Time to access microcode on the same TCM is 16ns.

Central Processor (CP)

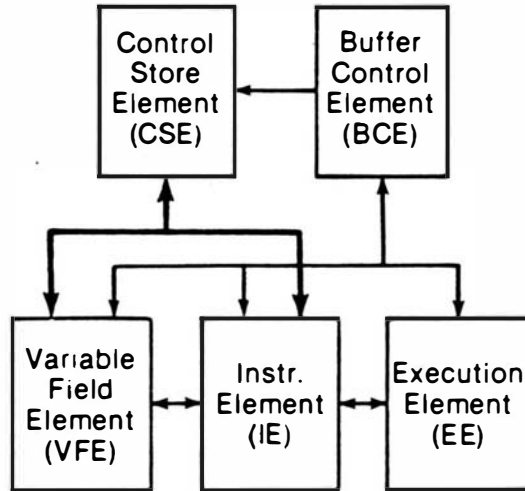


Figure 5

Each central processor has its own 32K high speed buffer, and management of the buffer is controlled by the system controller. Non-store-through (store-in) algorithms are used to provide a higher internal performance level than store-through when considered in conjunction with central storage access time. Access time of cache is 52ns or 2 cycles. Management of references to the same memory locations by a multi-central-processor system is controlled by the system controller. Due to store-in philosophy, if this happens, data can be transferred from one high speed buffer to the other by and through the system controller, and processing is resumed. The store-in algorithm provides and assists for a feature called hardware checkpoint retry. The basic concept is to establish a checkpoint at some particular instruction N, store the content of all registers in a backup set of registers and store changed cache values in a push down stack. Should an error occur, registers and cache will be returned to pre-error status to provide for integrity. The failing instruction is retried a specific number of times, each time the registers and cache are reset, and then before logging the error, the registers are reset once more.

IBM has currently announced 6 models of 308x processors:

- . 3083-E - uniprocessor 32K HSB - 3.5 MIPS
- . 3083-B - uniprocessor 32K HSB - 5.5 MIPS
- . 3083-J - uniprocessor 32K HSB - 7.5 MIPS

(The variances in speed of the above processors is attributed to the use or non-use of instruction pre-fetch and the number of translate lookaside buffer entries.)

- . 3081-G - dyadic processor 64K HSB - 10 MIPS
(2-3083-B central processors)
- . 3081-K - dyadic processor 64K HSB - 13 MIPS
(2-3083-J central processors)
- . 3084-Q - quadratic processor 128K HSB - 24 MIPS
(2-3081-K processors)

Sperry Univac

The Sperry Univac 1100/90 series of computers is a multi-processor providing a range of performance from 7 to 25 MIPS. One to four central processors, one to four main storage units, and one to four I/O processors

in any configuration can be used to design a system.
(Figure 6).

Each CPU consists of five separate components, each performing a specific function:

- . Instruction pipelining - a three level pipeline that provides for overlapped execution of three instructions. A "wraparound" feature is provided so that intermediate results from one instruction can be used in the next instruction in the pipe.
- . 32K high speed buffer for instructions, with an access time of 30ns per instruction.
- . 32K high speed buffer for operands, with an access time of 30ns per operand.
- . Arithmetic unit divided into 3 distinct, special purpose components:
 - . binary arithmetic component
 - . high speed multiply component
 - . decimal arithmetic component

(Each component is optimized to reduce execution time for its specific task)
- . Duplicate X file to accelerate operand and instruction address formation. This file provides two copies of the contents of each index register for internal manipulation.

Each component is hardwired, using no microcode, and achieves a 7.5 MIP uniprocessor with a 30ns cycle time.

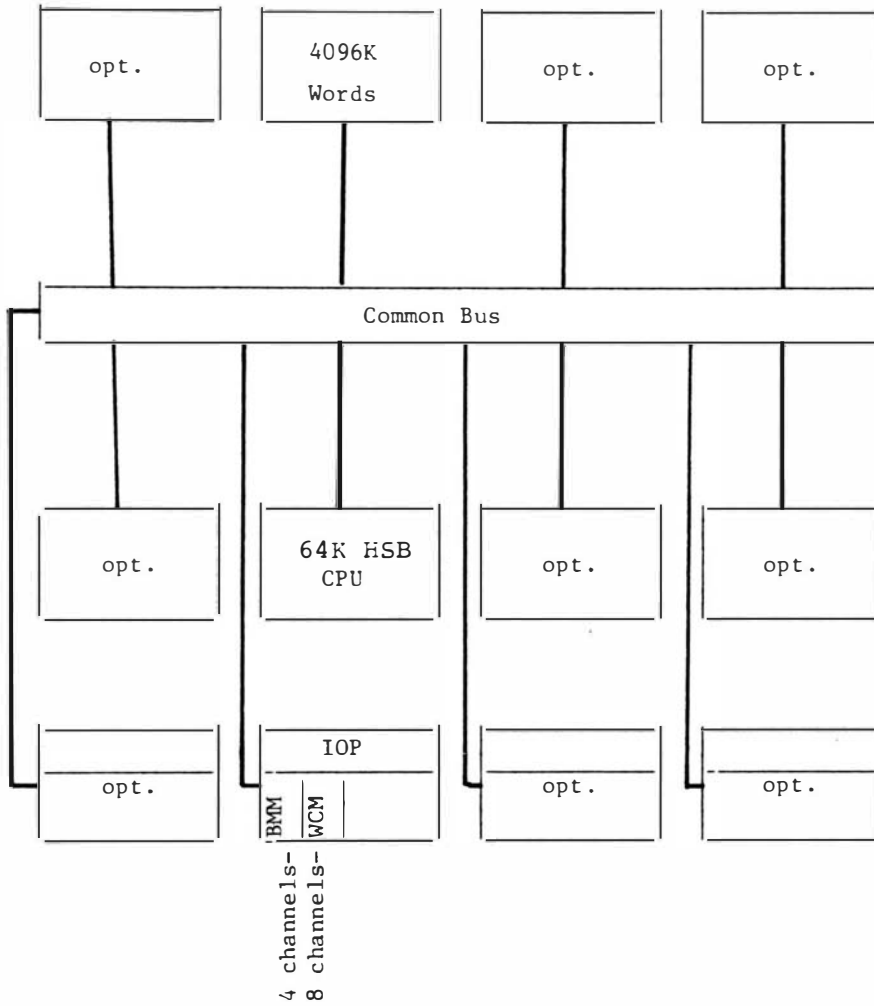


Figure 6

Communications between CPU, I/O processors, and main memory management is accomplished by a specially developed interrupt structure that uses a request/acknowledgement scheme for communication. Each component passes messages to each other using dedicated mailboxes, an implementation of using a bus design with little overhead of a system control mechanism. Optionally, a built-in performance monitoring set of hardware collection registers can be referenced to analyze system performance. Using this performance toll, overhead of a software monitor to gather performance data is eliminated.

Since no system controller feature is utilized, and communication is accomplished via mailboxing messages, configurations are free form, and no relationships exist between CPU's, I/O processors, and memory units to force specific configurations. As a result, typical CPU conflict and associated overhead is not a factor and system communications do not bottleneck on one system component.

Main Storage Systems

The primary function of main storage in a processor is to provide a storage area for instructions and data that will be acted on by the central processor. The activity of the main storage system is determined by instructions and data references required by the CPU. Therefore, the speed of the processor is somewhat dependent upon the speed of the memory accesses required.

Various methods have been implemented to accelerate fetch and store requests. Most notable is the concept of a high speed cache buffer. This buffer, usually considered part of the central processor, holds a certain amount of the last fetched data from main storage. Fabricated from high speed random access chips, access to the data within the high speed buffer is usually one machine cycle. Should the requested data not be in the high speed buffer, a storage request is sent to a storage control unit for fetching of the requested data. A storage fetch operation can take anywhere from one to twelve machine cycles, depending on a number of

factors such as type of memory used, number of simultaneous requests, and amount of data transferred per machine cycle. Once the data is fetched, it is moved to the high speed buffer where it is accessed by the central processor. Highest performance can be achieved if most fetch and store operations are done within the high speed buffer. Consequently, the greater the size of the high speed buffer, the better the machine throughput.

In order to provide fast access to memory, some designs have interleaved memory into arrays in which consecutive addresses are in different memory modules. For example, if a machine had four way interleaving and the central processor requested four consecutive memory locations, all four could be fetched simultaneously reducing memory fetch time by a factor of four.

Each system studied has incorporated certain attributes of main storage systems to optimize performance of the machine. Memory systems are generally standard within each vendor's product line which allows for easier upgradeability.

Amdahl

Amdahl's Main Storage system is composed of a Memory Bus Controller (MBC) and a Main Storage Unit (MSU). The MSU is a 16MB to 64MB unit, composed of 16K Dynamic NMOS chips with a cycle time of 288ns (12 machine cycles) fetching 8 bytes each cycle.

The MBC provides communication paths and message traffic control between components of the processor. A data integrity unit assures that the current version of a data line (32 bytes) is accessed where multiple copies of a data line can exist in the operand buffer, instruction buffer, I/O processors, or main memory. A component of the MBC, the Main Storage Controller (MSC), receives data requests for accesses to memory, sets the appropriate control latches and generates error checking and correction codes.

The MSU consists of 8 bit bytes arranged in 32 byte lines. Four way interleaving and four way quarterline multiplexing (each quarterline being 8 bytes in length) provide data bus paths of 72 bits (8 bit bytes plus 1

bit parity per byte). Each 8 byte message can be transferred every machine cycle. The 580 MSU has been optimized for main memory data fetches, which are the most common memory requests. A scenerio of a memory fetch is a follows:

- . Storage unit generates a read request
- . MBC receives read request
- . MSC unit receives request
- . MSC uses message opcode and address portion to create control signals for MSU
- . MSU accesses one of the four quarterlines from one of the four interleaves and latches them to a data-out register which is 8 bytes wide
- . Transfer of data is done to requesting unit.

CDC

The CDC 875 processor consists of 256K to 1024K words of 4K bipolar memory organized into 16 logically independent banks. The banks are phased so that successive addresses are in different banks, much like interleaving. One word can be fetched every 25ns, and having a three word transfer scheme, the memory cycle

time is 75ns, with a word size of 60 bits. Four memory interface ports, each with a three word buffer, are provided. Each central processor (maximum 2) is connected to a port, and each set of peripheral processors (maximum 2) is connected to a port. A data distributor services each of the memory interface ports on a priority basis and multiplexes data between the ports and memory. The data distributor also does error checking and correction. No cache memory is available since the bipolar memory speed is a very fast 75ns.

As an optional feature, Extended Semiconductor Memory (ESM) provides up to two million words of additional capacity. Data is transferred between main memory and ESM at the rate of 10 million words/second. When used with the Direct Extended Memory Access feature, where data is stored on disk, ESM functions as a buffer between disk and central memory. A Unified Extended Memory feature is standard which allows main memory to be partitioned into areas reserved for execution and areas reserved for large data storage.

IBM

The IBM 308x Processor contains between 8 and 64 MB (megabytes) of memory in 8 MB increments fabricated from 16K MOS technology housed within the central processing unit. Main storage access time is 312ns or 12 machine cycles, fetching 8 bytes per memory cycle.

Central storage is divided into Basic Storage Elements (BSE). Each BSE has either 8 MB or 16 MB of storage and contains logic for fetching doublewords (8 bytes) from or storing to data arrays in each BSE.

The BSE logic performs four functions:

- . Data storage and retrieval for the complex
- . Central storage communication with the CPU via the System Controller
- . Error Checking and Correction
- . Storage regeneration control

Central memory is 2 way interleaved by dividing the data arrays into Basic Storage Modules (ESM).

Interleaving of contiguous 2K blocks of storage provides for simultaneous access of separate Basic Storage Modules by multiple CPU processors, I/O processors, or combinations of processors. Key controlled storage protection is used for both store and fetch protection. Each 4K block of storage is protected by a 7 bit protect key, and is regulated by the System Controller. Storage requests can be queued, up to a maximum of 8, within the System Controller.

Sperry Univac

The Univac 1100/90 system consists of up to 4 Main Storage Units, (MSU), each composed of 4 million words of storage, fabricated from 65K chips. Each MSU has four storage banks, each of which acts independently, giving each MSU the capability to handle a total of four simultaneous requests from IOPs and CPUs. The storage system will interface to CPU buffers of up to 8 words, using an 8 word block transfer.

Two or four way interleaving is used and it

automatically allocates consecutive block (8 words) addresses to separate storage banks. Each CPU can stack up to 16 write requests if a busy condition is sensed for the MSU. Main storage can also be logically partitioned to each CPU, providing diagnostic and storage protection capabilities.

Input/Output Systems

An Input/Output system consists of components that transfer data between the processor and peripheral devices. These components generally consist of channels, which handle data transfer protocol and sense and status interpretations, and one or more I/O processors which interface between the processor and the channels. Communication channels exist between the I/O processor and the central processor for receiving commands for data transfer and between the I/O processor and the main storage system for the transferring of data into or from main storage and peripheral devices.

In order to provide flexibility of devices that are attached to a channel, the I/O processors are generally microcoded to allow for many different device types and configurations. The speed of the I/O processor and the number of channels it controls has a direct bearing on the aggregate data rate, the amount of data transferred through the I/O processor. With current high speed direct access devices, block multiplexor

channels (channels that interleave blocks of data from multiple devices on the same channel) are the dominant channel type. Byte or word channels function much the same as block multiplexor channels except bytes or words are interleaved rather than blocks. Byte or word channels are measured in the kilobyte or kiloword range whereas block multiplexor channels are measured in the megabyte range.

Redundancy, sharing of resources, and speed of I/O processors are the areas addressed by the processors studied here. Each has implemented a different design based on the type of workload, amount of built-in reliability, and relationship the I/O processor has with the other components of the system.

Amdahl

The Amdahl 580 provides 16 or 32 multiplexor channels implemented in LSI technology. Each set of 16 channels is controlled by an I/O processor (IOP) which is the primary interface between peripheral devices and the CPU.

The IOP consists of three components:

- . I/O controller (IOC)
- . Bus Handler
- . 16 interface handlers, one for each channel

An IOP including the IOC and Bus Handler is implemented on a single multiple chip carrier, MCC, and is shared by the 16 interface handlers. Each channel has 256 subchannels and can accommodate up to 6.0 megabytes per second data rate. The maximum aggregate data rate for the first 16 channels is 50 megabytes per second, with the second IOP increasing the total data rate to 80 megabytes per second.

Data paths in and out of an IOP are via the bus system. The Bus Handler is the interface and provides data buffering when required. The IOC performs the processing for the IOP and manages the IOC and the 16 Interface Handlers. Normal data transfer, including channel protocol and data buffering, is done by the Interface Handlers. Data and commands are fetched directly from the storage unit, via the data bus

thereby reducing contention between IOPs and the CPU.

Subchannel queuing provides the ability to hold I/O activities that were denied access to the system typically due to a busy device or channel. Once the desired device or channel is available, the held request is released for processing. Use of the feature reduces the load on the CPU.

One or two byte multiplexor channels are interfaced to the processor via the console interface. Each byte multiplexor channel has an Interface Handler and supports a data rate up to 200 kilobytes per second.

CDC

The CDC 875 contains an I/O unit to perform all control over external devices connected to the system. The IOU is composed of the following functional areas:

- . Peripheral Processors (PP)
- . I/O Channels
- . Central Memory access

- . Data Channel Connector (DCC)
- . Real time clock
- . Communications interface
- . Maintenance register

A basic IOU contains 10 PPs and 12 I/O channels and can be expanded to 20 PPs in groups of 5 with a maximum of 24 I/O channels. Each PP is an independent processor with its own memory and each set of 10 PPs comprises a multiplexing system which allows the PPs to share common hardware for arithmetic, logical and I/O operations without losing independence. Multiplexing is achieved by a set of 10 identical registers, one for each PP, rotated to provide time slicing of the set of PPs.

Each PP communicates with the central processor via central memory and flag settings, while PPs communicate with each other over an internal interface. Each PP has a 4K (12 bit words) memory, and each PP executes programs alone or in conjunction with other PPs to control data transfers. Requests from the central

processor are stored in central memory, fetched by a PP, translated into I/O requests, and scheduled and completed under the control of the PP program. The programs use the 4K memory as a data buffer between external devices and central memory. Any PP can access central memory directly via the central memory access function. Since data transfer is in 12 bit words, the central memory access function assembles 5 successive 12 bit words into one 60 bit central memory word, and likewise disassembles a 60 bit word for a write operation. Any PP can access any I/O channel and transfer at the rate of one 12 bit word every 500ns, and all channels can be active at the same time. A maximum of 20 PPs can simultaneously read central memory for commands, and a maximum of 5 PPs can write to central memory at the same time. Due to time slicing, each PP will be in a particular state of execution, so central memory bottlenecks are not a problem.

IBM

The IBM 308x series of computers contains an External Data Controller (EXDC), fabricated of LSI chips that perform channel functions. Provisions for up to 24 channels with a data rate of up to 3 megabytes per channel on each block multiplexor channel are available. Four of the 24 channels can be byte multiplexor channels with a data rate of up to 500 kilobytes/second. The EXDC contains the following logical elements:

- . Channel processing element (CPE)
- . Data server element (DSE)
- . Interface adapter element (IAE)

The CPE is a specialized processor for controlling I/O instructions and interrupts. It performs queuing of I/O requests, manages channel path selection, and communicates with the system controller. It is driven by pageable vertical microcode, has a two byte data path, and is packaged on one TCM. The CPE is shared by up to three DSEs, each DSE handling 8 channels. Each DSE is

a horizontal microcoded processor, housed on one TCM, with the microcode being shared by the 8 channels on an equal round robin basis. Each DSE has 256 bytes of data buffering capability per channel and controls the transfer of data to and from central storage. Each DSE port is connected to a non-LSI Interface Adapter Element, outboard from the processor. This hardwired IAE contains 8 bytes of data buffering and communicates and handles all data traffic sequences on the channel.

Sperry Univac

The Univac 110G/90 contains one to four I/O processors (IOPs), each with at least 8 word channels and 4 block channels, each set comprising a channel module. Each IOP can handle a total of six modules in any combination, with at least one being a block multiplexor module. Block multiplexor channels transfer up to 4.3 megabytes per second on input and 3.7 megabytes per second on output, with each block module having an aggregate data rate of 17.2 megabytes per second. Word channels, in sets of 8, have a maximum data rate of 3.7 megabytes per second per channel and 18 megabytes per second per word module.

The IOP receives commands via a Universal Processor Interrupt/mailbox system. The IOP does all channel processing without interrupting the central processor until completion. Queued I/O requests are kept in the IOP if device or path busy's are detected. The IOP communicates directly with main storage, completing the I/O request and signaling the requesting unit that I/O processing is complete via the mailbox system.

Reliability, Availability & Serviceability

The systems described in this paper all use some form of LSI technology as the basic building blocks for the processor. Compared to previous non-LSI technology of prior machines, these processors require an approach to error detection, fault isolation, and service different from their predecessors. But along with the technology comes an intrinsic reliability of a reduction in failures due to fewer physical connections and off-chip data paths, where most failures occur.

In the past, error re-creation strategy satisfied most diagnostic procedures because conventional tools and human intelligence were sufficient for problem isolation and repair. With the LSI packages used in these machines, a lesser number of individual components are available for repair or replacement, and diagnostics are structured to "call out" the failing field replaceable unit (FRU). Consequently, the less the number of FRUs, the easier it is to service.

Should major components, such as central processors, I/O processors, etc. be replicated, system availability can be enhanced, providing adequate real time problem diagnosis and built in system reconfiguration functions are available. The recording of any such actions, along with all intermittent problems for early problem detection, provides a vehicle for the vendor to follow what is or has happened within the processor, on parts such as LSI chips. This concept has removed the need to probe manually various circuits for problem diagnosis and provides real time capabilities to vendor personnel. Each manufacturer has implemented these features according to their processor design and components.

Amdahl

Amdahl has utilized the concept of fewer FRUs to improve the reliability of the processor, and by replacing discrete wiring with printed circuit boards for processor component interconnections. The high circuit density has allowed Amdahl to place an entire functional unit, such as the central processor, on a single MCC which provides fast fault isolation and repair.

Fault isolation is accomplished by having each MCC contain logic which records the identity of the circuit on the MCC detecting the error. Source data, control signals and latch contents are all recorded using the console processor. Built-in logic scan facilities can determine and/or set the contents of system latches for diagnostic purposes. Specialized RAMs are included on each MCC to maintain a microcode and bus transaction history for fault tracing. RAM is also used for maintaining a history of main memory correctable errors.

The console processor, implemented on one MCC, connects to a console complex which includes a floppy and hard disk, a system scan facility, a CRT/keyboard and an AMDAC control. Using AMDAC, any function done locally can be performed from a remote diagnostic facility via telephone connection to the console. Since the console processor is implemented in LSI and is part of the central processor, a separate microcomputer based support processor located in the console provides additional diagnostic capability. Component replication is done only in an attached or multiprocessor system.

CDC

CDC has implemented maintenance diagnostic facilities in a fault-analyze mode. No automatic reconfiguration is available, with the exception of an attached processor environment, where loss of one central processor will not cause total system outage. Failed peripheral processors can be configured out of the system and degraded I/O systems can be run.

Basic error detection is used and data is stored in maintenance registers for interrogation. Diagnostics are done down to the board level and the FRU is the failing board. RTA, Remote Technical Assistance, is available via telephone lines to allow for remote diagnostics in support of local service personnel.

IBM

IBM has implemented RAS in much the same form as Amdahl since the 308x includes much the same features, processors and facilities. IBM has also included many of the same concepts, such as real time logging,

diagnostic chips on each TCM, and remote access to system functions. The basic dyadic 3081 processor contains two central processors with automatic system reconfiguration should one fail. This is not true of the UP 3083 models.

There are also a limited number of FRUs per processor, between 18 and 54 TCMs depending on model, each which contains approximately 133 LSI chips. Each TCM contains diagnostic circuitry called Logic Support Stations to communicate using the Level Sensitive Scan Design capabilities of the Processor Controller to record the status of the machine down to particular latches. The facility can be used for loading of Engineering Changes to microcode as well as diagnostic capabilities of hardware functions.

The Processor Controller, a separate unit, contains necessary disk, modems and consoles for recording all activity within the processor complex.

Sperry Univac

Univac provides advanced processor availability through the use of multiple CPUs and IOPs providing continued system operation should a component fail. Partitioning of CPUs and IOPs can be accomplished so that a component can be removed from the system, diagnosed, repaired, and re-added to the system without an outage.

The system support processor, a micro-processor based unit, contains disk, memory and its own operating system that is used to communicate with each component of the CPU. Partitioning of components is done using this processor, and it also does recovery and isolation action when errors are discovered. Logging of all actions is done for later analysis.

Diagnostics are done down to the board containing the LSI chips. This board is the FRU of the 1100/90 system.

Applications

Each of the computer systems studied has been designed for a particular segment of the computer marketplace and components have been designed to provide high performance relative to the types of applications run on each system.

The ability of a system to perform work can be defined two ways:

- Performance : How fast can the work be accomplished
- Throughput : How much work can be done in X amount of time.

These concepts are applicable to the two major types of workloads:

- On-line : requiring fast response
- Batch : requiring high degree of overlap

Each area demands a specific type of design, yet most computer users have a mixture of both on their systems.

Specific performance in a commercial environment is generally measured in a benchmark, where a characteristic workload is run and measured on a particular computer system and re-run and measured on another system for comparison. Likewise the same procedure can be done within a family of systems to determine the improvement to be gained in upgrading from one system to another.

The applications stated here are based on manufacturer's claims and industry standards, along with conclusions based on information provided on internal processor organization and features. Each processor will perform differently on different workloads, but general trends can be seen based on architectural implementations.

Amdahl

The Amdahl 580 series of processors is designed around a high MIP (13) uniprocessor and independent I/O processors. Relative to other systems, it is the fastest commercial uniprocessor available. Because of the high MIP rate, the 580 lends itself very well to

on-line applications, where speed of turnaround of a transaction is of utmost importance.

Generally, in dealing with an on-line application, there is a single server queue mechanism, such as a control region or program, dispatching work. The speed of the on-line system is dependent on the speed of this control region. Since the 580 is a fast uniprocessor, the control region can then serve more items per unit of time than any other processor, thereby providing the fast on-line response time.

With regard to batch processing, the 580 provides fast turnaround of batch jobs due to a high degree of performance, that is, jobs will execute in less CPU time due to a high MIP rate. The effort placed on I/O processors support the batch and on-line application.

Amdahl has microcoded the execution unit for flexibility at the expense of performance of some instructions. Where other manufacturers have hardwired floating point instructions, Amdahl has utilized microcode at the expense of speed. Therefore, the

580 does not perform as well in a high computational environment where floating point calculations exist. Since the 580 was not designed at this marketplace, this has not affected Amdahl sales.

CDC

The CDC 875 machine has been designed for the scientific computing facet of the commercial data processing community. The hard-wired and specialized functional units provide high speed fixed and floating point calculations and the word size and register usage provide for a high degree of accuracy in computations. With its peripheral processors, the 875 provides good I/O overlap of CPU processing, and since I/O does not interrupt the system, more productive cycle time is achieved.

IBM

The IBM 308x processors have been designed for throughput. Since the maximum speed of a single processor is 7.5 MIPS, high speed units are made from

tightly coupled dual (dyadic) processors, providing 13 MIPs. Since 2 jobs can be executing simultaneously, a high degree of batch throughput is achieved. Independent I/O processors help to improve throughput also. With regards to on-line applications, the relatively slow MIP rate will tend to limit the speed of transaction turnaround.

IBM has microcoded all instructions except floating point, which are hard-wired. Therefore, the IBM 308x does very well in a scientific environment where multiple high calculation programs are needed to run. This also carries over into the commercial batch environment, providing a high degree of concurrency in batch processing.

Sperry Univac

The Unival 1100/90 series provides a high degree of redundancy at a moderate 6.5 MIP rate per processor. Multiple CPU configurations provide for a high throughput level, much the same as IBM. Since the 1100/90 system can easily be partitioned, dedicating processors to

tasks is available to provide a good overall throughput ratio. The independence of the I/O processors contributes to the throughput. The partitioning aspect lends itself to ease of testing of new applications, systems software, and maintenance/repair. It has obviously been designed for flexibility.

Futures

Each manufacturer described has ultimately the same goals in mind, to provide the highest performance, most reliable processor available. With the large software and peripheral investment each customer has made, these changes made by the manufacturer must be evolutionary rather than revolutionary. Within the next 5 to 10 years, computer manufacturers will concentrate on technology and its adaptability to reliability, serviceability, and performance.

Basic technological improvements will come in two categories, semiconductor technology and packaging, technology. With regards to semiconductor improvements, the use of other transistor technology (Figure 1) will provide needed lower gate switching speeds. Gallium Arsenide, GaAs, currently under development, provides a 4x decrease in signal delay time than silicon, the most commonly used media. TTL will more likely be phased out since it is a 1960's technology and its speed cannot compete with ECL. Josephson Junctions, though very fast and low power consumers, will not be practical

until 1990-1995 due to logistical problems with cooling, serviceability, and manufacture. Vendors will seek to reduce the delays per circuit with less regard to heat dissipation since it is easier to cool circuitry than it is to increase its speed.

With regards to packaging technology, the aim in the future will be to reduce the communication time between circuits and between chips. By placing multiple chips in a sealed package, chip to chip connections can be made without wiring, reducing signal delays. Since a ceramic substance slows down electrical signals, ceramic will be used only to hold a larger, more dense chip, perhaps a multi-layered cube of circuits.

Memory systems will be improved but not at the expense of cost. Cost per bit has and will remain the dominant factor in memory technology. Bits per chip exceeding 256K have been developed but yields are not high enough to warrant large manufacturing facilities. By 1985-1990, 256K chips should be commercially available. Manufacturers expect to have cycle time of RAM chips to be in the 2-3ns range, providing very fast memories

and ultimately affecting the cycle time of the processor.

Input/Output systems will be improved by taking advantage of technology developed by the communications industry, that is, the use of fiber optics for peripheral connections to the mainframe processor. Not only will this provide for faster data transfer, but more reliable and less bulky equipment will be required.

The effect of all these technological changes on the computers of 1985-1990 will probably be somewhere near the following figures:

- . Processors - 30 MIP uniprocessors will be available, and can be configured into 2 and 4 way multiprocessor systems
- . Memory - Memory size will be an average of 200-300 MB, fabricated from high density 256K chips. High speed buffer speeds will reach 2-3ns, lowering the machine cycle time.
- . Input/Output - Fiber optics will be used to provide for 20MB/second channel rates, a 7x improvement over current rates, with higher accuracy and less design constraints.

The processors of the 1985-1990 era will, as a result of these technological changes, be physically smaller, somewhat compatible with current systems architecture, inherently more reliable, and easier to service due to a reduced number of physical components that are field replaceable. Each manufacturer must anticipate the advantages of technological advances and evaluate them as to applicability to their processors, and use them to provide machines that meet the needs of the data processing community.

List of References

List of References

- Amdahl Corporation. 580 Technical Introduction, 1982
- Borgerson, B.R. and others. "The Evolution of the Sperry Univac 1100 Series: A History, Analysis, and Projection," Communications of the ACM, Volume 21 No 1, 1978
- Control Data Corporation. CDC Cyber 170 Computer Systems Models 865 and 875 Hardware Reference Manual, 1982
- Control Data Corporation. Cyber 170 Series 800 Computer Systems: Architectural Overview, 1982
- Credeur, K.R. Comparisons of Some Large Scientific Computers. NASA N82-11817/5, 1981
- Datapro Research Corporation. Amdahl 580 System. 1982
- Datapro Research Corporation. Control Data Cyber 170 Series 800. 1982
- Hamacher, V.C. and others. Computer Organization. New York: McGraw Hill, 1978
- IBM Journal of Research and Development. Vol. 26 No.1 January 1982
- IBM Corporation. IBM 3083 Functional Characteristics. 1982
- Reddi, S.S. and Fevstel, E.A. "A Conceptual Framework for Computer Architecture," ACM Computing Surveys Vol. 8 No. 2 1976, pp. 277-300
- Satyanarayanan, M. Multiprocessors: A Comparative Study. Englewood N.J.:Prentice Hall, 1980
- Scannell, T. "Univac Caps Large-Scale 1100 Line," Computerworld, (July 19, 1982)

Siewiorek, Daniel P. and others. Computer Structures:
Principles and Examples. New York: McGraw Hill 1982

Sperry Univac. Sperry Univac 1100/90 Facts and Figures.
1982

Wu, Lin C. VLSI and Mainframe Computers. IEEE COMPCON
78 Spring, San Francisco, 1978

Vita

