Notes from Interview with Joel Karp
Regarding Development of Intel 1K and 4K DRAMs

Jeff Katz 04 June 2007

SUMMARY

We spoke in Joel's home for about 2 hours, and discussed his recollections of the development of Intel's first Silicon-Gate 256 bit MOS RAM, the 1102 and 1103 1K DRAMs and the 2107 4K DRAM, for all of which he was either the principal designer, or a key contributor. Joel was at Intel from shortly after its founding in mid1968 through early 1973.

JOEL KARP CAREER SUMMARY

After graduating from MIT in the early 60s, Joel took a job at MIT's Instrumentation Laboratory (now Charles Stark Draper Labs) working on the electronics for inertial guidance systems for the US Navy. In 1966 he was recruited to General Microelectronics, a Fairchild spin-off. He joined Intel in 1968 as its first MOS chip designer, concentrating on memories and developing the first DRAMs. Leaving Intel in 1973, he worked for a short time at Intersil, and then formed a successful DRAM consultancy with former Intel colleague, John Reed. They developed DRAM designs for most of the major computer makers in the 1970s and early 80s. Since then he has had high profile jobs at several companies including Visic (Founder, CEO, venture backed semiconductor startup), Samsung Electronics (Senior Vice President) and Rambus (VP of Intellectual Property). He is now semi-retired but still consults on IP matters.

HOW JOEL GOT TO INTEL

As of 1962, he was one of the few engineers at the MIT lab who had studied solid state electronics, having taken one undergraduate course in the subject. (The course was taught by **Paul Gray**, who later went on to become President of MIT. A classmate was **Mike Callahan**, who later held executive positions at Motorola, Monolithic Memories, AMD, and was CEO of Wafer Scale Integration.) As part of his job at the Instrumentation Lab, Joel frequently visited California to evaluate capabilities of various semiconductor companies to support the Navy program, becoming friendly with **Bob Widlar** and several others at Fairchild. In 1966, some of these acquaintances who were working at Fairchild spin-off General Microelectronics recruited Joel. Two years later GME's new owner, Philco-Ford, tried to consolidate their semiconductor operations into their Bluebell, PA facility, and Joel resigned to avoid a move.

While wondering what to do next in mid-1968, he learned from an Electronic News article that Fairchild founders **Bob Noyce** and **Gordon Moore** were leaving that company to form a new company (no name for the new company was yet known). The new company was to be housed in the soon-to-be-vacated Union Carbide facility at 365 Middlefield Rd in Mountain View, but it wasn't known what would be its products. Having previously met Gordon Moore during visits to Fairchild, Joel applied to work for them. He became Intel's first MOS chip designer, employee number 20, in September 1968, not learning of the company's Silicon Gate MOS product plans until after he started work, reporting to **Leslie Vadasz**. He remained Intel's only MOS designer for a year and a half, until **Bob Abbot** was hired as employee number 200 in early 1970, and shortly later **John Reed** was hired.

WHO DID WHAT, IN EARLY INTEL MEMORY DEVELOPMENT?

1101 256 bit SRAM

Joel's first assignment was to work on a 256 bit SiGate PMOS SRAM, under supervision of engineering manager Les Vadasz. At this time Core memory was mature and dominant, and was priced fairly profitably for the makers at about 5 ¢ per bit. Intel's marketing director **Bob Graham** calculated that the 256b RAM could compete with Core if it could be priced at $12, or ~4.7¢/bit. This seemed feasible. Joel's design was completed and the first silicon almost worked, only requiring pull-up resistors on the bit lines, in less than a year. During the project Andy Grove provided the design rules and process parameters, and taught Joel how to use the IBM Service Bureau's mainframe to run FORTRAN to use Grove's MOS device equations. There were no other modeling tools available in the company. Design rules were 10 micron minimum line width and spacing. Partway through the project, Gordon Moore had to come up with a fix to the SiGate process, whose polysilicon was coming out too brittle and developing micro-cracks in the metal overlying the sharp steps. Moore had the idea to heat the material almost to the melting point to reflow it. But the word "reflow" never appeared in any process documentation; in order to preserve its secrecy, a euphemism term (anneal) was used. Joel remembers the day he and Vadasz discovered and demonstrated the fix which used bit line pull-ups was July 20, 1969. They were listening to the radio in the lab, learning of the landing of Apollo 11 on the Moon. Ten months after joining the company, Joel and all Intel employees celebrated the working 1101 SRAM chip at the Wagon Wheel tavern down the street on Middlefield Rd.

With Intel and other semiconductor makers threatening to compete with Core, the Core makers were quick to spend some of their fat profits on reduced pricing in order to keep their market. The price of Core dropped rapidly from 5 ¢/bit to 1 ¢/bit, forcing the semiconductor makers to sell at painful losses if they wanted to compete. A natural way to approach the problem was to develop denser, more cost effective semiconductor memory devices.

INTEL'S FIRST DRAMS

1102 Honeywell 1Kb DRAM

Late in 1969, Honeywell engineering manager **Bill Regitz** was making the rounds to semiconductor makers with his idea for a 1Kb Dynamic RAM. With only 3 transistors per DRAM cell, the data capacity of individual chips could be as much as 4 times that of 6-transistor per cell SRAMs, making it a more attractive option to replace Core memory. He came to Intel and met with Vadasz, Karp and **Ted Hoff**. The Intel team agreed to develop the DRAM for Honeywell. Joel was assigned to work on the Honeywell device, named by Intel the1102. As conceived by Honeywell, the device had a difficult-to-design analog circuit for the single read/write line, which had to operate over a very wide voltage and current range. Hoff suggested to the Intel team that the circuit would be much easier to design and operate more reliably if the read and write were separated. Using conventional design techniques of the time, this simplification would have greatly enlarged the cell size. But Vadasz and Grove approved a covert attempt to use the newly emerging buried-contact design technique to make a more robust yet small-cell-size 1Kb device with separate read and write lines, to be called the 1103.

The 1102 came out and, though it worked poorly, it was shipped to Honeywell. A second version with better placed decoder circuits worked somewhat better, but at Honeywell it exhibited some pattern sensitivities and marginal behavior. The Honeywell and Intel teams eventually discovered that the chip substrate being attached to the grounded package lid was causing the pattern

sensitivity marginalities. By cutting the ground connection to the package lid and adding external bias-voltage wires soldered to the package lids, Joel finally got the 1102 to work. The team presented the 1102 at the Feb 1970 ISSCC

Meanwhile during Joel's efforts to get 1102 working, Marketing learned of the 1103 and requested it to be offered to other customers.

1103 FIRST COMMERCIALLY SUCCESSFUL 1Kb DRAM

When Marketing pressed for the buried contact, better performing, more cost effective 1103 DRAM a project was started, to run in parallel with the 1102 project. There was an evening meeting at Vadasz's home in which Vadasz and Joel worked out how to architect Hoff's suggestion. Then new hire Bob Abbot was given the assignment to complete the design. He did bring the project to tape-out phase, whereupon it was given slightly newer but more experienced hire John Reed to debug and bring into production. Reed got the chip working, enjoying much better performance due to the elimination of the tricky analog circuit. But it had poor margins, with a very critical external timing signal overlapping with another timing signal.

During chip characterization, which Joel was assigned to do, capitalizing on his experience in characterizing the 1101, he proposed that a narrow metal than design rules permitted might solve the critical timing overlap problem. There developed a lot of tension between Fab, Engineering and Marketing. Then suddenly, while Vadasz was out of town Fab manager Gene Flath phoned Joel to say the Fab would try to make the narrower metal. When Vadasz returned he was surprised to find fab had finally agreed to the narrower metal, and moreover that the parts came out and the timing overlap problem was solved.

Ultimately Marketing VP **Bob O'Hare** and Noyce were able to convince Honeywell to abandon their 1102 design and use the 1103. Because of its larger capacity and much smaller cell size than traditional SRAMs, the 1103 DRAM became the first semiconductor memory able to compete profitably with Core memories. [Ed. Note: It was at this point that Marketing launched its campaign declaring that Core memory had lost the price war against semiconductor memory.]

THE PIN-OUT WARS IN 4K DRAMs

In late '70 and '71 Joel worked on Intel's 2107 4Kb DRAM. It proved to be another difficult analog problem. But ultimately it worked, and Joel and Regitz, by then an Intel employee, presented the device at the ISSCC in Feb '72.

With 4 times the capacity of the 1103, the 2107 had 12 address lines, forcing the device into a larger 22-pin package than the popular 16-pin 1103 and other 1Kb DRAMs. Thus the system board area required was twice that required for the same number of 16-pin devices. Competitor Mostek, with engineer **Bob Proebsting**, had invented a multiplexed-bus 4Kb DRAM that used only 6 address pins which timeshared the 12 signals. Thus it could fit in the smaller 16-pin package. The Mostek device was an instant hit, in spite of the fact that not only Intel but TI had compatible 22-pin devices. (Joel recalled that TI had originally copied Intel's 2107 device, but arranged the address pins backward from what Intel had done. Ultimately the 22-pin customers prevailed on Intel to make a version whose pinout matched TI's.) The clear advantage in system board density made the Mostek device inexorably the winner, even though it didn't perform as well as Intel's larger device. Ultimately all 4Kb suppliers had to adopt the Mostek multiplexed address bus, which continued to be the standard for more than 20 years.

KARP'S DEPARTURE FROM INTEL AND LATER CAREER

Joel left Intel in '73, feeling disappointed in the company culture that focused on improving negative aspects. He pointed out that by that time his stock options were fully vested. With job offers from several valley companies, he took a job with Intersil, who was trying unsuccessfully to make 1103s. His first assignment was to work on a 4 KB device. After he hired some former colleagues from Intel, he recalled that Intersil's president, **Marshal Cox**, received a call from Intel's Bob Noyce, threatening legal action if further hiring of Intel folks occurred. As an unprofitable company, Intersil had little funds and resources to invest in R&D. Eight months after joining the company a frustrated Joel left Intersil.

Word got out that Joel was a free agent and he began receiving recruiting calls from several companies. He decided instead of working for one of them, he would let *all* of them pay him in a consulting role. He recruited former Intel colleague John Reed to join him in a consultancy. Together they executed DRAM designs for many system and semiconductor companies that were in the process of designing and/or developing semiconductor memory devices. Joel and John initially charged $50 per hour to design 4 KB DRAMs for Burroughs, and National Semiconductor. As they expanded their practice, they consulted on various matters with Data General, RCA, Lockheed, Hughes, Control Data, Mitsubishi, Siemens and others.

After successful careers at Visic (later acquired by VLSI Technology), Samsung Electronics, and Rambus, Joel is now semi-retired and consults with several companies on intellectual property matters.

OTHER POTENTIAL SOURCES FOR STORIES ON EARLY SEMICONDUCTOR MEMORY

Joel suggested we talk with **Tom Rowe**, an Intel technologist known to have saved lots of potential museum artifacts. He also once shared an office with Intel's **Dov Frohman**, around the time Frohman discovered EPROM technology. An important DRAM development, the sense amp used for one-transistor cell designs, was invented by **Karl Stein** at Siemens, and presented by him at the '72 ISSCC.

OTHER ITEMS

Joel showed me several original documents and many chip samples of early DRAMs from Intel and other companies for whom he consulted, all in his personal collection. He was amenable to the suggestion that the SIG would like to borrow documents for archival scanning.

Also he owns original circuit diagrams for the Apollo guidance computer that was developed while he worked at MIT, and which the Museum displays in Visible Storage. Joel would consider donating these to the Museum's collection.

Should the CHM Semiconductor SIG decide to an Oral History panel discussion on early DRAMs, Joel would be a candidate participant. He can be reached at 650-322-4388 or jakarp@alum.mit.edu.