THE DEVELOPMENT OF MONOLITHIC INTEGRATED CIRCUIT MEMORY CHIPS AT IBM CORPORATION

by

Michael A. Rappa

Strategic Management Research Center
University of Minnesota
271 19th Avenue South
Minneapolis, Minnesota
55455

Preface

The following paper is a description of the events which led to the development of the first monolithic integrated circuit(IC) memory devices at IBM, and the process by which the new technology came to replace magnetic core devices in the computer's main memory unit. The story spans a decade, beginning with the formation the Components Division in 1961, and ending with the announcement of the System/370--the first commercial computer to implement an IC main memory in 1971.

The account presented here is primarily based on interviews with a number of IBM employees who were involved with IC memory development during the 1960's. They include: Ben Agusta, Paul Bardell, Erich Bloch, Paul Castrucci, Edward Davis, Jack Gersbach, William Harding, Robert Henle, Ed Hee, and Gerry O'Rouke. Other important sources included company newsletters, and testimony gathered in the Justice Department's antitrust case against IBM and subsequently published in IBM and the Data Processing Industry, by Franklin M. Fisher, et al.

IBM in the Early 1960's

During the late 1950's, IBM emerged as the industry leader in the manufacture of electronic data processing equipment. The company's second-generation computer machinery, the 1400 and 7000 series, gained wide popularity among users. The IBM 1401, in particular, just two years after its announcement in October 1959, became the most popular computer system yet produced by the industry. However, despite IBM's early success the company's management did not rest easy in its position. It was clear the pace of technological change would quickly ensure obsolescense if management was content with its current products.

By 1961, IBM's product mix included sixteen different processors comprising seven different lines of computer systems. Manufacturing responsibilities were split among two main divisions: the General Products Division(GPD) and the Data Systems Division(DSD). GPD was responsible for small and intermediate systems, including the 1400 series; while DSD had responsibility for high-end systems, such as the 7000 series. Each division conducted its own development programs to extend their current product lines.

Given the need to keep up with technological advances, serious debate arose over the company's ability to properly fund and coordinate development efforts among the two divisions, and their several product lines. During an executive committee meeting in April 1961, IBM's Chairman, Thomas J. Watson, Jr., expressed his concern:

> "Our competition is getting stiffer all the time....The best way to meet this competition is to keep our prices competitive. Prices involve costs and earnings....We need constantly to spend large sums in research and development of new products which will not produce revenue for some years to come. Without funds for this vital expense, competition will eventually surpass IBM."[1]

One executive who shared Watson's concern about IBM's ability to stay on the forefront was DSD's director of systems development and planning, B.O. Evans. It was Evans' responsibility to evaluating the leadership potential of the 7000 series extension, the 8000 series. With the 8000 program already nearing completion in 1961, Evans took issue over whether or not IBM should begin manufacturing the series. Evans argued that the 8000 series was a collection of different processors with loose ties in their structure; therefore, the series would provide users with little flexibility as their data processing requirements changed. Furthermore, Evans felt strongly that the transistor technology being used would be rendered obsolete by competitors not long after the 8000's market debut.

Evans outlook on the leadership potential of the 8000 series alarmed DSD's top management. As a result, T.V. Learson, IBM Vice President and group executive overseeing DSD and GPD, placed the 8000 program on hold. Learson then proceeded to assemble a task force to come up with a well-reasoned plan for IBM computer systems for the 1960's. Known as the SPREAD committee, Learson appointed GPD's Vice President of Development, J.W. Haanstra as chairman, and Evans as vice-chairman.

In their recommendations to Learson in December 1961, the SPREAD committee called for the "...termination of the proliferation of IBM products and the development of a family of compatible processors which

would employ a common component technology,... a compatible set of peripherals, and a compatible operating system."[2] According to the committee, the family concept would permit IBM to focus its resources and creative effort on the development of a leading-edge computer that would provide users with a maximum degree of flexibility.

In a bold and controversial decision, IBM's executive committee adopted the SPREAD committee recommendations early in 1962. Now, all of the company's development efforts were to be concentrated on a single program, called the New Product Line(NPL). The goal of NPL was to produce a family of computer systems, the System/360, that would become a price/performance leader in all user applications.

As Fisher explains, "the magnitude of the commitment--the devotion of vitually the entire business to that concept--carried with it a risk of staggering proportions. Both internally and externally the IBM System/360 program came to be referred to as a 'you bet your company' undertaking."[3]

The Formation of the Components Division

When Erich Bloch joined IBM in 1952, he was one among a small group of electrical engineers hired to develop the company's first computer system, the IBM 701. Later, as engineering manager of IBM's STRETCH project, Bloch was involved in the development of a method of packaging discrete semiconductor logic components called Standard Modular Systems(SMS). SMS was a significant advance in packaging technology that resulted in greater uniformity, reliablity, serviceability, and ease of

manufacture. And although the STRETCH computer, a supercomputer for its time, was not a financial success for the company, the SMS components became a standard feature in IBM's second-generation processors.

Early in 1961, while manager of systems development in DSD, Bloch was appointed chairman of the Advanced Technology Study(ATS) committee. The ATS committee was responsible for recommending the logic component technology for future IBM computer systems, and to set schedules and cost objectives for implementing their plan. At that time, the committee was faced with choosing among three alternatives: (1) an enhanced version of the SMS technology currently used; (2) monolithic integrated circuits under development at Texas Instruments and Fairchild; and (3) a form of hybrid circuits (between discrete and integrated cicuits) being developed within IBM, called Solid Logic Technology (SLT).

After careful evaluation of development efforts being conducted by both systems and components manufacturers, the ATS committee released its report in April, 1961. It was the committee's recommendation that IBM develop SLT components for use in its next-generation processors, and that the SLT program be given "high priority" status within the company.

The committee rejected further attempts to extend SMS technology because members felt its limits had been reached. This left much of the debate for deciding between integrated circuits and SLT components. The committee was briefed by representatives from Texas Instruments and Fairchild regarding the recent developments in semiconductor integrated circuits(ICs) at their firms. Bloch and other ATS members were convinced of the potential of ICs to meet performance specifications, and this fact sparked a wider concern: it was clear that breakthroughs in

component technology would be the driving force behind innovation in the next generation of computer systems. Indeed, this had shown itself to be true over the past decade, with the impact of semiconductor transistors.

In light of this, the ATS committee felt strongly that IBM needed to become more intimately involved with the component aspect of the computer system. The committee recommended that IBM begin in-house manufacturing of components used in its computers, and rely less on vendor procurement.

The desire to bring component manufacturing in-house affected the committee's decision between IC and SLT components. With no in-house experience in the development of ICs, the committee doubted whether the company could produce the new circuits at a reasonable cost and within the time-frame desired. On the other hand, the SLT technology was already under development within IBM, and the accumulated experience to date suggested SLT components would meet both performance and cost criteria.

The committee's recommendation for in-house component development was by no means readily accepted among the company's senior executives. Indeed, some executives were still learning to cope with IBM's involvement in computers (EDP revenues at this time were still less than half the firm's total sales.) But the proposal was appealing in a number of respects: it would allow greater coordination between component-level and system-level development efforts; it would enable the company to produce proprietary components which would give IBM computers a competitive edge; and it would reduce component costs by eliminating middleman profits.

But success in the components business was not a foregone conclusion. Afterall, the company had little or no experience in component manufacturing. Moreover, many new engineers would have to be hired, and large investments made in plant and equipment; and all in a very short time. Even with these hurdles, Bloch and other ATS committee members remained persistent in their belief that IBM's leadership in the computer industry required the company to become more directly involved in the component business.

After a few months of deliberation, IBM' executive committee adopted the recommendations of the ATS committee. On July 13, 1961, the company's chief executive placed John Gibson as President of the newly formed Components Division(CD). CD's mission: to develop and produce or procure logic and memory circuits for all IBM computers. The first objective on the division's agenda was to develop and manufacture SLT components in accordance with the recommendations set down by Bloch's ATS committee.

Although some executives would continue to question IBM's decision to form CD, there was no turning back. When Haanstra and Evans released the SPREAD committee report four months later, SLT components were at the core of the NPL program. Bloch's specifications for SLT cost/performance and availablity were essential to meeting NPL's overall cost/performance goals. The company's commitment to concentrate its efforts on a single family of compatible computers, predicated on SLT, meant there could be no second-guessing the CD decision. Thrust into a central role in the company's success, CD very quickly had to become the world's largest manufacturer of semiconductor components for computer

systems. (In the span of ten years CD grew from 1,000 employees and 140,000 square feet of plant, to 14,500 employees and more than 4 million square feet of plant worldwide.)

The demand for technical manpower to support the development and manufacturing activities of CD was tremendous. Engineers and scientists trained in semiconductor technology were needed to staff not only the SLT program, but also development efforts in integrated circuits. When making their recommendation to form CD, ATS members anticipated the company's immediate need for semiconductor-trained personnel, and expressed concern over IBM's ability to meet its manpower requirements. The problem was made more difficult by the fact that projected output of electrical engineers from the nation's universities would not satisfy industry-wide demand during the next few years.

In response to the critical human resource shortage resulting from the CD decision, IBM initiated the Resident Graduate Study(RGS) program to re-educate company engineers in fundamental knowledge in semiconductor technology and allied fields. Employees accepted into the RGS program were offered the opportunity to pursue a M.S. or Ph.D. degree as a full-time student at a major U.S. university. Program participants were given up to three years to complete their degree, during which time they received full-salary from IBM. Besides easing the staff requirements in CD, the program provided IBM engineers the means to make the transition from a maturing technology to an emerging one.

The Emergence of the IC Memory Concept

During the spring of 1964, Ben Agusta was one of several IBM employees who received an advanced degree in electrical engineering as a participant in the RGS program. Having completed his doctoral dissertation on gallium arsenide semiconductor devices at Syracuse University, Agusta returned to IBM to join CD's integrated circuit(IC) logic development organization.

The IC logic development organization, located in East Fishkill, New York, was led by William Harding. Known as East Fishkill's "Mr. Semiconductor," Harding's work in germanium and silicon semiconductors began back in 1950 after receiving his M.S. degree in physics from the Polytechnical Institute of Brooklyn. Harding joined IBM in 1957 after leaving his position as director of semiconductor research and development at the Radio Receptor Company.

The Harding organization was comprised of two separate groups, led by John Riseman and Dave Dewitt. Although each group had its own set of on-going responsibilities, both groups were actively engaged in a competitive rivalry to design and fabricate IBM's first operational IC logic circuit; one which would provide significant improvements over the SLT components being used in the NPL project to build the System/360. Each group had laboratory facilities and a pilot line to develop prototype circuits.

It was not long after Agusta settled into Riseman's group as a device-circuit engineer, when he began to question why the Harding organization's focus should be limited to the development of IC logic devices. Instead, Agusta suggested that IC technology could also be used to develop memory devices. Having worked as a circuit designer on magnetic

core memory devices before his departure to Syracuse, Agusta understood memory circuit design and was very comfortable with the notion of an IC memory device. Agusta felt strongly, moreover, that the relative simplicity of memory circuits would make the process of integration less complex than with logic circuits. In Agusta's view, not only were IC memories a possibility, they would be easier to develop than IC logic devices.

The decision to focus on the development of IC logic devices, at IBM and elsewhere in the industry, was less the result of conscious decisionmaking than it was a natural outcome of the progression of semiconductor technology. During the 1950's, when discrete semiconductor devices (such as the transistor) began to replace the vacuum tubes which performed the computer's logic function, semiconductors were not found to be a practical replacement for magnetic core memories. Therefore, the first semiconductor IC's developed in the early 1960's were designed for logic applications, with very little consideration was given to IC memory circuits.

Motivated by his idea, Agusta set out on his own time to design a IC memory device. His first design was a 3-by-3 array of memory storage cells, based on the Schmitt trigger circuit architecture, capable of storing nine bits of information.

Design in hand, Agusta now had to convince his colleagues in Riseman's group of the worthiness of his idea. But this proved especially difficult because the group's objective, afterall, was to develop an IC logic circuit and not a memory circuit--that was the responsibility of other engineers working on magnetic core technology. And with the intense competitive rivalry between Riseman's and Dewitt's

group to produce the first IC logic device, many people were openly reluctant to allocate time or resources to anything else.

Although his colleagues were unreceptive to his IC memory idea, Agusta was not totally discouraged. As he and others recall, the rivalry between the two groups was a serious matter, and it created an atmosphere not unlike the "space race" which permeated the U.S. technological community after Sputnik. Under this condition, Agusta understood the sentiments of his colleagues, and believed the group did indeed lack the manpower and technical resources to experiment with the IC memory concept. Rather than continue to push his idea on Riseman, Agusta decided to approach the Dewitt group.

Agusta was first attracted to Dewitt's pilot line manager, Paul Castrucci. Castrucci joined IBM in 1956 after receiving a B.S. degree in physics from Union College. Unlike Riseman's pilot line, Castrucci had a good deal of both human and technical resources in his operation. When confronted with Agusta's IC memory circuit design, Castrucci had no hesitations. Immediately after their first conversation, Castrucci arranged to have the necessary resources put at Agusta's disposal--everything he would need to make a prototype device.

Agusta was encouraged by Castrucci's receptiveness and support, but being a device-circuit engineer, he knew very little about how to practically fabricate a prototype device. Agusta needed assistance. Castrucci on the other hand, understood the manufacturing techniques necessary to build integrated circuits, and Agusta knew his help would be critical to make things happen.

There was something about the idea of IC memory devices that had a strong intuitive appeal to Castrucci. His experience in manufacturing

semiconductor devices led him to believe, as did Agusta, that memory circuits are well-suited for IC technology. Although Castrucci was responsible for manufacturing the experimental designs being generated by the IC logic race, slowly he began to feel personally committed to the IC memory idea. Both he and Agusta began collaborating on building a prototype device, and soon they found themselves consuming a good deal of time in the effort, working evenings and bootlegging time on the pilot line whenever possible.

Agusta and Castrucci were for the most part unable to construct a fully operational device in their initial experiments. However, word of what the two engineers were doing quickly spread around the East Fishkill facility, and eventually came to the attention of Robert Henle. Henle was then a third-level manager in CD, who joined IBM in 1951 after receiving his M.S. degree in electrical engineering from the University of Minnesota. Over the past twelve years, Henle acquired a highly regarded reputation for his outstanding technical achievements in the development of SMS components for project STRETCH and the 7000 series. Henle's inquisitive nature and unique knack for invention prevented his upper-level management position from isolating him from technical developments in the labs. As did Agusta, he too, had begun to conceptualize the notion of IC memory devices. And upon hearing of the experiments on Castrucci's pilot line, Henle contacted Agusta to offer some suggestions.

It occurred to Henle that Agusta and Castrucci had the right idea, but the wrong cell design. Recalling some development work done a few years earlier at IBM, Henle suggested replacing the Schmitt cell architecture with a Farber-Schlig cell in a four-by-four array (thus, holding

sixteen bits of data). The Farber-Schlig cell was conceived in an attempt to create a memory circuit using discrete semiconductor devices. Although Farber and Schlig were technically successful in building an operational circuit, their efforts were discontinued when the device proved too costly relative to magnetic core memory circuits.

Having brought together the IC memory concept and the Farber-Schlig cell architecture, Henle rekindled the efforts of Agusta and Castrucci. Because the Farber-Schlig cell was less complex and tailor-made to semiconductor technology, progress was made quickly. Now the prospect of operational 16-bit IC memory devices in the near future appeared promising. It soon became obvious to Henle that more time and resources should be dedicated to the project that Castrucci was able to bootleg from his pilot line.

The IC memory project needed formal recognition within the Harding's IC logic development organization. Only then could it be certain adequate resources would be available. However, the idea of IC memories deviated significantly from the organization's original mission, so approval had to come from CD's top executives. Futhermore, DSD executives also had to be consulted, since the new memory components would affect system-level design.

It is at this time when Henle played a most critical role as spokesman and proponent of IC memories in his capacity as a newly appointed IBM Fellow. One of the very first individuals to receive this honor, Henle was given the opportunity to pursue any research project of his interest for five years time (until 1969). This provided Henle the chance to dedicate his time and energy to exploring the possibilities for IC memories, and spreading the word within the company.

Henle began by authoring a white paper on the future potential of IC memory devices, and their performance relative to other memory circuit technologies. In the paper he discussed the development efforts in IC technology both within the company and in the industry. Soon after, Henle set out to give presentations on the subject of his paper. As an IBM Fellow, he was given the respect the position conferred, and quickly gained the attention of key IBM executives in CD and DSD.

After a series of more than twenty presentations during the autumn months of 1964, Henle's efforts began to payoff. In December, DSD's vice president for development, Erich Bloch, spoke with Henle about a project which called for high-speed memory circuits. Bloch had known Henle from their mutual involvement in the development of SMS components for the STRETCH project. And as chairman of the ATS committee, Bloch had become fully aware of the potential of IC technology.

The discussions between Bloch and Henle resulted in an agreement to develop 16-bit IC memory devices to be used in the storage-protect memory of a super-fast System 360 computer for NASA's Goddard Space Flight Center in Maryland. The storage-protect memory is a small unit that prevents unauthorized access to information stored in the computer's main memory. The critical requirement for memory circuits in the storage-protect unit is high-speed functioning, and both Bloch and Henle believed the IC memory devices might provide significant improvement over magnetic core memory devices. However, the IC memory effort was to be a back-up program, and would compete with the on-going efforts of other engineers working on high-speed magnetic core memories.

Late in December a meeting was arranged to initiate the SP-95 program to develop IC memory devices for the storage-protect units of

the System/360, Model 91 and 95. The program was to be a collaborative effort between engineers in CD and DSD. Attending the meeting were Agusta, Castrucci, Paul Bardell and Ray Pecorara from CD; and Ed Slobyzinski, Jack Shortle, and Ed Hee from DSD.

Bardell, like Agusta, had received his Ph.D. earlier in the year as a RGS program participant. The two now had circuit and device design responsibilities for the 16-bit memory. Agusta was transferred from Riseman's group, and now reported to Bardell. As pilot line manager, Castrucci was responsible for developing a manufacturing process and a means for testing the new devices. Both Castrucci and Bardell reported to Ray Pecorara, who was Dewitt's special devices development manager. Slobyzinski, Shortle and Hee were responsible for circuit board and system-level design.

Aside from the very high speed requirement, the basic objectives of the SP-95 program were conservative. The technical groundrules were set to shown feasibility of the concept--not optimum capability. But to term the program's objectives as conservative is not to say they were unchallenging. Although the experience accumulated from the SLT and IC logic development efforts was applicable to building IC memories, there were still new areas of technology to explore and many problems to solve. Over the next eleven months the project's engineers worked to isolate the problems and design experiments to test possible solutions.

From a methodical process of exploration and integration of new knowledge came a greater understanding of the basic properties of semiconductors. This understanding was essential to devising the process technology and manufacturing techniques necessary to produce an operational device; a device capable of attaining the required production

yields to make IC memories cost-effective.

In December 1965, there was cause for celebration among members of the SP-95 team, as IBM became the first computer manufacturer to publicly announce the usage of IC devices for high-speed memory applications. On a minute chip of silicon only 70-thousands-of-an-inch-square, were fabricated 148 electronic components forming sixteen circuits. By being in an "on" or "off" static condition, each circuit stores a bit of information; thus allowing each device to store sixteen bits.

The 16-bit IC memory chips compared very favorably with the alternative solution. A storage-protect unit employing IC memory components and having a 128-word, 20-bit array, had a 300-nanosecond cycle time. This proved to be faster while providing greater storage capacity than a magnetic core unit, having a 64-word, 20-bit array and 350-nanosecond cycle time.

The SP-95 program generated a good deal of excitement over IC memories in East Fishkill. For the engineers involved with the program, the success of the 16-bit chips was viewed as just the beginning. They recognized the potential of IC technology to fabricate chips with far greater storage densities and faster speeds. Rather than merely prove feasibility of IC memories, the engineers were now anxious to explore the technology's wider capabilities.

The Phase-2 Program

By 1966, the East Fishkill plant had become a very hectic place, due much to the tremendous success of the NPL program. Since the time the first System/360 was shipped in 1965, the marketplace demand far

outstripped IBM's production capacity. During 1965 and 1966, IBM booked orders for more than 9,000 System/360's, but was only able to ship 3,800 units. As the supplier of SLT and core memory components used in the 360, CD could hardly keep pace. During 1965, monthly production of SLT modules increased by a factor of ten, from 1 million to 10 million modules.

The feverish pace in East Fishkill to support the System/360 made it hard for the IC memory project's activities to attract serious attention. Certainly the new developments were considered interesting and exciting; but given the circumstances, few people not closely attached to the project could imagine implementing IC devices in something other than small, special-purpose memory units any time soon. The prospect of IC memories replacing magnetic cores in the computer's main memory was unimaginable to CD executives. For the time being, life in CD was complicated enough just trying to meet delivery schedules for SLT and core memory components. The ambitions of the IC memory development team were confronted with the stark reality that change would take time.

The development of IC memories at IBM was at a critical juncture, and it was unclear precisely which direction should be taken. At this stage, however, the problem was less one of selling the idea and more one of finding a use for the new devices. Soon a need did surface; and in retrospect it should be no surprise. Afterall, the faster cycle speed and greater circuit density made IC memories a superior alternative to magnetic cores, and this was quickly recognized by memory system designers. In fact, the flexibility offered by IC memories opened the door to new concepts in memory system design.

The buffered cache memory was one such new concept. The buffered cache memory unit was conceived by IBM systems engineers as a means to deal with the problematic trade-off between high-speed performance and extensive memory storage capacity. The memory system designer wants both; but, with greater memory capacity comes slower cycle speeds. As a solution, the buffered cache memory was designed to hold only that portion of the computer's main memory content currently in use. This was accomplished by a "least-recently-used" algorithm implemented in the hardware. A small cache memory unit functioning at high speeds and storing most of the immediately needed information, meant the systems engineer could design larger main memory units without slowing down the computer's processing speed. It also meant there was a need for high-speed memory circuits to be used in the cache unit. The high-speed characteristics of IC memories made them exceptionally well-suited for this purpose.

Early in 1966, CD executives gave the SP-95 team a new goal: to develop IC memory circuits to be used in the buffered cache memory unit of the System/360 Model 85 computer. The new memory devices were to be faster and have greater storage density than their predecessor, the 16-bit chip. This was important because the team would be competing with other development teams working on different technologies, as had been the case in the past. The new program was named Phase-2.

The design of the Phase-2 chip used an improved diode-load version of the Farber-Schlig cell architecture that included on-chip circuit drivers, developed by Agusta in collaboration with DSD system engineers. The chip layout was an 8-by-8 array forming 64 cicuits; thus, capable of storing 64-bits of information.

After a year-long effort, the Phase-2 chip was qualified in September, 1967. The new devices exceed the expectations of the Phase-2 team. Memory system cycle time was reduced from 300-nanoseconds with the 16-bit chip to 80-nanoseconds with the new 64-bit chip. Moreover, the production yields for the new device had shown remarkable improvement over its predecessor. It appeared that as experience in manufacturing IC memories accumulated, many process innovations were occurring and contributing to the overall success of the devices.

Once again, the IC memory development team had cause for celebration, when the 64-bit chips were adopted for use in buffered cache memory applications. At last, it was felt, the concept of IC memories had emerged out of the "experimental" realm. The new devices proved to the industry that IC technology could be economically feasible for wide-scale commercial use.

The success of the 64-bit chip meant the team's efforts could no longer go unnoticed by CD's executive management. The IC memory was quickly becoming a serious contender to replace the magnetic core technology used in the computer's main memory unit. The people involved in IC memory development were very excited about their work, and they began to press the division's management to clarify where the future of IC memories stood within the company.

The ambitions of the IC technology people were a strong force with which to contend. There was something about the technology that gave engineers involved with the project an overwhelming sense of accomplishment. For most, building IC devices was quite unlike anything they had ever done before, and there was an undeniable mystique about what they were creating. The IC chip is one of those rare technological feats the

very sight of which challenges the imagination. Like the images of an astronaut stepping on to the surface of the moon, the "chip" was quickly to become a frequently used symbol of modern technology.

For circuit designers, the thought of 1,000-bit or 4000-bit devices in the not-to-distant future created an unbounded optimism. It was clear that as line geometries were scaled down, the theoretical potential of chip density left many convinced that the accomplishments so far had only begun to scratch at the surface of IC technology. (Since the late 1960's, chip densities have increased to the point where 64,000-bit devices are now commonplace, and 256,000- and 1 million-bit chips are soon to be also as common.)

For engineers with a manufacturing orientation, IC memory chips represented a never-ending challenge to produce the next generation of higher-density IC devices. With each success came huge leaps in factory productivity, continually reducing the cost of computing at a time when it seemed the cost of everything else was going up. The IC memory chip was nothing short of a "once-in-a-life-time type of technology achievement."

The Phase-2i Program

The pressures created by competing memory technologies, and the prospects of an impending "technology" decision by IBM's corporate officers put CD and the rest of the organization under great strain. The question of which component technology would be employed in the company's next generation of computers, rested uneasy in the minds of many engineers connected with memory development. Careers were at stake; undoubtedly, the future stream of rewards, recognition and status would

go to those who had their training in the technology most valued by the firm.

Fears spread quickly, and an intense and sometimes dysfunctional competition stirred among the various technology factions taking hold in CD. Although competition between technical groups was considered essential and always encouraged by division executives, under these conditions it sometimes produced exagerated claims of the capabilities of new technologies. Such claims were dangerous input to the decisionmaking process.

Organizational politics began to serious hinder development activities in CD laboratories. As one engineer recalls, the organization was something to overcome; not something to be used as a tool. Political motivations, personality differences, and egotism meant that many good innovative ideas were dead-ended because it was not the merit of the idea that counted, it was who you knew and who you supported. Viewed from the lab, the growing political conflict made it hard to perceive just what was happening within management's ranks, and this created an atmosphere of mistrust. "Management-by-conflict" began to take its toll.

It was obvious a decision had to be made, but it was by no means an easy one. The company's commitment to magnetic core memory technology spanned nearly two decades and was very profitable. During this time an extremely large investment had accumulated in both physical and human capital. Semiconductors and magnetic core technology were sufficiently different in their nature, so that moving away from magnetic cores would create a major dislocation within the organization. Furthermore, it was unclear exactly which technology to chose: on the one hand, IC memories held much promise for the future, but then significant advances were also being made in magnetic cores.

The first move came toward the end of 1967, when IBM's executive committee approved a reorganization of CD. The central purpose was to consolidate responsibility for memory component development and manufacturing under one executive. This meant removing the IC memory group from the Harding logic organization, as well as shifting memory systems engineers from the Systems Development Division(SDD, formerly DSD) into CD.

The corporate executive committee placed Edward Davis in charge of the new memory organization. Davis received his Ph.D. in solid-state physics from Stanford University in 1958, and joined IBM shortly thereafter. He had been involved in the development SLT components, when his managerial competence earned him a position on the fast-track in CD. Eventually, he was placed as an administrative assistant to IBM's chief executive officer--a position referred to as being "knighted" within the company.

As head of the memory component organization, the first task on Davis' agenda was to decide on a technology that would become his candidate to replace magnetic cores. Here, he knew he could not move too fast. The technology factions which developed within the organization made certain that any decision he decreed would not receive a wholehearted commitment from all sides. Davis understood the organizational politics and the challenge that confronted him; that is, the need to rally the support of key individuals around a particular technology, in a situation where people's self-interest had pitted different groups against each other.

Davis first set up a series of technology task forces to gather information on the relative performance capability of the of the primary

candidates to replace magnetic cores: IC memories and magnetic films. In addition, opinions were sought from an ever increasing number of sub-factions blossoming among IC supporters. Now there were opposing groups proposing the use of bipolar versus FET (field effect transistor) IC's, and a still further split within the latter group between n-channel and p-channel FET devices.

While the task forces gathered valuable comparative data on the various technology options, they also gave Davis a clear understanding of where individuals stood with respect to technologies. This was important because now he had a better idea of who he could rely on and who he had to convince, given the decision to be made. It was important that those individuals who could make things happen were behind his decision.

In January 1968, Davis initiated the Phase-2i program to develop IC memories to be used for main memory applications. Phase-2i was an extension of the Phase-2 program and involved many of the same individuals responsible for IC memory development within IBM since SP-95. However, the product to be developed in Phase-2i was in many ways different from its predecessors, the 16-bit and 64-bit chips. The new devices would be based on the Harper cell architecture and not the Farber-Schlig cell previously used. Futhermore, each chip would store 128-bits of data, and for the first time include pulse-powered driver and decode circuits.

The basic linkages between the Phase-2 and Phase-2i programs were the people involved and the decision to continue using bipolar process technology. Davis' decision to proceed in this manner had little to do with issues of relative performance between the various technologies. Instead, Davis had shown greater concern over the abililty of his can-

didate technology to gain acceptance among the SDD executives who would decide when and in what computers the new memory devices would be employed. The Phase-2i program had the appeal of being the latest phase of a well-established development effort within IBM--with a track record in which SDD people could place some confidence in CD's ability to deliver a quality product. To switch from bipolar to FET devices at this point in time would only serve to to fuel the arguments of magnetic core proponents; that IC memories were still experimental devices and thus, too risky. If and when bipolar IC's were successful in replacing magnetic cores, then Davis could evaluate the relative merits of bipolar versus FET technology.

Phase-2i was in progress for nearly a year when in November, 1968, the company was stunned by the overnight exodus of twenty employees, many of whom were engineers working on the Phase-2i program. The departing group was led by an upper-level manager within the Davis organization who joined with George Cogar (formerly with Mohawk Data Systems) to form a new company to build and market IC memory devices.

The reasons of those who left were many and varied. However, there was a common feeling among "technology people" that they were not receiving due recognition for what they believed were significant contributions to IBM's success. From their perspective, the opportunity for hefty financial rewards offered to marketing people was incongruent with the gold-plated pen and pencil set reward system reserved for them. This feeling of inequity resulted in a deep-seated bitterness among engineers.

As Bloch recalls, the rewards for technical personnel did differ from marketing, but for good reason: the risks facing each are substan-

tially different. Compensation in marketing is directly linked to individual performance, whereas compensation to technical personnel is tied to group performance and represents a long term investment in human capital.

For whatever reasons the group departed, IBM's corporate headquarters moved quickly to stabilize the situation. The loyalties between engineers in the memory development group threatened to cause a snowballing effect among those who still remained. Indeed, the Cogar group actively sought to enlarge the exodus to their new company in the months that followed.

The Cogar incident had a serious impact on Phase-2i. Not only did it deplete the program of some key participants, it left those who remained somewhat dismayed. The sudden exodus of friends and colleagues led many to question their commitment to IBM--and the company's commitment to them. The uncertainty over whether one made the "right" decision by staying with IBM lingered on for many months. The fact that many of those who chose to stay and those who departed, lived in the same communities, and had close personal relationships, made the incident a difficult one with which to cope.

If it could be said that some good had come from this incident, it would be the greater awareness among corporate officers about the work being done in CD. On several occassions, the company's executive officers made their way through development labs and manufacturing lines to gain a better understanding of IC technology. Efforts also were made to bring technology and marketing personnel together in day-long conferences to foster communication between the two groups, and build a sense of

the of the contribution made by each.

Eventually, the Phase-2i program was back on track, and as determined as ever to put IC memory technology into the main memory units of IBM computers. The proponents of magnetic core memories were not silent over this prospect, and battles over this issue became increasingly frequent at corporate headquarters in Armonk, New York. Davis' staff was preoccupied most of the time in meetings countering the allegations made against a movement toward IC memories. Meanwhile, Davis concentrated his efforts on convincing SDD executives of the strategic importance of going with IC technology. Since SDD was responsible for system-level design, acceptance of IC memory components by SDD executives would go a long way in convincing Armonk to approve Davis' plan.

The efforts of the Phase-2i program came to fruition in August, 1969, when the first 128-bit memory chips passed qualification. Now with a product in hand, the struggle to convince SSD to design-in the new components intensified. Finally, agreement was reached to employed the chips in the company's next-generation computer, the System/370.

It was referred to as the "biggest technology push in IBM's history," when in June, 1971, the first industrial and commercial computer with an IC main memory was shipped. IBM's early commitment to IC memories took the industry by surprise. It was just a week earlier that industry observers reported IC main memories were at least four years in the future--not until 1975.

Notes

[1]Fisher, Franklin M., et al., IBM and the U.S. Data Processing Industry, (New York: Praeger Publishers) 1983, p. 102. See chapter five.

[2]Ibid, p. 103.

[3]Ibid, p. 104.

IBM East Fishkill Today

EXPERIMENTAL MONOLITHIC CHIP TO BE ANNOUNCED

IBM integrated circuits for memory units will be announced in Washington today when Dr. Benjamin Agusta, Dept. 911, presents his paper on "Sixteen-bit monolithic memory array chip" at a meeting of the IEEE Professional Group for Electronic Devices.

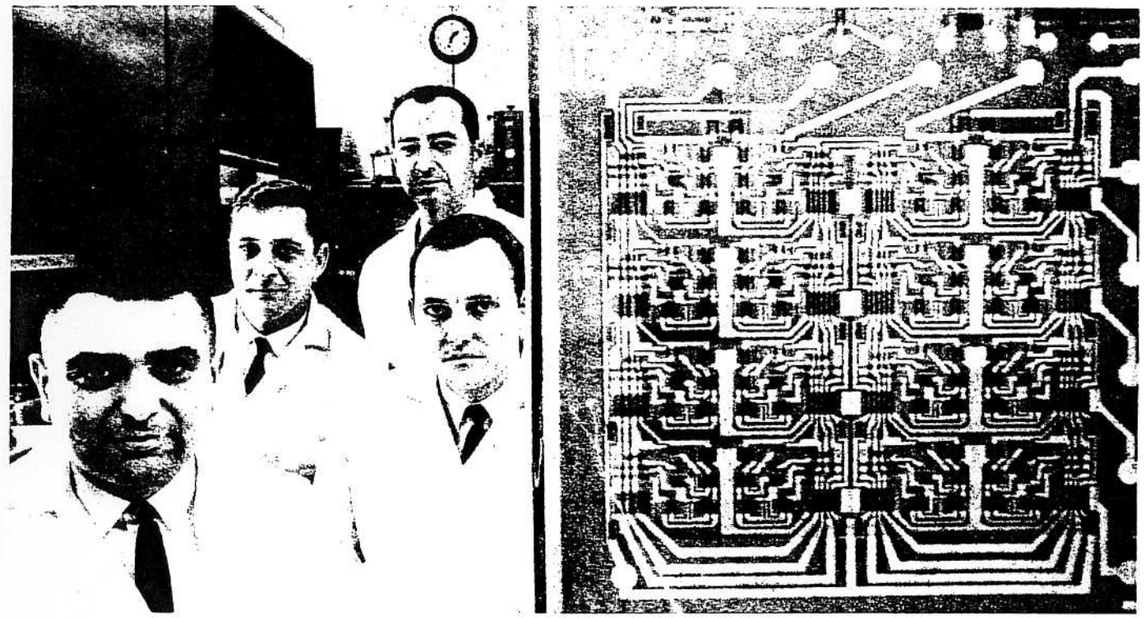

Front to back: Castrucci, Bardell, Pecoraro, and Agusta. Chip is smaller than letter o.

The paper will describe experimental circuits fabricated on tiny chips of silicon 70-thousandths-of-an-inch square. The chips represent the first reported use of monolithic semiconductors for high-speed computer memories.

Coauthors are Dr. Paul Bardell, Dept. 911, and Paul P. Castrucci, Dept. 210. Authors are shown above with Raymond P. Pecoraro, project manager. The chip, shown greatly magnified in a mock-up, will fit into the typewritten lower-case letter o.

One chip has 16 circuits, with 4 auxiliary diodes, and each circuit has 5 transistors and 4 resistors. With this density, nearly 7500 components could be placed side-by-side in an area no larger than a fingernail. The chips will be used in experimental three-dimensional, scratch-pad memories. They will make possible faster speeds than such memory elements as ferrite cores.

Circuits are called "integrated" when they cannot be broken down into separate components. (An example of a separate component is a transistor in a portable radio; it can be changed without affecting any other component.) Integrated circuits are called "monolithic" when all electrical elements are fabricated within a single chip of silicon.

Monolithics can mean major advances in memory speed, design, and packaging.

10-21-65