Product development profile

The Intel 1103: The MOS memory that defied cores

A young company banked on a silicon-gate p-channel dynamic MOS random-access memory to replace cores in mainframes: a Honeywell team aided in the design, which has become an industry standard

by George Sideris, *San Francisco bureau manager*

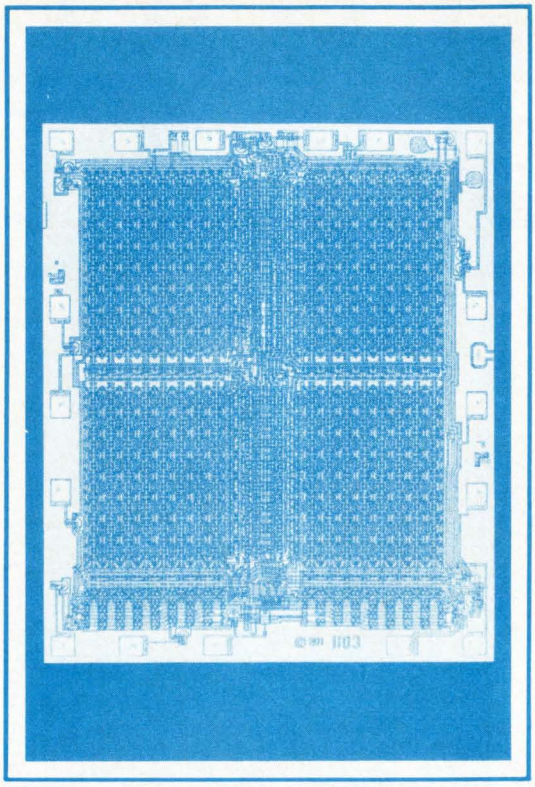

☐ Next month, Intel Corp. will start reducing production of the 1103 random-access memory and phase in an easier-to-use replacement, the 1103A. The new chip, which can be slipped directly into 1103 sockets, has simpler timing procedures and higher speed, and it is TTL-compatible. When the original 1103 design is retired in a year or so, the frontier days of solid-state memory development will end. But the trends set by the 1103 in 1970 will probably continue through the decade.

The 1103 showed that dynamic MOS RAMs could compete with cores in the computer mainframe and peripherals markets. It firmly established semiconductor-memory organization in multiples of 1,024 bits. And it made silicon-gate MOS the dominant process technology throughout the world of non-IBM memories. Few contenders for the market created by the 1103 depart radically from that concept.

Fairchild's 709 amplifier and Texas Instruments' 7400 TTL gate, two other landmark circuits, have been used by more engineers than the 1103. However, Intel's RAM is unquestionably the greatest economic success. Intel's second-source list reads like a "Who's Who in the Semiconductor Industry."

Costs drop

Since the 1103 was introduced at $60, cost and volume curves have quickly intersected like two parabolas. This year, 6 million 1103s will be shipped at an average price of $4—0.4¢ a bit—estimates A.C. "Mike" Markkula, Intel's sales manager.

True, 6 gigabits is small, compared with core volumes. But huge numbers of solid-state RAMs will be used in new computers. Moreover, 4,096-bit versions of the silicon-gate dynamic RAM built with a high-density n-channel cell structure—costing still less per bit—are already being aimed at that market. And with the 1103-type RAM, megabyte solid-state add-on memories are becoming commonplace in the IBM System/370 peripheral market.

Perhaps the 1103's true measure is that system designers have been willing to struggle with its idiosyncracies for months at a time because it is economical. "They hate it, but they use it," Markkula jokes. Others have used more pungent phrases to describe the 1103. But whatever phrase is used, the fact is the 1103 was the first widely used device that put 1,024 bits of memory on a chip in a form that was economical for the user.

The Intel and Honeywell Inc., design teams who developed it accepted the system design difficulties in order to cram 1,024 bits on a manufacturable chip. Timing with the 1103 is difficult because it requires overlapping cycles that are timed to narrow windows. However, the quirks will be retired with the 1103, Intel promises.

Early explorations

In 1968, when Intel was spun out of Fairchild Semiconductor, it was widely predicted that semiconductor memories would find a $500 million market in the 1970s and perhaps a $2 billion market in the 1980s. Robert Noyce, formerly Fairchild's general manager, and Gordon Moore, who was engineering manager and R&D director at Fairchild, started Intel with that market as their goal.

However, a review of memory technology convinced them that none of the then-popular concepts could catch up quickly with the falling cost curve of core memories. They decided to pursue Schottky TTL for a

quick entry into the small, high-speed memory market and silicon-gate MOS for the longer haul into the mainframe market.

The silicon-gate technology was then developmental. The first definitive paper was published in 1968 by Bell Telephone Laboratories researchers, but some of Intel's founders had begun their silicon-gate work in 1967 at Fairchild. Fairchild's initial silicon-gate product was a multiplexer, though [*Electronics*, Sept. 15, 1969, p. 67].

Noyce and Moore, already contemplating 1,024-bit RAMs, foresaw silicon-gate MOS as the most likely solution to chip-density and yield problems. In addition, the silicon gate's low threshold voltage offers high speed and can be made bipolar-compatible.

Intel's MOS staff cut its teeth on silicon-gate design with the 1101, a 256-bit static RAM introduced in 1969. Initial yields of good chips were pleasantly high—10% or better, compared with an anticipated range of 2% to 5%. Metal-gate yields were then running around 5%.

The 1101 was a significant advance. It could be operated directly by system logic because it had TTL-compatible address decoders and sense amplifiers on the chip. Low logic overhead made it useful in small systems, but the complex static cell structure (Fig. 1a) made the 1101 too slow and costly for mainframe memories.

A new trail

Static RAMs were the style in 1968. In the proceedings of the 1968 Fall Joint Computer Conference, for example, the only paper that discussed solid-state memories covered static RAMs. The author, W. B. Sander, of Fairchild, accurately predicted: "Both (bipolar and MOS) will be developed, and the final edge of one over the other will require some dramatic development in one of the technologies."

At Intel, M. E. "Ted" Hoff Jr., a young Ph.D. from Stanford University, devised a simple dynamic storage cell. It needed only three transistors, compared with the conventional four (Fig. 1b). More importantly, the intraconnections, which in conventional cells occupied more room than the transistors, were sharply reduced.

Since the cell promised three to four times the density of a static design, Noyce, and Moore started a team developing dynamic RAMs. Leslie Vadasz, Joel Karp, and Hoff were assigned to design circuits with variations of the cell. (Vadasz, the team leader, is now Intel's engineering manager, Hoff is applications research manager, and Karp, long a mainstay in Intel's design staff, has recently joined Intersil Inc.)

They worked mostly on three versions: 1102, 1103, and 1104. The 1102 (Fig. 1c) and 1103 (Fig. 1d) were 1,024-bit designs (Fig. 2 is the chip diagram) that could be paralleled to form 1,024 words with any number of bits per word. The 1104, a 512-bit device, was soon rejected as not cost-competitive (it was later produced experimentally to test n-channel processes).

At first, the 1102 was deemed the most promising design. Hoff wrote an article about it [*Electronics*, Aug. 3, 1970, p. 69]. Karp and William Regitz, then a Honeywell engineer, described it in a paper at the 1970 International Solid State Circuits Conference. But at the conclusion of system trials in 1970, the 1103 got the nod.

Another bullet that had to be bitten during the conceptual design stage was the fact that, unlike static RAMs, dynamic operation would raise system overhead and design costs. However, core memories also have

The RAMs to come

The Intel Corp. 1103 merely blazed the trail to mainframe computer memories. Coming closely behind are several more efficient MOS random-access devices that are larger by multiples of 1,024 bits. Two with quadrupled capacity are already nearing production.

Microsystems International Ltd., on its own, has enlarged the 1103 to 4,096 bits with n-channel substrates, doubling the chip size. And Intel has started pilot-line production of an n-channel 4,096-bit RAM. Conceptually, this RAM is similar to the 1103A, retaining a three-transistor cell and single-clock operation. However, the n-channel design is TTL-compatible and has an on-chip sense amplifier—features that will enable it to go further than the 1103A in reducing memory overhead.

The n-channel design makes it only about 50% larger than the 1103. Its chip measures 137 by 164 mils, compared with the 1103's 113 by 139 mils. The chips will probably cost about four times as much as the 1103 to manufacture, but the overhead reduction will make them less costly at the systems level (a rule of thumb is that, once a chip reaches 100 mils on a side, each 15% increase in area approximately doubles processing cost.)

Intel will introduce a medium-speed version of the 4,096-bit RAM this summer, following up with a high-speed version later in the year. Whether or not these, too, will become industry standards is still moot. This time, memory-system designers will have a choice of many RAMs. But each one builds on the basic 1103 concept—a simple dynamic cell.

Undoubted originals. Robert Noyce and Gordon Moore, founders of Intel, join in admiring some of the first 1103s produced.

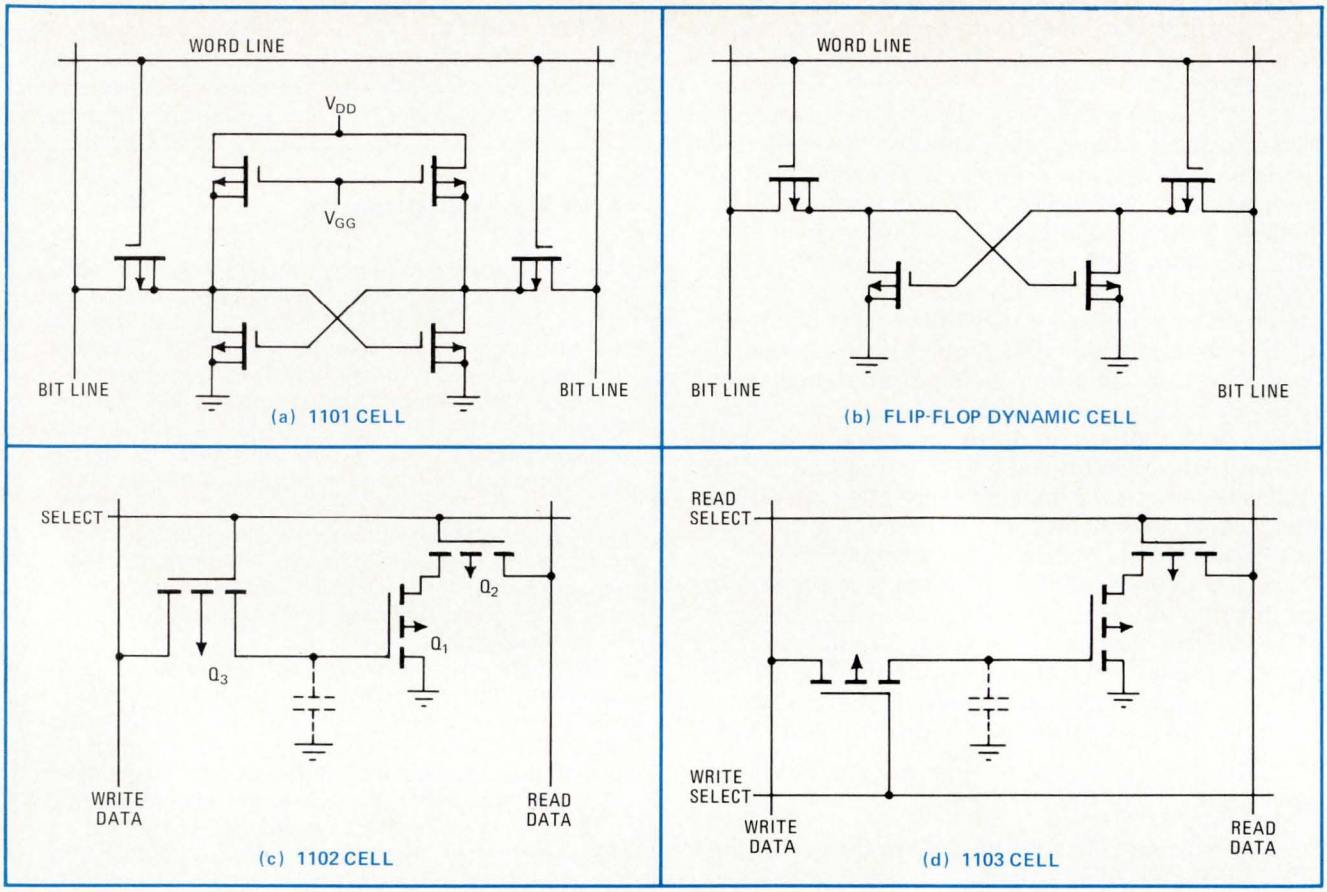

1. Shrinking cells. Static cell used by Intel in 1101 RAM was too slow and costly for mainframe memories (a). Leapfrogging conventional dynamic cell designs (b), the 1102 and 1103 1,024-bit RAMs contained three-transistor cells (c and d). Though smaller in area because it had one select line, the 1102 cell was rejected in favor of the easier-to-make 1103 design.

high overhead costs—divided among many bits. The team went for broke.

To increase speed, the team opted for large logic swings in hopes of winning additional cost-performance tradeoffs in the inevitable comparison with cores. The decisions were to require multiple clocks (Fig. 3), powerful drivers, level shifters, sense amplifiers, timing and control subsystems, and other support circuitry shown in Fig. 4.

The dynamic design promised brutal noise levels and noise-related problems in system control, timing, and sense circuits. Heavy capacitive loads would have to be driven at rates around 1 volt per nanosecond, assuring large current surges in the printed-circuit traces. Then, there were such complications as cell-refresh and volatility (loss of data) to be considered.

The 1103 requires three clocks with carefully timed overlays: precharge, cenable (chip-enable) and write. Precharge was not really an innovation, Vadasz remarks. Precharge and cenable are comparable to an MOS shift register's 0_1 and 0_2 clocks—one charges up the MOS capacitances so that the nodes can be discharged very rapidly or not be discharged in the following logic operation.

Of more lasting importance, Vadasz thinks, is a technique of building into the MOS transistors varactor diodes that bootstrap the logic levels. Once Intel's secret method of speeding up decoders, varactor bootstrapping is periodically rediscovered, he notes.

Uncertain whether or not the designs would work well in a memory system and reluctant to choose between the 1102 and 1103 arbitrarily, Noyce and Moore sought the opinions of potential customers. Here the risk that industry gossips might cause damage by broadcasting the design ideas was less immediate than the risk of expending large amounts of time and money on a design that might fail.

A new partner

Honeywell's interest in developing a standard product to compete with cores coincided with Intel's. If either the 1102 or 1103 evolved into an industry standard, Honeywell would be assured a supply of low-cost components. Other semiconductor manufacturers would surely second-source any circuit chosen by major computer manufacturers.

Honeywell assigned to the project a group at the computer plant in Framingham, Mass. William Jordan was group leader, William Regitz the principal components designer, and Henry Bodio the systems engineer (Regitz went to Intel as manager of MOS-memory engineering in 1971, then Jordan became manager of Intel's new Systems division, and Bodio joined the division as engineering manager).

In effect, Jordan's group acted as the non-IBM-computer industry's semiconductor-memory steering committee. They helped Vadasz' team firm up timing specifications, skew tolerances, control configurations, and a

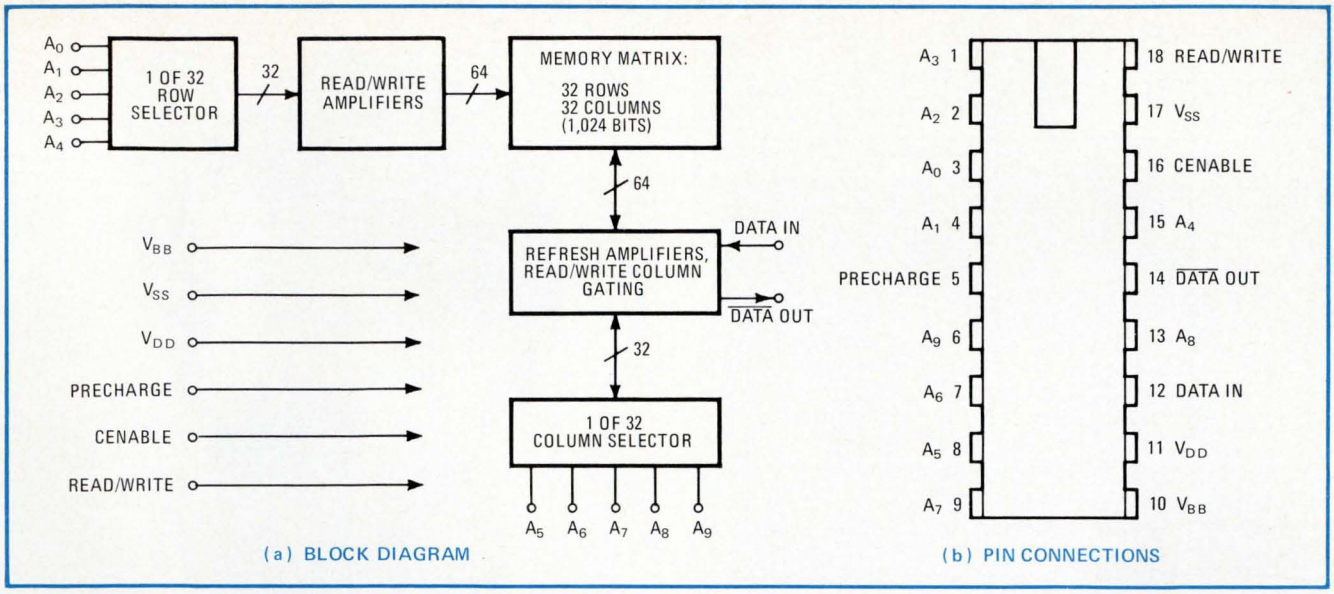

2. Chip block. Outline of the 1103 chip is identical with the 1103A, except for the latter's elimination of precharge.

host of other system design aspects during 1969 and 1970.

But the 1102 and 1103 were largely paper designs when the work started, Jordan recalls. "We worked right along with Vadasz's group." First, they assembled 16-bit prototype cell arrays. When those were debugged, 1,024-bit arrays were prepared. Finally, prototype systems were assembled.

As the quirks became known, they were compensated for by changes in chip or system design. For instance, certain address sequences that "bombed out" stored data forced chip changes. Timings sensitive to process or temperature variables might be compensated on the chip and in the system logic. The 1103 is sensitive to certain system operating conditions, involving, among other things, certain combinations of line-charging rates, addressing patterns, and timing patterns. Determining how to avoid these problems through system design took much of the group's time.

At Intel, thousands of chips were "wrung out." Known also as characterization, wringing-out means that chips, produced in batches under varying processing conditions, are tested under every conceivable combination of operating conditions. One objective is to find out what process controls give the best yield to the customer's specifications—and in this case, the customer was to be the computer industry. Thousands of performance curves were laboriously prepared.

The showdown

The 1102 emerged from development first. It was the more advanced concept and fitted on a smaller chip, which could eventually mean higher yields and lower cost. Early in 1970, Jordan presented Honeywell's engineering manager with an H516 minicomputer with an 1102 memory instead of the standard core memory. Soon afterwards, the 1103 design was completed.

"We chose the 1103 because it was the more conservative," Vadasz reports. Jordan indicates that Honeywell decided to buy the 1103 because of fears that 1102 might have yield problems that would impede second-sourcing (the 1103 has two operating voltages, while the 1102 cells operated at three voltage levels and required tighter tolerances).

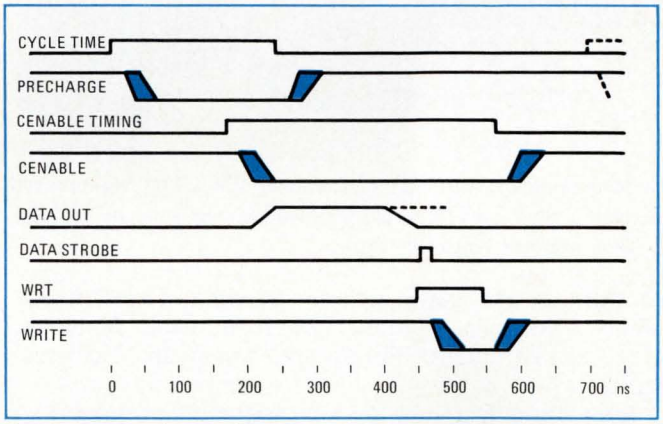

3. Cycle slowdowns. The 1103's complex timing diagram is now world-famous. This simplified version shows how system skews and tolerances add nearly 200 nanoseconds to the basic cycle time.

"We decided on the part with the highest confidence level," says Jordan. "Today, we would probably choose the 1102, but that is academic now. The whole key was to settle on a standard part in order to realize the economies of scale in high-volume production."

The 1102 was also somewhat slower than the 1103. Was that a factor? No, says Jordan, because both were fast enough to compete with 18-mil and 20-mil cores. Besides the speed difference became slight, once system delays and skews were added to basic operating times (see Fig. 4).

Intel quickly established a second source. Microsystems International Ltd. took a license. In return, Intel helped MIL buy equipment and set up a plant, and Intel also supplied mask sets and taught MIL to make the 1103. In a short time, the Canadian company was more than a second source—MIL became Intel's chief competitor.

Markkula says of companies that have copied the design without the formalities of licensing, "I stopped

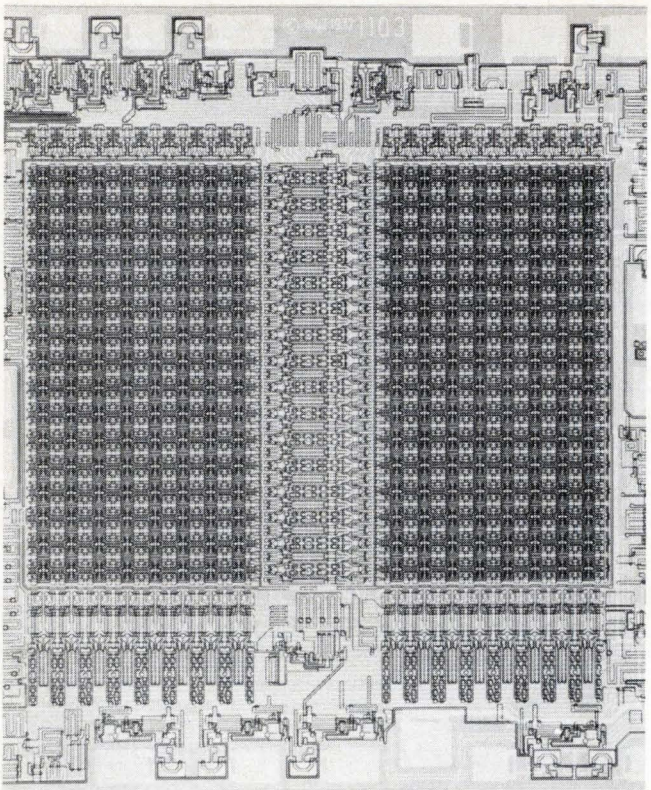

4. New standard. Retiring the venerable 1103 will be this improved version. Called the 1103A, this improved version is faster, TTL compatible, and has simpler timing.

counting at 18." They include Fairchild, Texas Instruments, Motorola, National Semiconductor, Signetics, General Instrument, Philips, and American Micro-systems.

Even though a number of small-systems manufacturers adopted the 1103 quickly ($3.9 million worth of 1103s were sold in the latter part of 1970), bigger companies hung back to make evaluations and design studies. Mainframe computer shipments started building up in mid-1972. Among major systems containing 1103s, Markkula lists the Digital Equipment Corp. PDP 11/45, Burroughs 7400, Texas Instruments Advanced Scientific Computer, Univac 9480, and the Honeywell 5800. Intel itself is penetrating the IBM market—Jordan's Systems division makes a 9-megabit add-on memory for IBM Systems/370-155 and 370-165.

The 1103 has also generated a kind of sub-industry of producing special clock drivers, address latches, and other support circuits. These have helped bring overhead costs down to a small fraction of a cent per bit.

Round two

When the dust of development settled, Hoff wrote a 28-page note to explain the operation of the 1103. It was sprinkled with warnings about operating conditions that could cause the 1103 to malfunction, and it recommended ways to solve those problems.

To spare system designers further grief—and to make the 1103 more competitive—Intel began redesigning the 1103 late in 1971. The result was the 1103A, scheduled for introduction in May of this year after only six months of wringing out and system trials. Among the

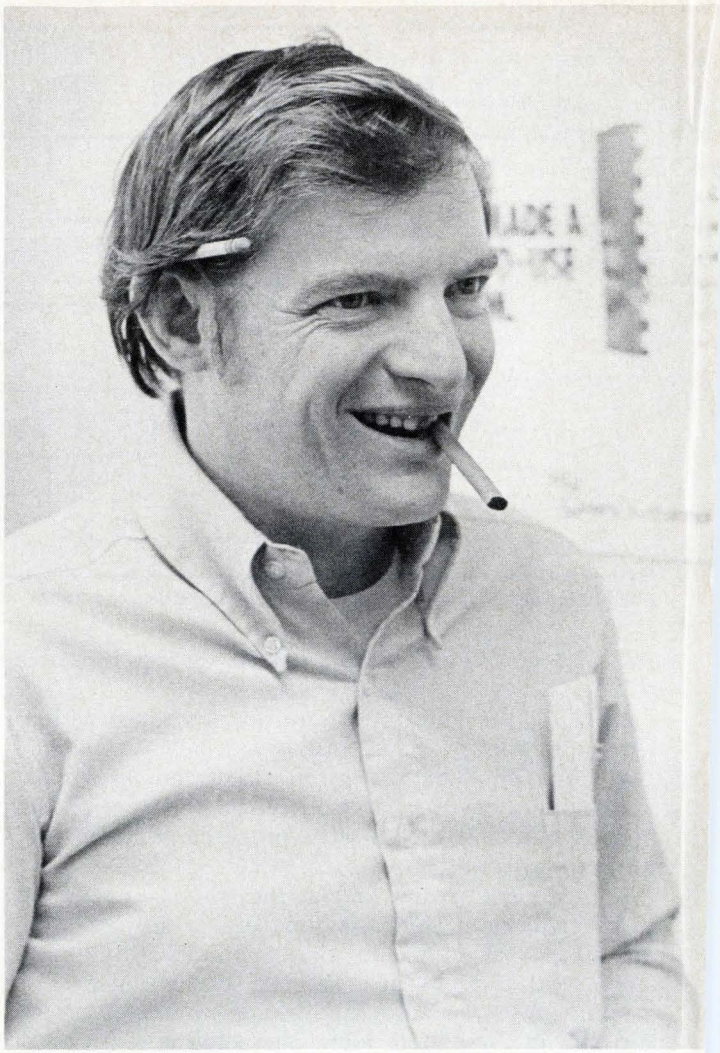

Circuit fixer. At Honeywell, William Regitz helped debug the 1103 and built prototype array assemblies.

changes cited by the designers are:

- Precharge and critical timing problems are gone.
- Sensitivity to process, temperature, and timing variations is reduced.
- Access time is almost halved in system operation.
- Address buffers are on the chip.
- Standby power is 1/40th that of the 1103, and operating power drops more rapidly with clock frequency.
- The chip is smaller, promising higher yield and lower cost than the 1103 might achieve.

But, Regitz stressed, the 1103A was developed as a direct "socket replacement" for the 1103. Existing system designs need not be changed because:

- Package pinouts are the same (the precharge pin is not connected to the chip).
- Cycle time is the same.
- Cell design and refresh methods are the same.
- Clock formats are the same (except that precharge can be removed).

Regitz calls the 1103A a single-clock RAM. On the leading edge of cenable, one-shots and other chip delays take the place of an 1103 system's timing overlaps. The chip itself controls timing sequences. When the read pulse or write clock arrives, the chip is all set up for

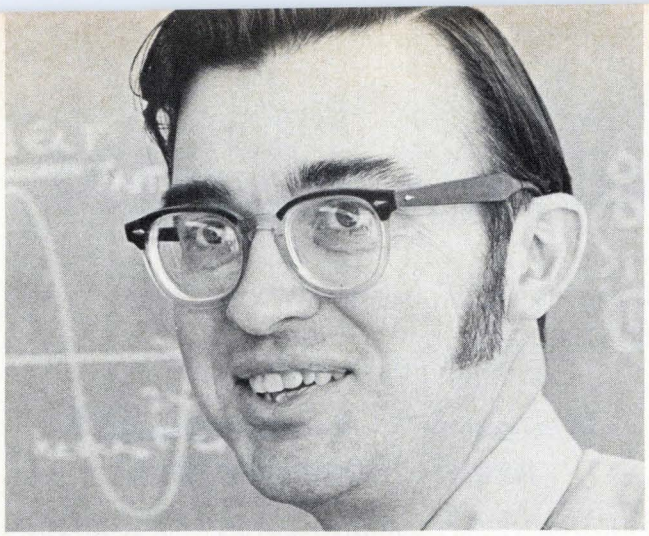

Cell mate. Ted Hoff is credited with inventing the three-transistor cell that sparked the 1103 dynamic design.

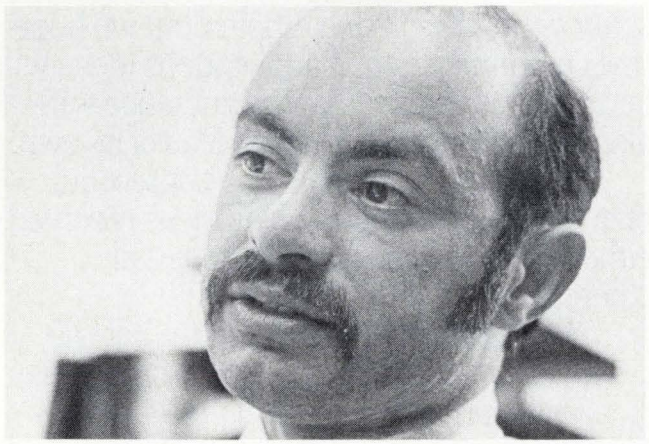

Happy Hungarian. Leslie Vadasz, who was head of the 1103 design team, now manages all Intel engineering.

Systems man. William Jordan led the Honeywell group that tried out early 1103 designs in computer memories.

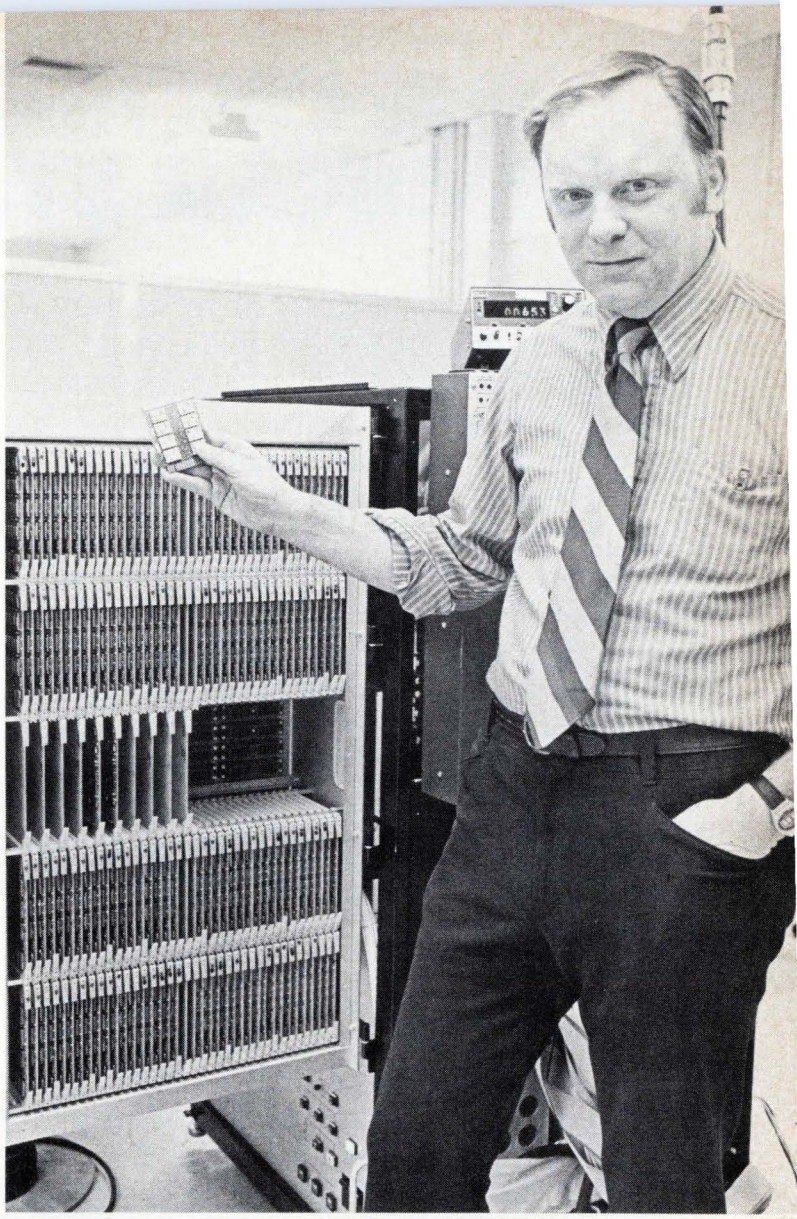

the operation. When cenable goes off, inverters on the chip keep the logic nodes charged (cells are refreshed as in the 1103, however).

For the older chip, the address input had to be stabilized with latches throughout the cycle time. Now, inputs are picked up by a buffer register on the chip. The buffers are isolated within 100 nanoseconds of the onset of cenable. After that, the addresses can vary without affecting decoding.

A system with an 1103-type timing controller need not be redesigned. But if the cycle is tightened to reflect elimination of precharge, overlaps, and skew tolerances, access time will drop from about 450 to 250 ns and cycle time from about 675 to 650 ns, Regitz estimates.

To conserve power in an 1103 system, precharge had to be decoded (addressed to selected memory segments) and unselected segments switched into a power-down mode. Not counting driver and control-logic dissipations, power consumption averaged 300 mW per chip without precharge decoding and 100 mW with decoding.

In operation, the 1103A consumes full power only during cenable transitions. The chip is driven through dynamic buffers that dissipate little dc power. The on-chip logic is all made of low-power dynamic circuits controlled by cenable. When cenable turns off, the chip drops into a low-power mode automatically.

Even though the 1103A is more complex than the 1103, the chip is smaller. Familiarity with silicon-gate processing allowed the designers to shrink the cells to 1.8 square mils—exactly the same area as single-transistor cells that are proposed for 4,096-bit RAMS, Vadasz points out.

He ruled out single-transistor cells since they require special sense amplifiers on the chip and would have made the 1103A incompatible with system designs based on the 1103.

Regitz adds that the 1103A is more tolerant of processing and temperature ranges than the 1103 because of the on-chip timing controls. The timing circuits tracks better in monolithic form. "We wouldn't produce the 1103A—we'd stick with the 1103—if yields were not improved," he asserts.

Although the 1103 is near retirement, Markkula doubts that it will fade away before 1980. The 1103 concept is not yet fully developed, he points out. But Intel's original goal of developing a standard RAM has merely been amended to one of providing a smooth transition from one standard part to the next. □