Electronics

Which digital voltmeter? page 84

Sensors for celestial guidance: page 94

Transistors can multiply, divide: page 109

April 4, 1966

75 cents

A McGraw-Hill Publication

Below: Scratch-pad memory increases computer capability: page 118

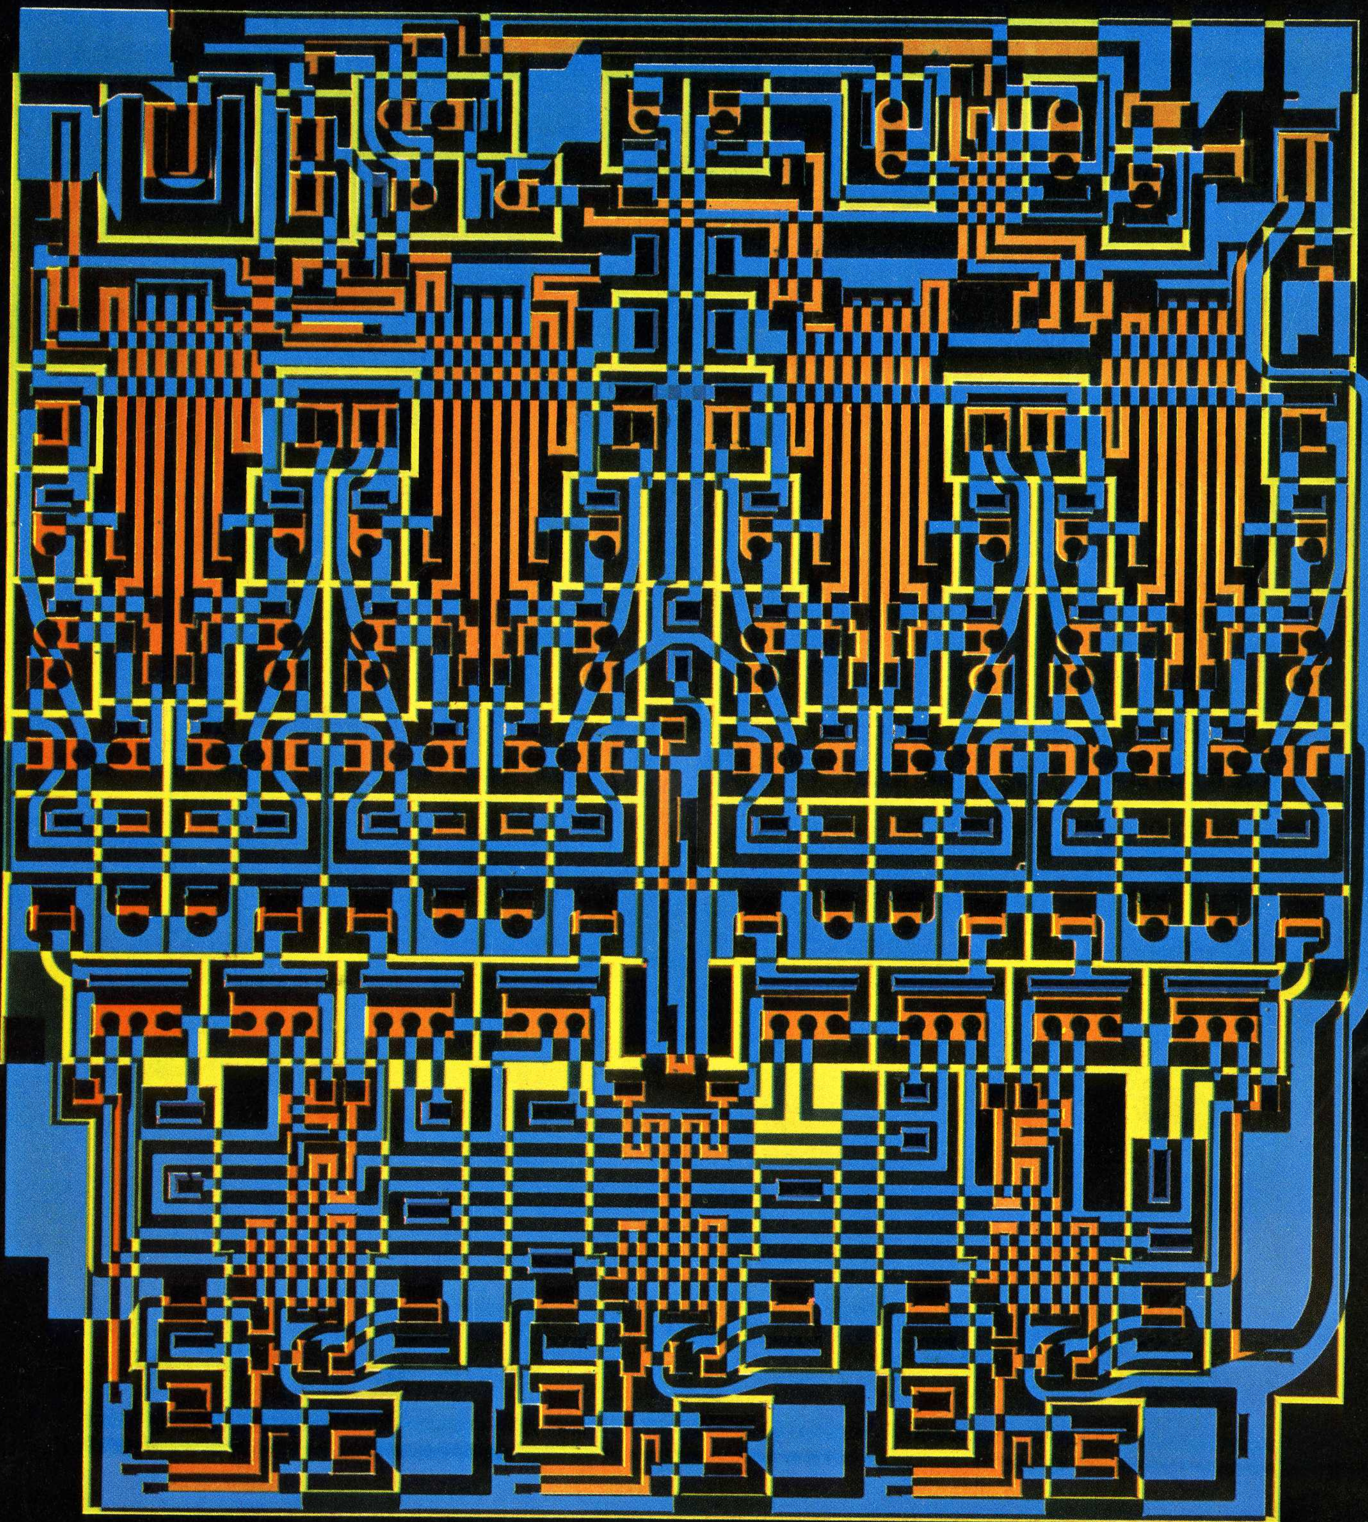

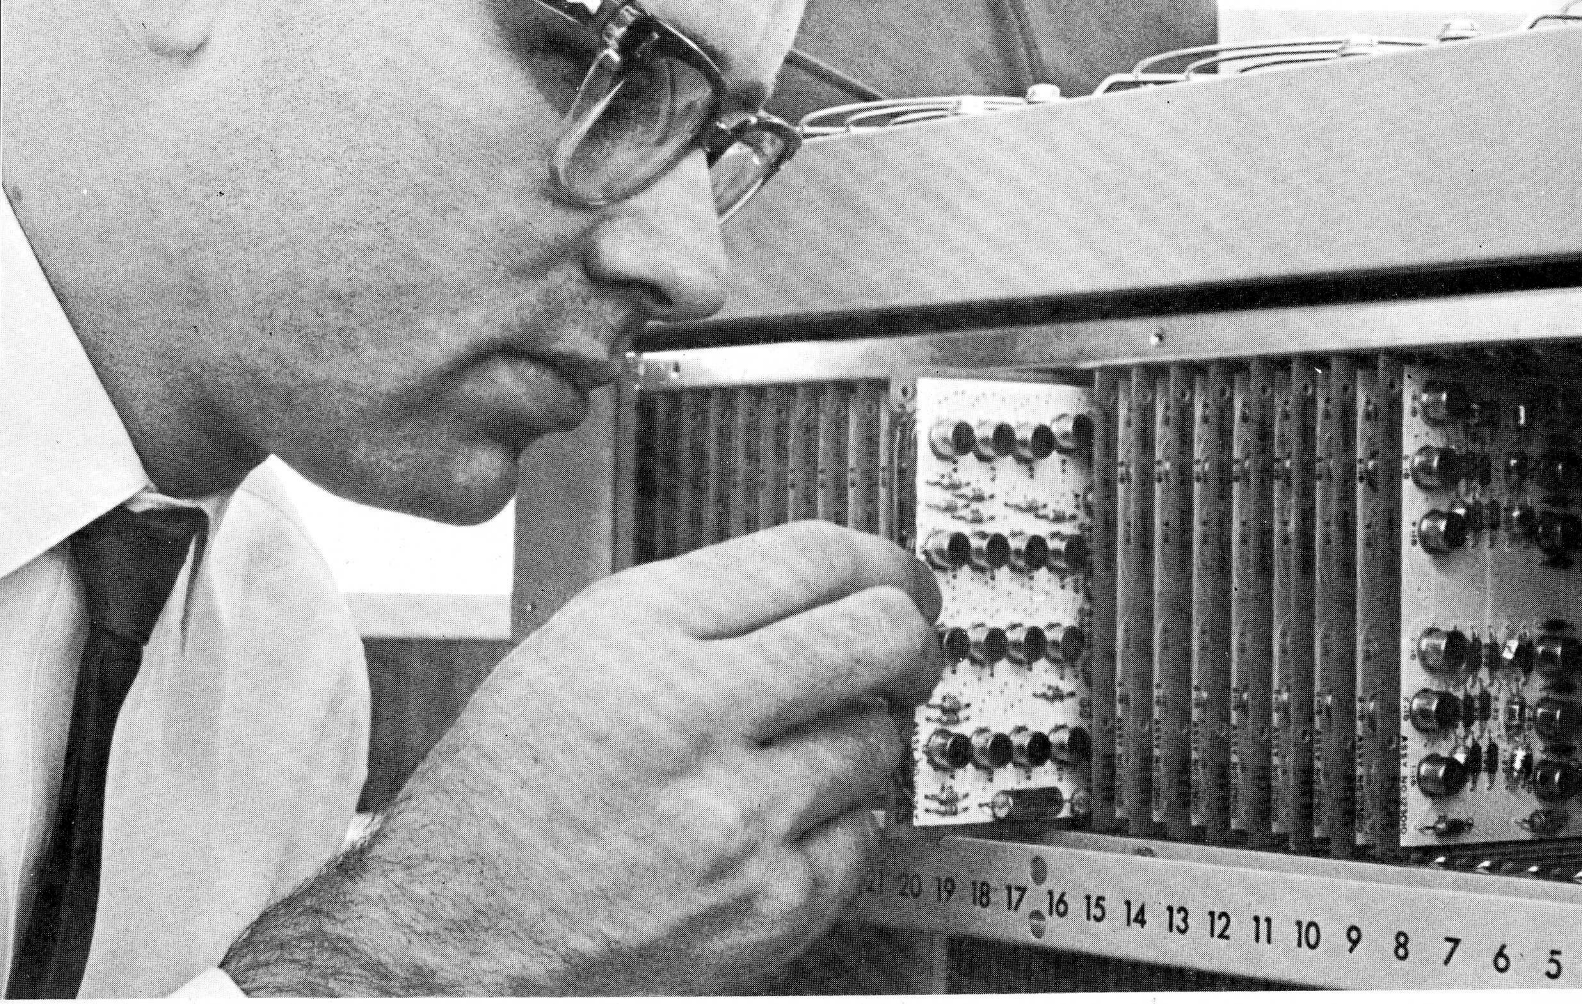

Gene Potter plugs a 16-byte memory module into one of the standard logic-module cages of a computer. Each cage can contain 128 words, each 32 bits (4 bytes), of scratch pad memory.

Computers

Integrated scratch pads sire new generation of computers

When a single chip of silicon carries eight bits of memory and their buffered decoding network, numerous high-speed memories go into a computer and time sharing becomes economical

By Gene B. Potter and Jerry Mendelson, Scientific Data Systems, Inc., Santa Monica, Calif.
Sam Sirkin, Signetics Corp., Sunnyvale, Calif., a division of Corning Glass Works

The mass production of integrated circuits containing several bits of memory—complete with decoding and output circuitry—makes practical a new generation of computers. With the IC's, many small memories, called scratch pads, operating at the speed of the logic circuitry, can be built economically into workaday data-processing systems.

The scratch pads enable moderately priced computers to provide the speed, capacity and versatility once available only in huge, expensive systems. As a result, the cost of processing large amounts of data can be halved.

An example of the efficiency that can be obtained when scratch pads are the keystone of the computer organization is in the Sigma 7 family of computers, which Scientific Data Systems, Inc. (SDS) introduced March 15 [Electronics, March 21, 1966, p. 37]. A typical Sigma 7, costing about $500,000, has an input-output rate of some 160 million bits a second.

The high efficiency results from a new control memory and an arithmetic organization based upon the use of from two to a dozen scratch pads. The scratch pads are modular plug-in units, ranging in size from 16 to 512 words. The memories are made with the monolithic integrated circuits pictured on the cover and on page 122. The IC's were developed in a cooperative effort of SDS and the Signetics Corp. and are being produced by Signetics for SDS. Each circuit contains eight flip-flop storage bits and their encoding, decoding, control and buffer circuits.

A printed circuit card, carrying 16 IC's plus the drive circuitry pictured on page 121, forms a 16-byte module (8 bits per byte) that is the basic building block for all Sigma 7 scratch pads. For example, four such modules make up a 16-word scratch pad. If more capacity is needed, more boards are plugged into the computer's address and data buses. The modules are self-contained, having address-decoding circuits for up to 128 words.

The main job of the scratch pads is to quicken the response of the system to multiple inputs by reducing its dependence on the main memories for many operations and by allowing it to react instantly to priority requests. The Sigma 7 can service up to 200 users at remote consoles while simultaneously performing several routine chores such as inventory control and while operating peripheral equipments. The functions of the scratch pads are outlined in the block diagram on page 125.

The multiple-register architecture of the system, with its indexing, mapping (program relocation) and memory-write protection, requires multiple accesses to scratch pads within the main memory cycle. Therefore, to be effective, scratch pads must operate at least five times as fast as the main memory. Otherwise their operation would not mesh effectively with logic and control circuitry.

The Sigma 7 scratch pads typically have a writing speed of about 90 nanoseconds and a reading speed of 60 nanoseconds. The main memories, built of ferrite cores, have a cycle time of 1.2 microseconds and an effective cycle time of 700 nanoseconds when the operation of the main-memory modules is overlapped.

The cost problem

The costs of scratch pads appeared to be a crucial problem when development of the system began two years ago. Designers of time-sharing machines have long known that unrestricted use of scratch pads would immeasurably improve system performance, but the expense limited their use in commercial systems.

Except for a few systems, new at the time, which had scratch pad-like registers composed of logic IC's, scratch pads were made of magnetic storage elements and transistor circuitry. In contrast to large, ferrite-core memories, scratch pads were inefficient in space used, power consumed and access time per bit. The cost was more than 50 cents per bit compared to a few pennies per bit for mass memories.

At that time, these problems were being attacked through the development of batch-fabrication methods, among them monolithic circuits, various forms of thin-film matrices, the weaving of magnetic wires and the molding of multibit ferrite elements.

IC's were the most attractive solution. Production processes were established. Functional density was high and could be made higher. Power and voltage requirements were compatible with the other IC's planned for the system. The ability of the IC's to operate at normal logic levels and provide useful logic drive simplified application. The eight-bit design selected provided the best trade-offs among power, packaging and system-design constraints and semiconductor-manufacturing capabilities.

Today, IC's containing 16 bits of storage with internal decoding, or 32 bits without it, probably can be made. If the system constraints on power dissipation and packaging can be eased, they may be used. One advantage of the present IC design is that it can be packaged in a 10-lead, TO-5 can, which is low in cost, easily assembled and reliable.

The design selected has 178 components on a silicon die measuring 104 by 95 mils. The components are bipolar transistors and diodes, and diffused resistors. Considerably higher functional density could have been obtained with circuits composed of field effect transistors made by the metal-oxide-silicon technology. But MOS circuitry was less developed at the time of design selection and is still too slow for scratch pads in the Sigma 7.

How big a scratch pad?

As the system design took shape, it became apparent that scratch pad modules could be used with significant gains in several places in the system, most of which are discussed on pages 124 to 126. To take full advantage of the savings possible with high-volume production of IC's, all the scratch

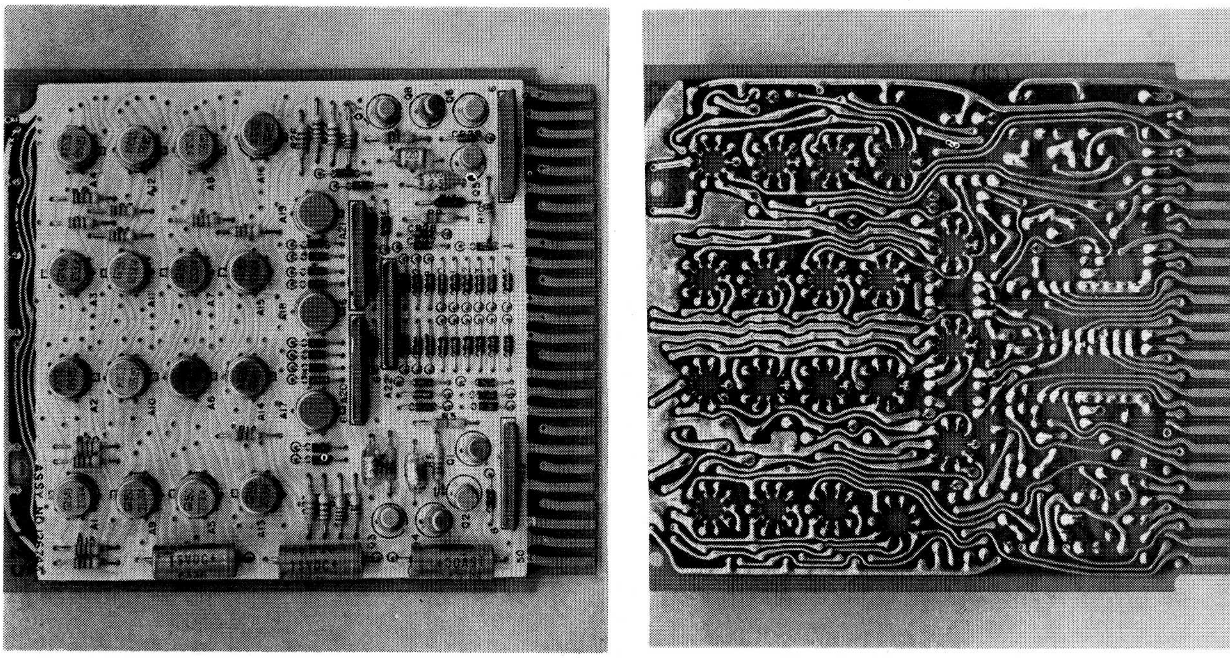

Scratch-pad module, front and back. On the front (left), the wiring—coated with an insulation—runs in a direction generally vertical to plug-in position. On the back, wiring is generally horizontal. The design avoids the need for expensive multilayer printed circuits. The bar-shaped objects are thick-film resistors.

pads should use the same IC and be assembled in similar configurations.

Since even large machines need only a few one-word registers, a scratch-pad size of 16 bytes appeared ideal. This number meshed with system word lengths, nominally 32 bits composed of four 8-bit bytes. It also meshed with the main memory, which can be addressed in bytes, half words, words, and 64-bit double words.

A linear-select (word-organized) form of memory was required. This is an advantageous form for IC storage because bits are positioned within a regular, two-dimensional wiring network. The coincident (bit-organized) form requires three wiring planes.

The two-dimensional matrix is easily provided by printed circuit cards, as illustrated above and below. The wiring runs in the X direction on one side of a card and in the Y direction on the other side. In a register memory, four 16-byte cards are placed side-by-side to form a 16-word block. These are organized so that a single set of control lines are shared by each row of IC's. Without that sharing, packaging would be expensive.

The memory module operates in a parallel manner, obtaining one bit from each of eight separate IC's, as part of a linear-select memory scheme.

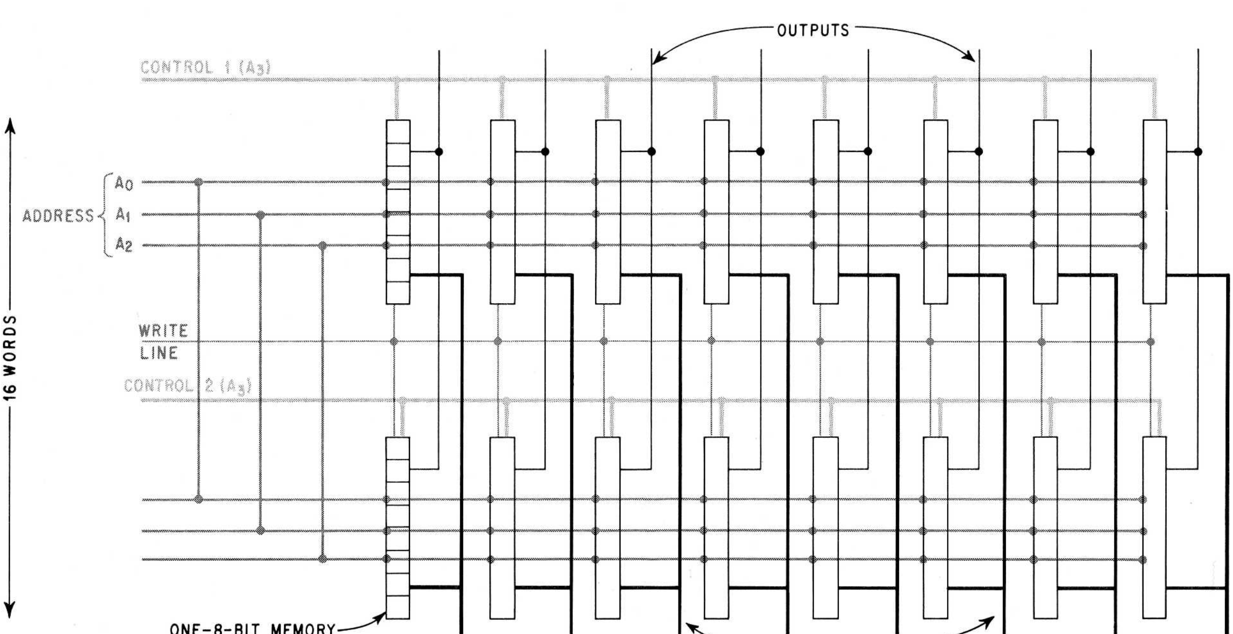

Module wiring configuration that is one byte wide by 16 words high. Ground and voltage lines are not shown.

The monolithic choices

IC-processing and scratch pad design constraints pointed toward an eight-bit or 16-bit IC design. For economy and reliability, the circuits had to fit into the 10-lead, TO-5 cans, which had also been chosen for the other IC's in the system. Four-bit IC's were rejected as too small and not economical; they would require a doubling up of the number of circuits and the amount of processing and assembly.

If 16 bits were used, the buffered decoding network and the group of storage flip-flops would have to be made as two types of IC's, costing some of the mass-production advantages of one type. With external decoding, the individual flip-flops would present a heavy current load, requiring more and higher-current drive circuits. Also, the bits would be sensitive to the noisy environment they would create, due to the fast switching of heavy current and capacitive loads. Since these and other design factors would increase over-all cost, eight bits appeared optimum.

Volatility and compatibility

In a large IC scratch pad, power consumption tends to be high and heat dissipation becomes a problem. Unfortunately, IC scratch pads are volatile, so that power cannot be removed from those portions not in use. A unique circuit design minimized the difficulty. Standby power is provided to only the actual storage flip-flops of an IC when that IC is not in use.

This was accomplished by using one of the required address lines as a control line. It switches a transistor, in the IC, which in turn gates power to all internal logic except the elementary storage flip-flops. The technique saves 60% of the power that would ordinarily be consumed.

IC scratch pads are extremely convenient to use in a computer because they can be matched electrically to the rest of the system. The machine's normal logic-voltage and current levels can interface directly with the inputs and outputs of the memory IC's. The nondestructive readout (NDRO) operation of the scratch pad makes character regeneration unnecessary, producing highest system speed. Writing is simple since the memory need not be cleared or reset; when the voltage on the write-enable line rises, the address lines steer the data to the proper locations and force the storage flip-flop to the proper state.

Electrical specifications for the IC's are given on page 124 and the schematic is on page 122. The circuitry is transistor-transistor logic, which is the fastest type for the logic-voltage levels employed.

Note that the tenth lead of each IC is not used; this allows addition of a mass-reset function. The other nine leads are for an address-control line, three address-decoding lines, write-enable line, input, output, supply-voltage and ground lines.

Power-saving circuit

A data bit pulse enters a given IC, is gated by the write line and directed to the selected bit by the three address circuits (three variables will identify any location). The control line deprives the decoding and write circuitry of power, when they are not in use, suppressing noise and saving up to 60% of the power. It also provides the option —essential, since the outputs of a column of IC's are tied together—of not altering or looking at the state of any flip-flop.

The address-control line switches a transistor Q_1 (in the IC schematic), which feeds supply cur-

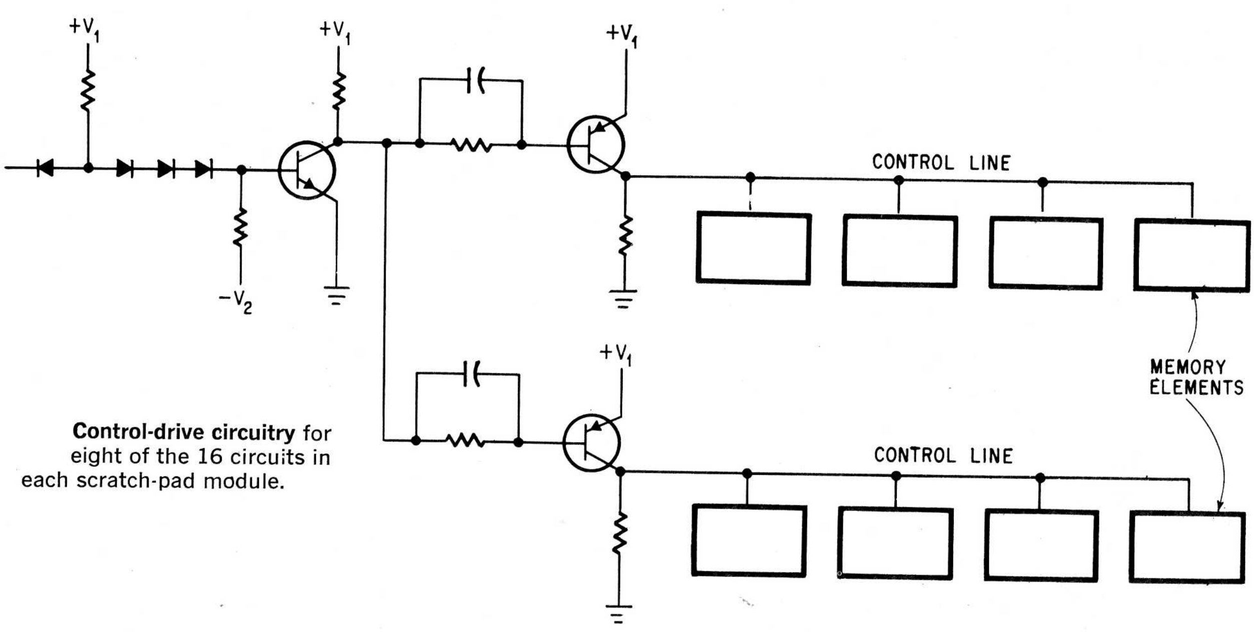

Control-drive circuitry for eight of the 16 circuits in each scratch-pad module.

Reading, writing and eight bits on a chip

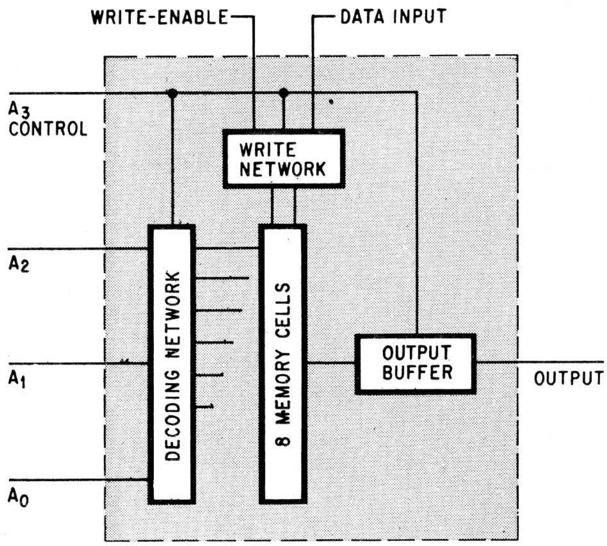

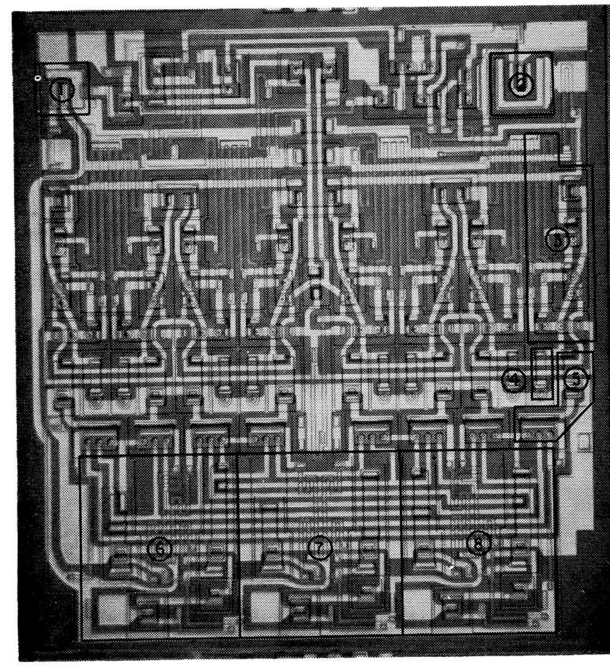

Layout of memory circuit, shown schematically at the left, is rearranged to pack the 178 components compactly on a single chip measuring 104 by 95 mils. The circuitry is transistor-transistor logic. Key areas of the chip, outlined on the photograph, are: 1, 2. Control and output transistors
3. One of the eight flip-flop storage cells in a band across the center of the chip
4, 5. Cell's clamp diodes and three-input AND gate
6, 7, 8. Address-decoding circuits.

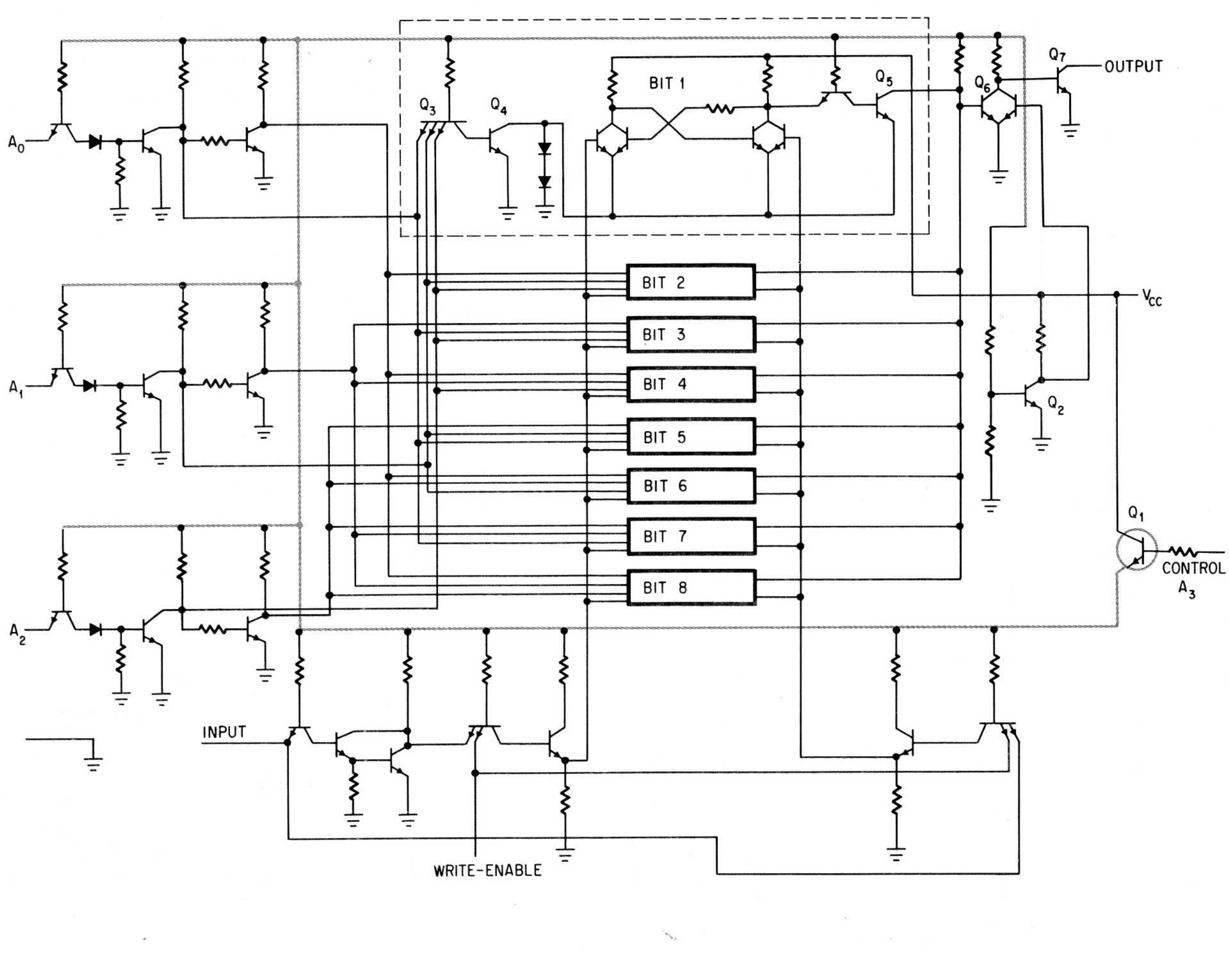

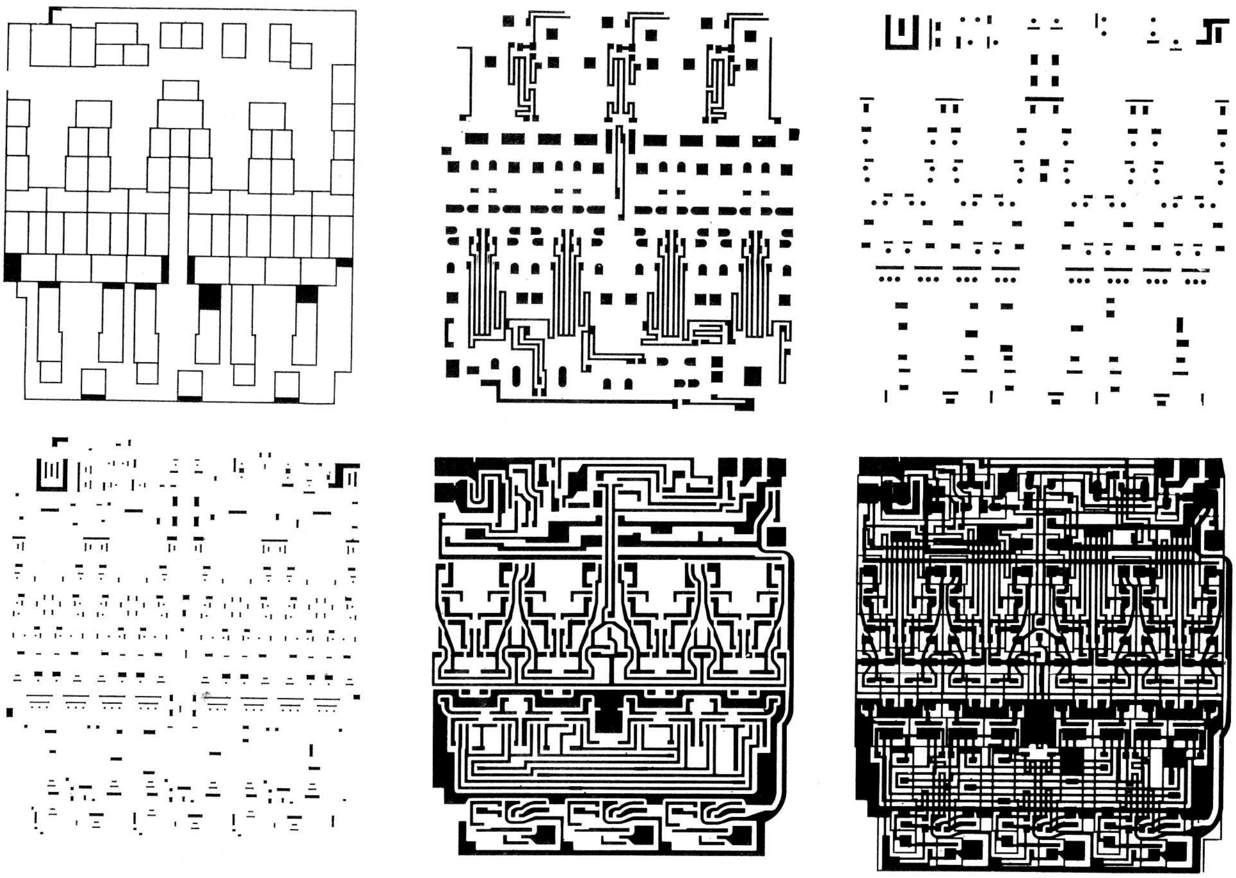

Processing stages that produce the monolithic memory circuits can be traced in these patterns. They are a series of masks for etching the oxide passivation coating of the silicon crystal and the thin-film metal layer. The patterns etched in the oxide bare the silicon for diffusion of isolation moats and the resistors, diodes and transistors. The next to last pattern produces the thin-film interconnections. These patterns were superimposed to create the composite pattern at the lower right and the color illustration on the cover.

rent to 20 other transistors. When Q_1 is on, current flows to the decoding and write logic and to the output buffer. When Q_1 is off, these circuits are deprived of power. The flip-flops draw power directly from the supply-voltage lead V_{cc}. Signals on the data-input line cannot be inserted by the write circuitry, since that is not functioning.

The output buffer is also deprived of power. When the control line is down, Q_2 is deactivated. This holds the buffer off because the right-hand side of transistor pair Q_6, which feeds Q_7, is turned on. The voltage swing for the control line is between 6.6 volts and ground.

Q_1 and Q_7 are the biggest and most powerful transistors in the IC. They are oversized so they can carry large currents without excessive loss and overheating of the IC.

Low-power flip-flops

So that decoding can be done on the chip, the address circuitry subjects only the selected bit to a write-line input. Likewise, only that bit controls the output of the memory. The decoding design also reduces the unselected bits' power consumption about 30%. The power saved is used to speed up the address and write circuitry.

If it is assumed that a diode drop equals approximately 0.75 volt, a transistor's $V_{be(sat)}$ equals 0.70 volt and $V_{ce(sat)}$ equals 0.20 volt, the operation of the circuit can be described as follows:

The base reference of a selected bit is only $V_{ce(sat)}$ above ground, while that of all the unselected bits is equal to two diode drops above ground, thus reducing their apparent supply voltage 1.5 volts. The selected bit clamps the write lines at $V_{be(sat)} + V_{ce(sat)}$, or 0.9 volt. The other bits cannot be written into because a write pulse has to be at least three times the diode voltage drop or 2.25 volts, to get into a flip-flop.

For example, if bit 1 in the schematic on page 122 is not selected, its base reference will be about +1.5 volts. That is the voltage drop across the two clamp diodes that are connected in series from the collector of Q_4 to ground. However, if bit 1 is selected, the base reference is only about +0.2 volt, the collector-to-emitter saturation voltage of Q_4. All the emitters of Q_3 are high.

The state (on or off) of the selected flip-flop is reflected by Q_5 and detected by Q_6. Q_5 can turn Q_6 off only if Q_5's emitter has been grounded through Q_4. Q_6 is normally on. It can be switched off by Q_5 only if bit 1 has been selected and bit 1 is in the off state. Since the system already knows which bit it is looking at, the presence or absence of an output from the IC (through Q_7) tells the state of the bit.

If bit 1 has not been selected, the collector of Q_5 will be at a high enough level in either state to tend to keep Q_6 on. With Q_4 shut off, the emitter of Q_5 is at 1.5 volts and Q_6 is on regardless of the state of Q_5. Thus, Q_6 has to look elsewhere for a turn-off signal.

Third generation

The IC's qualify Sigma 7 as a third-generation computer since the design goes beyond that of tube and transistor systems. In another sense, it is a "third-and-a-half-generation" system because the scratch pads represent a move toward the expected organization of future generations. Many computer designers feel that when huge arrays of circuits can be batch-fabricated cheaply, it will be practical to make self-programing and self-organizing systems, or systems composed of many small computers working in parallel, or even machines composed of thousands of cells in which the system's control, logic and memory functions are combined.

There are strong resemblances between Sigma 7 and the "fourth generation." It can do any combination of time-sharing, multiprograming and multiprocessing [Electronics, Nov. 29, 1965, p. 71]. It can operate in real time, run several programs asynchronously, correct errors when programs are being run and automatically adjust its internal operation to changing user demands.

The scratch pads helped make this possible by avoiding the need to build a large quantity of fixed-capability hardware into the system. In each of their four major applications, described below and outlined in the block diagram on the facing page, the scratch pads have various optional applications and sizes.

IC specifications

V_{cc}	4.0 volts± 10%
I_{cc}	50 ma maximum
A_3 "1"	6.6 volts maximum
Output current	60 ma at 0.5-volt maximum
Operating temperature	5° to 70° C

Scratch pad I

The first application of the scratch pads is as general register blocks that implement the Sigma

Other scratch pads, other systems

The integrated circuits of today do not end the search for the ultimate component for scratch-pad memories. As the authors of the Sigma 7 article point out, it should be possible soon to double or quadruple the storage content of each IC. Furthermore, competing forms of high-speed memories are being sought by researchers.

The breadth of the development effort was indicated in February at the Solid State Circuits Conference when four companies gave reports on four distinctly different scratch-pad designs. A common theme—high-speed memories produced by low-cost techniques — emerged:

- Bell Telephone Laboratories, Inc. is developing a diode-selection matrix, with performance comparable to that of a transistor matrix, for a 1,024-word magnetic memory. The magnetic-storage elements are wires plated with magnetic alloy.
- Univac division of the Sperry Rand Corp. and Motorola, Inc. described a 64-word memory built by vacuum-depositing magnetic films and interconnections on a small glass plate and then attaching to the plate integrated circuits and other semiconductor components. An earlier report described the fabrication techniques [Electronics, Oct. 4, 1965, p. 102].
- Nippon Electric Co., Ltd. of Japan, is preparing an associative memory of integrated circuits made by the metal-oxide-semiconductor technique (MOS). The active elements in the storage cells are p-channel MOS transistors. With p-n-p transistor sense amplifiers, the cycle time is 150 nanoseconds, unusually fast for MOS circuitry.
- The Radio Corp. of America is also developing a high-speed MOS memory. The storage cells, constructed of complementary n-channel and p-channel MOS transistors, operate in 20 nanoseconds. A 16-word array containing 1,080 MOS transistors, including decode and drive circuitry, is expected to read in 50 nanoseconds and write in 75 nanoseconds.

Bipolar IC's. Several companies, in addition to Signetics Corp., are making scratch pad IC's with bipolar transistors. An IC containing 16 bits of storage and their associated circuitry will soon be introduced commercially by Texas Instruments Incorporated [Electronics, March 21, 1966, p. 144]. Sylvania Electric Products, Inc., a subsidiary of the General Telephone & Electronics Corp., will soon offer a 16-bit chip. Transitron Electronic Corp. is making them on custom order. Signetics also has a 16-bit project. Motorola, Inc. is developing storage arrays.

General Micro-electronics, Inc., is making large-capacity MOS IC's on custom order and GMe and other companies are also making off-the-shelf circuits, such as shift registers [Electronics, Oct. 4, 1965, pp. 84 and 96]. One GMe custom circuit stores 64 bits in a flatpack. One of its main features is low standby power, typically 10 microwatts per bit; the read-write cycle time is 500 nanoseconds.

For future use in its own computers, the International Business Machines Corp is readying monolithic memory circuits that are compatible with the hybrid IC's now used in its System 360 computers. One research IC, described

7's instruction set. Each set of registers, termed a register block, has 16 registers and is made up of four 16-byte modules. One such 16-word scratch pad is standard; there can be as many as 32 blocks, or 512 words.

A five-bit block-pointer register designates which of the 32 blocks is active: that is, which block is to be employed in the execution of an instruction.

These registers greatly improve the speed with which the computer responds to interrupts. Interrupt requests can be honored instantly by aborting long instructions in midexecution without loss of information, because operands come from storage into active circuit registers for instruction execution. Moderately long instructions (10 to 30 microseconds) are aborted and restarted upon return from an interrupt. Instructions longer than 30 microseconds are designed so they may be aborted almost instantly and then resumed from the point of interrupt upon return from the interrupt routine. Even though some instructions take 500 microseconds, it takes only 2 to 20 microseconds to initiate an interrupt response. The total response time is never more than 25 microseconds and it may be as short as 6 microseconds.

In conventional systems, interrupt processing requires the content of the general registers to be stored in core memory to preserve the operating

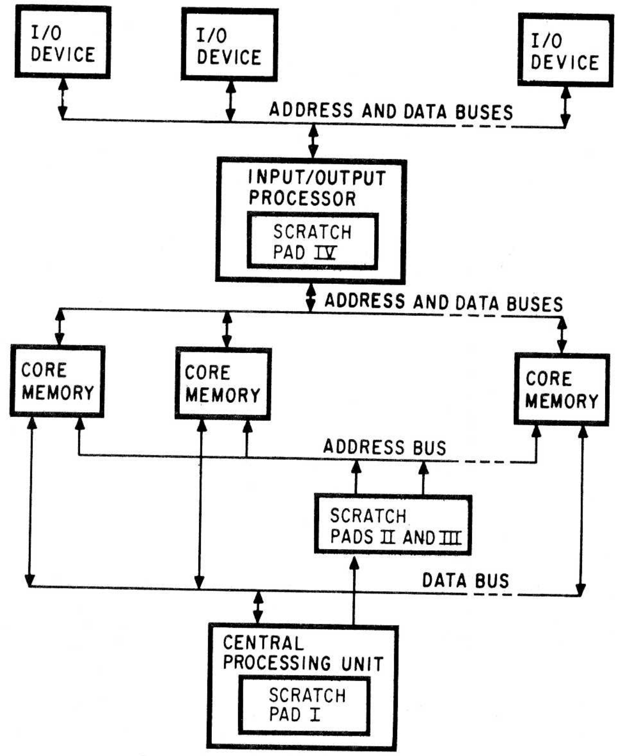

Main applications of scratch pads: I, as registers; II and III, memory mapping and protection; IV, addressing and local storage.

last fall, contains 16 storage cells and their control circuitry on 0.07-inch-square chip of silicon [Electronics, Nov. 1, 1965, p. 31].

The Fairchild Semiconductor division of the Fairchild Camera & Instrument Corp. has been working for more than a year on IC scratch pads [Electronics, March 8, 1965, p. 46]. Its latest model, a test bed for circuit, system and packaging design, is a mix of the company's standard logic circuits and special 36-bit storage cells. Four printed circuit cards, each carrying 160 IC packages, form a memory that stores 256 words that are 72 bits long.

Other systems. Scratch pads have been employed in small quantities in computers for years. Some small airborne computers make do with simple delay-line storages. Others have high-speed magnetic storage assemblies. Univac and the Burroughs Corp., for example, have been putting high-speed thin-film memories into their computers. The plated-wire type of thin-film memory is also in commercial use. Woven-wire scratch pads, resembling a piece of cloth, are built into France's newest computers, the Bull-General Electric Gamma 140 and 141 [Electronics, March 7, 1966, p. 304].

Computer manufacturers began using conventional logic IC's as scratch-pad elements about two years ago. One of the earliest applications was in Electronic Associates Inc.'s EAI 8400. The central processing and input-output control sections are made with discrete components, but built into the circuitry are stacks of IC registers that also handle several scratch-pad functions.

The EAI 8400's central processor has eight 32-bit registers, split into 16-bit halves. Four of these are "save" registers that act as bookmarks in computer programs. They permit interruptions and store intermediate results of programs. Another register stack in the input-output section aids in memory control and buffers variations in data-transfer speeds.

To improve the speed of the larger models of the IBM 360 computers, the register and scratch-pad circuitry is scattered throughout the logic circuitry. Instead of making separate plug-in modules, IBM puts storage circuits a few at a time onto the logic-circuit cards to shorten lead lengths and control wiring delays. The arrangements are designed by computers [IBM's design-automation techniques and the organization of the System 360 are described in the April, 1964, issue of the "IBM Journal of Research and Development"; the wiring techniques are detailed in Electronics, Nov. 1, 1965, p. 90].

The Models 65 and 75 of the System 360 have 25 registers that are each 36 bits long (four bytes of eight bits plus a parity bit for each byte). Among the functions which qualify the registers as scratch pads are working storage and local provision of instructions. In the larger Model 90, there are 76 one-word registers. Scratch pads in three smaller models are separate ferrite-core arrays. Twenty-four of the registers in each can be manipulated by the computer users.

IBM has been using hybrid integrated circuits. It is expected to use monolithic circuits, such as the 16-bit memory chip, as these emerge from the IBM labs and go into production.

George Sideris

environment of the interrupted program. The registers must also be loaded from the core memory to establish the operating environment of the interrupt process. Later, the interrupted program has to be returned to the registers.

The Sigma 7's multiple-register blocks preserve the status of the interrupted program in register by simply changing the content of the block-pointer register. The change takes place within the six-microsecond execution time of a single instruction that stores in memory the control state of the interrupted program and obtains from memory the corresponding control states for the interrupting program. In effect, the program-switching instruction causes the system to walk away from one block of registers and step into another with no additional time cost.

Scratch pad II

Scratch pad II stores constants, assigned by the supervisory program, for program relocation (main-memory mapping) and memory protection.

Mapping is essential for time-shared or multi-programed operation. It allows programs to flow continuously in and out of the main memory, occupying different fragments of the memory each time it is entered. The user's instructions generally call for a contiguous block of memory at one location, but honoring such reservations for room wouldn't be efficient. Instead, the system files the program parts as though it had a loose-leaf notebook. At one time a word may appear on page 8, word 9, and at another time at page 88, word 9. A memory of 131,072 words will have room for 512 words on each of 256 pages.

The system keeps track of the program's actual addresses in core by storing address-change constants in the scratch pad. The system refers to these each time it initiates an access to memory. The validity of each reference is controlled by 2-bit access-protection codes stored in the scratch pad.

The access-protection codes prevent one program from interfering with another. They also prevent snooping. One user can't get at another's proprietary information, such as the pay rates for key personnel, since the codes can be set to prevent it.

The memory-mapping scratch pad is provided by 16 cards, each carrying 16 bytes' worth of IC's. That works out to one byte per page for the 131,072-word addressing range. Core memory size is expandable to eight 16,384-word modules. If main memories faster than those now used become available, they can be plugged into the system.

The 512-bit scratch pad for 256 pairs of access protection bits is provided by four 16-byte modules. The map scratch pad and these modules are addressed concurrently.

Scratch pad III

Another 512-bit scratch pad stores 256 2-bit "locks" that provide write-protection control for a main memory of the maximum size, 256 pages. Each program is given a 2-bit "key" when it is given control of the computer. Write access is permitted a program when its key opens the lock assigned to a page (512 words). The key and lock codes must have values of 00 or some nonzero value.

The speed of the scratch pads allow the access-protection and the write-protection techniques to be performed sequentially without slowing down system operation.

Scratch pad IV

In the multiplexing input-output processor (IOP), the scratch pad speeds up I/O processing and avoids idling of the central processor unit (CPU) while data is read into or out of relatively slow peripheral equipment. Each IOP can handle 32 peripherals and there can be eight IOP's in a system.

The IOP gets its initiating instructions from the CPU, but operates independently thereafter. The CPU specifies the peripheral, the direction of data transfer and where transfer should begin in memory, and the number of words to be transferred. The CPU proceeds with the program, while the input-output control information remains in the IOP scratch pad until the data arrives.

The basic IOP scratch pad provides for eight channels. It is built with five cards, providing 16 words of 40 bits. Generally, 10 bytes control a channel. The scratch pad can be expanded in sets of sixteen 40-bit words to a maximum size of 64 words. That size will control concurrent operations on 32 channels.

The authors

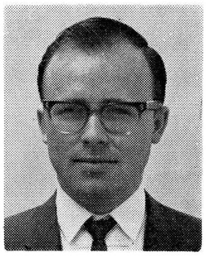

Gene B. Potter is responsible for semiconductor applications, including custom integrated circuits, at Scientific Data Systems, Inc. He is an honor graduate of the University of California and joined SDS in 1964.

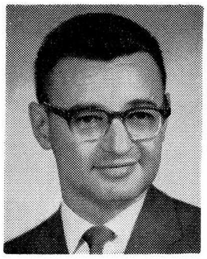

Jerry Mendelson, another University of California honor graduate, directs the planning and design of all SDS products, including the Sigma 7 computer. Until 1964, he was manager of the advanced programs staff at Litton Industries, Inc. He has eight computer patents.

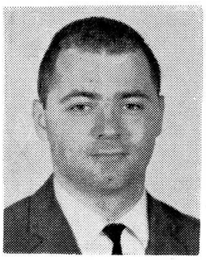

Sam Sirkin, a naval aviator from 1957 to 1962, joined Signetics, Inc., in 1964 after graduation from the University of California. He was responsible for the design of the SDS memory circuit and is now working on ultrahigh-speed designs for digital integrated circuits.

УДК 681.142.65: 621.382.82

Новое поколение вычислительных машин с буферными запоминающими устройствами на интегральных схемах[1]

Поллер и Мендельсон

Фирма „Сайентифик дейта системз"
(Санта-Моника, шт. Калифорния)

Сиркин

Отделение „Сайнетикс" фирмы „Корнинг гласс уоркс"
(Саннивейл, шт. Калифорния)

Если на одной пластинке кремния размещается 8-разрядная ячейка памяти вместе с буферными обслуживающими цепями, то в вычислительную машину можно ввести большое количество быстродействующих запоминающих устройств, в результате чего автоматическое распределение машинного времени становится экономичным.

Серийное производство интегральных схем (ИС), содержащих законченную ячейку памяти на несколько двоичных разрядов вместе с цепями записи — считывания и выходными цепями, делает возможным создание нового поколения ЭВМ. В обычные системы обработки информации можно ввести большое количество малых запоминающих устройств на ИС, так называемых буферных запоминающих устройств (БЗУ), которые работают со скоростью логических цепей ЭВМ.

Применение БЗУ в ЭВМ невысокой стоимости позволяет получить такое быстродействие, емкость памяти и универсальность, какими ранее обладали только огромные дорогостоящие системы. В результате стоимость обработки больших массивов информации может быть снижена вдвое.

Примером эффективности применения БЗУ в качестве основного блока структуры ЭВМ может служить серия ЭВМ «Сигма 7» фирмы «Сайентифик дейта системз» («СДС»)[2]. Типовой вариант ЭВМ «Сигма 7» стоимостью около 500 тыс. долл. имеет скорость ввода-вывода информации около 160 млн. *дв.зн./сек*.

Такая высокая производительность получена благодаря новой организации памяти и структуре арифметического устройства, построенного с применением от двух до 12 БЗУ. БЗУ представляют собой сменные блоки емкостью от 16 до 512 слов. Ячейки памяти, выполненные на монолитных ИС, показаны на стр. 2 обложки и на фиг. 7. Эти ИС были разработаны при сотрудничестве фирмы «СДС» и «Сайнетикс» и выпускаются фирмой «Сайнетикс» для «СДС». Каждая ячейка содержит 8-разрядный регистр на триггерах с цепями записи, считывания, управляющими и буферными цепями.

Печатная плата, на которой располагается 16 модулей ИС вместе с цепями управления (см. схему, фиг. 4), представляет собой блок памяти на 16 символов (один символ — 8 двоичных разрядов). Этот блок является основным блоком, использованным для построения всех БЗУ машины «Сигма 7». Например, четыре таких блока образуют БЗУ на 16 слов. Если требуется бóльшая емкость БЗУ, то к адресным и числовым шинам ЭВМ подключаются дополнительные блоки. В блоках содержатся все необходимые для работы цепи, в том числе дешифратор адреса на 128 слов.

Главное назначение БЗУ состоит в том, чтобы сократить время реакции системы на многочисленные запросы путем уменьшения ее зависимости от работы основных ЗУ для многих операций, а также благодаря возможности немедленной реакции на требования приоритета. ЭВМ «Сигма 7» может обслуживать до 200 абонентов, связанных с машиной с помощью дистанционных оконечных устройств, при одновременном выполнении нескольких обычных программ (например, составление описей) и работе периферийного оборудования машины. Функциональное назначение БЗУ иллюстрируется на блок-схеме фиг. 8.

Многорегистровая архитектура системы, индексация, переадресация (перемещение программ в памяти) и защита памяти требуют многократных обращений к БЗУ в течение одного цикла основного ЗУ. Поэтому, чтобы получить наилучшие результаты, БЗУ должны работать по меньшей мере в пять раз быстрее, чем основная память; в противном случае они не будут эффективно взаимодействовать с логическими и управляющими блоками системы.

БЗУ ЭВМ «Сигма 7» в среднем имеют время записи около 90 *нсек* и время считывания 60 *нсек*. Основные ЗУ, выполненные на ферритовых сердечниках, имеют длительность цикла 1,2 *мксек* и эффективную длительность цикла 700 *нсек*, если работа блоков основного ЗУ происходит с перекрытием по времени.

Экономическая эффективность

Два года назад, когда начиналась разработка системы, казалось, что решающей проблемой будет стоимость БЗУ. Разработчики ЭВМ давно знают,

---

[1] *G. A. Poller, J. Mendelson, S. Sirkin* Integrated scratch pads sire new generation of computers, pp. 118-126.

[2] *Электроника*, № 6, 1966, раздел «Обозрение электронной техники».