APPLYING MICROPROCESSORS

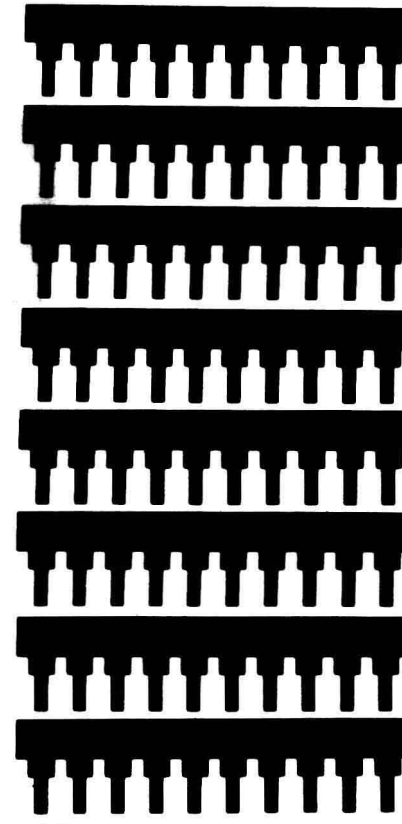

Edited by Laurence Altman and
Stephen E. Scrupski, Senior Editors,
Electronics

New hardware, software and applications

Electronics® Magazine Book Series

ELECTRONICS BOOK SERIES

Also published by *Electronics*
- **Microprocessors**
- **Large scale integration**
- **Data communications**

Library of Congress Catalog Card No. 76-30685

McGraw-Hill Publications Co.
1221 Avenue of the Americas
New York, New York 10020

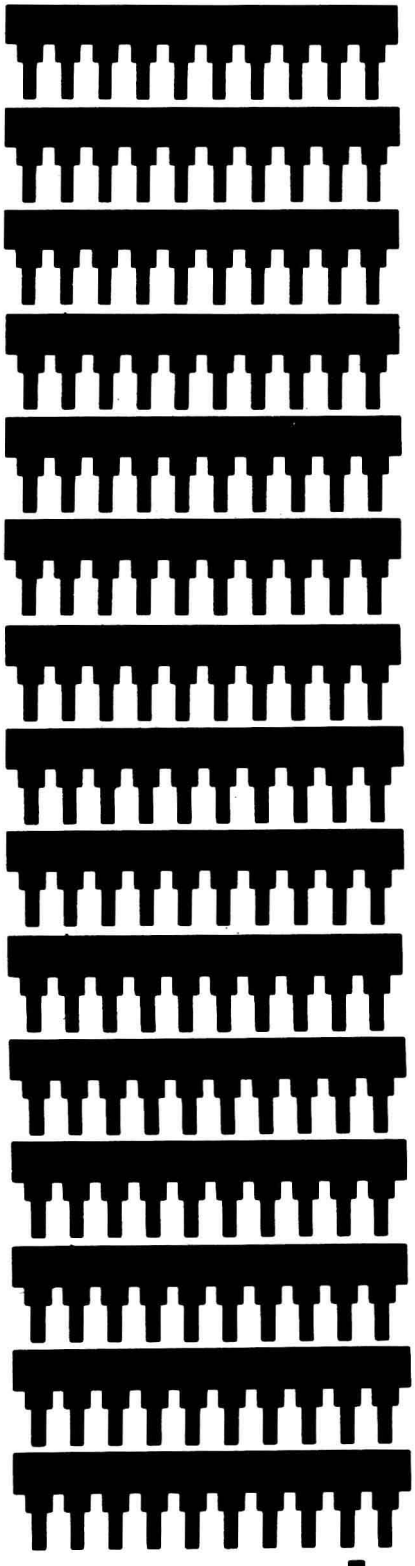

introduction

Microprocessors are the most versatile and powerful tools electronics engineers have ever had. Right from the start, these large-scale-integrated circuits promised to spread into many diverse applications, being both programable and—in terms of what they can do for the money—very low-cost. They are rapidly fulfilling their promise. Manufacturers of electronics equipment everywhere are using them to improve the performance of their systems, add new features, and even stimulate the development of altogether new equipment.

This volume is intended to ease the engineer's transition from old methods of electronic design to the new world of microprocessor engineering. It picks up from where *Electronics*' first book on the subject, "Microprocessors," left off. That volume, a collection of articles from the pages of *Electronics*, analyzes the inner workings of the earliest microprocessors up through such devices as the Intel 8080, Motorola 6800, and the Fairchild F8.

But much has happened since its publication just one year ago. The field is already producing second- and third-generation devices, which have larger instruction sets, operate faster, and pack truly significant amounts of memory onto the same chip as the arithmetic/logic unit. There are also new peripheral control devices that complement the central-processor chips to form complete circuit families.

Part 1 covers the hardware—the microprocessor and peripheral devices themselves—and leads off with an extensive overview of the rapidly expanding variety of devices now available to designers. Then follow descriptions of the latest generation, ranging from the Intel 8048, which puts an erasable read-only memory on the same chip as an 8-bit processor, to the Texas Instruments TMS9900, a 16-bit minicomputer-like device.

Part 2 presents the ins and outs of programing and prototype design. Surveys of the available software and development systems precede discussions of specific development systems, among them Intel's Microcomputer Development System, the first to include the in-circuit emulation capability. The computer-programer approach to designing microprocessor systems is contrasted with a hardware-oriented approach by Ed Lee, president of Pro Log Corp., and the often overlooked problem of testing microprocessors also comes up for extensive discussion.

Part 3 proceeds to the payoff—actual applications. An assortment of many short articles tells how today's most popular devices are already controlling all types of systems, from engine-temperature monitors and display terminals to blood analyzers and weighing systems.

Laurence Altman
Stephen E. Scrupski

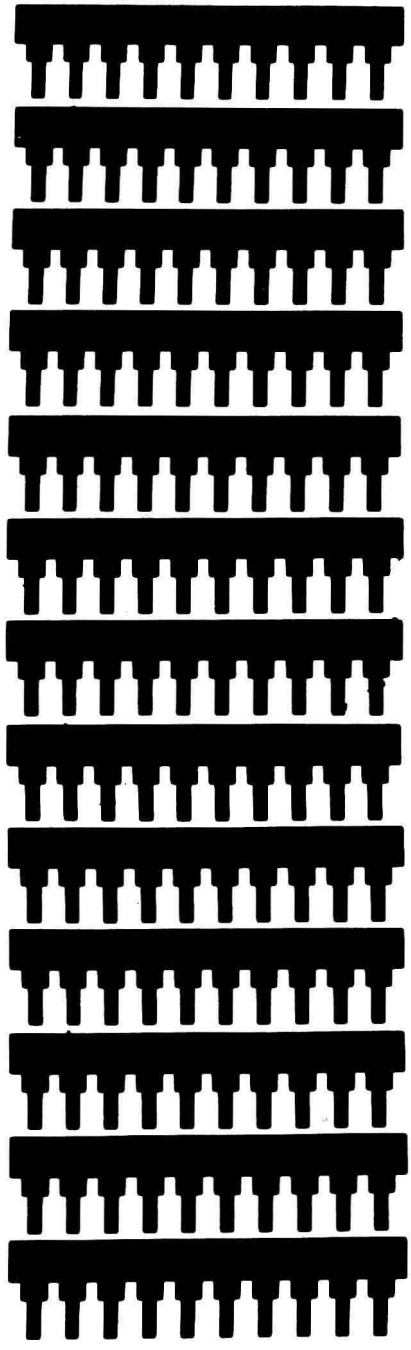

Table of contents

Part 1

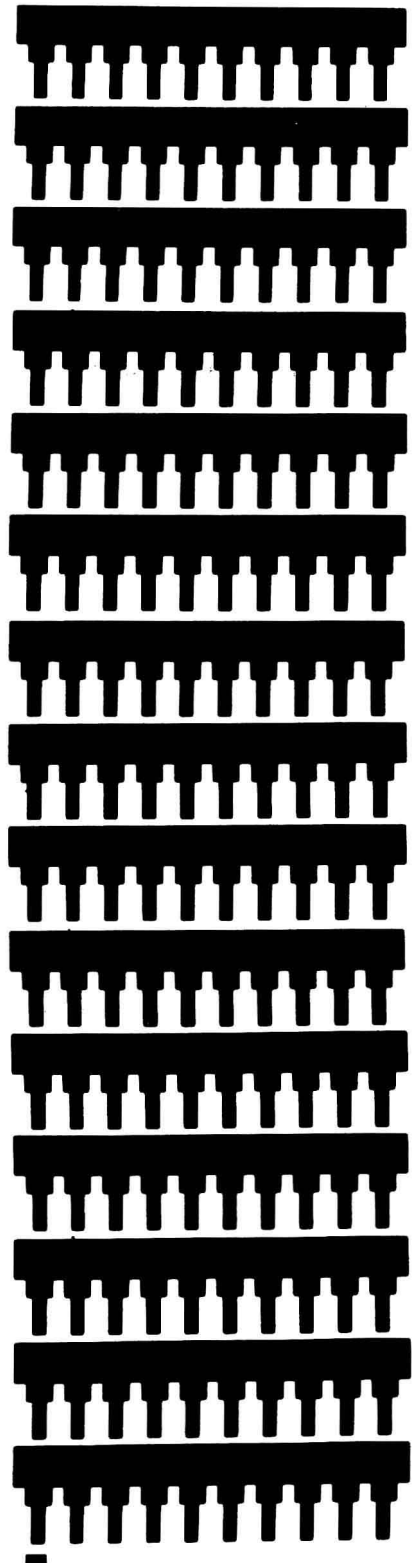

Hardware

Designers gain new freedom as options multiply

Enhancements add lustre to the capabilities of established families, and new devices are extending the performance range in both directions. Low-cost microcomputer boards offer another alternative.

After 18 months of calm, in which microprocessor manufacturers have consolidated the designs of their first general-purpose families, a second big wave of activity has begun. This time, however, it reaches a more sophisticated level.

Urged on by savvy users, whose concern is with system design rather than chip architecture, the manufacturers are introducing second- and third-order refinements aimed at boosting microcomputer capacity while lowering system cost. At the same time, they are rushing new devices to market to extend the microprocessor performance range both at the low end, where existing chips present an overkill solution, and at the high end.

Four trends are emerging. First, established families are being enhanced. System throughput is being increased and instruction sets enlarged as manufacturers turn to new metal-oxide-semiconductor processing and improved central-processer architecture. Input/output power, too, is being increased with new sets of programable I/O chips.

Second, the new 16-bit single-chip processing units are heading upwards. What they are aiming for is the high-performance end of the microprocessor market, where precision arithmetic and large memories must be accommodated.

Third, the one-chip controllers, as their name implies, contain enough computing power to handle many stand-alone controller functions on their own. On the same chip as the central processing unit sit control read-only memory for program storage, random-access memory for data storage, and input/output registers for system manipulation.

Finally, there's a host of single-board microcomputers, beguiling alternatives to the do-it-yourself approach of buying just the chips.

All these developments are changing the microprocessor universe. In order to graph this change, Fig. 1 charts the various family types against the applications spectrum.

Clearly, the 8-bit system covers the most ground, being used in many more different designs than either the 4- or 16-bit devices. Indeed, the 8-bit word seems just about right for most of today's microcomputer systems, in contrast to the 16-bit words that are the staple of minicomputers.

How much overlap there will be between powerful 8-bit general-purpose systems and the 16-bit high-performance systems is still to be determined, especially in large-memory process-control applications. The current wisdom is that the 16-bit families will remain primarily on the high-performance end of the application spectrum for several years, since the 8-bit families are so well established.

Besides, enhanced 8-bit microprocessors are coming along that are faster and can handle 16-bit word data anyway, making it easy for a user to upgrade his 8-bit system to a 16-bit design without much additional investment in software. Moreover, the use of single-chip controllers in distributed processing systems boosts the performance of 8-bit designs by taking much of the burden off the central processing unit and, in many cases, by making it unnecessary to move up to a higher-capacity 16-bit CPU system.

Meanwhile the multichip 4-bit systems—the earliest microprocessors to appear—are feeling increasing pressure from the minimum-chip system designs. Rockwell's PPS-8/2 and PPS-4/2 two-chip systems, Fairchild Semiconductor's F-8, National Semiconductor's SC/MP, and Electronic Arrays' 9002 can all handle many of the jobs formerly done by the 4-bit Intel MCS-4 or Rockwell PPS-4 but often with fewer packages and at lower cost. Moreover, the single-chip 4- and 8-bit microcontrollers already mentioned will increasingly eliminate the need for multichip 4-bit designs.

Puts it all together. Activity in microprocessors is fast and furious, as manufacturers make available a wide range of products, from low-cost microcontroller chips to powerful, general-purpose families and boards. This 16-bit microcomputer is from Data General.

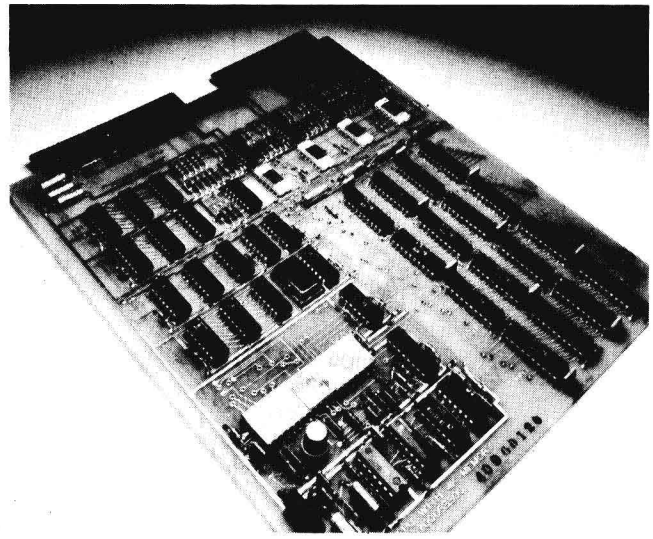

The 8-bit mainstream

In the 8-bit microprocessor applications spectrum, Intel Corp.'s 8080 family, with its enhanced 8080A CPU, Motorola Semiconductor's 6800 family, with its enhanced 6800D CPU, and Rockwell's PPS-8 family currently rank one, two, and three in popularity among users. The 8080 system is being used in a wide range of industrial process controls, games, intelligent data terminals, and so on. The 6800 has found its greatest penetration in data-communications terminals and instrumentation. The PPS-8 has found strong acceptance in skid-control automotive designs, as well as in other high-volume systems. All are second-sourced—the 8080 by AMD, TI, NEC, and Siemens, the 6800 by AMI, and the PPS-8 by National.

What makes these chip families so suitable for general-purpose applications is the centralization of their computing capabilities—an orientation borrowed from minicomputer architecture. Unlike many newer designs, such as the F-8, which distributes its computing power among its family of devices, the 8080, 6800, and PPS-8 concentrate that power all on a single chip. In effect, their central processing units act as their own peripheral controllers, using generalized bus lines to manipulate external memories, interface chips, and input/output chips.

These CPU chips are well equipped for their job. Both the 8080A and 6800D have a 16-bit address bus, an 8-bit bidirectional data bus, and fully TTL-compatible control outputs. Besides supporting up to 65 kilobytes of random-access memory, they can address a large number of peripheral devices, providing for practically unlimited system expansion.

Moreover, both CPUs show a considerable improvement in architecture over the preceding generation of 8-bit designs. The 8080, for example, contains a 16-bit stack pointer that controls the addressing of an external stack located in memory. The proper instructions can initialize this pointer to use any portion of external memory as a last-in/first-out stack, so that almost unlimited subroutine nesting becomes available. The stack pointer in addition allows the contents of the program counter, the accumulator, the condition flags, or any of the data registers to be stored in or retrieved from the external stack.

The 8080's stack control instructions also permit multilevel interrupts. The current program or "status" of the processor can be pushed onto the stack when an interrupt is accepted, then popped off the stack after the interrupt has been serviced, and this can be done even when the interrupt service routine is itself interrupted.

Where the two families differ is in several of the system requirements. The 6800's single 5-volt power supply contrasts with the 8080's ±5 V and 12 V supplies. The 6800 timing is quite simple. All instructions are executed in two or three cycles, which are identical in length. Control outputs are real-time signals instead of look-ahead instructions. Moreover, in the 6800 system, separate I/O instructions are unnecessary since memory locations can house either I/O or memory data.

On the other hand, the 8080 has more powerful instructions, with stronger branch and interrupt capability. It can interface with a wide variety of peripheral devices. It has tremendous software support, such as an in-

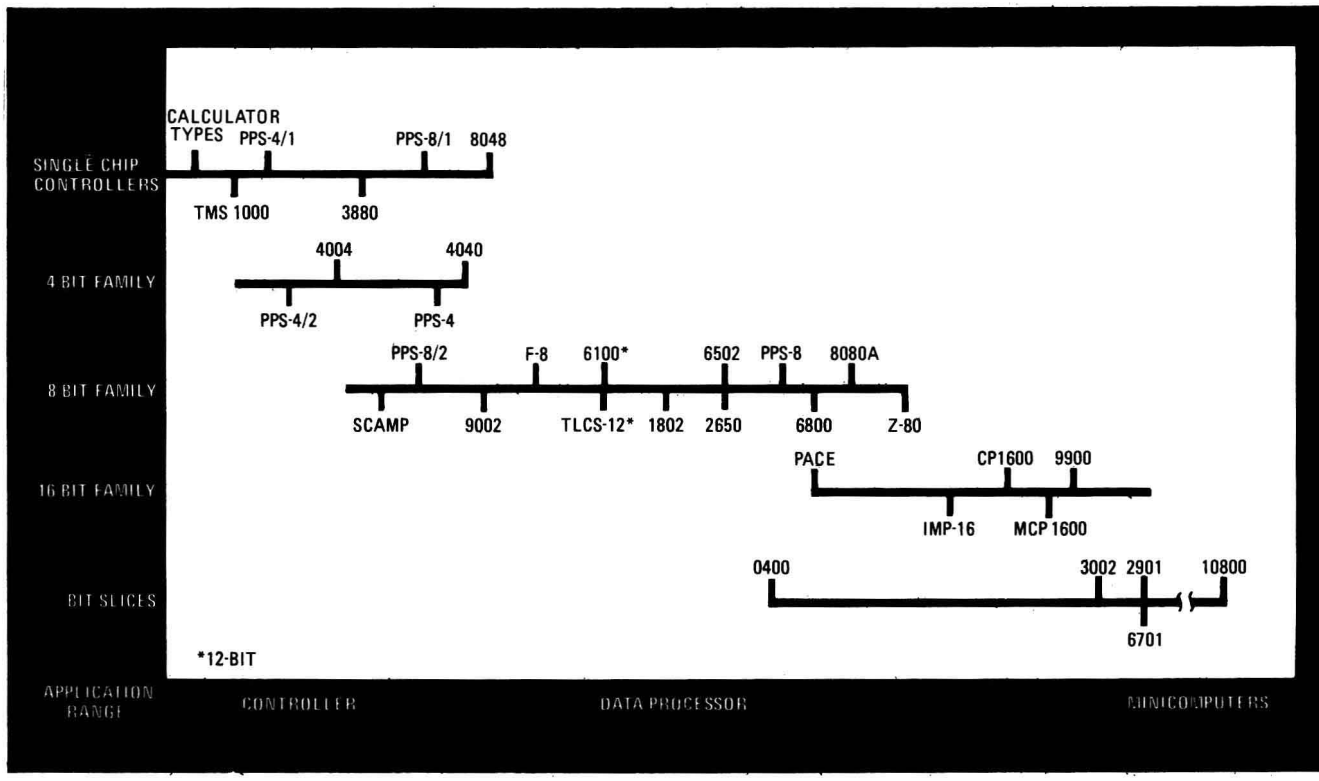

1. The universe. The 8-bit microprocessor families cover most ground, being found in everything from peripheral controllers to powerful data-processing systems. The 4-bit systems handle the smaller controller jobs, while the 16-bit chips and bit slices rise to minicomputer work.

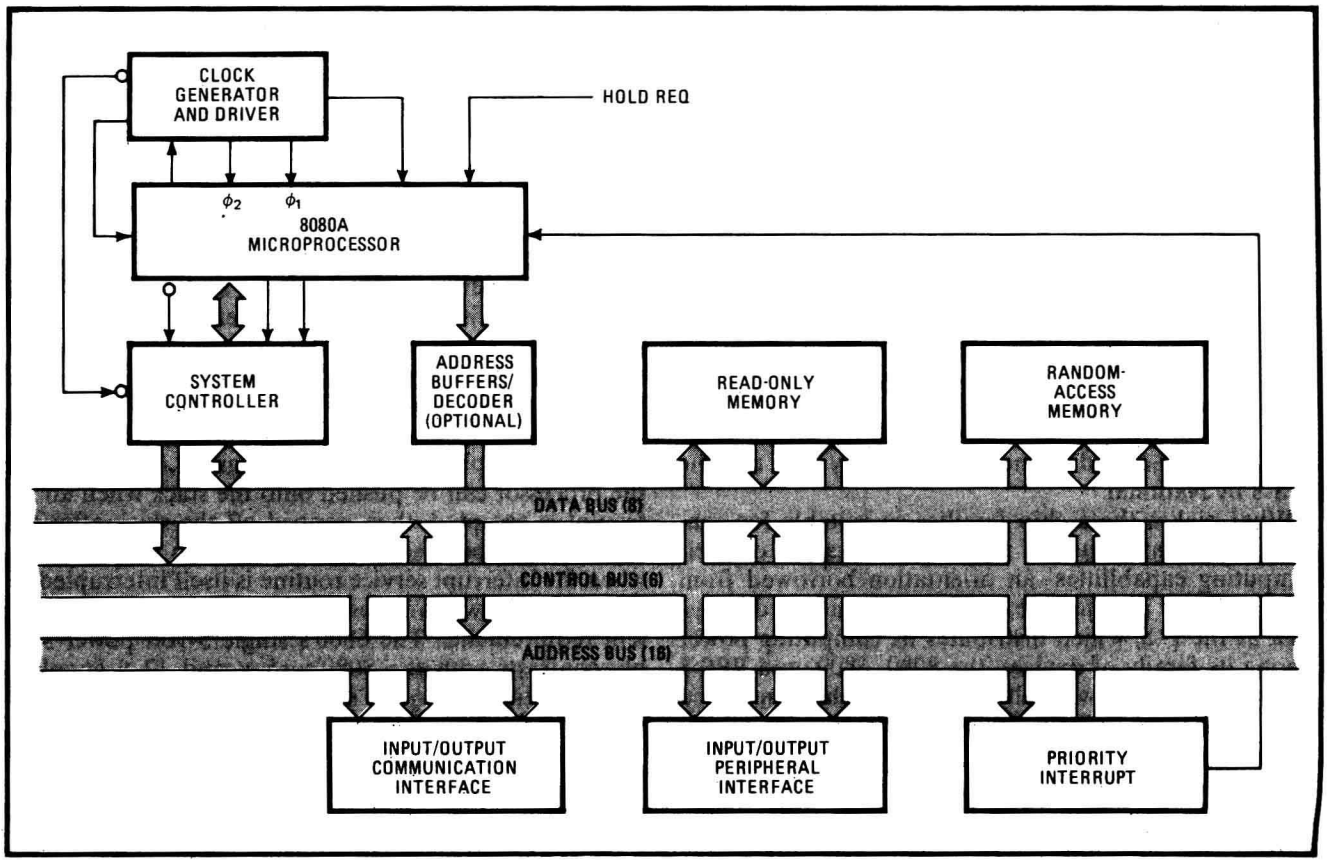

2. The 8080. Intel's 8080 system can produce dozens of peripheral and I/O device configurations. Because all devices connect to a system bus, hooking up parts into complex configurations becomes quite straightforward once the instruction program has been developed.

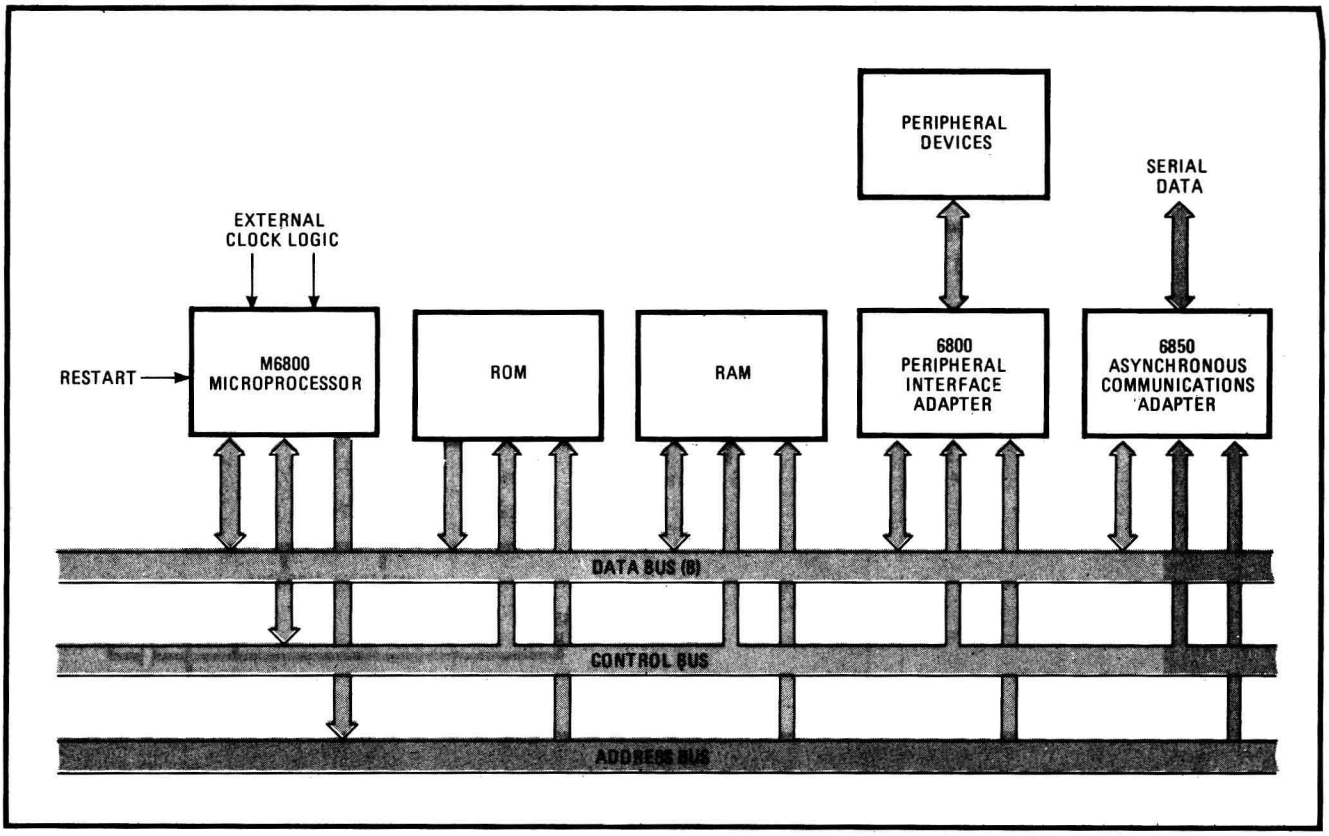

3. The 6800. A Motorola 6800 system, also designed around a central bus, keeps the package count to a minimum by using powerful peripheral devices, such as the peripheral interface and communication adapters. System works off standard ROMs and RAMs.

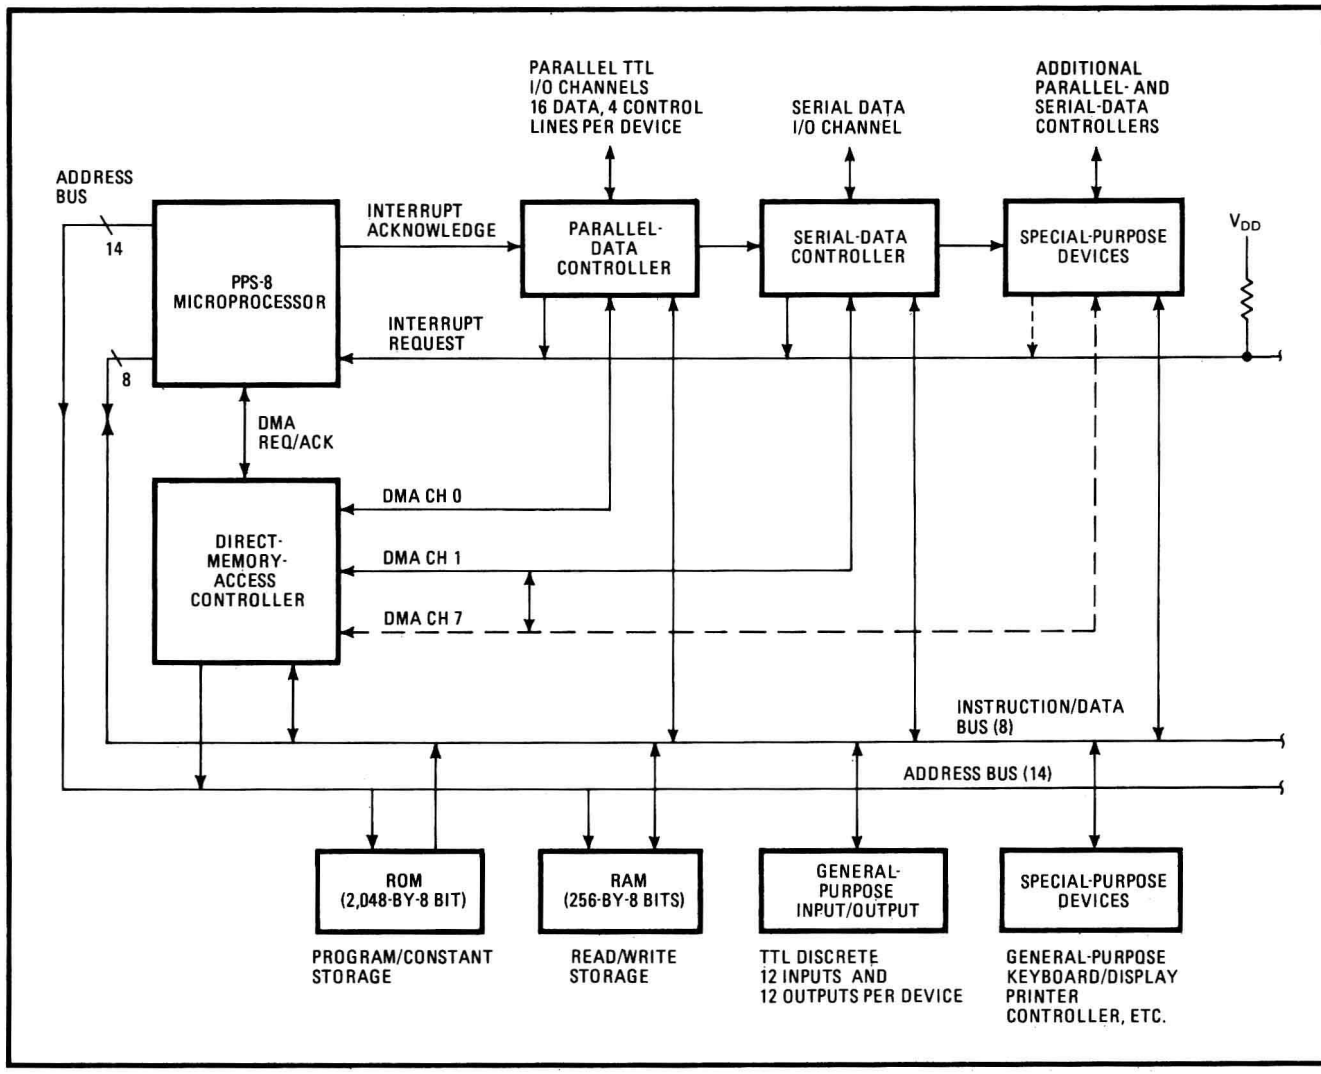

4. The PPS-8. In Rockwell's PPS-8 microcomputers, general-purpose input/output, ROM, and RAM chips again hook up directly to address and instruction data buses. Family contains a direct-memory-access controller and parallel and serial I/O data channel modules.

circuit prototype developing system and a large library of user-generated instruction programs. All this accounts for the 8080 system's overwhelming success.

Rockwell's PPS-8 differs from both the 8080 and the 6800 in the basic architecture of its CPU. Its CPU chip has an arithmetic/logic unit, a control unit, accumulators, and address registers, laid out much as in the 8080 and 6800, but interlinked quite differently from either the 8080 and 6800 families (see Figs. 2, 3, and 4). For example, the program and data memories in a PPS-8 microcomputer system have completely separate and parallel address spaces, so that a memory address may legitimately identify two different memory locations—maybe a byte of program memory containing an instruction code and a byte of data memory containing binary data. This double-duty memory address accounts for the PPS-8's high throughput, even though it is built with p-channel MOS technology.

A unique feature of the PPS-8—its clock signals—adds to its flexibility. These signals serve for both synchronization and control. The four-phase clock generator transmits two clock signals to every device in a system, and every device contains logic to decode them, interpreting the contents of the data and address buses in different ways during different phases of a machine's cycle. As a result, the CPU can program a wide variety of peripheral devices—if the user is willing to learn his way around the clocking scheme.

While the CPU or microprocessor chip is the central controlling element in a microcomputer system, just as vital are the ROMs, RAMs, and various input, output, and interface circuits that make up the balance of the design. Figures 2, 3, and 4 illustrate typical system configurations for the 8080, 6800, and the PP-8: these, the most established general-purpose microcomputer systems, have highly-developed system components that hook up directly with the CPU on simple bidirectional bus lines.

The bus line configuration of the 8080 and 6800 families generally has become the model for most 8-bit systems. Its three bus lines—a data bus, a control bus, and an address bus—handle all elements of a system. Standard ROMs, which store the program data in the form of lookup tables, are hooked up by connecting the ROM's output lines to the data bus and input lines to the address bus. RAMs get their data written from CPU com-

AVAILABLE MICROPROCESSORS

Type No.	Technology	Address Capacity (bytes)	Manufacturers* and Comments**
4-bit			
4004	p-MOS	4-k	Intel
4040	p-MOS	4-k	Intel (National)
PPS-4	p-MOS	4-k	Rockwell (National): SV
PPS-4/2	p-MOS	8-k	Rockwell: CC, SV
PPS-4/1	p-MOS	–	Rockwell: CC, SV, RAM on chip
TMS-1000	p-MOS	8-k	Texas Instruments: SV, MP
8-bit			
EA 9002	n-MOS	65-k	Electronic Arrays: SV
F-8	n-MOS	65-k	Fairchild (Mostek): CC
8008-1	p-MOS	16-k	Intel
8080 A	n-MOS	65-k	Intel (AMD, TI, NEC, Siemens)
8048	n-MOS	2-k	Intel: 512-bit RAM on chip
6502	n-MOS	65-k	MOS Technology – other versions are available with lower address capacity
5065	p-MOS	32-k	Mostek
6800	n-MOS	65-k	Motorola (AMI): SV
SCAMP	p-MOS	65-k	National: CC, SV
1801	C-MOS	65-k	RCA: 2-chip CPU
1802	C-MOS	65-k	RCA
PPS-8	p-MOS	32-k	Rockwell (National): SV
PPS-8/2	p-MOS	32-k	Rockwell: CC, SV
2650	n-MOS	32-k	Signetics: CC, SV
300	TTL-S	8-k	Scientific Micro Systems
Z-80	n-MOS	65-k	Zilog: SV
12-bit			
6100	C-MOS	4-k	Intersil (Harris): SV, CC
TLCS-12	n-MOS	4-k	Toshiba: MP
16-bit			
CP1600	n-MOS	65-k	General Instruments: MP
MCP-1600	n-MOS	65-k	Western Digital: MP, MC
IMP-16	p-MOS	65-k	National: MP, MC
PACE	p-MOS	65-k	National: MP
PFL-1600A	n-MOS	65-k	PanaFacom: MC
TMS-9900	n-MOS	65-k	Texas Instruments: SV, general-purpose registers in memory
Bit slices			
2901	TTL	65-k	Advanced Micro Devices (Motorola, Raytheon): MP
9400	TTL	65-k	Fairchild: MP, SV
3002	TTL	512	Intel (Signetics): MP, 2-bit slice
6701	TTL	65-k	Monolithic Memories: MP
10800	ECL	65-k	Motorola: MP, CC, ECL
SBP0400	I^2L	65-k	Texas Instruments: CC, MP

NOTES

*Developing manufacturer listed first.

**Key: MP – microprogramable; ECL – emitter-coupled logic; TTL – transistor-transistor logic; I^2L – integrated injection logic; SV – single voltage; CC – clock on chip; MC – multi-chip central processing unit

Assessing microprocessors

Making comparisons between the available microprocessors on the basis of data sheets is a very tricky business. Even a simple specification like cycle time can be highly misleading. In most cases cycle time by itself tells you practically nothing—you must know how many cycles are needed to execute what instruction. For example, some microprocessors boast cycle times as low as 1 microsecond but require multiple cycles to execute even the simplest instructions. Others list longer cycle times but require fewer cycles to do the same instruction.

Nor does it help too much to compare execution times of simple instructions. Often the time to do a fetch or a register-to-register ADD has little relation to the time required for executing more complex instructions, like calling in a subroutine on the basis of various bit settings.

Even more misleading is ranking CPU complexities in terms of numbers of registers, or I/O ports, or whether the chip has built-in direct memory access, and so on. Many powerful microprocessors, such as TI's 9900 expel all the general-purpose working registers from the CPU chip and locate them in external RAM. But the chip is more powerful than most CPUs with multiple general-purpose registers. Likewise, a minimum-chip system design, such as the F-8, has computation logic distributed over two or three matched chips, so just looking at the CPU doesn't begin to show the capability of the system.

The instruction set is another area that lends itself to vendor specmanship. Repertoire size alone has little meaning, unless you know how the supplier is counting instructions. Are multiple, closely related instructions counted as one or as many? What instructions are included? And the various types of instructions that differ only in their "if" conditions, how are they counted?

That's why this microprocessor chart is kept fairly simple. Breaking down the chips by word length gives an idea of a processor's range—but only a rough one, since the efficiency of doing anything certainly does not depend on word length alone. The technology is broken out only because knowledgeable users feel more comfortable knowing what's in the device, but it too has to be related to design—whether processing is done serially or in parallel, and so on. (All things being equal, devices built with n-channel MOS are faster, smaller, and easier to interface than those built with p-channel MOS, whereas devices built with bipolar technology are faster and can do more but are larger and cost more to build than MOS LSI devices.)

As for address capacity, obviously the more memory that a chip can access, the larger the system that can be implemented. But again, watch out. Some processors can access large amounts of memory directly. Others need external devices to reach large bytes of memory.

Another area that concerns users is alternate sourcing In general, the alternate sources of microprocessors are proving well able to satisfy customers' demand for multiple-sourced devices. For instance, AMD's 9080A series claims speed and power specifications that in some respects exceed Intel's 8080A specifications, and AMI undertook considerable process development in building its version of Motorola's 6800 family. Mostek Corp. has done a nice job supplementing Fairchild's F-8 support and applications effort. Then too, there is the National/Rockwell technology exchange that made their respective microprocessor families available to each other.

mands that travel on the address bus, and their data is read out to the CPU on the data bus. Peripheral interface circuits receive their inputs on the control and address bus and return their data outputs on the data bus. Thus, this bidirectional bus system is the conduit serving all members of the microcomputer family. New interface and peripheral chips, regardless of their complexity, will use these bus lines, ensuring a user a simple and well-formulated method of upgrading his basic system with more powerful I/O and peripheral chips.

The dedicated 8-bit types

Unlike the general-purpose 8-bit systems, which generally use at least a dozen chips, Fairchild Semiconductor's F-8, also supplied by Mostek Corp., and National Semiconductor Corp.'s SC/MP families were designed to realize controller-type systems with the fewest possible chips at the lowest possible cost. Both families can dish up useful designs with just two chips, although the F-8 is a more powerful system, readily expandable into memory-rich designs.

The dissimilarities of these two devices stem from dissimilar design philosophies. The Fairchild F-8 designers chose a configuration that is quite unlike the CPU orientation of minicomputers. Instead they distribute process and memory control throughout the system. The F-8 therefore works best where two or three of its powerful family members can do the job standing alone, without a large number of external memory (although they can be added if necessary).

Because they interact so intimately, a designer must be familiar with the functions of the F-8 chips. Besides the CPU chip, there's a programable storage unit, which provides read-only memory plus various logic functions—it combines with the CPU to form a complete microcomputer if so desired. A dynamic-memory interface chip links the first two chips to either dynamic or static RAMs storing data, and a static-memory interface is for use with state RAMs only. Finally, a direct-memory-access chip implements the direct-memory-access logic in conjunction with the dynamic-memory interface chip.

Because various logic functions are distributed among the four peripheral chips, the CPU contains only the arithmetic/logic unit, the control unit and instruction register, the logic associated with interfacing the system bus with the I/O control signal, and the accumulator register. It does not contain memory-addressing logic, memory-addressing registers, stack pointer, program counter, and data counter, all of which reside in the companion memory and memory interface chips.

This configuration has both advantages and disadvantages, the chief advantage being that fairly powerful systems can be implemented with remarkably few chips.

Moreover, the lack of memory-addressing logic on the CPU chip itself means that no address lines are needed on the system bus, and the 16 address signals, which CPU-oriented systems need to interface with the bus, can instead be used for two 8-bit I/O ports on each

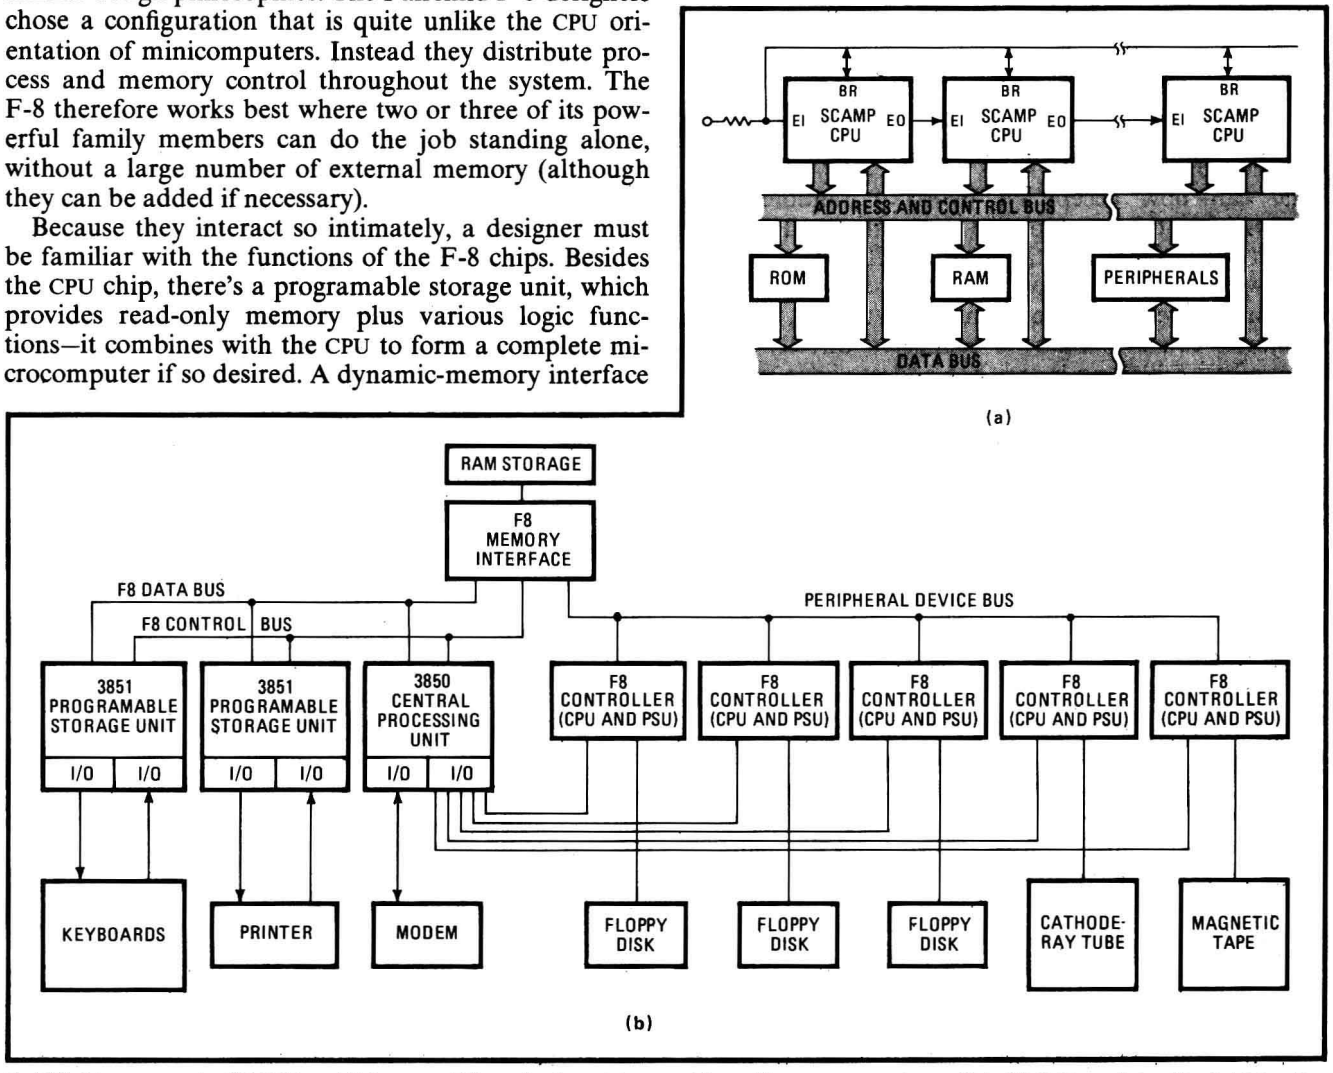

5. Efficiency experts. SC/MP and F-8 carry off the prize for minimum-chip multiprocessor systems. This SC/MP configuration (a) from National daisy-chains several microprocessors along one bus. The Fairchild F-8 disk controller has five F-8s operating as distributed processors. F-8 systems can readily be expanded to handle a wide variety of peripheral control and complex data communications applications.

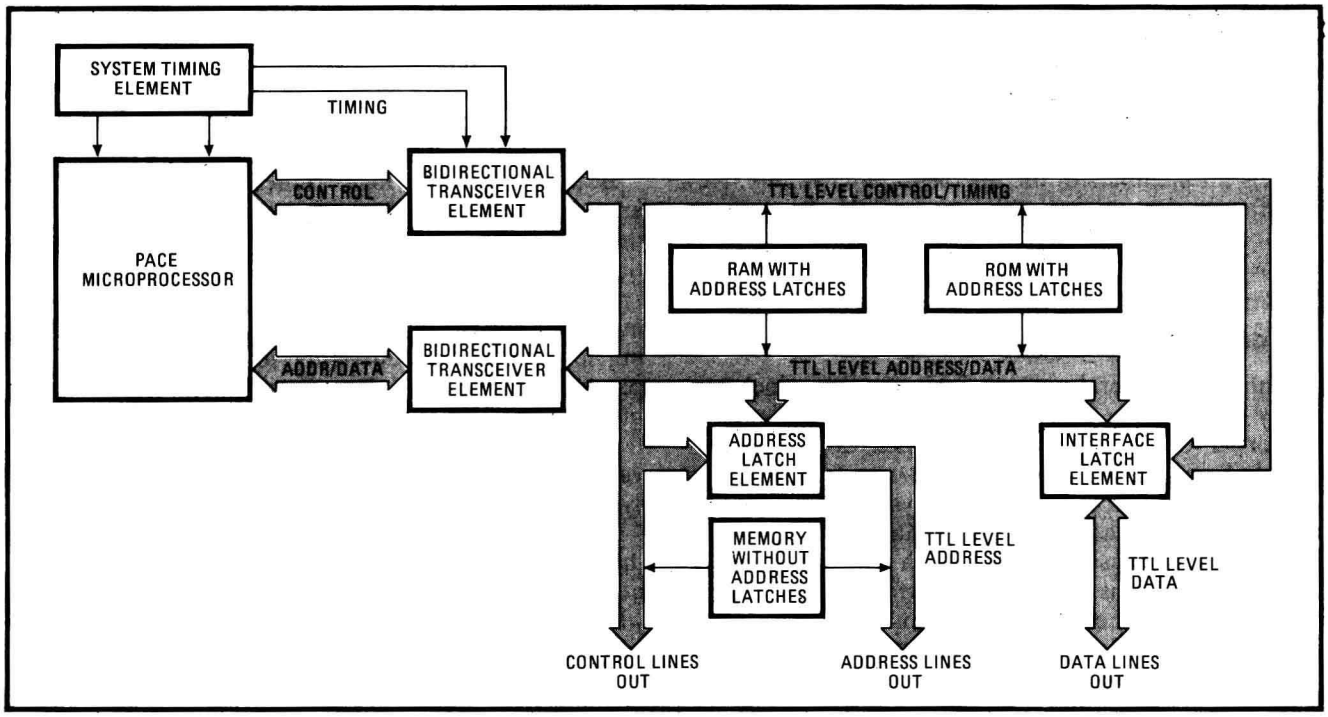

6. Relating. National's 16-bit PACE system reaches the TTL world with transceivers that interface directly with any RAMs or ROMs designed with on-chip address latches. Standard memories can also be used if an 8-bit address latch element is included.

device. Better yet, the place on the CPU chip formerly occupied by address registers and memory-addressing logic can now accept 64 bytes of random-access memory. It is this on-chip RAM that makes an F-8 minimum two-chip configuration functionally useful.

Now for the disadvantages. Because of the removal of memory-addressing logic from the CPU chip, external memories can no longer be connected directly to the system bus, which no longer has address lines, and the family's other devices must be used. Of course, this is easily done with the memory-interface devices, but the extra packages do add to the cost of the design. Worse yet, this memory-addressing logic must be duplicated if more than one memory device is present.

On the other hand, SC/MP (pronounced Scamp) centralizes its computing capabilities in the CPU, just like the 8080 and 6800 families, so that systems can be configured with standard memories directly. The SC/MP chip can handle up to 4 kilobytes of memory with no additional logic or interface packages. Systems requiring more memory are also possible: a five-chip system, handling up to 65 kilobytes of RAM, would consist of the SC/MP, a two-chip bidirectional transceiver, an address latch, and a buffer.

Internally, SC/MP is a programable 8-bit parallel processor. It contains one 8-bit accumulator, four 16-bit pointer registers (one of which is dedicated to the function of program counter), an 8-bit status register, and an 8-bit extension register. On-chip timing circuits eliminate the need for external clocks, and TTL compatibility allows easy interfacing with other system components.

Architecturally, SC/MP, again like the 8080 and 6800 families, employs a unified bus system, to which the central processing unit, memory, and peripheral devices are each connected. The common data bus enables memory-reference instructions to reference peripheral devices. In addition, SC/MP architecture provides serial data and control streamlining under software control and has built-in programable delay.

Both the SC/MP and F-8 families lend themselves to multiple processor systems. In SC/MP, the bus configuration is responsible, allowing many SC/MPs to be tied to the bus for daisy-chain operation (Fig. 5a). When one SC/MP stops transmitting or receiving, it notifies the next SC/MP in line that it may take over.

The F-8 CPU chips can serve either in a multiple processor system or in two-chip peripheral controllers subordinate to a multichip processor-based system, such as large point-of-sale terminals. The floppy-disk controller shown in Fig. 5b contains five F-8s working in conjunction with floppy disks, a magnetic-tape unit, a cathode-ray-tube display, a keyboard, printer, and modem. While the low-speed devices (the keyboard, printer, and modem) can be adequately handled by the programed I/O structure, the other, high-speed devices require separate F-8 CPUs and programable storage units.

The 8-bit newcomers

While established suppliers of microprocessors have recently come out with upgraded products—most notably the 8080A and 6800D—newcomers to the field are trying to gain entry with still higher-performance versions of the earlier devices. A good example is Zilog Inc.'s Z-80 chip. In a tribute to the success of the 8080, designers at the Los Altos, Calif., company have based their design on it, but have added more data-processing and instruction-handling capability. At the same time, they have tried to simplify the system configuration along the lines of the 6800.

For example, the Z-80 is heavily CPU-oriented, like

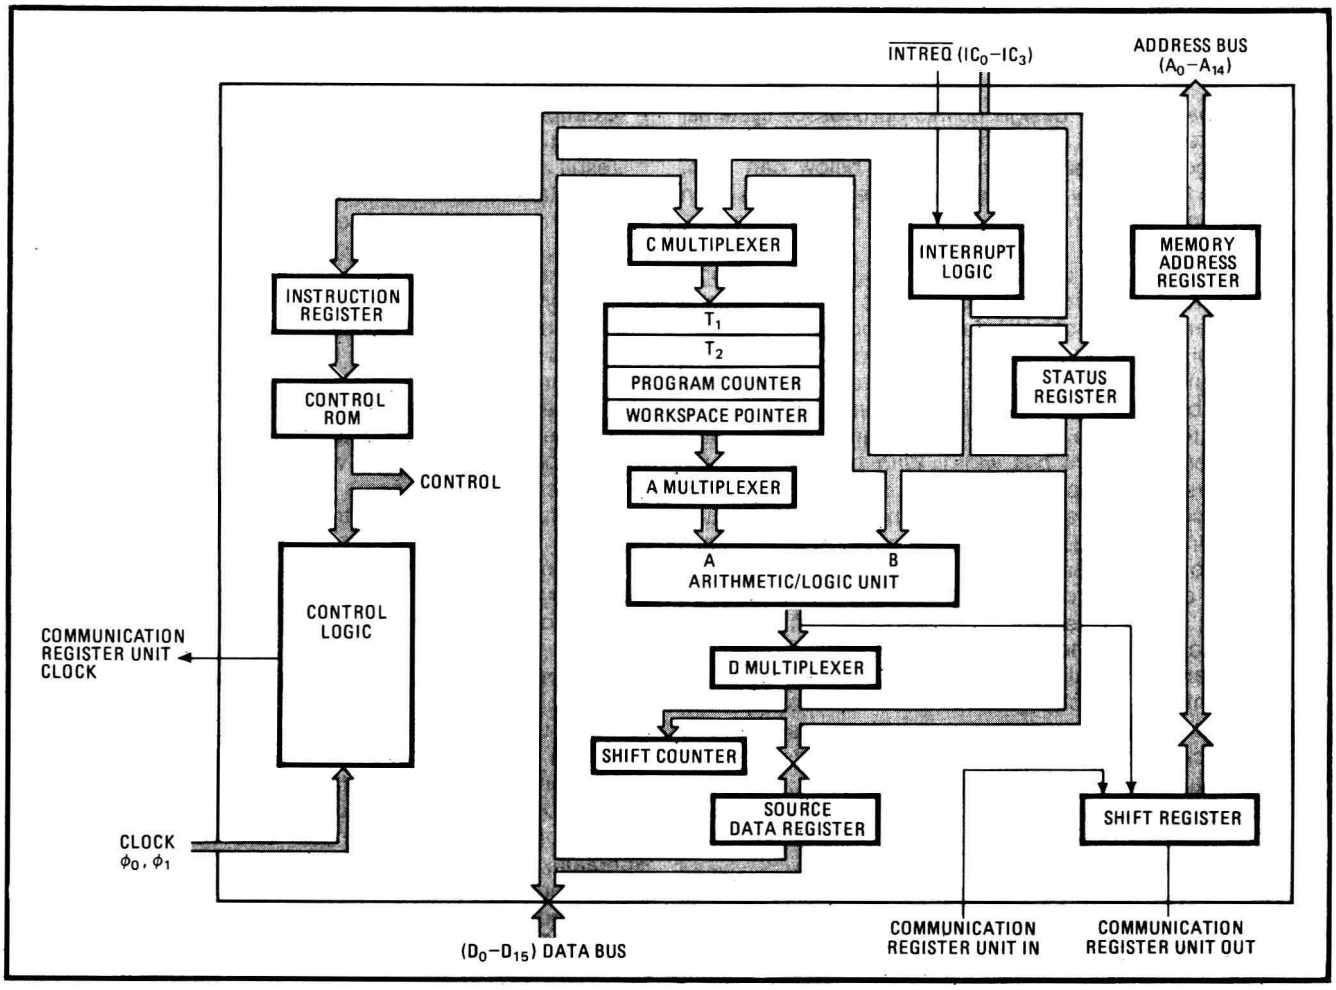

7. Thinking big. Texas Instruments' 9900 16-bit microprocessor has full 16-bit data bus and 16-bit ALU on chip but exiles all general-purpose registers to external RAM locations. A wide range of interrupt capability is included, making the chip very flexible.

the 8080A, and is completely compatible with 8080A software. But thanks to depletion-mode technology, it, like the 6800, has a single-phase clock on the chip and requires only a single 5-V power supply.

The Z-80 can handle 158 different instructions and, like the 8080A and 6800D, has an internal 16-bit-wide data bus. Unlike them it contains both an 8-bit and 16-bit external address bus, so that it can process either 8- or 16-bit words in one cycle.

Architecturally, the Z-80's CPU chip resembles the 8080A, where general-purpose registers perform basic computation operations and special-purpose registers perform various program operations, such as program counting and stack pointing. Also as in the 8080A, the CPU contains the accumulator and flag registers.

The Z-80's block of general-purpose registers has a distinctive feature: it consists of two matched sets of six 8-bit registers. Now a programer can use them individually, as 8-bit registers, or in tandem as 16-bit register pairs, depending on whether he is handling 8-bit or 16-bit words. Moreover, the programer may select one set of registers for a single exchange command while using the other set for the rest of the sequence. This saves interrupt time—and is especially useful in systems that require a fast interrupt response—because there's no need to transfer the register contents to an external stack during the fast-cycle interrupt or subroutine processing.

As for the Z-80's special-purpose registers, the program counter and stack pointer function much as they do on the 8080A. The program counter holds the entire 16-bit address of the current instruction just fetched from memory, and the stack pointer keeps track of the 16-bit address of the current instruction. An external stack memory, organized as a last-in/first-out file, allows simple implementation of multiple level interrupts, unlimited subroutine testing, and simplification of many types of data manipulations.

Like the Z-80, another recently introduced microprocessor that takes cognizance of established 8-bit general-purpose designs is the 9002 from Electronic Arrays Inc., Mountain View, Calif. But unlike the Z-80, which is being supplied with its own set of dedicated support devices, the 9002 has been conceived as a stand-alone digital process controller that can interface with standard peripheral chips through an 8-bit parallel TTL-compatible data bus.

The 9002 timing and control signals allow a user to bring the chip together with any bus-oriented peripheral devices he may choose. Examples are: Motorola's 6820 peripheral interface adapter for general-purpose controller applications; the asynchronous communications interface adapter and low-speed modem for

Learning about microprocessors

In approaching the design of microprocessor systems, the first requirement for the novice is to learn the specialized jargon. The basic terms defined below can help. They are followed by some advice on the next steps in an education in microprocessor theory and applications.

Central processing unit: a group of registers and logic that form the arithmetic/logic unit plus another group of registers with associated decoding logic that form the control unit. Most metal-oxide-semiconductor devices are single-chip CPUs, in that the registers hold as many bits as the word length of the unit (the 8080 and 6800, for example, are 8-bit devices and thus the basic registers are eight bits wide). With bit-slice devices, however, central processing units of any bit width can be assembled essentially by connecting the bit-slice parts in parallel. Externally, a bit-slice device will appear to be a coherent single CPU capable of handling words of the desired bit length.

Register: logic elements (gates, flip-flops, shift registers) that, taken together, store 4-, 8-, or 16-bit numbers. They are essentially for temporary storage, in that the contents usually change from one instruction cycle to the next. In fact, much of the microprocessor's operation can be learned by studying the registers, which take part in nearly all operations.

Accumulator: a register that adds an incoming binary number to its own contents and then substitutes the results for the contents.

Program counter: a register whose contents correspond to the memory address of the next instruction to be carried out. The count usually increases by one as each instruction is carried out, since instructions generally are stored in sequential locations.

Instruction register: storage for the binary code for the operation to be performed. Usually this instruction represents the contents of the address just designated by the program counter. However, the contents of the instruction register or the program counter may be changed by the computations. This, of course, represents one of the key ideas of a stored-program computer—instructions, as well as data, can be operated upon and subsequent operations will be determined by the results.

Index register: some memories are organized by index number (the contents of the index register). The address of the next instruction may be found by summing the contents of the program counter and the index register. Increasing the index register by one will cause the processor to go to a new section of memory.

Stack pointer: a register which comes into use when the microprocessor must service an interrupt—a high-priority call from an external device for the central processing unit to suspend temporarily its current operations and divert its attention to the interrupting task. A CPU must store the contents of its registers before it can move on to the interrupt operation. It does this in a stack, so named because information is added to its top, with the information already there being pushed further down. The stack thus is a last-in first-out type of memory. The stack-pointer register contains the address of the next unused location in the stack.

Flag: usually a flip-flop storing one bit that indicates some aspect of the status of the central processing unit. For example, a carry flag is set to one when an arithmetic operation produces a carry. A zero flag is set when the result is zero. These flags aid in interpreting the results of certain calculations. Others are sometimes provided to permit access by interrupt request lines—for example, if a CPU is engaged in the highest priority of calculation, it may set all status flags to zero—which, loosely translated, means "don't bother me now." If only some of these flags are set, then only certain interrupt lines will be able to get through according to their priority.

Direct memory access: a technique that permits a peripheral device to enter or extract blocks of data from the microcomputer memory without involving the central processing unit. In some cases, a CPU can perform other functions while the transfer occurs.

In going beyond these definitions, an engineer will probably find that there's not an abundance of good basic information on microprocessors. However, the gap is filling.

Certainly, a first source on the details of a particular product is the manufacturer's product descriptions. Some of them are quite readable. Most provide easily understood introductions to the microprocessor, with just enough information to get started. Best known are Intel's "8080 Users Manual," Motorola's mammoth "Microprocessor Applications Manual" and 6800 System description, Fairchild's "F-8 Circuit Data Book," Signetics' 2650 manual, Rockwell's microprocessor family descriptions, and the descriptive literature National puts out on SC/MP, PACE, and IMP-16 families.

Independently produced sources also are available, but they're of varying quality. A useful one is a monthly publication on a variety of microprocessor subjects called "New Logic Notebook," edited by Jerry L. Ogdin of Microcomputer Technique Inc., 1120 Reston International Center Office Bldg., Reston, Va. 22091. A monthly compilation of microprocessor news and product introductions is a newsletter called "Microcomputer Digest," P.O. Box 1167, Cupertino, Calif. 95014.

One of the best books is a paperback called "An Introduction to Microcomputers," from Adam Osborne and Associates, 2950 Seventh St., Berkeley, Calif. 94710. It has a compact tutorial section on basics, followed by good comparisons of key families.

Then there's "Microprocessors," first volume in the Electronics Book Series. It is a compilation of all the original articles on major microprocessor designs that appeared in this magazine—from the first 4004 to today's complex 8- and 16-bit designs. It also contains detailed design and application material. It's available for $8.95 (see page vi).

A good source of basic information is the independent seminars that are becoming widely available. One of the most successful is Integrated Computer Systems' three-to-five-day courses held across the country. A schedule is available from David Collins at ICS, 4445 Overland Ave., Culver City, Calif. 90230.

An opportunity for hands-on experience is the suppliers' seminars. These are manned by applications specialists who travel around regularly, offering a good review of a particular microprocessor line. Finally, there are the courses offered by the IEEE and the universities.

communications-controller applications; or any of Intel's new programable interface devices, such as the programable peripheral or communications interfaces. This means that a system designer can use the 9002 as a powerful controller chip, managing the operation of any TTL-compatible peripheral device.

The 9002 designers also picked the best features of existing processor designs. The CPU combines the on-chip 64-byte scratch-pad RAM of the F-8, the push-pop subroutine stack of the Intel 4040, the simplified timing concepts of the PPS-4, the straightforward peripheral addressing techniques and single 5-V-supply requirement of the 6800, and the general-purpose registers of the 8080.

To these borrowed features the 9002 adds some purely its own. It contains a seven-level subroutine stack for multiple interrupt capability and eight 12-bit general-purpose data registers. With its 64-byte scratch-pad memory it can handle many stand-alone controller jobs without requiring additional RAM. Moreover, one of the 9002's internal flags allows the user to perform either 8-bit binary arithmetic or packed binary-coded-decimal arithmetic (dual 4-bit operands) with built-in, automatic decimal correction. To choose, he simply sets the flag in one state or another. This is useful for peripheral controllers where CRT displays need BCD data.

With all this computing power, ample control signaling, and on-chip RAM capability, the 9002 can realize many fairly powerful designs with only two or three packages. For example, a controller can be built with the 9002, a 1,024-by-8-bit ROM, such as the EA 4700, and two Intel 8212 peripheral interface chips, or else it can be built with the 9002, a 2,028-by-8-bit ROM, such as the EA 4600 and Motorola's 6820 PIA chip.

C-MOS: another choice

Another enhancement of an existing device is RCA Solid State division's single-chip version of its 8-bit C-MOS microprocessor. Designated the 1802, the chip is three times faster than the old two-chip design, has one third more instructions—a total repertoire of 91—and costs less. This came about thanks to RCA's new silicon-gate process that yields C-MOS devices almost half the size of metal-gate designs and also increases transistor switching speed. As a result, a C-MOS microprocessor becomes as fast, cost-effective and flexible as today's p- and n-MOS microprocessors.

To illustrate, the 1802 has a cycle time of 1.25 microseconds and takes only one or two cycles, plus a fetch cycle, to execute any instruction. This gives it an instruction time of either 2.5 or 3.75 microseconds that puts it well in the speed range of either the 8080 or 6800. Moreover, with its 91 instructions, it is as powerful and as flexible. Yet RCA designers were careful to retain the architecture of the two-chip design, so that the 1802 is software-compatible with its predecessor.

What distinguishes the 1802 CPU from other 8-bit designs is its separate instruction and address registers. The address data is placed in an array of sixteen 16-bit scratch-pad registers, each of which can point to either data or program. That means that a user is not forced to provide an address with each memory reference instruction—something he must do with other processors.

As address pointers, individual scratch-pad registers in the array are selected by any one of three 4-bit registers, so that the contents of any address can be directed to any one of three destinations. As data pointers, the 16 scratch-pad memories are equally flexible. They can be used either to indicate a location in memory or as pointers to support a built-in direct-memory-access function.

The only other C-MOS microprocessor is Intersil's 12-bit device. By using the same software as the PDP-8A, the device lets users of that popular computer implement their systems in low-power easy-to-use C-MOS technology. The 40-pin package has an instruction capacity of about 40, can access 32-k bytes of external memory, and can control 64 I/O parts. For the 1600, Intersil plans to supply a complete set of C-MOS peripheral devices, such as C-MOS ROMs, RAMs, and UARTs.

Two n-channel 8-bit microprocessors that have begun to make headway for general-purpose applications are the Signetics Corp. 2650 and the MOS Technology Inc. 6500 family of microprocessors. The Signetics part, available only in sample quantities about a year ago, lately gained momentum—especially in Europe, thanks in part to Philips' recent acquisition of Signetics.

The 2650 is a single 5-V parallel 8-bit binary processor capable of performing 75 instructions in a machine cycle time of 2.4 microseconds, which puts it in the same general class as the 8080 and 6800 families. The chip can address up to 32 kilobytes of external memory (compared to 65-k for the others). But its ability to execute variable-length instructions makes it somewhat more efficient, since a one- or two-byte instruction may often be used for memory addressing. Moreover, most instructions require only 6 of the first 8 bits, so the remaining bits can be used for the register field.

MOS Technology's family is unique in that it includes a number of software-compatible microprocessor chips differing primarily in the amount of memory they can address. The 40-pin 6502 can handle 65-k bytes of

What it costs

Like all semiconductor chips, microprocessor prices are coming down. Here's a rough guide to how much it costs to do some typical jobs at today's prices (in appropriate volumes):

Job	Number of packages	Cost
General-purpose minicomputer emulation	30 and up	$1,000
Dedicated minicomputer	20–30	600
Process controllers	15–20	400
Smart terminal (i.e. communications, etc.)	10–15	300
Complex general controllers (i.e. traffic lights, medical, machine tool, etc.)	10–15	200
Complex peripheral controllers, industrial	10	150
Point-of-sale terminals	10	150
Games, instruments, etc.	5–10	75
Simple controllers, hobby gear, appliance control	1 or 2	10