




OCTOBER 1974

HEWLETT-PACKARD JOURNAL



A User-Oriented Family of Minicomputers

HP's minicomputer section manager discusses the philosophy behind the design of this new computer series.

by John M. Stedman

WHAT DO MINICOMPUTER USERS want? In setting design objectives for HP's new 21MX Series minicomputers, we tried to make the objectives conform as closely as possible to the answers to this question, as we saw them.

Minicomputer applications have broadened tremendously in the last few years. One finds minicomputers today solving problems that only a few years ago would have required a large expensive computer system or a dedicated system designed to solve one particular problem. In more and more cases a minicomputer turns out to be the best solution to a problem.

What Do Users Want?


In general, a minicomputer user wants the most cost-effective solution to his problem. He would like to have the solution as quickly as possible, and not be required to design special hardware to do the job. In addition, he wants a solution closely matched to his specific application, and doesn't want to pay for capabilities he doesn't need or want.

Minicomputers should be able to match closely to the number of peripheral devices required by the particular application. If the user needs only four he should not have to pay for twelve I/O slots. However, it is desirable that the minicomputer be extendable, allowing a user to expand the number of peripherals at a later date if this need arises.

Physical size is also important to some users. They don't want to have half a rack filled up with the CPU, power supply, I/O system, and so on, especially if they can do the same job in just a few inches of rack space.

Different systems require different amounts of main memory, and again users don't want to be paying for capability they don't really need. A dedicated system application may require only up to 32K words of main memory. Another information management

system application may require 128K words of main memory today, with the capability of expanding as additional needs arise. And it shouldn't be necessary to trade off physical memory space for I/O controller space.



Cover: *The HP 21MX Series is a family of advanced minicomputers featuring modular design, a choice of semiconductor memory systems, user-microprogrammable processors, and customized instruction sets, and a power system that has unusual immunity to substandard electrical conditions. The memory systems use the new 4K RAMs, a few of which are shown here.*

In this Issue:

A User-Oriented Family of Minicomputers, by John M. Stedman **page 2**

Microprogrammable Central Processor Adapts Easily to Special User Needs, by Philip Gordon and Jacob R. Jacobs **page 7**
Testing the 21MX Processor, by Cleaborn C. Riggins and Richard L. Hammons, page 10.

All Semiconductor Memory Selected for New Minicomputer Series, by Robert J. Frankenberg **page 15**
The Million-Word Minicomputer Main Memory, by John S. Elward, page 19.

A Computer Power System for Severe Operating Conditions, by Richard C. Van Brunt **page 21**



Fig. 1. The 21MX Series now consists of two computers—M/10 and M/20—with different levels of capability. Both are user-micro-programmable and have semiconductor main memory.

Software for many minicomputer systems is being written in higher-level languages, such as FORTRAN, ALGOL, and BASIC, so it's essential for these applications to have a good set of compilers or interpreters and associated support software. For other applications the best solution is to write the applications software in machine or assembly language. In this case, it is desirable to have a powerful, yet easy to use instruction set.

Capabilities like decimal arithmetic instructions for applications such as information management, and extended arithmetic and floating point arithmetic instructions for computational requirements, should be easy and inexpensive to add to minicomputers, if they are not standard. Again, a user should not have to trade memory or I/O space for these capabilities.

In some specialized applications the user's program spends much of its time doing a certain set of calculations or a certain operation over and over. If overall system response time is important and proportional to this, such as in a real-time measurement or control system, it is desirable that the user be able to add instructions easily, thereby replacing whole operations or subroutines and making them run much faster than the equivalents in assembly language. In other words it is desirable to be able to tailor the minicomputer to the user's custom requirements.

Reliability is of importance to any minicomputer user. It is especially important in such areas as tele-

phone switching or measurement/control applications where minimum down time is essential. In all other areas, however, it is certainly desirable to have a reliable minicomputer. And in the event that a component does fail, it should be easy to replace the failed subsystem and get the minicomputer back on the air as soon as possible.

For many applications, reliable operation even under abnormal power-line conditions is also important. The minicomputer should be able to operate normally if the power line voltage dips as much as 25% indefinitely, and be able to withstand complete loss of line power for short periods of time. Some applications also require power-fail-auto-restart capability; an example is a minicomputer monitoring data in an unattended location. In addition to these considerations, the computer should be able to run normally over a wide range of temperature, humidity, and vibration conditions.

The New Family

As a result of these considerations, our project team created not just one new minicomputer, but a family of processors, the 21MX Series. All of the members of this family are compatible with assembly-language programs written for the earlier HP 2100A,¹ and are also I/O compatible with previous HP machines. This protects the large amount of investment in these earlier minicomputers—over 2000 man-years of software and over 70 peripheral interfaces supported.

The Value of User Microprogrammability

User microprogrammability can be an extremely valuable feature. It allows the user to customize the computer, dramatically increasing performance, increasing software security, or adding features that are important to a particular application but are not offered in the base instruction set. For example, benchmark programs run on an HP 21MX Computer with the HP 12977A Fast FORTRAN Processor, an optional microcode package, show performance increases of up to 28 times over the same programs in software. Typical FORTRAN programs run four to six times faster using the 12977A. Similar dramatic performance increases are expected in user applications.

Input/output is another area that can benefit from special microprogramming. Since I/O is under direct microcode control in 21MX Computers, the application of microcoding for higher throughput in I/O-intensive applications should be very profitable.

Several companies concerned about easy transportability of their sophisticated application programs are now considering changing their software programs to microcoded subroutines executed as machine instructions from fixed control store. The 21MX has 512 macroinstruction codes reserved for new instructions, so this approach is viable for a large number of routines. Implementing routines in firmware is not a foolproof protection mechanism against software pirating, but it does make the job much more difficult.

Microprogramming of features useful in specific application environments, but not offered in the base 21MX instruction set, can speed execution, simplify programming, or provide useful new architectural features. For instance, certain advantages of a stack-oriented computer, such as subroutine linking and parameter storage, may be easily and inexpensively incorporated into the 21MX family by user microprogramming. Two instructions called PUSH and POP use the A accumulator as the source and destination register to and from the memory stack. For simplicity, the valid stack range is defined by the X register (the lower limit) and the Y register (the upper limit). Accumulator B points to the last valid stack entry. An overflow bit is set if the stack is full and a PUSH is attempted, or empty and a POP is attempted. B will be updated to point to the new top of the stack if no overflow condition occurs. Implementing a stack capability of this caliber requires only 12 words of microcode, leaving 244 words in a standard 256-word user control store module for other useful architectural enhancements.

Other features that might be useful in certain applications are a microcoded DMA that searches for key characters, specialized pattern recognition instructions, special arithmetic operations like complex arithmetic, or queue manipulation operations.

Beyond this, many new features and capabilities are incorporated in the new family, thereby opening the way for new applications.

There are now two minicomputers in the family, designated M/10 and M/20 (Fig. 1). Each has a different amount of capability to match different combinations of immediate and long-term user needs.

The 21-M/10 contains four powered I/O slots and has space in the mainframe for 32K words of main memory. Like all 21MX mainframes, it has standard

extended-arithmetic and single-precision floating-point instructions. It takes only 5¼ inches of rack space, and its applications are expected to include such dedicated system areas as satellite navigation or processing of oil exploration data.

Designed for larger systems applications, the 21-M/20 contains nine powered I/O slots and can hold 65K words of memory in the mainframe, extendable much further via memory extenders. The dynamic mapping system, optional on the 21-M/20, is one of the many significant contributions in the 21MX family, allowing the 21-M/20 to be used in such applications as real-time measurement or control where quick access to large amounts of data stored in main memory is essential.

Design Contributions

Lowering the cost of these processors significantly was a key goal of the design team, along with increasing their capabilities in several areas. Solid-state memory, specifically the 4K MOS RAM, was chosen as the key component of main memory for several reasons, including lower cost, lower power, higher density, and potentially great increases in reliability. This decision had to be made before these RAMs were available in production quantities, so close relationships were established with several vendors of these components while they were still in development. This allowed design of the minicomputer family to proceed while vendors were still building up their capabilities to produce reliable parts in volume.

Modular design of the minicomputer family was another major design goal (see Fig. 2). One advantage of this approach during development was that memory subsystems could be designed around memories available at that time (1K RAMs) and it was then possible to adapt the designs in a very short time to 4K RAMs as they became available. Other advantages to this approach will be realized as enhancements to the family are developed. It isn't necessary to redesign the entire computer, but only a subsystem, to increase performance in specific areas such as memory, CPU, power system, and so on.

Modular design also plays a key role in the serviceability of the family. All major subsystems are built to be easily removed and replaced in case of malfunction; for instance, the entire CPU and control store is one assembly, and can be easily replaced by removing a few screws. Another advantage of modularity is commonality of subsystems between members of the family, allowing lower manufacturing cost by economy of scale. For instance, the CPU assembly is the same for all processors, so only one automatic test system is needed.

The CPU or control processor, totally micropro-




Fig. 2. Modular design helps tailor the new computers to any application. It also facilitates future improvements, since specific subsystems can be upgraded without redesigning the entire machine.

grammed and user-microprogrammable, is the key to allowing these general-purpose minicomputers to be tailored to custom, specific requirements without requiring special hardware design. Users can custom-tailor the machine to their own requirements easily using support hardware and software available from HP. HP will continue to take advantage of this capability also, offering as standard features or inexpensive options capabilities such as extended-arithmetic and floating-point instructions.

The 2100A has a good record of reliable operation, often in severe environments. Another goal of the 21MX project team was to increase the reliability of the new family significantly, lowering the maintenance costs in all applications, allowing its use in new market areas not previously addressed, and enhancing its capabilities in current markets. Examples of high-reliability applications are data acquisition systems in aircraft, process control systems in mills, and computational/data monitoring systems in oil exploration applications. In addition to significant potential gains in reliability achieved with the 4K RAM, specific areas where reliability increases are achieved are described in several of the following articles.

Reliable operation under abnormal power-line conditions was also achieved by an efficient power supply which achieves greater than 70% efficiency under most load and line conditions. Its wide tolerance of voltage and frequency variations (47-66 Hz, 88-132 volts or 176-264 volts) allows its use in applications

where the power source is a poorly regulated motor-generator, or under brownout conditions on conventional power sources. Battery backup is provided to sustain the contents of memory in the event of complete line failures, allowing auto-restart capability sometimes required in remote unattended systems for control or monitoring applications.

Acknowledgments

Those involved in developing the original concepts of the 21MX Family included Fred Coury, Stan Mintz, Jake Jacobs, Jim Toreson, Dick Hackborn, and Larry Peterson.

Jake Jacobs, Lam Dang, Jim Toreson, Dave Tsang, Larry Peterson and Al Kinney developed the first prototype. Dave Tsang continued on development of memory subsystems, and was later joined by Don Cross, Gordon Goodrich, Don Bowman, and Bob Frankenberg. Bob managed development of the 21MX memory subsystems and is now project manager of the 21MX program.

Al Kinney worked on the 21MX power supplies in the early phases of development and was joined later by Dave Baker, Dave Kuykendall, Ron Villata, Ken Check, Jack Elward, and Dick Van Brunt. Dick's previous experience on the 2100 power supply plus help and guidance from Dick Crawford and Greg Justice of Hewlett-Packard Laboratories were key elements in developing superior power supplies for the 21MX family.

Jake Jacobs continued design and management of


the control processor and options, and was joined by Jack Elward, Phil Gordon, Gordon Matheson, Ron Villata, Walt Lehnert, Don Bowman, and André Schwager. Pat Mulreany helped in innumerable ways with software aids and development tools. Darlene Harrell, Joe Dixon, Jim Fouts, Karl Balog and Dave Willis have also contributed in many ways during development of this family. The I/O Extender was developed by Phil Gordon and Earl Kieser.

Larry Peterson was joined by Dave Horine on the product/industrial design of the 21MX family. Rose Carson, Gary Thomsen, Gerry Priestly, Dennis Silva, Gary Koopman, Dave DeMoss, and Roger Wilder also were among the many that contributed in this area.

The Dynamic Mapping System for the 21-M/20 was developed by Ron Matsumoto, Janelle Bedke, Bob Frankenberg, Bill Gimple, Jim Kasson, George Anzinger and Jack Elward. Eric Ha and Gordon Matheson developed the EIG. Dave Crockett provided both overall guidance and specific inputs on the instruction set and marketing aspects. Walt Lehnert developed the Memory Extender and assisted in many other areas.


I wish it were possible to mention all of the others who had a part in this program. Inputs from other divisions and customers aided in making tradeoffs during design phases. Bob Jones and his printed circuit group showed great patience and skill in laying out complex printed circuit boards under continued tight schedules, and the manufacturing division provided excellent response in manufacturing complex printed circuit boards.

Cle Riggins and his production engineering group were involved from the start of the project to ensure an easily manufacturable product. The CPU tester was developed by Kirby Miller, Russ Scadina and Rich Hammonds. Dave Boone was involved in many production aspects of the product, and the guidance of Bob Cornell, Dick Ellis and Paul Hammel have ensured a smooth introduction into manufacturing. Ex-

tensive tests by Jim Gillette's quality assurance group and Bernie Levine's material engineering group have contributed to a quality machine. Many others from marketing assured a proper product introduction including sales literature and manuals by Ray Ahrens, Downey Overton, Jim Higgins, Ed Hayes, Jim Lally, Bob Kadarrauch, Larry Lotito and Wayne Gartin. 

Reference

1. F.F. Coury, "Price, Performance, Architecture and the 2100A Computer," Hewlett-Packard Journal, October 1971.



John M. Stedman

John Stedman is engineering section manager for 2000-series computer products at HP's Data Systems Division. With HP since 1969, John has contributed to the design of the 5401B Multichannel Analyzer and several nuclear instrument systems: the 5402A, 5403A, 5406A, and 5407A. He designed the 21MX memory protect system and managed the initial memory and power supply design efforts, then served as 21MX project manager before taking on his present responsibilities. John received his BSEE degree from Walla Walla College in 1966 and his MSEE from San Jose State University in 1969. He's a member of IEEE and the Sierra Club, and enjoys hiking, skiing, swimming, gardening, and golf. John is married and has two small daughters. The Stedmans live in Campbell, California, not far from John's native San Francisco.

SPECIFICATIONS HP 21MX Computers

The 21MX Computer family is a combination of 21-M-Series Microprogrammable Processors (M/10 and M/20) and 21-X-Series Semiconductor Memory Systems (X/1 and X/2).

HP 21-M/SERIES PROCESSORS

User-microprogrammable processor—optimizes performance by tailoring the CPU to each application.

Microprogrammable by programmable Read-Only Memory or by Writable Control Store.

DATA REGISTERS:

- 2 Accumulators
- 2 Index Registers
- 12 High-Speed Scratchpad Registers

BASE SET INSTRUCTIONS:

- 128 Instructions in Base Set
- 32 Index Register Instructions
- Bit, Byte, and Word Manipulation Instructions

Extended Arithmetic Instructions including double-word integer arithmetic, shifts, stores, and loads are standard.

High-Speed Floating Point.

Software compatible with previous 2100 Series Computers.

Power-Fail Interrupt allows for orderly system shutdown.

Standard bootstrap loader in read-only memory and space for four loaders.

Flexible Input/Output—I/O configuration totally independent of memory configuration and up to 32 additional devices added with I/O extenders.

Vectored Priority Interrupt.

HIGH-SPEED CONTROL STORE:

Bipolar LSI ROM Semiconductor
4096 words of addressable control store space (only 1024 used for standard instructions)

ROM cycle time of 325 nanoseconds

175 microinstructions

POWER SUPPLY:

Voltage: 110 or 220V ±20%

Frequency: 47.5 to 66 Hz

Power requirements, maximum: M/10, 400 watts

M/20, 525 watts

2½ line cycle power loss toleration

Switching regulator power supply

Crowbar input overvoltage protection

Heat dissipation: M/10, 1365 BTU/hr

M/20, 1795 BTU/hr

ENVIRONMENT:

Operating temperature: 0° to 55°C

Altitude: Non-operating, 25 000 feet

Operating, 15 000 feet

Shock: 30 g for 11 milliseconds, ½ sine wave shape

Vibration: 1 g at 44 Hz

PHYSICAL:

Height: M/10, 5½ in; M/20, 8½ in

Width: M/10, 19 in; M/20, 19 in

Depth: M/10, 23½ in; M/20, 23½ in

Weight: M/10, 39 lb; M/20, 45 lb

HP 21-X/SERIES SEMICONDUCTOR MEMORY:

650-nanosecond system cycle time

16-bit word, 17th bit for parity

Uses 4K Metal-Oxide-Semiconductor Random-Access Memory components

PLUG-IN MEMORY MODULES:

X/1 high-density memory in 8K or 16K word modules

X/2 medium-density memory in 4K or 8K word modules

Each module is one board, 7-3/4 by 8-11/16 inches

Only one memory controller required per processor, regardless of how many memory modules are used.

MEMORY VOLATILITY PROTECTION:

Ac standby mode selectable from operator's console

Memory sustained through loss of 10 line cycles

Power-fail recovery system available

Memory Protect allows a programmable boundary to be set, protecting all memory below the boundary.

DUAL-CHANNEL PORT CONTROLLER:

Assignable to any two I/O channels

Allows maximum transfer rate of 616 666 words per second

Maximum block size 32 768 words

PRICES IN U.S.A.:

M/10 Processor: \$4150.

M/20 Processor: \$5300.

X/1 Memory: Controller, \$650; 8K, \$2500; 16K, \$4600.

X/2 Memory: Controller, \$500; 4K, \$1300; 8K, \$2150.

Writable Control Store: \$1000.

Dual-Channel Port Controller: \$750.

MANUFACTURING DIVISION: DATA SYSTEMS DIVISION

11000 Wolfe Road
Cupertino, California 95014 U.S.A.

Microprogrammable Central Processor Adapts Easily to Special User Needs

The 21MX processor maintains program and I/O compatibility with its HP predecessors, but has a new microinstruction format that makes it easier to extend the instruction set.

by Philip Gordon and Jacob R. Jacobs

AMONG THE DESIGN GOALS for the 21MX Computers was compatibility with their predecessors in both programs and input/output. It was up to the design team to create a computer family that would build on the software and peripheral base that HP had established over the years.

The 21MX central processor, like that of the earlier 2100A Computer,¹ exploits the flexibility of microprogramming to the fullest possible extent. This permits even the smallest member of the family, the M/10, to have standard hardware multiply and divide and automatic bootstrap loader, and an enhanceable instruction set in which floating-point and bit, byte, and index instructions are standard. Available extensions to the instruction set include a fast FORTRAN processor and decimal arithmetic instructions.

A microprogrammed processor is really a computer within a computer. The lowest-level computer in the 21MX is a 24-bit microprocessor that cycles at 325 ns. This microprocessor emulates the instruction set of earlier HP computers, controls the front panel in the halt mode, operates the automatic bootstrap, and implements the enhanced instruction set.

Like the 2100A, the new 21MX family supports user-generated microprogramming. However, although the microprocessor word size (24 bits) is the same as that of the 2100A, the formats and fields are different and microprograms are not compatible. The decision to change the microinstruction format was a big one and was based on a number of reasons, including lower cost from new technology, easier microprogramming, and larger address space.

For example, to reduce part counts and cut costs, 64-bit integrated-circuit random-access memories (RAMs) were incorporated as scratch registers and working registers. These RAMs were unavailable when the 2100A was designed. To incorporate these new circuits, changes in the internal architecture and

therefore the microinstruction format were mandatory. It was decided that compatibility at the base instruction set level was far more important than at the microprogram level. Microprograms are typically small compared to applications programs, and our experience has shown that it is relatively easy to convert these small programs from the 2100A microcode to the 21MX microcode.

Easy Microprogramming

Another reason for changing the microinstruction format was to gain more precise, more versatile control of the microprocessor. Ease of microprogramming, especially by the user, was a worthwhile goal. In most earlier microprogrammed machines the microprogrammer must have an intimate knowledge of the internal gate structure of the computer and therefore microprogramming is done by the design engineer and purposely proscribed for the user. The 2100A was a pioneer in the field of user-microprogrammable machines. The new family continues this trend and is even easier to microprogram.

What is easier about the new microinstruction format? First, conditional branching is performed by a conditional jump rather than a conditional skip. Skips are satisfactory if, upon some condition, only one instruction need be inserted into the instruction stream. For example, to guarantee that a number is positive one might test its sign. If the sign is positive, one could skip over the complement instruction; otherwise the complement instruction would be executed. But what happens if two instructions must be conditionally inserted into the instruction stream? Now the conditional skip will not work, since it will skip only one instruction. However, a conditional jump will work, for upon some condition a jump of two instructions beyond the current instruction can be executed. Also, backward jumps can be program-

med for iterative loops.

In the 21MX, the conditions upon which one can jump include carry, zero, least-significant-bit, most-significant-bit, flag, extend-bit, overflow-bit, run-mode, halt-mode or interrupt, and most of the front-panel buttons. There are also many special jump conditions, not normally accessed by the user, that facilitate emulating the 2116 instruction set.

Twelve scratch registers are available to the 21MX microprogrammer, eight more than in the 2100A. These registers may be used to hold temporary or intermediate variables during execution of a microprogram. More registers results in fewer accesses to main memory and hence faster execution.

Microinstruction Formats

There are actually four microinstruction word types and four formats. The four formats and examples of typical microinstructions are:

Word Type 1

23	20	19	15	14	10	9	5	4	0
OP CODE		ALU	S-FIELD	STORE-FIELD	SPECIAL-FIELD				

SOP INC T A RTN
 SOP means standard operation. This microinstruction increments the contents of the memory transfer (T) register and stores the results in the A-register. Also, return from subroutine (RTN) is executed because this is the last microinstruction in a subroutine.

LGS PASS B B L1
 The B and A-registers, considered as a single 32-bit register, are left shifted one place, logically.

Word Type 2

23	20	19	18	17	10	9	5	4	0
OP CODE		P/C	U/L	OPERAND	STORE-FIELD	SPECIAL-FIELD			

IMM P L #7 S3
 The value "7" is passed (P) into the lower byte (L) of scratch register 3 (S3).

IMM C U #7 S3
 The ones-complement (C) of the value "7" is placed into the upper byte (U) of scratch register 3.

Word Type 3

23	20	19	15	14	13	5	4	0
OP CODE		CONDITION	RJS	ADDRESS (8-bits)			SPECIAL-FIELD	

JMP OVFL #123 CNDX
 Jump to control store location 123 only if the overflow (OVFL) bit is set, otherwise continue to next microinstruction.

JMP RUN RJS *+3 CNDX
 Jump to control store location three addresses beyond the current one only if the computer is not (RJS) in the RUN mode.

Word Type 4

23	20	19	17	16	5	4	0
OP CODE			ADDRESS (12 bits)			SPECIAL-FIELD	

JSB #1234 UNCD
 Jump to subroutine (JSB) in control store location 1234 and save the current address plus 1 in the SAVE register for later use with the RTN instruction.

Common to all formats is the SPECIAL field, which controls single and double-word shifts and rotates, setting and clearing of the flag, run-mode, and overflow flip-flops, and enabling the various jump tables.

Expanded Control Store

With the changing of the microinstruction format, the addressability range of the microprocessor was expanded. The 21MX computers support up to 4096 words of control store space, four times the space available in the 2100A. This space is divided into sixteen 256-word modules. The basic instruction set, consisting of all instructions standard on the 2100A including multiply and divide, is implemented in the first module. The second module holds the program for controlling the front panel and implementing the automatic bootstrap loader.

Two additional modules implement 2100-compatible floating-point operations and a new extended instruction group, which includes 42 versatile operations. Index registers X and Y are introduced with over thirty supporting instructions. This package also provides instructions to access, manipulate, and test bits or bytes, plus the ability to move or compare up to 32K bytes or words.

Besides these four standard modules, there are two others that are optional. These implement the fast FORTRAN processor and decimal arithmetic.

Microprocessor Operation

Unlike the 2100A, the 21MX family has no phase logic. Phase logic is what transfers the machine from one operational phase to another, such as fetch, indirect, execute, and interrupt. In the 2100A, flip-flops and combinatorial next-state logic are used to establish the current phase and the next phase. To minimize costs, it was decided to eliminate these circuits from the new computers. The burden was picked up by the microprogram for the basic instruction set. "Phases" are now merely microcoded subroutines.

The fetch phase is emulated by a three-word micro-routine, the essentials of which are as follows:

	OP	SPECIAL	ALU	STORE	S-BUS
(1)	READ		INC	PNM	P
(2)			PASS	IR	TAB
(3)	READ	JTAB	PASS	CM	ADR

Statement (1) does three things in parallel. It begins a memory read operation, loads the memory address register with the contents of program counter (P), and increments the program counter. Fig. 1 shows the internal bus structure of the machine as the old value of P is stored into the memory address register and the new incremented value is stored in P. The P register contained in the RAM is gated onto the S-bus. Since the memory address register (M) is loaded from this S-bus, it receives the old nonincremented value of P. However, P is loaded from the T-bus which is an incremented copy of the S-bus.

Statement (2) stores the memory data from the T-register into the instruction register (IR). The data path is on the S-bus.

Statement (3) takes the address portion of the instruction register (ADR) and *conditionally* loads the memory address register (CM) if the instruction is a memory reference instruction. The READ in statement (3) causes a read from memory even if the instruction is not a memory reference instruction. This may seem strange, but there are good reasons for it. More than 50% of the instructions executed by typical programs require an additional memory read (e.g., load accumulator, add to accumulator, increment memory and skip if zero). For these the memory read operation

is started early, during the instruction fetch phase. For instructions that do not require a memory read, the fact that a memory reference was started is of no consequence; it merely goes unused.

JTAB in statement (3) completes the phase, causing a jump to a microcoded routine that implements the instruction contained in the instruction register. This begins the execute phase. The n-way jump is accomplished by mapping the eight most significant bits of the instruction register into a 256-word read-only memory (ROM). This ROM is called the main Look-Up Table (LUT); it is different from the ROM containing the microprogram.

In the execute phase the microprocessor executes the instruction contained in the instruction register. In some cases, the instruction register contents are used during execution (for example, in the alter/skip and shift/rotate instructions). In other cases, the mapping through the main LUT is sufficient to define what action is to be taken by the microprocessor, and the contents of the instruction register are no longer needed.

As an example, consider the assembly listing for the microprogram for the ADA or ADB instruction add to A register or add to B register, depending on bit 11 of the instruction).

	LABEL	OP	SPECIAL	ALU	STORE	S-BUS
(4)	AD*I	JSB				INDIRECT
(5)	AD*			PASS	L	CAB
(6)		ENVE	RTN	ADD	CAB	TAB

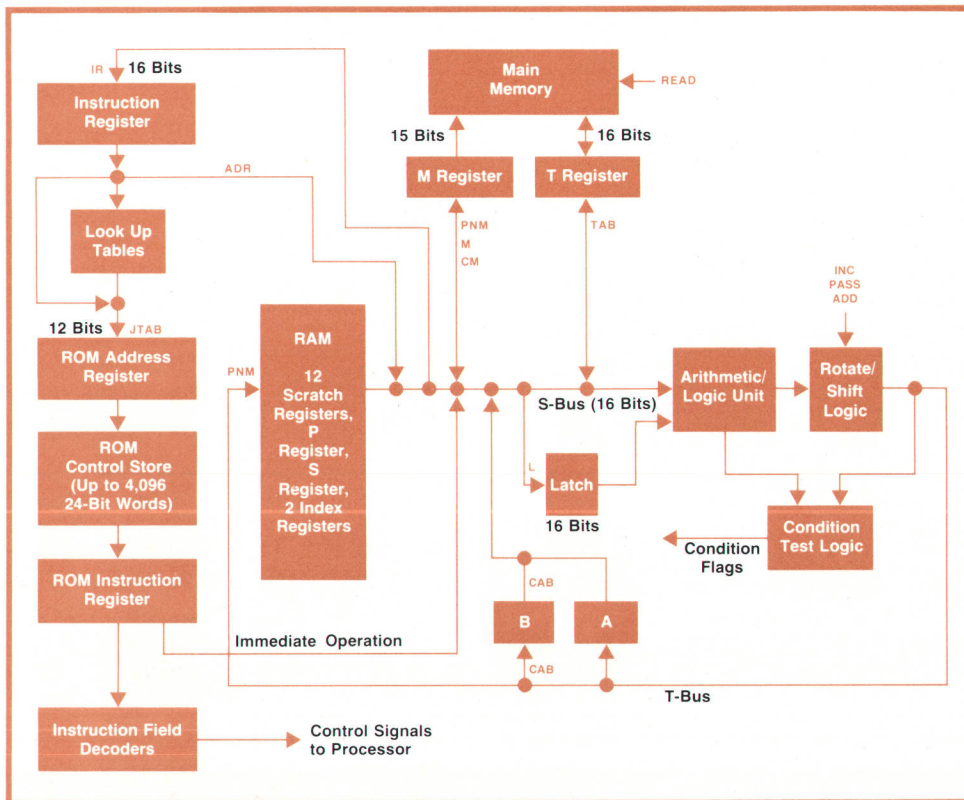


Fig. 1. 21MX microprocessor block diagram, showing the bus structure and the contents of the buses as a typical user instruction is fetched and executed. Control store is four times as large as the 2100A's, providing more room for special instructions and user-generated microprograms or instructions.

Testing the 21MX Processor

by Cleaborn C. Riggins and Richard L. Hammons

To test the 21MX Processor, which is housed on a single printed circuit board, the special automatic test system shown in Fig. 1 was developed.

During the definition phase of the product, the design team established objectives that were considered necessary to have a competitive product. The cost objective called for a labor content that represented a factor of four improvement over the present HP minicomputer, the 2100A. From the labor content, the allowable test time was extracted. After considering the various alternatives, it became obvious that general-purpose test systems would not meet the test time goal. To meet this target, a tester would have to give the processor a complete test at CPU speed in ten minutes or less, and troubleshoot failures to the failing component 85% of the time in less than three minutes.

How the Problems Were Solved

The ten minute limitation was relatively easy to meet. The 21MX is controlled by the microprocessor control store, a ROM. The ROM was replaced with a high speed RAM, the HP12908B Writable Control Store, and a set of microdiagnostics were written to exercise the processor and detect faults. To make the test a complete one, a memory subsystem is used in conjunction with the software diagnostics to check the processor with memory. An I/O simulator was designed to test the I/O control and logic.

Testing at CPU speed was a more difficult requirement to meet because the control store is located some distance away electrically. This problem was solved by using terminated twisted-pair wires and by selecting RAMs for the writable control store that will operate at faster than average speeds.

The troubleshooting requirement was the toughest. The microdiagnostic will detect a problem in a few seconds, but locating the failing component in less than three minutes is difficult. A troubleshooting tree is impractical because of the size of the tree. Any circuit changes would require a change in the tree




Fig. 1 Special test system completely checks 21MX CPU in ten minutes or less and troubleshoots failures to the component level 85% of the time in three minutes or less.

and a lengthy rewrite of the diagnostics.

For the system to correctly isolate a faulty device at least 85% of the time, it had to be able to distinguish between equivalent faults. For example, to a downstream device, an inverter that always outputs a logic one is equivalent to a device driving that inverter always outputting a logic zero. The solution to this problem was to provide the system with visibility to every node in the unit under test.

Because a unit as complex as the 21MX CPU would be likely to undergo several production changes in its lifetime, particularly in the first few production runs, the system had to be able to accommodate changes in the design of the CPU with a minimum of reprogramming effort (less than one man-week for a typical minor circuit change). This was accomplished by having the system acquire the data needed for the fault isolation routines by memorizing the responses of a known-good unit to the test stimuli. This allows most circuitry changes to be accommodated in just a few hours.

The System

A block diagram of the system is shown in Fig. 2. The basic system is an HP S310 Data System consisting of a 2100S Computer with 32K memory; a 7900A Disc Drive, a 12960A Magnetic Tape Drive, a 2600A CRT Terminal, and a 12925A Photoreader. Also included are two 12908B Writable Control Store (WCS) units and a specially designed card that controls operation of the UUT (unit under test).

The special interface card provides all clock pulses to the UUT and can be programmed to generate a single clock pulse, a specific number of clock pulses in a burst mode, or a continuous string of pulses. This board also contains circuitry to monitor the ROM address register (RAR) of the UUT and can be programmed to break, or turn off all clock signals to the UUT, at a specific address.

The UUT is connected to the system by means of a vacuum-operated "bed of nails" fixture that makes contact with the UUT in approximately 1400 places, about 400 of which are the normal inputs and outputs of the CPU. The other 1000 contacts are used to connect each individual node in the UUT to the node state registers. The node state registers allow the system computer to examine the state of any node in the UUT as it is currently, as it was at the end of the preceding CPU clock cycle, or as it was at the end of the CPU cycle before that.

The peripheral hardware unit contains a standard 21MX power supply and a 4K memory subsystem with dual-channel port control and memory protect options. Also included is an I/O simulator that is used to test the UUT's I/O functions.

Also shown on the block diagram, although not part of the CPU test hardware, is the 21MX battery pack tester. This tester is controlled by the CPU test system computer and all of its control inputs and status reports go through the CPU test system console. However, the operation of this tester is independent of the CPU test system software and transparent to it.

The Software

The CPU test system software, which was written in HP Assembly Language and HP ALGOL, runs under control of the HP Real-Time Executive operating system with the RTE file management package.

The major part of the testing is done via microdiagnostics. A

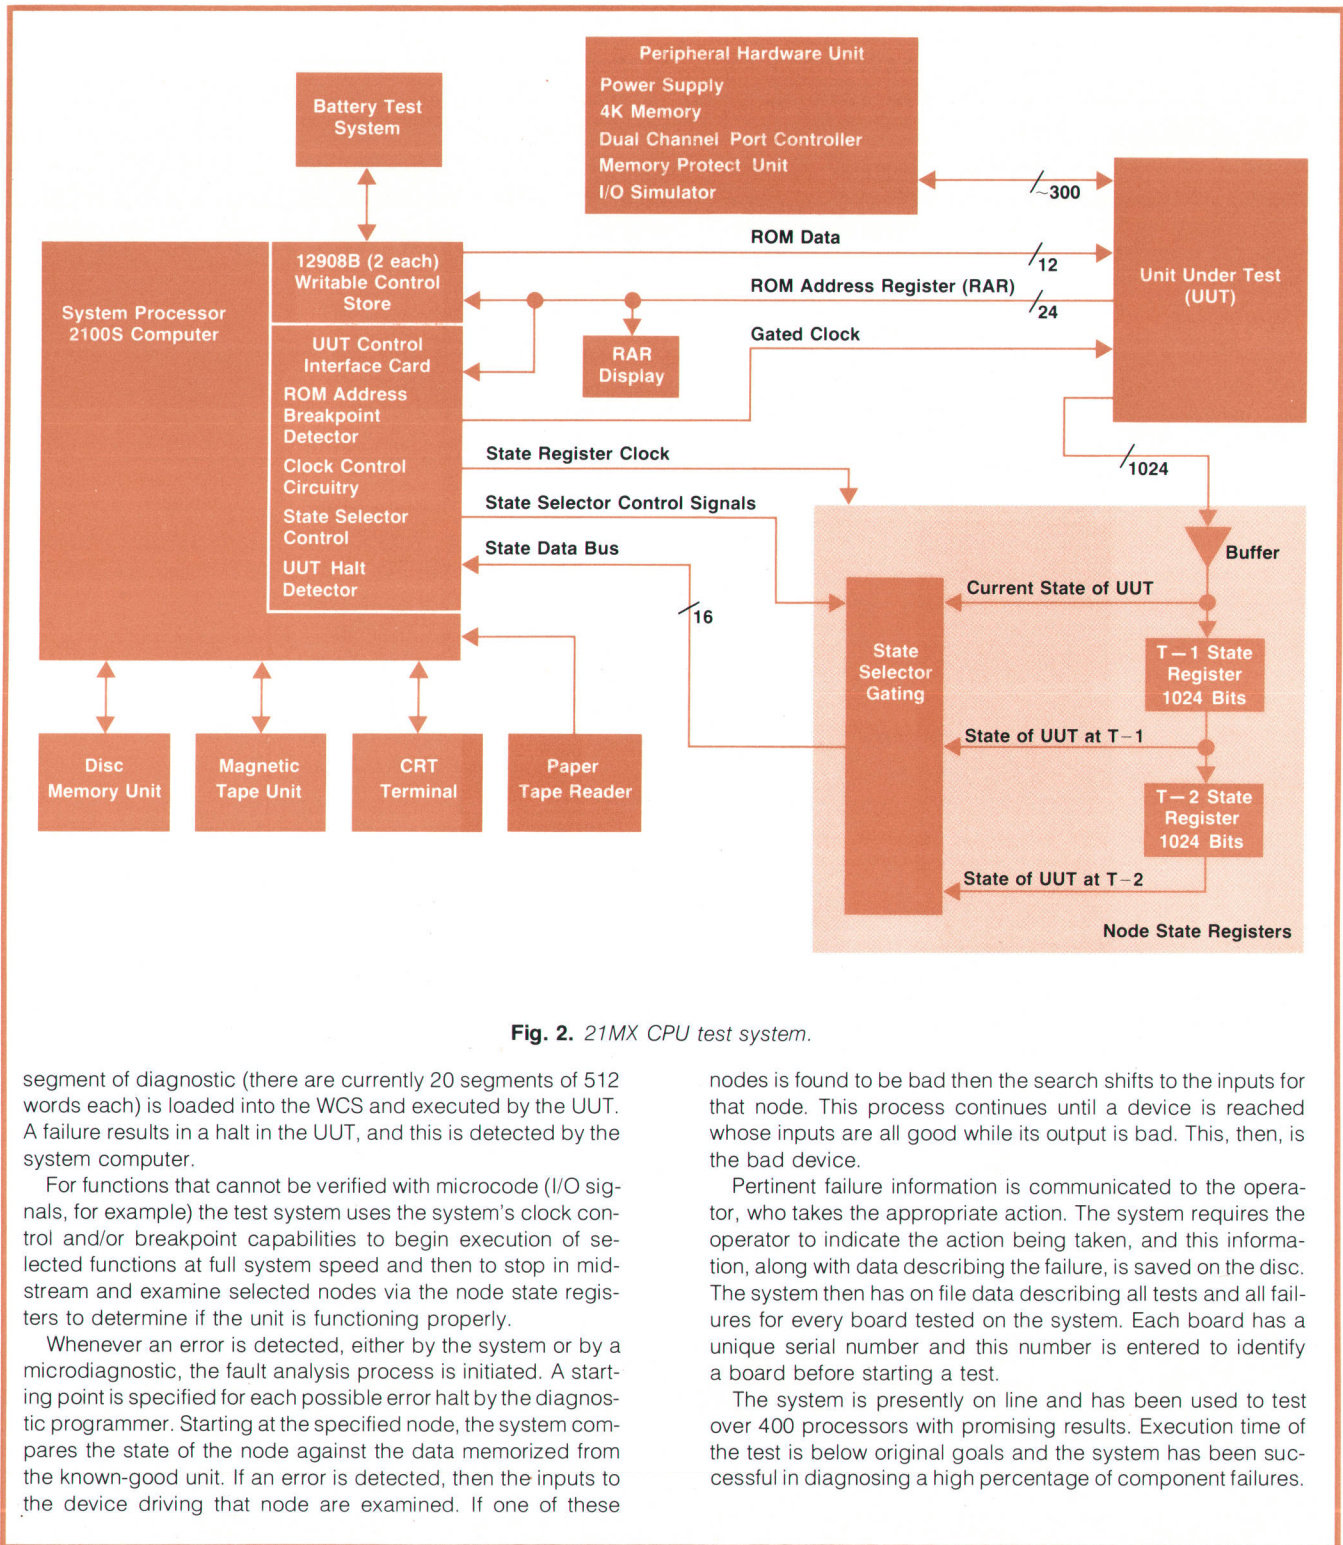


Fig. 2. 21MX CPU test system.

segment of diagnostic (there are currently 20 segments of 512 words each) is loaded into the WCS and executed by the UUT. A failure results in a halt in the UUT, and this is detected by the system computer.

For functions that cannot be verified with microcode (I/O signals, for example) the test system uses the system's clock control and/or breakpoint capabilities to begin execution of selected functions at full system speed and then to stop in mid-stream and examine selected nodes via the node state registers to determine if the unit is functioning properly.

Whenever an error is detected, either by the system or by a microdiagnostic, the fault analysis process is initiated. A starting point is specified for each possible error halt by the diagnostic programmer. Starting at the specified node, the system compares the state of the node against the data memorized from the known-good unit. If an error is detected, then the inputs to the device driving that node are examined. If one of these

nodes is found to be bad then the search shifts to the inputs for that node. This process continues until a device is reached whose inputs are all good while its output is bad. This, then, is the bad device.

Pertinent failure information is communicated to the operator, who takes the appropriate action. The system requires the operator to indicate the action being taken, and this information, along with data describing the failure, is saved on the disc. The system then has on file data describing all tests and all failures for every board tested on the system. Each board has a unique serial number and this number is entered to identify a board before starting a test.

The system is presently on line and has been used to test over 400 processors with promising results. Execution time of the test is below original goals and the system has been successful in diagnosing a high percentage of component failures.

The LUT maps the microprocessor to statement (4) if the INDIRECT bit is set in the instruction; otherwise statement (5) will be executed. Statement (4) is labeled AD*,I, which has no meaning to the microprocessor and serves only as a label. Somewhere else in the program, this statement might be referenced by this label. For example, JMP AD*,I would execute a

jump to statement (4).

Statement (4) does a jump-to-subroutine to a label called INDIRECT, which is the start of the indirect phase simulated in microcode. Upon completion, the INDIRECT subroutine returns control to statement (5).

Statement (5) says PASS the contents of CAB into register L, the holding latch (see Fig. 1). But what

does “the contents of CAB” mean? CAB means CONDITIONAL A OR B depending on bit 11 of the instruction, which specifies whether the add is to take place in the A-register or the B-register. The result of statement (5) is that the contents of the A-register or the B-register are put into the L-register.

Statement (6) actually performs three operations in parallel. The ADD CAB TAB portion of the instruction says to add the memory data (TAB) to the contents of the latch and put the result into the A-register or B-register. Recall that this memory data was requested earlier in the fetch phase. The ENVE will enable the setting of the overflow and extend flip-flops depending on the result of the add portion of the microinstruction.

Finally, the RTN in statement (6) means “return from subroutine.” This causes the microprogram, with hardware assistance for speed, to ask whether the machine has been halted, and whether there is an interrupt pending. If the answers to both of these questions are “no”, the microprocessor then goes to the fetch phase. If either answer is “yes”, then a further firmware test is made to determine which question has the “yes” answer. Normally the answers will be “no” and no time will be lost answering the question, “which one?”

If the machine is not in the halt mode, it is assumed that an interrupt is pending and a microroutine that services interrupts is entered (interrupt phase). Unlike the microprocessor of the 2100A, which is used only in the run mode, the microprocessor in the 21MX remains alive and well even in the halt mode. In this mode it controls the front panel, or programmer’s console.

The flow of control through the various micro-coded machine phases is shown in Fig. 2.

Microprogrammed Front Panel

In the halt mode, the front-panel microroutines continuously scan the switches on the panel to determine which button has been depressed by the operator or programmer. The microprocessor then jumps to a routine to carry out the desired function. Because the front panel is controlled by the microprocessor, it contains only a minimal amount of logic, thereby enhancing reliability and minimizing cost.

In addition to this small amount of logic, the front-panel switches and light-emitting diodes are mounted on a single printed circuit board. The switches are an adaptation of a novel design developed for Hewlett-Packard’s Model HP-35 Calculator.² Strips of beryllium-copper are raised at each switch location over a printed circuit trace. Pressing a switch pushes the raised strip down and makes contact between the strip and the trace. The “oil can” or “cricket” principle provides positive, tactile feedback to the user that contact has been made.

Of special interest is the front panel IBL (initial binary load) function. On the processor printed circuit board, up to four loader ROMs may be mounted. Each loader ROM contains, in packed form, a 64-word binary loader. In contrast, the loader in the 2100A is resident in the highest 64 words of main memory. Although these 64 words are protected by a switch on the front panel, many times through operator error this loader is destroyed and has to be loaded into memory again manually.

To use the IBL feature the operator places the de-




Fig. 2. 21MX microprocessor controls the front panel as well as program execution. This diagram shows the flow of control through the various machine phases.

vice number of the loading peripheral into the front panel switch register and presses the IBL button. The microprocessor scans the memory and finds the uppermost 64 words. Next, the loader ROM, which contains 256 4-bit words, is unpacked into 64 16-bit words and placed into the uppermost 64 memory locations. The microprocessor scans the 64-word loader program, and finds all I/O instructions, patches in the correct device, and sets the program counter to the first word of the loader program, all automatically. Loader ROMs for up to four different peripherals may reside in a CPU, each immediately available to the operator via front panel switch selection.

User Microprogramming

The new microprocessor with its easy-to-use microcode and expanded address space, along with the assembler and debug software packages, offer the user a strong invitation to write his own microprograms. There are several reasons why a user may want to do this. Performance may be improved by increasing the speed of frequently used software routines, and memory space may be saved as a consequence. New instructions may be invented to take advantage of internally available registers. It may be desirable to customize the computer for certain applications. Because the front panel is under microprogrammed control, the user can design his own front panel and offer a specialized machine for a specific application.

Hewlett-Packard continues to support the user in developing his own microprograms. The 21MX has been made compatible with writeable control store (WCS).³ Each WCS card may be dynamically loaded with up to 256 words (one module) and up to four WCS cards may be used. Microprograms can also be placed in nonvolatile programmable read-only memories (PROMs). Hewlett-Packard offers a high-speed PROM-writer subsystem with full software support to allow the user to convert his code easily into PROMs. The user generates and assembles his code with the HP microassembler, then enters the special mask tapes into any 2100-Series Computer equipped with the PROM-writer system. Modular control-store assemblies, accepting up to twelve PROMs (two modules) are available. These mount underneath the processor board along with the basic microprograms (Fig. 3).

The Hardware

The three machines in the 21MX family use many common subassemblies, thereby decreasing the quantity of different parts to be built and tested. The common parts include the CPU printed circuit assembly, the front panel printed circuit assembly,




Fig. 3. Entire 21MX CPU is on a single eight-layer board. User-generated microprograms may be put into programmable read-only memories in modular control-store assemblies, which mount underneath the processor board along with the basic microprograms.

the memory protect option, the control store assemblies, the dual-channel port controller, memory assemblies, and many mechanical parts. Custom to each of the three machines are the printed circuit backplanes for memory and I/O, the power supplies, and some sheet metal and mechanical assemblies.


The entire CPU is on a single eight-layer printed circuit board measuring approximately 33 cm by 43 cm. The eight-layer processor board contains approximately 240 integrated circuits. An eight-layer board is more costly than a two-sided board but the added cost is outweighed by the inherent noise immunity and reliability of the multilayer approach. Two planes are dedicated to power distribution; this

provides very low inductance and large distributed capacitance. The remaining six layers are signal layers. Because there are six layers, the trace width and separation is much larger than in two-sided boards, thereby providing a more reliable package.

A dual-channel port controller, program-compatible with DMA in 2100-Series Hewlett-Packard minicomputers, is available for all three 21MX Computers. All required logic resides on a single plug-in board.

In contrast to the 2100 Series, several system protection features have been moved from within the standard CPU and placed on a separate optional plug-in board. The protection package, called Memory Protect, can interrupt and protect a programmable portion of memory from alteration, prevent certain I/O instructions from executing, and keep the CPU from processing an instruction with a parity error. All three features are program-compatible with the 2100 Series and are available in both the M/20 and M/30 processors.

Time Marches On

The breadboard of the 21MX Computer was built in a 2100A chassis and used the 2100A's power supply and core memory. A 2155A I/O Extender housed the peripheral controllers. The logic was mounted on eight solderless wrapped printed circuit boards which, along with the backplane, were wrapped on semi-automatic machines. Ironically, these solderless machines, which were instrumental in the development of the new computer family, have been pushed close to obsolescence because 21MX Computers have printed circuit board backplanes and no solderless wrapped connections are used. 

References

1. C.T. Leis, "Microprogramming, ROMs, Firmware, and All That," Hewlett-Packard Journal, October 1971.
2. T.M. Whitney, F. Rodé, and C.C. Tung, "The 'Powerful Pocketful': an Electronic Calculator Challenges the Slide Rule," Hewlett-Packard Journal, June 1972.
3. F.F. Coury, "Microprogramming and Writable Control Store," Hewlett-Packard Journal, July 1972.




Philip Gordon (left)

Phil Gordon received his BSEE degree from the University of California at Berkeley in 1972 and joined HP the same year. As a project engineer working on the 21MX family, he has contributed to many aspects of the design, particularly the CPU. Phil was born in Huntington Park, California, near Los Angeles, and now lives in Santa Clara. When he takes a vacation he usually heads for the mountains, backpacking or hiking if it's summer, skiing or snowshoeing if it's not.

Jacob R. Jacobs (right)

Jake Jacobs received his BSEE and MSEE degrees in 1965 and 1966 from the University of California at Berkeley. Since coming to HP in 1969, he has worked on the I/O system for the HP 3000 Computer, designed the controller for the 7900 Disc Drive, and been responsible for the 21MX CPU design, front-panel design, and basic instruction set microprogramming. He's now an engineering section manager at HP's Data Systems Division. Jake was born in New York City. He's married, has two children, and lives in Mountain View, California. His hobby is photography, but he now spends much of his spare time studying for his MBA degree at the University of Santa Clara. He's a member of ACM.



Cleaborn C. Riggins (left)

Cle Riggins is production engineering manager at HP's Data Systems Division. Born in Dill, Oklahoma, Cle served in the U.S. Air Force for four years, then enrolled at Oklahoma State University, graduating in 1960 with a BSEE degree. He joined HP the same year as a development engineer, later becoming electronic tooling supervisor and then production engineering supervisor. In 1970 he received his MSEE degree from Santa Clara University. He's a member of IEEE. Cle enjoys fishing, woodworking, and suburban living in San Jose, California with his wife and three children.

Richard L. Hammons (right)

Rich Hammons was project leader for the 21MX CPU tester. He joined HP in 1962 after two years at the University of California at Berkeley, starting as an electronic assembler. He later served as a production supervisor, a production test technician, and an electronic tooling engineer. Self-taught in digital electronics and software design, Rich has been involved in numerous electronic tooling projects for 2000 and 3000 computer systems. He and his wife and two children live in Morgan Hill, California, and are currently busy putting the finishing touches on a home they built themselves. To get away from it all, they favor camping and waterskiing.

All Semiconductor Memory Selected for New Minicomputer Series

Considerations of cost, reliability, power, density, and speed all pointed to the 4K RAM as the best choice.

by Robert J. Frankenberg

TO SATISFY THE ANTICIPATED needs of users, 21MX computers had to have a memory system that met stringent requirements. The memory had to be inexpensive, because memory is still the most expensive single hardware element in a minicomputer. It had to be very reliable, because there are more circuit elements in one large memory module than in the processor, power supply, I/O interfaces, and front panel together. The memory system had to be fast, because memory speeds are still a major limiting factor in minicomputer performance. The memory had to be extremely dense (the 21-M/10, for example, may contain 32K 17-bit words of memory on three 18×23-cm circuits boards, a total volume of about 1230 cubic centimeters). The memory had to consume a minimum of power to reduce power requirements to the point where the power supply could, after significant innovation, be economically produced in the allotted space and at a reasonable cost. The memory had to be expandable from a minimum of 4K words to at least 32K words in the 21-M/10, and from 4K to 65K words in the 21-M/20. Memory expansion could not be allowed to degrade memory performance, because larger systems require good processor performance as much as smaller systems, and perhaps more.

Core versus Semiconductor Costs

Let's consider these constraints one at a time and compare the two major contenders, core and semiconductor, for a new minicomputer memory design. First, let's look at cost. At today's 4K random-access memory (RAM) prices, taking into account the system savings (fewer boards, less overhead circuitry, easier testing and debugging), the semiconductor RAM system is about 10% less expensive to build than the equivalent HP core system. It can be argued that a less expensive core system can be built today using a 16K

core stack. This is true. However, the cycle times of many of these core systems have been in the 1-to-1.2 microsecond range, which is relatively slow. Also, the minimum expansion increment of such systems is 16K words, forcing users to pay a higher system price to get a lower cost per bit. This is often not a good trade-off.

Another important cost consideration is that 4K RAM manufacturers are now at the very beginning of their learning curve. Decreases in part cost by a factor of two to four are almost certain within the next two to three years. Core, on the other hand, has been experiencing price decreases for over 20 years and it is not very likely that costs will decrease by as much as a factor of two within the next three years.

Reliability and Power

Second, let's look at system reliability. The average minicomputer core memory system on the market today displays a mean time between failures (MTBF) of about 9% per thousand hours.* The equivalent 4K semiconductor RAM system by all present indications should be as good or better. Furthermore, HP has been producing n-channel MOS LSI parts for over three years with a demonstrated part MTBF of less than 0.13% per thousand hours. RAM parts are the major determining factor in memory system reliability, and RAMs with MTBFs similar to these would produce a memory system nearly fifteen times as reliable as the equivalent core system. As RAM manufacturers learn how to produce even more reliable parts, we expect to see demonstrated long-term system reliability somewhere between our tested value and the value that would be achieved given 0.01%/khr parts, an acknowledged long-term reliability target of 4K RAM manufacturers.

The third major consideration is power, not only

*Computed using Rome Air Development Center methods.

the amount of power consumed, but also the regulation required, and the effects of both on system cost, size, and complexity. Reliable core memory systems require large temperature-compensated drive currents and high-gain sense amplifiers. This combination forces the power system to supply well regulated multivoltage power to a varying load. Semiconductor memories, by comparison, require about one sixth the power of equivalent core systems and this power need only be regulated to $\pm 5\%$.

Density, Speed, Expandability

Density is the fourth major consideration. Using any core technology that is economical, it has not been shown to be possible to put 32K 17-bit words of core memory including all control and interface circuitry in the 1230 cubic centimeters that this amount of memory occupies in the 21MX.

The fifth major consideration is performance. Memory systems made with 4K n-channel MOS RAMs have a cycle time of about 650 ns, whereas core systems with comparable prices and lower densities cycle at about one microsecond or more. These 4K RAMs are the first marketable version of these parts, and all manufacturers anticipate faster versions (some as much as two times faster) as more is learned about processing and as designs are improved to increase performance and decrease cost. Here again, core cannot hope to reach comparable performance levels at a competitive price.

The last major consideration is expandability. Because of their high densities and low power consumption, 4K semiconductor memory systems are easily expandable. Large core memories require greater volume, larger and mechanically weaker circuit boards, high-current power supplies, and better cooling, and because of long buses must often be operated at a slower speed when expanded.

The Problem of Volatility

The major problem with semiconductor memory is its volatility. Unlike a well designed core memory, when the power is removed from a semiconductor memory, the information stored in the memory is lost. The solution to this problem is a system solution. Dick Van Brunt's article in this issue describes how this problem was addressed in the power system. In the memory system, significant design effort was required to control and guarantee consistent refreshing. However, the cost of refreshing was far outweighed by the semiconductor memory's advantages in cost, reliability, power, density, performance, and expandability.

4K RAM Chosen

Taking all of the constraints into account, it became obvious that semiconductor memory was the best

choice and that 4K n-channel MOS semiconductor RAMs were the best part type for the new memory system. We came to this realization in 1972, long before any vendors were delivering parts or even reasonable samples. To get a breadboard of the 21MX working, a 2100A core memory was first interfaced to the system. This was only temporary; a system based on 1K semiconductor chips was soon developed and used on the first prototype versions of the 21MX. It appeared for a time that it would be necessary to introduce the 21MX with a 1K chip memory, because 4K parts appeared to be too far away from production. To determine just how far away, a campaign was initiated to contact every MOS manufacturer. Some had not started 4K RAM projects, but most were considering it, and a few, notably Intel, Mostek, Motorola, and Texas Instruments, were deep into 4K RAM designs with design goals which met our requirements. When the stage of development of these projects was compared with similar complex MOS LSI design projects conducted in the past by HP, and after production capacities and price projections were analyzed, we came to the conclusion that production volumes of 4K RAM parts would be available in mid-to-late 1974. This coincided very well with the 21MX's completion date. We then committed the product to 4K RAMs.

Implications of the Choice

Committing the 21MX to 4K RAMs implied several things about the system design. We were reasonably certain that parts would be available in large volume by mid-to-late 1974 but did not know which of the four leading manufacturers would deliver them first. This problem, coupled with the need for multiple sources of parts and the design constraints mentioned earlier, helped determine the system design. The memory control functions were separated from the CPU and memory arrays, creating a separate memory controller. This allowed the flexibility required to interface different RAMs to the same CPU without changing the CPU for different memory parts.

With one version of the memory controller we were able, with minor additions, to interface both Texas Instruments and Motorola parts to the 21MX. Although Texas Instruments and Motorola parts have separate module boards, they can be mixed in a single system. This memory system, the 21-X/2, is a medium density memory that comes in 4K and 8K modules.

The Mostek 16-pin parts were sufficiently different to require a different controller and module. Modules hold up to 16K words each, allowing 32K in the 21-M/10 processor. This high-density memory system is the 21-X/1.

The basic elements of these two memory systems

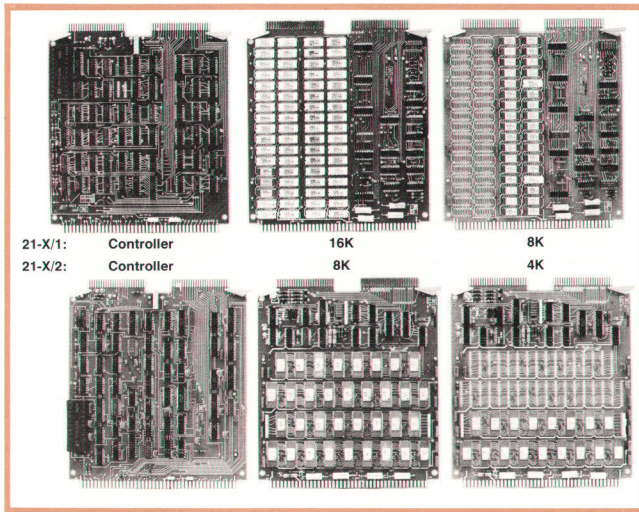


Fig. 1. Two memory systems allow 21MX Computers to use 4K RAMs from many manufacturers.

are shown in Fig. 1. Both memory systems were designed to run semi-asynchronously with the CPU, synchronizing only to refresh. This approach makes it possible to speed up the entire computer without major changes to the processor when faster memories become available from any source.

While the variations in RAM parts provided some of the design constraints, the most important considerations were user-based. For example, parity was included as a standard feature, as in all HP computers. Memory problems are rare, but in the event of an error, systems without parity detection will cause many days of grief for the user who tries to solve what appears to be a software problem when the real problem is a memory fault that would have been easily detected by parity circuitry. When a system is designed to include parity as a standard feature, the added system cost is very low, especially when compared to the potential cost to the user without parity protection.

Printed circuit boards used in the memory systems are multilayer boards. These boards are sturdy and

have withstood many severe vibration and shock tests. This is especially important, because the large 22-pin ceramic parts are susceptible to cracking when large boards are allowed to flex. Cracked packages don't seal the chip cavity and early failure is inevitable. The multilayer boards also provide good power distribution, more than adequate density, and increased static discharge immunity.

Memory expansion is an extremely important aspect of the memory system designs. Jack Elward's article in this issue explains how the 21MX memory was logically expanded above 32K words. Physical expansion does not cause any decrease in speed up to a 196K-word memory size, so users of relatively large systems can still have small-system performance. Both memory systems are designed as one-million-word memories to avoid the necessity of redesign should further expansion be required in the future.

Design Details

As Fig. 2 shows, the memory system is logically and physically separated from the rest of the computer. The data, address, and control lines enable the memory to communicate with the processor, the dual-channel port controller, and the input/output system. Refresh information consists of synchronizing signals from the processor and port controller, and a refresh status signal from the memory to the processor. Power status signals from the power system via the I/O system provide the memory controller with information required to refresh memory in power-up and power-down situations as well as transitions between the two states.

A block diagram of the memory controller is shown in Fig. 3. Because the Mostek RAM has an output data latch, a memory data register is not used in the 21-X/1 memory. The address multiplexer, used only for the Mostek 16-pin part, is not required in the 21-X/2 memory. The MEMORY MANAGER ACTIVE signal, when present, commands the memory to look at a 20-bit

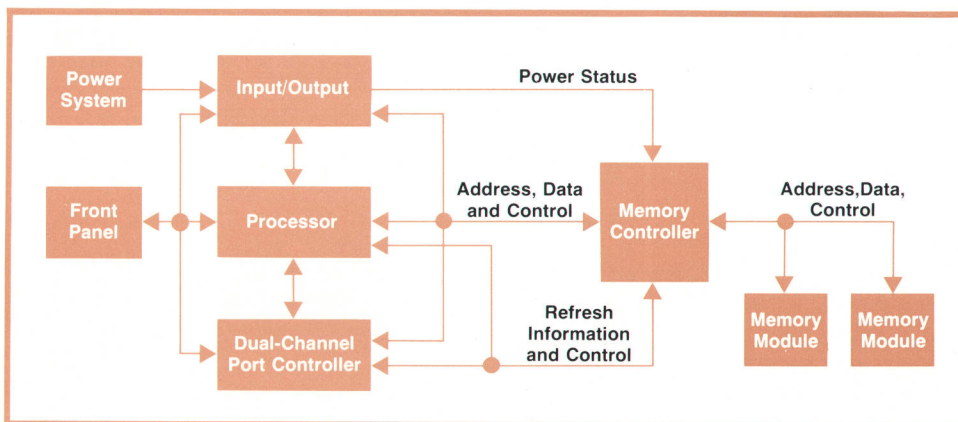


Fig. 2. Memory system is logically and physically separate from the rest of the computer. Memory runs asynchronously with the processor, synchronizing only to refresh.

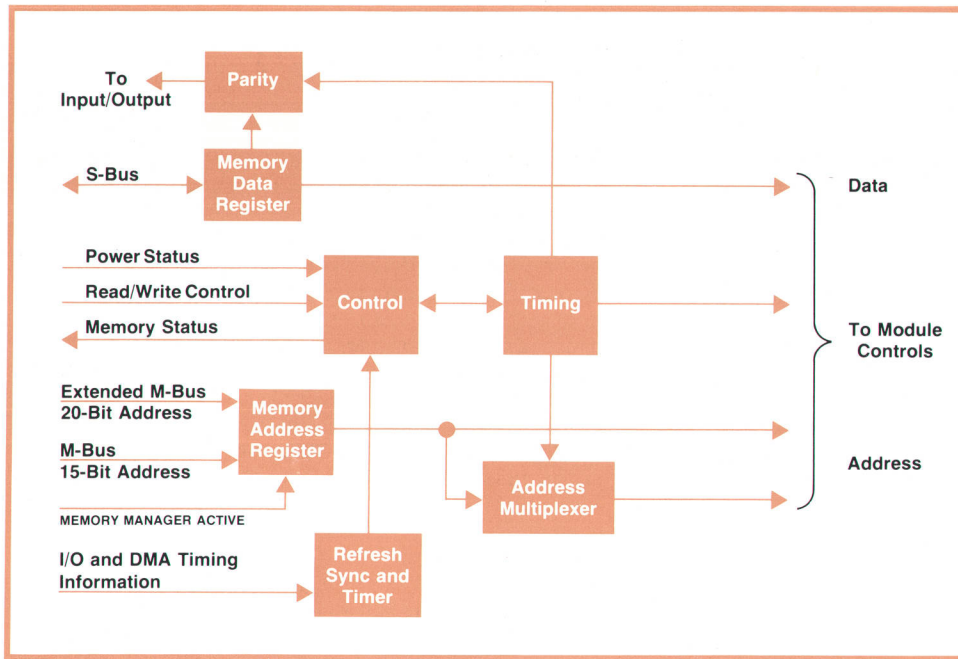


Fig. 3. Memory controller block diagram. MEMORY MANAGER ACTIVE signal tells the memory to look at a 20-bit address bus instead of the normal 15-bit bus. This expands the memory address space from 32K words to 1M words.

address bus instead of the standard 15-bit address bus. This expands the memory address space to 1M words from 32K words.

Memory controllers for the 21-X/1 and 21-X/2 address 1M words of memory, but physical space availability and bus lengths limit the expandability of the 21-X/1 to 32K words in a 21-M/10 and 65K words in a 21-M/20. The M/20 can add another 131K words to its physical address space by means of a 12990A Memory Extender. Bus lengths are kept to a minimum by mounting the memory extender directly below the computer and routing the bus through the bottom of the computer and the top of the 12990A.

Memory is extended without a decrease in performance by taking advantage of a characteristic of dynamic MOS RAMs. These memories require a small amount of inactive time between operations (or clocks). This time is sufficient to allow address translation and its associated delay to take place without system performance loss when the memory is equipped with the 12929A Dynamic Mapping System.

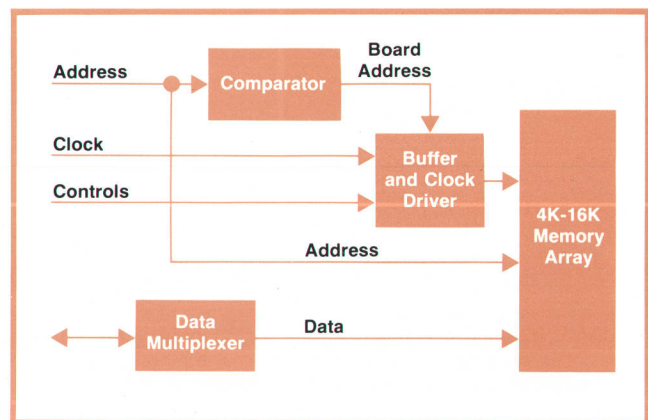
Refreshing of the entire physical address space is accomplished by the refresh timer, control circuitry, and memory modules. The refresh timer is an oscillator with a period of 31 microseconds. Every 31 μ s a pulse is generated and, if the processor is powered, the memory is synchronized with the processor, the I/O system, and the dual-channel port controller and a refresh is performed on 1/64 of the entire memory. The oscillator is allowed to free run, assuring a refresh every 31 μ s, on the average. If the processor is not powered, no synchronization information is required and refresh occurs without synchronization.

Memory modules (Fig. 4) decode their location with-

in the address space and can therefore be put in any slot in the memory backplane. Since modules vary from 4K words to 16K words, decoding varies with each module type. All control information is transferred through one connector on the top of the board. Only power is obtained from the backplane. This makes the 12990A Memory Extender much less expensive, because it doesn't have to provide control and busing, but merely power and the mechanical enclosure. The 12990A also makes a good powered enclosure for user-designed applications hardware.

Reliability Assurance

Because many memory parts are used in each system, the reliability of individual memory chips plays an important role in the overall system reliability. Therefore, it was important to properly qualify the



The Million-Word Minicomputer Main Memory

by John S. Elward

Traditionally, Hewlett-Packard minicomputers have been designed with a word length of 16 bits. The new 21MX Series, which is compatible with its predecessors, is no exception. Memory is organized into pages of 1024 words each. Each memory location is given a 15-bit address, the five most significant bits specifying the page number and the ten least significant bits specifying the location or displacement within that page. The maximum number of pages, therefore, is $2^5 = 32$, and the maximum memory size is 32,768 words.

The 12929A Dynamic Mapping System (DMS) alters the architecture of 21MX Computers to expand the memory address size to 20 bits, thereby providing a maximum main memory size of 1,048,576 16-bit words.* It does this without adding to the 650-nanosecond cycle time, so the user pays no penalty in speed when he adds more memory. The DMS is transparent to any user who wishes to ignore it, but is easily enabled and controlled when needed.

How It Works

Fig. 1 shows how the DMS fits into the 21MX architecture. It receives 15-bit addresses on the M-bus from the central processor (CPU) and dual-channel port controller (DCPC) and translates them into 20-bit addresses on the ME-bus. The CPU and DCPC operate normally, and user programs are unchanged. In an operating-system environment, the instructions that control the DMS are executed under a privileged mode and can be called only by the operating system.

Fig. 2 shows how the DMS translates 15-bit addresses into 20-bit addresses. The user program works, as before, within a logical address space of 32 pages, and the DMS performs a one-to-one mapping of this logical address space into the larger physical memory space by means of a translation table, or "map," consisting of 32 12-bit registers. The five most significant bits of the incoming 15-bit memory address are used to select one of these map registers. Ten bits from the map register then become the most significant bits of the new 20-bit address. The ten least significant bits are the same as those of the incoming address. The remaining two bits from the map register are used to implement the memory-protect features.

There are actually four maps in the DMS. Two are dedicated to the CPU and two to the DCPC, one to each channel. The two CPU maps, designated the system map and the user map, are used for program execution; one or the other is selected under program control. The two CPU maps are enabled by software commands and remain active until specifically disabled. (An exception occurs when an interrupt is acknowledged; the system map is immediately enabled to return the program to

* Physical limitations may restrict actual memory size to less than 1M words, depending on processor and memory type.




Fig. 1. Dynamic mapping system (DMS), when present, expands the maximum physical address space of 21MX Computers from 32K words to 1M words.

an environment for processing interrupts.) The DCPC maps are automatically enabled when the DCPC requests a memory cycle. The four maps may address entirely separate or overlapping areas of physical memory.

It is the responsibility of the operating system to load the translation tables. When the computer is first turned on the contents of the maps are meaningless because the DMS memory is volatile. Until the maps are loaded and the DMS enabled under program control, the DMS hardware is completely disabled and the memory accepts addresses from the M-bus. Thus the DMS can be ignored if not needed. In case of power failure, one map can be saved in about 50 microseconds and all four can be saved in 175 microseconds; this is well within the capabilities of the 21MX memory and power systems.

It was possible to insert the DMS into the existing 21MX architecture without affecting the design of the CPU or DCPC mainly because of the multi-bus 21MX design. A hybrid package consisting of hardware and firmware, the DMS can be installed in the field. The 21MX memory system has been designed to recognize the presence of the DMS and respond to the ME-bus.

DMS firmware implements over 45 special assembly-language instructions, including the fast FORTRAN group. DMS commands may also be included in user-generated microprograms.

System Enhancement Features

The DMS was designed to be a powerful tool for the system




Fig. 2. DMS maps 32-page logical address space into a 1024-page physical address space. Translation tables are loaded by the operating system.

programmer. Individual memory pages may be protected against writing and/or reading. Many DMS instructions are privileged and may not be executed with memory protect enabled.

A key feature is the ability of the programmer to segment the base page (addresses 0 to 2000_a). A ten-bit "base-page fence" stored in the DMS separates this page into a portion that will be mapped and a portion that will reference physical memory directly. This allows common parameters to be shared by two or more maps. The ten bits of the base-page fence are stored in a 16-bit status register along with other information about the state of the DMS.

Another 16-bit DMS register, the violation register, is activated whenever a program attempts a privileged operation in a non-privileged mode or when an attempt is made to access

a protected area of memory. A program interrupt is generated to signal the violation and the violation register stores specific information about the violation.

Several instructions are provided for exchanging data between program areas and to move words or bytes from one memory location to another. When the system map is enabled the operating system can, for example, move a block of words into the user's memory space or from one place to another within the user's space.


If a program needs more than 32K words of memory space a process similar to virtual memory may be used. However, instead of moving program segments to and from a disc, all that is required is a simple map swap, taking microseconds instead of milliseconds.

memory parts, provide adequate testing capability, and establish good manufacturing screening processes.

Memory parts were qualified initially by individual part tests. This was accomplished by the use of special testers and temperature-controlled ovens. Some parts were then opened and inspected for workmanship and design tolerances to determine the manufacturer's quality control level. All memory systems and manufacturers' parts were then sent through HP's stringent environmental tests, which include extended heat, cold, humidity, shock, and vibration, as well as system performance checks.

Manufacturing process screens were then established, including an extended dynamic burn-in at

125°C, a system heat run for at least 24 hours at 55°C, power cycling for 48 hours, and several quality assurance checks at various points in the process. These screens are aided by a randomly sampled complete characterization of each manufacturer's parts to monitor the manufacturers' process control. A quick-response (demand reporting) failure reporting computer system covers any in-process part failures, and an overall failure reporting computer system reports any part failures from the receipt of the parts to the end of HP's repair service on the computer.

The screens and the instantaneous failure reporting give an excellent means of screening out weak parts, detecting the presence of bad part lots, and assuring manufacturer quality control. Such efforts are essential for the manufacture of reliable, high-quality memory systems. 



Robert J. Frankenberg

21MX project manager Bob Frankenberg joined HP in 1969 with two years of college and four years in the U.S. Air Force under his belt. Starting out as a test technician, he evolved into a computer designer, contributing to the development of the 2116C and designing the memory systems for the 2100A and 21MX Computers. Author of several patent applications and instructor in computer fundamentals, Bob received his BS degree in computer engineering this year from San Jose State University. Bob was born in Wisconsin and now lives in San Jose. He's married and has a small son. Now that he has his degree, he hopes to have more time to indulge his interests, which include fishing, painting, restoring a 1935 Chevrolet, and inventing consumer electronic devices.



John S. Elward

Jack Elward received his BS degree in electrical engineering from the University of Michigan (his home state) in 1972, and came to HP the same year to work on the 21MX. He helped with the CPU and power supply and designed the dynamic mapping system for the M/20. A sports-minded bachelor, Jack plays tennis, racquetball, squash, paddleball, football, and softball. He's also studying for his MS degree in computer engineering at Stanford University, and expects to receive it in December. He lives in Cupertino, California.

A Computer Power System for Severe Operating Conditions

The power supply system of 21MX Series Computers differs in many respects from the power supplies of other minicomputers. It is less vulnerable to poor ac line conditions than the supplies of most minicomputers of similar size and cost.

by Richard C. Van Brunt

AS MINICOMPUTERS ARE USED in increasingly diverse applications the physical operating environments to which they are subjected can be extremely demanding. Typical is a minicomputer-based navigation system aboard an oil tanker. It is subject to salt spray, the vibration of the engines, and a power-line voltage that varies dramatically as massive electric cargo pumps are turned on, suddenly straining the ship's electrical system.

Other HP minicomputers have been used aboard airplanes and on factory floors, near arc welders in an automobile assembly plant, and controlling giant electric motors in paper mills amid steam, chemical fumes, and load-induced line-voltage fluctuations.

Experience with these types of environments was one factor that significantly influenced the design of the 21MX Computer and especially its power supply. Other factors that influenced the power supply design were the wide range of power-line voltages in different parts of the world, projections of increasing power shortages, brownouts, and blackouts and, more significant, the use of a volatile semiconductor memory system instead of core memory.

Computers have always been more seriously vulnerable to power interruptions and voltage variations than many other types of electrical equipment. Often large and expensive uninterruptable power installations are required to protect computer systems because a several-millisecond line-voltage transient, hardly enough to make the lights flicker, might wipe out hours of work or thousands of dollars.

Power-Fail/Auto-Restart

One approach to making computers immune to power failures is called power-fail/auto-restart. With this system the computer must sense a power failure

while there is still enough energy stored in the power supply to permit the computer to store certain key information in a memory that will not lose its contents when the power is removed. When power is restored the computer can pick up where it left off or perform any special restart procedure programmed by the user. Power-fail/auto-restart has always been available on HP computers.

Although power-fail/auto-restart is adequate in many situations it has the disadvantage that computer operation is interrupted during the line-power fault. In many real-time operations this can be bothersome or even unacceptable. Many of the "power failures" that cause a computer to enter the power-fail/auto-restart mode are only transients; for example, the line voltage may drop out of tolerance for a few cycles, or an automatic switchover may disconnect the power but then restore it in just the time required for mechanical contacts to open and close.

In real-time systems, it is often in these very situations that the computer is most required to perform adjustments or acquire data. If the computer goes momentarily out of service during such an interval it may be necessary to restart a complex process or operation, possibly at great expense. However, if the computer could continue operation through such a transient and keep the system under control much time and expense might be saved.

21MX Power Supply Characteristics

The power supply system of Hewlett-Packard 21MX Series Computers is less vulnerable to poor line environments than any other minicomputer of similar capability and cost. The line voltage rating is 110/220 Vac $\pm 20\%$, which means the computer can be turned on for line voltages as low as 88 Vac on a

nominal 110-volt line or 176 Vac on a nominal 220-volt line. Once the computer has been turned on it will continue to operate indefinitely even if the voltage dips to 70 Vac or lower. The computer is unaffected by line frequency variations between 47 and 64 Hz.

When operating at nominal line voltage the power supply's internal energy storage is sufficient for the computer to continue operating for a complete voltage cutoff lasting between two and ten line cycles at 60 Hz (32 to 160 milliseconds), depending on how heavily the system is loaded. A battery powered standby system is provided for the memory, so in the event of a power failure lasting longer than ten line cycles the contents of the semiconductor memory will remain valid for a period of at least two hours. The battery supplied is nickel-cadmium. Longer standby periods may be obtained by using a larger external nickel-cadmium battery.

The power supply operates in three different modes: operate, line standby, and battery standby. The computer does not distinguish between the line standby mode and the battery standby mode, but these two states are entirely different within the power supply.

In the operate mode, all output voltages are present and current is available up to the full capacity of each output. In the two standby modes, only those voltages necessary to permit the semiconductor memory to retain its contents are present. As the two names imply, the line standby mode receives input power from the line, while the battery standby mode operates on power from the storage battery.

Protection for Critical Circuits

Due to the relatively high replacement cost of the circuits using the output voltages of the power supply, especially the memory modules and the CPU board, all output voltages are protected against over-voltages, whether caused by a power supply failure or by a short to an external source of power. Similarly, the outputs of the supply are protected against over-current, which might be caused by a component failure or a misplaced screwdriver.

Unlike many power supplies, which current-limit to a maximum value when an output is shorted, the 21MX power supply completely shuts off the voltages involved in an overload until manually reset from the front panel. This reaction helps prevent fires caused by inner-layer PC board shorts and other insidious forms of damage.

The supply output voltages and their current ratings, in amperes, are as follows:

OUTPUT TERMINAL VOLTAGE	21-M/20		21-M/10	
	OPERATE CURRENT	STANDBY CURRENT	OPERATE CURRENT	STANDBY CURRENT
+5V (CPU,I/O)	35	*	25	*
-2V (CPU,I/O)	5	*	5	*
+12V (I/O)	2	*	1.5	*
-12V (I/O)	2	*	1.5	*
+5V (mem)	5	5	5	5
+12.5V (mem)	1.0	0.5	1.5	.5
-12.5V (mem)	1.0	0.5	1.5	.5

*Indicates that this output voltage is zero in standby mode.

The total output power available is 300 watts in the M/20 and 235 watts in the M/10. The efficiency is approximately 70%, depending upon which outputs are most heavily loaded.

Basic Principles

The design of the 21MX power supply is based on several physical properties:

- The energy storage in a capacitor is proportional to the capacitance and the voltage squared, while the physical size of an aluminum electrolytic capacitor is proportional to the capacitance and the first power of the voltage rating. This results in a value of energy storage per unit volume which, within certain physical limits, increases directly in proportion to the voltage to which the capacitor is charged. By using higher voltages it has been possible to extend the holdup time and reduce the volume of capacitors over that required for a more conventional approach.
- The size and weight of power magnetic components are reduced as the operating frequency is increased. An operating frequency of 20 kHz permits all of the power transformers of the 21MX power supply to be mounted directly on printed circuit boards, resulting in reduced fabrication and assembly costs.
- For a given power level, control at a high voltage and low current is inherently more efficient than at a low voltage and high current. This is because of the power lost by the relatively fixed forward voltage drop across a semiconductor junction.
- Within certain upper limits on frequency, a switching regulator is inherently more efficient than an active-region regulator such as shunt or series pass. Furthermore, this efficiency is relatively insensitive to input/output voltage differentials. As a result, a very wide input voltage range is permissible without significantly reducing efficiency. This aspect, coupled with the high storage voltage, allows a voltage drop of nearly 50% of the voltage on the energy storage capaci-




Fig. 1. 21MX power supply efficiency is approximately 70%. Optional battery preserves memory contents in case of line power failure.

tor, thus permitting nearly 75% of the stored energy to be extracted in the event of a momentary line-voltage interruption.

- The frequency at which the switching losses of appropriate power transistors begin to approach and exceed the static losses is somewhat above the frequency at which acoustic noise becomes inaudible to the human ear. Thus, an operating frequency of 20 kHz is a good compromise between the switching losses in the power transistors, the physical size of magnetic components and the undesirability of audible acoustic noise.

Power Supply Design

A simplified block diagram of the 21MX power supply is shown in Fig. 1. Power from the ac input line passes through an RFI filter and a turn-on surge-limit circuit. It is then rectified (in the case of 115 Vac operation the voltage is also doubled) and stored in the input ripple-filter capacitors at a voltage ranging between 180 and 380 Vdc depending upon the line voltage and load.

From this variable dc voltage at the input capacitors the power passes through a 20-kHz transistor switching-regulator circuit. This circuit uses a pulse-width-modulated transistor and diode switch with a filter choke to step the input voltage efficiently down to a regulated intermediate voltage of approximately 160 Vdc.

The efficiency and reliability of the preregulator circuit are enhanced by the use of a slow-turn-off circuit, which diverts current away from the switch-

ing transistor into an auxiliary circuit during its turn-off transition. This turn-off period in a switching regulator is almost always the period of greatest instantaneous power dissipation.

The effect of the slow-turn-off circuit is to reduce the collector current to approximately 25% of its initial value before the collector-emitter voltage has reached 25% of its final value. This reduces the peak and average power dissipation during the turn-off interval by a factor of approximately 16, to 125 watts peak and 1.25 watts average, and results in a preregulator-stage efficiency of greater than 96%, thereby improving the reliability of the circuit and reducing cooling problems. As a result of this high efficiency the preregulator switching transistor can be mounted directly on a printed circuit board with a heat sinking area of less than 20 square centimeters.

Dc-to-Dc Converters

At the output of the preregulator stage are two dc-to-dc converters, which step the relatively high voltage and low current down to the final low-voltage high-current outputs and provide transformer isolation between the computer circuits and the ac power line.


One of the dc-to-dc converters is used mainly to power memory-related circuits while the other converter powers CPU and I/O circuits. This allows the high-current +5 volt supply for the CPU to be shut off without affecting memory voltages. Thus a line standby state can be provided, so semiconductor memory contents can be kept intact when shutting off all other power to change I/O interface cards

or to reduce power consumption when the computer is not in active use. This line standby feature is also coordinated with the protection circuits so that if an overvoltage or overcurrent should occur in the CPU or I/O area those voltages will be automatically shut off separately from the memory voltages, thereby maintaining the memory contents until the cause of the problem can be determined or the misplaced tool removed.

The two-converter approach also allows memory operation to be extended in the event that the optional battery standby system, or "power fail recovery system," is not used. The CPU and I/O voltages are automatically shut off before the preregulator input voltage becomes too low for the preregulator to regulate. With this high-power load now removed from the preregulator, the remaining energy still stored in the input capacitors will provide memory voltage regulation for an additional 100 ms or more in the event of a total line loss. In the event of a voltage drop only, the reduced loading in this state will allow the preregulator to maintain the memory voltages indefinitely for line voltages as low as 50 Vac.

The voltage at which the CPU converter is turned off is factory-preset to allow the computer to be turned on for a line voltage of 88 Vac. However, this voltage level can be readjusted to increase the memory hold-up time at the expense of low-line CPU operation.

If the optional power fail recovery system is used the battery operated power supply will maintain the memory contents during power outages of two hours or more.

The power supply package consists entirely of printed circuit board assemblies. It requires a minimum of hand wiring in assembly and can be removed easily from the computer for field replacement. Repairs can be handled at the factory as part of the regular board exchange program. 



Richard C. Van Brunt

Dick Van Brunt joined HP in the summer of 1969 after graduating from Massachusetts Institute of Technology with a BSEE degree. He was extensively involved in the development of the power supply for the 2100 Computer, and was responsible for development of the 21MX power supply. Originally from Connecticut, Dick is an avid skier, figure skater, and swimmer, but has set these interests aside for the present to build a house that he's designed on a 47-acre parcel of land in the Santa Cruz mountains. The house will have a power supply, of course: Dick plans to generate and store his own electricity.

Hewlett-Packard Company, 1501 Page Mill Road, Palo Alto, California 94304

HEWLETT-PACKARD JOURNAL

OCTOBER 1974 Volume 26 • Number 2

Technical Information from the Laboratories of
Hewlett-Packard Company

Hewlett-Packard S.A., CH-1217 Meyrin 2
Geneva, Switzerland

Yokogawa-Hewlett-Packard Ltd., Shibuya-Ku
Tokyo 151 Japan

Editorial Director • Howard L. Roberts

Managing Editor • Richard P. Dolan

Art Director, Photographer • Arvid A. Danielson

Illustrator • Sue M. Perez

Administrative Services, Typography • Anne S. LoPresti

European Production Manager • Kurt Hungerbühler

Bulk Rate
U.S. Postage
Paid
Hewlett-Packard
Company

CHANGE OF ADDRESS: To change your address or delete your name from our mailing list please send us your old address label (it peels off). Send changes to Hewlett-Packard Journal, 1501 Page Mill Road, Palo Alto, California 94304 U.S.A. Allow 60 days.