


OCTOBER 1971

HEWLETT-PACKARD JOURNAL



Price, Performance, Architecture and the 2100A Computer

Here are the why's behind the design of HP's new minicomputer.

By Fred F. Coury

HEWLETT-PACKARD'S LATEST MINICOMPUTER, MODEL 2100A, is an entirely new design with the two notable exceptions of its basic instruction set and its input/output structure, which are identical to those of the earlier HP 2114, 2115, and 2116 computers. The unchanged I/O structure makes the 2100A compatible with all existing HP peripherals, and the identical instruction set allows it to work with the large existing library of HP software.

The 2100A's price is well below that of its lowest-cost predecessor, yet the new mini is faster than its fastest predecessor—980 ns cycle time versus 1.6 μ s—and includes as standard features many things that used to be options. Everything fits in a 12-inch-high mainframe which has room for 14 I/O interfaces and up to 32,768 words of core memory.


The other articles in this issue describe how the 2100A was implemented. As project leader, I'd like to say a few words about why.

Price/Performance

Price/performance ratio is often used as a means for comparing computers, and rightly so. However, I think the terms are often used in too narrow a way. Price, for example, is much more than the cost of a computer mainframe. It is the total amount a customer pays to solve his problem. In many applications, the price of the mainframe is as small as ten percent of the total cost.

Performance is difficult to measure because it is hard to define. People often use addition times,

memory cycle times, or other computer-related parameters as measures of a computer's performance. But these parameters are only indicators of potential performance. Often actual performance is limited by the application. For example, if a com-



Cover: Inside the new 2100A Computer is another computer, the microprocessor, most of which fits on the two boards shown here. The article on page 4 describes how the microprocessor's read-only memory can be expanded to extend the instruction repertoire of the 2100A.

In this Issue:

- Price, Performance, Architecture, and the 2100A Computer,*
by Fred F. Coury **page 2**
- Microprogramming, ROMs, Firmware, and All That,*
by Charles T. Leis **page 4**
- A Lot of Memory in a Small Space,*
by Robert J. Frankenberg **page 10**
- A Bantam Power Supply for a Minicomputer,* by Richard D. Crawford and Gregory Justice. . . **page 13**
- UTC Time Scale to Change in 1972.* **page 16**

puter is transmitting ten characters per second to a teleprinter, doubling the speed of the computer will not increase the printing rate. It will only make the computer wait faster.

In evaluating a computer's price/performance ratio, then, price must be the total price paid for the solution of a problem, and performance a measure of how well the system solves the problem or how large a problem the system can solve. Here's how these considerations affected the design of the 2100A.

Architecture

Whether or not to change architecture received a lot of serious consideration in the early planning stages of the 2100A project. What is architecture, really? There are two elements involved: the logical structure of the machine and its implementation.

Logical structure is implementation-independent and can usually be described in terms of a block diagram, a list of instructions, and a set of rules of behavior. Implementation is the transformation of the logical structure into actual hardware. It is the task of a computer designer to arrive at a combination of a logical structure and its implementation which provides the optimum price/performance ratio for a given range of applications at a particular point in time.

After analyzing and evaluating several alternatives we arrived at the following decisions:

Keep the same logical structure. We found that in our applications areas, radical changes in logical structure would not provide startling increases in performance per se. However, considerable systems, application, and support software and peripheral development efforts would have been necessary to support a new logical structure, and these costs would have had to be reflected in the price a customer pays. Also, many of our customers have themselves developed extensive systems using earlier HP computers and their investment in software development and peripherals should not be wasted. *Although we could have designed a logical structure capable of exhibiting better computer-related performance than earlier designs, the net effect on price/performance as defined above would have been negative.*

Keep the design open-ended. We found that, although radical changes to the instruction set did not necessarily produce significant improvements, there are certain areas where extensions to the instruction set can make a big difference. Therefore we wanted to keep the door open for additions and extensions,

such as floating-point instructions, that would really increase performance at a low incremental cost. To satisfy this requirement we decided on a micro-programmed implementation, as described in the next article.

Concentrate on implementation. The 2116A is five years old and technology has changed a lot in five years. Although there were several modifications made during these years to bring the design more up to date, we found that we could make significant improvements in implementation by using state-of-the-art components and by making contributions in CPU, memory, and power supply design. These contributions are described in the articles which follow.

Acknowledgments

In a team effort such as the 2100A project, it is very easy to start a list of acknowledgments but very difficult to stop. Instead of trying, let me simply thank all those who can look at a 2100A and get a feeling of satisfaction for having contributed in some special way to a job well done. 



Fred F. Coury

Fred Coury is minicomputer section manager at HP's Cupertino Division. He was the 2100A project manager. A University of Michigan graduate with a B.S. degree in science engineering (1963) and an M.S. degree in systems engineering (1967), Fred has written and taught extensively in the fields of minicomputer design and applications, digital systems design, digital laboratory equipment, and computer education. He's the author of a forthcoming book on minicomputer applications and a member of the COSINE Task Force on Digital Laboratories. He came to HP in 1969 after several years of research and teaching at the University of Michigan. He's a student of woodcrafts, bowhunting, and judo, and he drives to work in a 1934 Ford pickup which he's 'not restoring, just recycling.'

Microprogramming, ROMs, Firmware and All That

Behind the scenes in HP's new mini is a sophisticated microcomputer with its own instruction repertoire and an expandable memory.

By Charles T. Leis

MODEL 2100A COMPUTER is much more than a faster, more powerful version of its predecessors. The big difference is in how its instruction set is implemented.¹ In earlier HP computers, the software instructions written by the user were decoded and executed by permanent hardware. In the new 2100A, these same instructions are decoded and executed partly by hardware and partly by 'firmware,' the latter being a microprogrammed read-only memory (ROM). While the new mini may look like other HP machines to the software programmer, behind the scenes is a sophisticated microprocessor in the control section of the computer which is running an entirely different set of instructions out of its own memory.

The microprocessor can be thought of as a special-purpose computer within a computer. The microprogram stored in the microprocessor's read-only memory controls the data transfers that cause the 2100A to execute the user's instructions. Fig. 1 gives an overview of the 2100A system and shows the relationship of the microprocessor to the other machine elements.

One big benefit of microprogrammed implementation of instructions is expandability. It takes six bipolar integrated-circuit ROM packages (Fig. 2), which together have a capacity of 256 24-bit words, to implement the basic instruction set, including extended arithmetic instructions (hardware multiply/divide, etc.). However, enough addressing and board space was built into the 2100A for 1024 24-bit words of ROM. Thus there is plenty of room for enhancements to the instruction set. Hence the name 'firmware': the ROM chips are permanently programmed when manufactured, so they can't be called software, but since the ROM can be ex-

panded to extend the instruction set, it isn't hardware in the traditional sense.

In the future HP plans to design, manufacture, and support several extensions to the 2100A instruction set. One six-package extension is already available as an option. It provides floating-point double-word arithmetic instructions. It's also expected, of course, that some users will want to design, manufacture, and support their own extensions of the instruction set.

A Special-Purpose Emulator

If all user instructions were decoded and executed at the micro level, the 2100A computer would be a general-purpose emulator. In fact, user instructions are decoded by hardware and the required data transfers are implemented at the micro level. The partitioning between hardware and firmware is optimized for the 2116 instruction set. Thus the 2100A can be thought of as a special-purpose emulator, optimized to emulate the earlier HP 2116, 2115, and 2114 computers.

These earlier machines use 16-bit instruction and data words. There are two accumulator registers. The instruction set consists of three groups: memory reference instructions, register reference instructions, and input/output instructions. With additional plug-in boards, double-length instructions and hardware multiply/divide (extended arithmetic) can be executed by these machines. There are four machine phases or states: fetch, execute, indirect, and interrupt. There is an automatic phase change—from fetch to execute, for example—after each memory cycle (1.6 microseconds in the 2116) except when executing extended arithmetic instructions.




Fig. 1. Block diagram of Model 2100A Computer. The micro-processor is a computer within a computer, running micro-programs stored in its read-only memory (ROM). There's a microprogram for each user instruction.




Fig. 2. Basic 2100A ROM consists of six packages, each containing one 1024-bit bipolar integrated-circuit chip. (On this prototype board the ROM packages are mounted in sockets.) Empty spaces on the boards can be loaded with three more six-package ROM groups to extend the instruction set. A six-package floating-point option is now available.

The microprocessor in the 2100A was designed to complete any of the four machine phases in one 980 ns memory cycle, except when executing the increment-memory instruction or the extended arithmetic group. This in effect optimizes the machine for executing memory reference instructions. Memory cycles are tightly packed, that is, there is no idle time between one memory cycle and the next, so the machine gets the most out of each cycle.

Optimizing for memory use was justified by research into 'standard' instruction mixes, which revealed that memory reference instructions can be expected about 80% of the time.^{2,3} Since the 2100A cycle time is about 40% faster than that of earlier machines and memory is fully used 80% of the time the new machine runs 'standard' instruction mixes roughly 32% faster. Extended arithmetic instructions also run faster in the 2100A, and the tightly packed memory cycles give a very high DMA transfer rate.

There were some tradeoffs. Register reference instructions, for example, which were single phase instructions in the earlier machines, now take two phases. This means it takes about 23% longer to execute these instructions. But while some applications use these instructions extensively, the 'standard' or average variety of program uses relatively few of them.

The Microprocessor

When the RUN button on the 2100A front-panel is pressed, it is the inner computer within the computer, the microprocessor, that is started. The

microprocessor fetches and executes microinstructions out of ROM at the rate of one every clock period of 196 ns. Each microinstruction is 24 bits long. The basic clock period of 196 ns is determined by data path propagation delays.

When running a user program the microprocessor, under control of the microprogram in ROM, fetches a 16-bit word from main (core) memory, stores it in the instruction register, then does a microjump to a routine in the ROM microprogram that will enable the data paths and special hardware needed to execute this user instruction. The microprocessor has access to six more 16-bit registers than the user; this gives it a great deal of flexibility in executing user instructions.

The microprocessor controls the timing of the phase changes (e.g., from fetch to execute) so that once in the execute phase, for example, the microprocessor can do complicated operations taking many microseconds. Only when these are complete will it allow a return to the fetch phase. It is this semisynchronous feature which made it easy for the extended arithmetic instructions to be included as a standard feature of the 2100A. All that was necessary was a sophisticated enough set of data path controls, or micro-operations.

Fig. 3 gives a closeup view of the microprocessor. During each basic clock period, one microinstruction is executed. Meanwhile the ROM address register (RAR) is incremented and the ROM contents are presented to the ROM instruction register (RIR)

to be executed during the next clock period. Normal indexing through the ROM locations is altered in case of a microjump, in which case the RAR is parallel-loaded with a new address. A new address is also generated for the RAR when changing phases. In this case the new address points to the beginning of the microroutine to do what is required of the phase.

An Example

Here's an example to show what happens in ROM as the 2100A executes a typical 16-bit user instruction.

Assume the user instruction is ADB (M), a typical memory reference instruction which adds the contents of memory address M to the contents of the B accumulator and stores the result in B. The instruction format is

Bit Position	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
Contents	D/I		Instruction				Z/C		Memory Address							

D/I = direct or indirect address
Z/C = zero page or current page

For a direct address this instruction takes two machine phases to execute, a fetch phase and an execute phase. An indirect address would require at least one indirect phase between these two. Each phase takes one 980 ns memory cycle, and in each memory cycle are five 196 ns clock periods. Each clock period, one 24-bit ROM instruction is executed. Here's the sequence for a direct address.




Fig. 3. The microprocessor fetches a 16-bit user instruction from core memory, stores it in the instruction register, then jumps to a routine in the ROM microprogram that will cause the 2100A to execute the user instruction.

Clock Period	ROM Address	ROM Instruction Operation
Fetch Phase		
1	000	Load Memory Address Register with Contents of Program Counter (P); Start Read/Write
2	001	Wait
3	001	Wait
4	001	Load Instruction Register with Contents of Memory Location Specified in Memory Address Register
5	002	Load Scratch Pad* Register 1 with Memory Address (M); Check Whether Address is an Accumulator Register (A or B); Change Phases.
Execute Phase		
1	144	Load Memory Address Register with Contents of Scratch Pad Register 1; Start Read/Write; Increment Program Counter ($P \leftarrow P + 1$)
2	145	Wait
3	145	Wait
4	145	Load Scratch Pad* Register 2 with Contents of Memory Location Specified in Memory Address Register
5	146	Add Contents of B Accumulator to Contents of Scratch Pad Register 2; Store Result in B; If Carry Out, Set E Register; Change Phases.

Older HP computers executing the same instruction would perform basically the same operations. The difference is that in these hard-wired-logic implementations, time and condition signals are generated throughout the machine and are combined to form control functions. In a microprogrammed machine like the 2100A, the generation of control lines is concentrated in the ROM.

Powerful Microinstructions

The microprogramming language that's used to write the microprogram stored in the 2100A's ROM is roughly comparable to the assembly language of earlier machines. Fig. 4 shows the format of the microinstructions. There are six fields, each of which is fully encoded. With few exceptions the operations specified within each field produce only one action (e.g., set up the AND function in the logic unit).

A typical instruction for the microprocessor might be

R-Bus	S-Bus	Function	Store	Special	Skip
A	S1	AND	A	-	TBZ

This would put the contents of the A accumulator on the R-bus, the contents of scratch pad register #1 on the S-bus, AND the two, and store the result in the A accumulator. If the result were zero the

* Four 'scratch pad' registers are among the six extra 16-bit registers the microprocessor has access to.




Fig. 4. 24-bit, three address microinstructions are more powerful than the 16-bit user instruction set they are designed to implement. Format of microinstructions makes microprograms easy to write.

next microinstruction would be skipped.

Four of the six fields—the R-bus, S-bus, function, and store fields, control primarily data flow. The skip field controls the sequencing of microinstructions, although the function field shares some of this responsibility.

The special field is used to decode and execute the register reference instructions. This group had to have a special hardware implementation, the reason being that these instructions have some bit-per-function encoding, and to decode them in ROM would have required too much ROM and too much time. Most of the logic for these instructions is external to the microprocessor, and commands from the special field enable this hardware. Memory read/write operations are also initiated in the special field, as are arithmetic unit output shift signals and a 'reverse skip sense.' The special field provides the principal hardware control functions of the 2100A.

Notice that three of the microinstruction fields contain addresses—the R-bus, S-bus, and store fields. This, plus the fact that the microprocessor has access to six more registers than the user instructions, makes the microprocessor language more powerful than the assembly language it is designed to emulate. A measure of this power was evident when the floating-point package was written for the optional second 256-word ROM module. This package adds double-word floating-point add, subtract, multiply, divide, fix, and float instructions to the basic set. The routines are comparable with their software counterparts but require fewer words and execute 20 to 30 times faster.

Vertical Microprogramming

The 2100A microinstruction format has sometimes been called 'vertical microprogramming,' as opposed to 'horizontal microprogramming,' which would feature a much longer instruction word with each bit causing some specific action to occur. Ver-

ROM Simulator Tests Computer in Five Minutes

Production and field testing of the 2100A minicomputer posed some new problems and represented an opportunity to develop some unique test methods. To test earlier models, technicians used the front-panel switch register to check out the registers and small portions of the hardware when the computer was first turned on. Once it was verified that some of the hardware was operational, software diagnostics were used to verify that the entire computer was operational. With the 2100A, on the other hand, much of the hardware must be operational to exercise even simple commands from the front panel. Software instructions are decoded into read-only-memory (ROM) addresses and then the ROM instructions exercise the hardware. Thus the key to exercising the hardware in the 2100A is the ROM output lines. The 2100A was designed so the ROM can be disabled and the ROM output lines can be easily accessed.


A ROM simulator was designed to use the ROM output lines to exercise the hardware of the 2100A. The ROM simulator contains a bipolar random-access memory (RAM) that operates much as the ROM does in the 2100A. A diagnostic microprogram, which is stored in the RAM, was written to exercise small segments of the 2100A hardware initially, and additional segments of the hardware as progress is made through the program. This allows much smaller segments of hardware to be exercised at any one time and reduces the amount of time it takes to isolate any troubles that occur.

The RAM in the ROM simulator is small (32 words). The diagnostic microprogram is written in 32-word segments and each segment is completely independent. Segments may be added at any time if changes occur or may be altered by means of a switch register. Each segment is loaded into the RAM through an interface in the ROM Simulator. In the HP computer manufacturing facility is a master computer which has the microdiagnostic and software diagnostic programs stored in its memory. The ROM simulator interfaces to this master computer to obtain the microprograms. Another model of the ROM simulator, designed for field service, obtains its microprograms via a photo reader.

During the execution of the diagnostic microprogram, the computer clock is supplied by the ROM simulator. The operator may run the diagnostic microprogram at normal computer speed or at a reduced rate. He may stop the clock and manually change its phase. Other features of the ROM simulator that aid the technician if troubleshooting is necessary include:

- a ROM address display indicating the ROM address to be executed next
- a ROM instruction display indicating the state of the ROM output lines
- a sync-on-ROM address that provides a sync pulse for an oscilloscope when the ROM address coincides with the address selected
- a break-on-ROM address that stops the clock to the computer when the ROM address coincides with the address selected
- a loop-on-segment mode that allows the operator to loop endlessly on the segment of the diagnostic microprogram that is stored in the RAM
- a segment number display that displays the current segment of the diagnostic microprogram that is stored in the RAM
- a self-check feature to assure that the ROM simulator is operational.

The field service model is designed to give the field service technician all of the tools necessary to diagnose a malfunctioning computer as quickly as possible. This model has all of the features of the manufacturing model plus a photo reader for loading the RAM. With this model a service technician can exercise all the logic in the 2100A in less than five minutes. The field service model is packaged for easy carrying (see photo). In the package are all the necessary service tools except for an oscilloscope.



tical microprogramming has two disadvantages: decoding hardware is required for each field, and execution of specific algorithms usually takes longer. These are offset in the case of the 2100A for two reasons. First, it takes only two relatively inexpensive integrated-circuit packages to decode the output of each ROM package. Second, any further increase in speed would be fruitless since the ROM would just be idle more of the time, waiting for the main memory.

The main advantage of the vertical approach is that it makes it very easy to write microprograms.

A bit-per-function ROM output which would simplify the hardware would make ROM programs more difficult to write. The expectation that some users would want to write their own microprograms seemed to require an approach that would allow someone unfamiliar with the hardware to write useful microcode, and the vertical approach satisfies this requirement. Microdiagnostics are also easy to produce and debugging is a matter of using the right tool (see above). Microinstructions are relatively independent of user instructions (i.e., very few microinstructions refer to the 2116 in-


struction set), so the amount of machine knowledge needed to write correct microcode is a minimum.

Acknowledgments

The microprocessor circuit design was done by Don Jenkins. Gary Mueller designed the I/O control and DMA hardware. Thanks go to Pat Mulreany for his work on the backplane. The ROM simulators are the work of Cle Riggins and his group. Technical support was provided by Ed Holmes. Arne Bergh was a source of many valuable suggestions during the early stages of design.

References

1. C. T. Leis, 'Microprogramming Features of a 16-Bit Mini,' Workshop on Microprogramming, University of California at Santa Cruz, September 13-14, 1971.
2. J. C. Gibson, 'The Gibson Mix,' IBM Technical Report TR 00.2043, June 18, 1970.
3. C. C. Church, 'Computer Instruction Repertoire — Time for a Change,' AFIPS Conference Proceedings, SJCC 1970, Volume 23, pages 343-349.



Charles T. Leis

Chuck Leis was the system designer of the 2100A. He's been with HP's Cupertino Division since 1969. His B.S. and M.S. degrees in electrical engineering came from the University of Florida in 1968 and 1969. Chuck is a member of IEEE and two honor societies, and has authored a couple of professional papers, one on magnetics and one on microprogramming. He's a four-year veteran of the U.S. Air Force.

SPECIFICATIONS HP Model 2100A Computer

MEMORY

TYPE: Folded planar core.
WORD SIZE: 16 bits with 17th parity bit.
PAGE SIZE: 1024 words.
DIRECT ADDRESSING: 2 pages.
INDIRECT ADDRESSING: All pages.
MODULE SIZES: 4K and 8K word memory modules provide 4, 8, 12, 16, 24 and 32K configurations all in the 2100A main-frame without additional power supply or cabinetry.
CYCLE TIME: 980 ns.
LOADER PROTECTION: Switch protects last 64 words.

REGISTERS

ACCUMULATORS: Two (A and B), 16 bits each. Directly addressable.
MEMORY CONTROL: Three (T, P, M), 16 bits each.
SUPPLEMENTARY: Two (Overflow and Extend), one bit each.
MANUAL DATA: One 16-bit switch register.

INSTRUCTION EXECUTION TIMES

MEMORY REFERENCE GROUP (14 total): 1.96 μ s (SZ: 2.94 μ s).
REGISTER REFERENCE GROUP (43 total): 1.96 μ s.
INPUT/OUTPUT GROUP (13 total): 1.96 μ s.
EXTENDED ARITHMETIC GROUP (10 total)
MULTIPLY: 10.7 μ s.
DIVIDE: 16.7 μ s.
DOUBLE LOAD: 5.9 μ s.
DOUBLE STORE: 5.9 μ s.

SHIFT/ROTATE: 2.9 to 7.8 μ s dependent on type and length.
INDIRECT ADDRESSING: 980 ns/level (1.96 μ s/level in Extended Arithmetic Group).

FLOATING POINT HARDWARE EXECUTION TIMES (Optional)

	Minimum	Maximum
Add	23.5 μ s	59.8 μ s
Subtract	24.5 μ s	60.8 μ s
Multiply	33.3 μ s	41.1 μ s
Divide	51.9 μ s	55.9 μ s
Fix	5.9 μ s	8.8 μ s
Float	9.8 μ s	24.5 μ s

INPUT/OUTPUT

MULTILEVEL AUTOMATIC PRIORITY INTERRUPT
Determined by interface location.
I/O CHANNELS IN 2100A COMPUTER: 14.
I/O CHANNELS IN 2100A COMPUTER PLUS 2155A EXTENDER: 45.
I/O COMPATIBILITY: HP 2114/2115/2116.

DIRECT MEMORY ACCESS (Optional)

NUMBER OF CHANNELS: 2.
REGISTERS PER CHANNEL: Word Count Register, Address Register.
MAXIMUM BLOCK SIZE: 32,768 words.
ASSIGNABLE: To any I/O channels.
TRANSFER RATE: Greater than 1 million words per second.

POWER FAIL INTERRUPT WITH AUTOMATIC RESTART (Standard)

PRIORITY: Highest priority interrupt.

MEMORY PARITY CHECK WITH INTERRUPT (Standard)

PRIORITY: Second highest priority interrupt (shared with Memory protect).

MEMORY PROTECT (Standard)

PRIORITY: Second highest priority interrupt (shared with Memory Parity).

PHYSICAL*

DIMENSIONAL
WIDTH: 16 3/4 in with adaptors for mounting in 19 in. rack.
HEIGHT: 12 1/4 in. (rack mounted).
DEPTH: 2100A—26 in. (23 in. behind rack mounting ears).
2155A—23 1/2 in. (23 in. behind rack mounting ears).

WEIGHT

MINIMUM: 91 pounds (41 kg).
MAXIMUM: 111 pounds (50 kg).

ELECTRICAL*

POWER REQUIREMENTS:

115V/230V \pm 10%.
47.5 to 66 Hz.
800 watts maximum.

ENVIRONMENTAL*

OPERATING TEMPERATURE: 0° to 55°C (+32° to +131°F).
RELATIVE HUMIDITY: To 95% at 40°C (104°F).

VENTILATION

INTAKE: Rear panel.
EXHAUST: Sides of front panel and cabinet.
HEAT DISSIPATION: 2700 BTU/hr maximum.

PRICES:

2100A COMPUTER: \$3,750.
Without memory modules. Includes operator's panel, extended arithmetic instructions, memory parity check with interrupt, power fail interrupt with automatic restart, memory protect, 14 I/O locations. Memory module(s) must be specified.
4,096 Word Memory, \$3,500.
8,192 Word Memory, \$7,000.
16,384 Word Memory, \$14,000.
32,768 Word Memory, \$28,000.
Multiplexed I/O, \$1,000.
12901A Floating Point Hardware, \$2,500.
2155A I/O EXTENDER, \$3,500.

MANUFACTURING DIVISION: CUPERTINO DIVISION
11000 Wolfe Road
Cupertino, California 95014

*Except as noted, applies to both the 2100A Computer and the 2155A I/O Extender.



A Lot of Memory in a Small Space

This new core memory is fast, rugged, reliable, and small enough to expand to 32K words inside the 2100A mainframe.

By Robert J. Frankenberg

A MINICOMPUTER LIKE THE 2100A needs a memory with special qualities if it's to be able to cope with the diversity of applications awaiting it. For one thing, it can't expect air conditioning, so its memory must be designed to operate flawlessly in relatively harsh environments. Many minicomputer applications have severe space and weight limitations. Others start with relatively small memory needs which grow as time goes by. This means a memory should be small, so the mini can start small and stay small as its memory expands. Low cost is important because a mini by definition is a computer that can do a job when a larger, more expensive machine wouldn't be economically feasible. And for the many applications which call for fast computations, a good memory must be fast, to allow quick access to the data it contains.

Even more important than all these qualities is reliability. Reliability stands above all other characteristics as the prime attribute a computer memory must possess. Related to reliability is the question of volatility, or what happens when a power failure occurs. Many applications don't require a non-volatile memory, but in others it's vital that the data stored in memory not be lost when the power fails. Non-volatility, therefore, plus reliability, high speed, low cost, small physical size, and ability to withstand harsh environments are all qualities a memory must have to be useful in the widest possible range of applications.

All these considerations made it apparent that core memory was the best kind of memory for a general-purpose minicomputer like the 2100A. Despite many impressive advances in semiconductor memory devices, the problem of volatility hasn't yet been solved. Also, the bits-per-package density

of semiconductor devices isn't high enough yet to permit a semiconductor memory to be made as small as a core memory of equivalent capacity.

For the 2100A a three-wire three-dimensional core memory system was designed to have all the desirable characteristics in a good balance. Fig. 1 shows its organization. The memory is expandable from 4096 words (4K) to 32,768 words (32K). In a 4K or 8K configuration the memory fits on five printed circuit boards. This increases to 12 boards for a 32K memory. The boards are installed simply by plugging them into existing slots in the 2100A mainframe and configuring the memory data control board. Cycle time of the memory is 980 nanoseconds.

Corestack and Sense Amplifier Volume Reduced

The huge volume usually occupied by the core array and sense amplifiers was drastically reduced in the 2100A by the use of HP-developed planar folded corestacks and hybrid sense amplifiers.

The memory corestack design was prompted initially by evolutionary developments in core memory technology. Smaller cores have become more economical, thanks to higher demands and improved manufacturing techniques. Continuous wiring capabilities have been developed that both increase reliability and decrease costs by decreasing solder connections 400:1 over older methods. The three-wire three-dimensional organization reduces the amount of drive and sense circuitry needed for the memory system. These trends established a direction for the 2100A corestack development, and other needs and constraints dictated the size and speed necessary.




Fig. 1. Core memory for 2100A Computer fits entirely on printed-circuit boards which plug into existing slots in mainframe. Memory expands from 4K on five boards to 32K on twelve boards.

One constraint was to maintain the same size printed-circuit boards used in earlier HP computers. This was desirable for reasons of economy. The problem then was how to take advantage of continuous wiring and reduce the area of an 8K array to fit on a printed-circuit board. The solution was to fold an initially planar core array so it could be accommodated on a board of the given size (see Fig. 2). Because the boards have only a limited number




Fig. 2. 8K folded-planar corestack with sense amplifiers and diode decoding matrix occupy only one printed-circuit board. Cores are 20-mil wide-temperature-range type.

of output pins, the sense amplifiers were put on the same board with the corestack.

The x-y decoding matrix for the corestack also required considerable compression. This was accomplished by using conventional diodes, fabricating them into 16-unit modules, and standing them upright to form a compact matrix.

Memory Data and Control

The memory data control board (Fig. 3) occupies only one-third the space required for memory control in earlier HP computers. Commercial availability of recently developed integrated parity checker-generators, quad D flip-flops, synchronous binary up-down counters, and eight-line two-way multiplexers made it possible to put a data register, a parity checker, and an S-bus/sense amplifier multiplexer in a very small area of the control board. New integrated circuits that don't source current even intermittently when the supply voltages are turned on and off made it possible to use a compact and simple power-fail protection circuit. Repartitioning of memory control functions reduced the number of board interconnections required and thus overcame the output pin limitations of a single board. The space needed for memory timing circuits was reduced by using recently developed precision delay lines in conjunction with precise timing logic. The timing circuits automatically compensate for the difference in the speeds of 4K and 8K core-stacks.




Fig. 3. Memory data control circuitry for up to 32K memory occupies only one board, one-third as many as in earlier HP computers. Virtually all circuits were redesigned.

8K/16K Inhibit Driver

The 2100A inhibit driver can be in either a 16K or an 8K configuration, depending on how the inhibit-driver board is loaded. The inhibit circuit is a simple circuit that requires very few parts. It is repeated 68 times on a 16K board and 34 times on an 8K board. An unusual aspect of the circuit is that one collector load resistor is shared by as many as eight transistors. This allowed the reduction of one whole board in a 32K machine.


Wide-Range Cores

The 2100A memory system is compensated for temperature changes and uses wide-temperature-range cores. Temperature sensing is done on the resistor load board and thus represents an average of the memory and the ambient temperatures. The sensed value of the temperature is sent to the power supply which adjusts itself in accordance with the core's requirements.

Acknowledgments

The compactness of the printed-circuit boards in the new memory is the work of Cupertino Division's fine printed-circuit layout group and of the HP corporate printed-circuit fabrication facilities. Joe Olkowski was the principal developer of the folded planar corestack. I would also like to acknowledge with gratitude the hard work and advice of Gordon Goodrich, Wayne Gladwin, Tak Watan-

abe, and Dennis Wong, and the help and encouragement of Fred Coury. Although too numerous to name here, many HP production people made indispensable contributions to the development of the memory. 



Robert J. Frankenberg

Bob Frankenberg was responsible for the design of the memory system in the 2100A. He's filed two patent applications as a result of that work. Bob came to HP in 1969 after four years in the U.S. Air Force working on computer-controlled radar systems. Before the 2100A he was involved in the design of the 2116C computer and in production testing of various computers and systems. Bob holds an A.A. degree in electronic technology and is working for his B.S.E.E. He likes to fish, he dabbles in oil painting, and he restores antique cars.

A Bantam Power Supply for a Minicomputer

How do you get a 500 watt power supply with several regulated dc output voltages into half a cubic foot of space?

By Richard D. Crawford and
Gregory Justice

A MINICOMPUTER WHOSE POWER SUPPLY IS BIGGER than the rest of the computer is a little like a sports car that needs a trailer to carry its gasoline. Yet today's minicomputers are faced with this problem, now that MSI, LSI, and smaller memories have helped shrink everything but their power supplies. We decided that the 2100A Computer would have to have a bantam power supply appropriate to the compactness of the rest of its circuitry. Small size, however, was not to be achieved at the sacrifice of reliability, good performance, a high degree of protection for the computer and its stored data, or low acoustic noise.

For a start, we aimed for better than 70% efficiency. High efficiency means low heat dissipation and therefore small heat sinks.

Second, we went to higher voltages to reduce the size of the filter capacitors. In the event of an ac power failure the supply needs sufficient energy storage for the computer to complete an orderly shutdown. This energy is stored in the filter capacitors. Energy stored is equal to $\frac{1}{2} CV^2$, where C is capacitance and V is voltage. Volume, however, is generally proportional to CV. Therefore, higher voltages mean smaller capacitors for a given amount of energy.

Next, there was no room for a 60 Hz power transformer. This meant we had to find alternate ways of providing ac power line isolation and easy conversion to 115 V or 230 V operation. Isolation was achieved by using inverters, which are dc-to-ac converters with transformer-coupled outputs. To regulate the dc output voltages of the power supply, we put a *preregulator* ahead of the inverters. The preregulator also provides energy storage and 115 V or 230 V operation, and helps make life easier for the inverters. Fig. 1 is a block diagram of the supply.

Preregulator

How the preregulator works is illustrated in Fig. 2. The preregulator is a choke-input phase-controlled SCR circuit which functions as a conventional controlled rectifier for a 230 V ac input voltage, or as a voltage doubler for a 115 V ac input voltage. Diodes D_1 and D_2 are 'free-wheeling' diodes which carry the current flowing in L_1 when SCR_1 or SCR_2 are not conducting.

For 230 V operation, switch S1 is open. If SCR_1 conducts for 180° and SCR_2 conducts for the other 180° , the voltages will be as shown in A of Fig. 2.

For 115 V operation, switch S1 is closed. If there were no mutual coupling ($M = 0$) between the two halves of choke L_1 , the voltage waveforms would be the same as in A, but the magnitudes of the voltages would only be half as large. Actually, however, there is tight coupling ($M = 1$) between the two halves of L_1 , and the waveforms are as in B of Fig. 2. The mutual coupling and closing of S1 has filled in the missing half sinusoids, with the result that the dc voltages V_+ and V_- are the same for either 115 V or 230 V ac line voltage.

V_+ and V_- can be varied for either the 115 V or the 230 V case by phase control of the SCR conduction angle. If the conduction angles are less than 180° the output voltages will be smaller than those shown in Fig. 2. The actual regulated dc output voltage is approximately 160 V.

Inverters

The main job of the inverters is to provide isolation from the ac line by means of transformers. A problem here was the choice of inverter frequency. Too low a frequency means big transformers. Too high a frequency causes the switching losses to grow outlandishly. Audio-frequency transformers




Fig. 1. Power supply for 2100A Computer provides 500 watts at seven regulated output voltages, yet occupies only 800 in³. Computer, software, and supply are fully protected.

are plagued by acoustic noise. We chose 800 Hz as a good compromise. There is some acoustic noise, but it's quieter than the gentle whirring of the fans, so it isn't objectionable.

As Fig. 1 shows, there are two inverters. They generate square waves with an overlap of one-quarter period. During the finite switching interval when one inverter is changing states, the other is carrying the full load. This prevents the full load current from flowing through the equivalent series resistance of the output capacitors during the switching interval, thereby minimizing ripple voltage. What's more, the inverters aren't switching under load, and this means better reliability. Another advantage is that the switching interval can be longer; we made use of the longer switching interval to prevent overlapping conduction of the transistors within each inverter.

Outputs

The square waves generated by the inverters feed several sets of rectifying diodes which produce these outputs:

+30 V at 100 mA	= 3 watts maximum
+12 V at 3 A	= 36 W
+4.85 V at 50 A	= 243 W
-2 V at 23 A	= 46 W
-12 V at 3 A	= 36 W

Regulation for all these outputs is provided by the preregulator. Space is saved because separate regulators for each output aren't needed.

Core memory requires +20 V and -20 V supply voltages with a higher degree of regulation, so additional regulators are used for these voltages. Because the optimum operating voltage of core memory varies with temperature, the +20 V and -20 V regulators are temperature compensated such that they always supply the optimum voltage to the memory. The +20 V regulator can supply up to 6A or 120 W and the -20 V regulator can supply up to 400 mA or 8 W.

Total power output capability of the 2100A power supply is therefore 492 W. We stretch it a little and call it 500 W. Volume is 800 in³.

Protective Circuits

The power supply could have been much smaller




Fig. 2. Preregulator produces regulated dc output of 160 volts for either 115 V or 230 V ac line voltage. Five of the seven power supply output voltages require no further regulation.

if we hadn't provided such complete protection for the software stored in core memory, for the hardware of the computer, and for the power supply itself. Software protection is the purpose of the Internal Power Up (IPU) and Power Up (PWU) signals shown in Fig. 1. If all is well with the power supply IPU and PWU are in a high logic state (+5 V). IPU for several 2100A computers can be connected in parallel so each computer monitors the well-being of all others. If IPU fails even momentarily for any reason, PWU drops to a low logic state for 0.5 second. When this happens the computer starts its shutdown procedure.

IPU will go low if the ac line voltage drops below 100 (or 200) volts, if any of the output voltages is below preset limits, or if the internal temperature-sensing switches open. The power supply will automatically recover from any of these conditions as soon as it is corrected.


Hardware protection is provided by overvoltage sensing. If any supply voltage exceeds preset limits, not only does IPU go low, but also all output voltages are very rapidly discharged (crowbarred). The power supply will not recover automatically from

this state; the presumption is that some part has failed, causing a hazardous condition for the integrated circuits in the computer.

Protection for the power supply itself is provided by output current limits and fuses at key points.

Acknowledgments

So what do we have? A bantam power supply. 500 watts in 800 cubic inches. Greater than 70% efficiency. Seven different output voltages. Regulation adequate for a computer. Energy storage. Protective circuits. A minicomputer of appropriate size and weight. We would like to express gratitude to all those who contributed to the success of the supply. Barney Oliver and Dave Hilbiber provided many creative solutions to the problems we encountered. Zvonko Fazarinc and Kent Stockwell provided us with valuable information from their computer simulation of the power conversion circuits. Dick Van Brunt designed the protection circuits. Bob Pierce and Bob McCaw did an excellent job of mechanical design. Technician support was provided by Frank Lee, Mike Farrell, Bill Newell, Ken Check, Joe Dixon and George Canfield. John Grimaldi ably directed the production design phase and provided many valuable suggestions.



Richard D. Crawford

Dick Crawford (left) is an 11-year HP veteran, a design engineer with HP Laboratories, and one of the principal developers of the 2100A power supply. He received his B.S. degree in electrical engineering from California Polytechnic Institute in 1960 and his M.S.E.E. from Stanford in 1969. He's a hi-fi bug and an amateur enologist.

Gregory Justice

Greg Justice left college after two years and taught himself electronics. Now a design engineer with HP Laboratories, he's been with HP for 15 years and has helped design a variety of advanced instruments. His latest project was the 2100A power supply. He holds several patents and has filed for half a dozen more.

UTC Time Scale To Change In 1972

On January 1, 1972, the time scale known as UTC (Coordinated Universal Time, also called GMT, or Greenwich Mean Time), the most commonly used time scale around the world, will be slightly altered. The present rate offset from the atomic time scale will be eliminated, and step adjustments in time will be made in increments of 1 second instead of 0.1 second. Thus, after January 1, 1972, UTC will accumulate time at the same rate as Atomic Time (AT), except that whole-second step adjustments, called leap-seconds, will be made as needed to maintain approximate agreement with a time scale based on the earth's rotation. UTC is broadcast by many standard time transmission stations, including WWV, operated by the Boulder Laboratories of the National Bureau of Standards, U.S. Department of Commerce, and is the basis for standard time in the U.S.

UTC, at present, is a compromise time scale arrived at by international agreement through the International Radio Consultative Committee. It is a time scale that has been adjusted, or offset in rate, from an atomic scale to approximately follow a time scale based on the rotation of the earth. The offset for the last few years has consisted of a continuous retardation of 300 parts in 10^{10} . In addition, step adjustments of 0.1 second are made whenever needed to keep UTC within 0.1 second of a time scale based on the rotation of the earth.

The need for the change arises from the fact that atomic standards of time and frequency, e.g., a Hewlett-Packard Cesium Beam Frequency Standard, provide a time scale that is much more uniform than the scale provided by the earth's rotation on its axis. The atomic time scale was chosen to agree in rate with the time scale based on the orbital motion of the earth about the sun in the year 1900. The time scale based on the present rotation rate of the earth on its axis differs from this by about 300 parts in 10^{10} or 1 second per year. That much variation in the time scale cannot be tolerated by many technical and scientific projects, so atomic clocks, which run at a much more uniform rate, are used to generate our time scales today.

However, even though the time scales used officially throughout the world are maintained by atomic sources, it is convenient to coordinate atomic time with the time indicated by the rotation of the earth. This is done so that navigators and others can use clocks to determine their position on earth.

At present this coordination, between the varying rate of the earth's rotation and the constant rate of the atomic clock, is done by slowing down the atomic clock rate by a certain fraction, and by subtracting or adding fractions of a second several times a year. All these adjustments can cause errors for the scientists and technicians involved.

The new system will be much simpler; the atomic clock rate will not be slowed down at all, and instead of adding or subtracting a fraction of a second every few months, everyone will add or subtract a whole second once in 6 to 12 months.

Of course clocks which characteristically lose or gain more than one second in a year won't need these tiny adjustments!

The standard time and frequency radio stations maintained by various countries will cooperate with the International Time Bureau in broadcasting the new time scale and in making the adjustments simultaneously. The adjustments should not be needed more than once in 6 to 12 months, and will be made preferably on December 31 or June 30. To provide a traditional service to navigators and astronomers, who need earth-related time, these stations will broadcast information concerning the difference between the transmitted time and the astronomical time that is relevant for navigation. The difference will not be more than 0.7 second, and will probably be broadcast with a resolution of 0.1 second.

In the U.S., therefore, the NBS standard broadcast services of WWV, WWVH, WWVL, and WWVB will all have zero offset from Atomic Time in their carrier and modulation frequencies and time signals. At 00 hours on January 1, 1972, UTC will be reset a fraction of a second, sufficient to give the new UTC scale an initial difference of an integral number of seconds (probably 10,000 seconds late) with respect to AT as maintained by the International Time Bureau. UTC is now about 9 seconds late, and during the next year the difference will probably grow to about 10 seconds; thus, the reset should be only a few hundred milliseconds. Thereafter, the difference between UTC and AT will always be an integral number of seconds.

Here's how the time-scale change will affect users of HP counters and frequency standards.

All HP counters and frequency standards will be calibrated to the new UTC offset when shipped from the factory after January 1, 1972. This is significant to users who are counting

All HP counters and frequency standards will be calibrated to the new UTC offset when shipped from the factory after January 1, 1972. This is significant to users who are counting frequencies to 3×10^{-8} accuracy or better or are using HP frequency standards as calibration references and clocks.

The 117A VLF Receiver/Comparator user will have little problem with the change. The received signal (60 kHz from NBS, Ft. Collins) has been on zero offset (to within better than 1×10^{-11} for several years. In order to calibrate to the present offset of -300×10^{-10} , a Time Scale Translator was offered as an option. The user can just switch off (internal slide switch) or remove the translator. He then can adjust his house standard for minimum frequency difference relative to the received frequency. He will then be on the new UTC time scale frequency.

Users with the 5061A Cesium Beam Standard and the 5065A Rubidium Standard can easily make the frequency change. The operation manuals provide instructions and no new parts are necessary.

5060A Cesium Beam Standard users will need to plan ahead. A plug-in crystal (HP #05060-8002, \$100 each) is required to make the change. The manual has all the necessary instructions so that the user can easily do it himself. A set of instructions will also be sent with each crystal.

HEWLETT-PACKARD JOURNAL OCTOBER 1971 Volume 23 • Number 2

FROM THE LABORATORIES OF HEWLETT-PACKARD COMPANY 1501 PAGE MILL ROAD, PALO ALTO, CALIFORNIA 94304 U.S.A.
Hewlett-Packard S.A. 1217 Meyrin - Geneva, Switzerland • Yokogawa-Hewlett-Packard Ltd. Shibuya-Ku, Tokyo 151 Japan
Editor: R. H. Snyder Editorial Board: R. P. Dolan, H. L. Roberts, L. D. Shergalis Art Director: Arvid A. Danielson Assistant: Maridel Jordan