Written by
Nick Tredennick

DynamicSilicon

Published by
Gilder Publishing, LLC

Vol. 2, No. 5 The Investor's Guide to Breakthrough Micro Devices May 2002

Goldilocks and the Three Memory Chips

Semiconductor memory chips are everywhere! It's a $28-billion market—and that's in real dollars, not in dot.con suppositions. Memory is used in computers, cars, digital cameras, cell phones, MP3 players, microwave ovens, and just about any other electronic device you can imagine. Today, three kinds of semiconductor memory predominate: DRAM, SRAM, and flash. But, unlike *Goldilocks and the Three Bears*, none of these is "just right." DRAM is cheap, but it's too slow, SRAM is fast, but expensive, and flash memory is non-volatile, but flash memory is slow and it wears out. *The industry needs semiconductor memory that is as dense and as cheap as DRAM, is as fast as SRAM, remembers its contents through power cycling as flash memory does, doesn't wear out, and is readily manufacturable with CMOS logic*. There's $28 billion at stake ($28 billion is memory *chips*, there's three times as much in semiconductors with on-chip memory), so contenders are coming out of the woodwork.

kb - 1,024 bits
Mb - 1,024 kb
Gb - 1,024 Mb

Memory makers are looking for a new semiconductor memory that does the job of all three. As we move to system-on-chip designs, a memory's compatibility with the manufacturing process used for CMOS logic becomes important. (CMOS logic is like the Windows operating system. You can use something else, but if you do you are going to miss most of the market.) SRAM and flash memory integrate well with CMOS logic, but DRAM doesn't.

The three most promising candidates are MRAM, FRAM, and OUM. Their promise is to replace memory as we know it. They see even the hard disk as prey. The companies behind these new memories are established semiconductor manufacturers, not startups. These candidates benefit both tethered (powered from a wall socket) and untethered systems, but untethered systems benefit more because these new candidates use less power and are smaller than current semiconductor memory chips.

The opportunity

DRAM, SRAM, and flash own today's CMOS memory market. Together, they collect more than 90% of the dollars spent on memory chips. Their share of the market is increasing. In spite of their huge markets, each memory cell has its flaws. Could a new memory cell combining virtues and lacking flaws displace one or more of the incumbents? Huge corporations back the three incumbents; none will cede market share without a fight. Even a perfect memory cell would take years to replace its entrenched competition.

When you turn your PC on, it takes a minute or so to get ready to work. When the microprocessor gets power, it looks in the flash-based boot memory for the initialization software. This software moves the operating-system software from the PC's hard disk to its DRAM and turns control of the PC over to the operating system. As it runs, portions of the operating system move from the slow DRAM to the microprocessor's faster on-chip SRAM caches. The PC uses each memory type for its strengths and abides its weaknesses. Imagine a memory type with the density of DRAM, the speed of SRAM, and the non-volatility of flash. It could replace DRAM, SRAM, flash, and maybe even the hard disk. No more elaborate startup and shutdown for the PC. It would be instantly on in exactly the state it was in when you turned it off. The PC would certainly benefit from a non-volatile DRAM replacement—and so would any other system that uses DRAMs today.

The world is going mobile with cell phones, cameras, personal digital assistants, MP3 players, DVD players,

In This Issue: Most of the major semiconductor makers are working to replace random access memory (RAM) as we know it. These new memories are "non-volatile," they remember their stuff even when the power is off. They consume little power, yet they operate as fast as today's RAM. The new acronyms to learn are MRAM, FRAM, and OUM. In this issue, I explain how they work, summarize their progress, and assess them.

notebook computers, watches, GPS receivers, and radios. All use memory chips, most use more than one type of memory, and some use hard disks. The mobile world is a lucrative target for a practical non-volatile memory cell.

Many of today's microcontrollers integrate three memory types onto the chip: read-only memory (ROM), flash memory, and SRAM. Consolidating the microcontroller's ROM, flash, and SRAM into a single type of non-volatile memory would reduce the manufacturing cost and it would increase the range of applications for each microcontroller type since programs that used to be in ROM could be loaded into the non-volatile memory *after* the microcontroller is manufactured.

Programmable logic devices that today use SRAM for configuration must load the configuration file into the SRAM each time the power comes on. A non-volatile equivalent to the SRAM configuration bits would eliminate the external configuration memory and shorten the

Today's memory types

■ **DRAM.** In mid-2002, leading-edge DRAM chips are 256 Mb and will soon be 1 Gb. These chips have access times of about forty nanoseconds and cost a few dollars. DRAM chips store a 1 or a 0 by the presence or absence of electrons on the plates of a capacitor that is isolated by a single transistor. DRAMs like high voltage so they can store lots of electrons in the small capacitor and they want low-leakage (stay stubbornly off) transistors so electrons stay in the capacitor longer. Since each bit uses only a single transistor and a single capacitor, DRAMs cram lots of bits onto one chip. But the transistor and the capacitor "leak," requiring a periodic "refresh" of the bits to replenish electrons that have leaked away. DRAM cells are literally tiny containers that hold a 1 or a 0 value. DRAMs are volatile—bit values stay only as long as the chip has power.

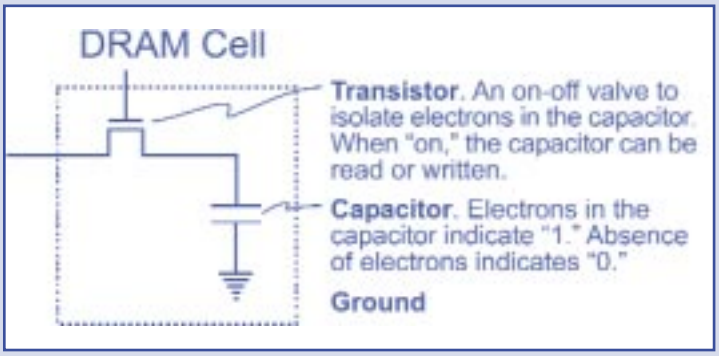

DRAMs have a "destructive" read property. To read, the valve is opened and a 1 or 0 is declared by whether electrons escape from the capacitor. Once the value on the capacitor has been read, the value must be restored. So each read is really a read followed by a write; this takes longer than a simple "non-destructive" read. DRAM's high-voltage, low-leakage transistors don't integrate well with CMOS logic's low-voltage, high-speed (flip on or off easily) transistors.

As their silicon process improves, semiconductor manufacturers first increase DRAM's storage capacity then do what they can to speed access time. In 2001, DRAMs accounted for more than 40% of memory chip revenues. And DRAM's share of memory chip revenues is growing.

■ **SRAM.** Leading-edge SRAMs are 16 Mb. Access times can be under two nanoseconds. A leading-edge SRAM chip costs about four times as much as a leading-edge DRAM chip. Since leading-edge DRAMs contain about sixteen times as many bits, SRAM bits cost sixty-four times as much as DRAM bits. The SRAM cell stores its information in a circuit made up of a pair of inverters. SRAMs are fast because the inverters use fast CMOS logic transistors. SRAM cells are active circuits constantly calculating their 1 or 0 value. Since the SRAM cells are made up of fast CMOS logic transistors, SRAMs integrate well with CMOS logic. SRAMs are "volatile" because they retain their information only as long as they have power.

As their silicon process improves, semiconductor manufacturers increase the SRAM's access time and then do what they can to increase its capacity. One or two levels of SRAM-based "cache" bridge the gap between fast microprocessors and relatively slow

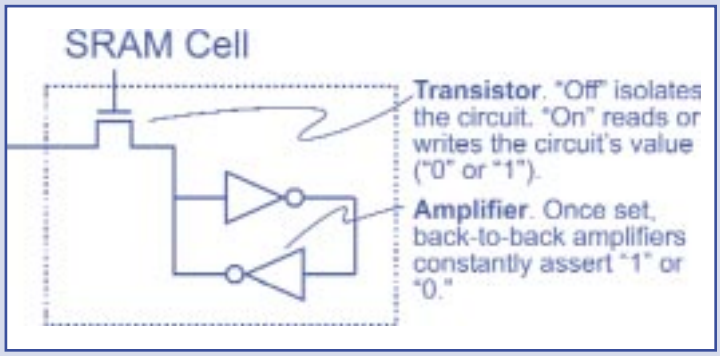

time to initialize the chip. The leading non-volatile candidates to replace SRAM cells in programmable logic devices have cell sizes as little as a tenth the size of an SRAM cell, which would increase the capacity of the programmable logic device or it would make it cheaper.

The leading candidates

Among a dozen or so contenders, three stand out: magneto-resistive random-access memory (MRAM), ferroelectric random-access memory (FRAM or, sometimes, FeRAM), and ovonic unified memory (OUM).

MRAM. Magneto-resistive random-access memory stores its information by polarizing the magnetic domains in "ferromagnetic" material. Ferromagnetic materials throw off a magnetic field. The alignment of the magnetic field flip-flops between two directions depending on the direction of a nearby electrical current. Once set though, the magnetic field stays aligned in one of its two

DRAM main memory. As the semiconductor process improved, the first and then the second level of SRAM cache migrated onto the chip with the microprocessor. In 2001, SRAMs accounted for more than 15% of memory chip revenues, but the SRAM's share of memory chip revenues will decrease as SRAM migrates on chip in system-on-chip designs.

■ **Flash memory.** Flash memory chips are as dense as DRAMs but flash memory costs twice as much. Read times for flash memory are about the same as they are for DRAM, but writes take twenty times longer. Like DRAM and SRAM, flash memory represents 1s and 0s with the presence or absence of electrons. The flash memory cell contains a "floating" gate. High voltage causes electrons to "tunnel" onto the floating gate, which is so isolated that the charge will stay there for ten years or so even if the chip loses power. Because it doesn't lose its information when it loses power, flash memory is "non-volatile." Flash memory is good for any number of read cycles, but loses its properties after a few hundred thousand to a million write and erase cycles. This is enough for things like cameras and GPS receivers, but not for PCs.

Because writing is slow and because it wears out the cell, flash is best in applications with frequent reads and infrequent writes. The PC's "boot" memory is a good example. The boot memory holds the software that starts the PC and helps it find and initialize the operating system. The boot memory is read every time the computer starts. The boot software is loaded into the flash memory at the factory and changed infrequently thereafter.

In 2001, flash memory accounted for more than 30% of memory chip revenues. Flash memory's share of the market is growing. The combination of its reasonable read time (as fast as DRAM, but slower than

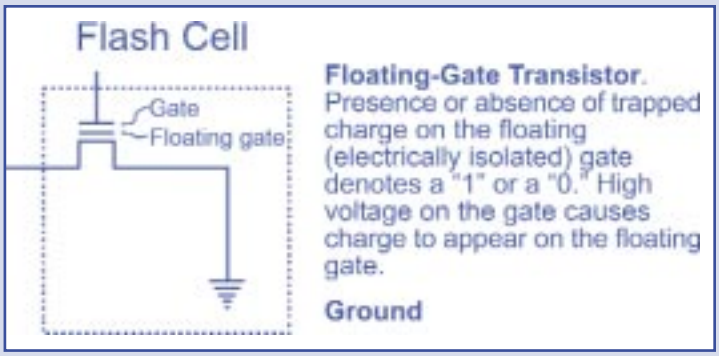

SRAM) and its non-volatile memory cell makes flash well-suited for portable devices. Flash memory is also good for portable devices because of its low power use. An MP3 player with a hard disk uses ten times the power of a flash-based player. MP3 players use hard disks for their capacity and their cost. Flash memory costs 250 times as much per bit; flash memories above a gigabyte aren't yet practical. Like SRAM, flash memory has migrated onto the chip with other functions. Flash memory is popular in microcontrollers, for example. The largest market for flash memories is cell phones. The fastest-growing market for flash memories is portable storage modules for consumer devices.

Flash memory cells, which contain only a single floating-gate transistor, can be smaller than a DRAM cell. The capacity of flash chips could outstrip the capacity of DRAM chips. Intel and AMD offer flash memories with two bits per cell. In addition to building the capacity of their chips, flash manufacturers, such as AMD, Hitachi, Intel, Fujitsu, and SanDisk, work to improve write/erase endurance so the cells won't wear out and they work to speed read and write cycles. As the market for mobile devices grows, the flash memory market grows. Flash manufacturers have an eye on these growing applications and they have an eye on the DRAM market. ■

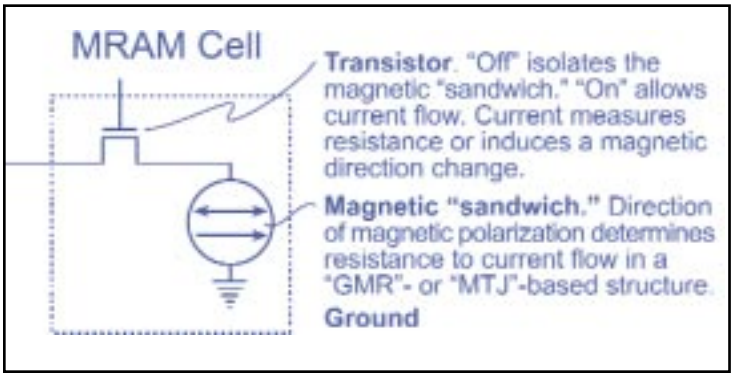

directions even if power is turned off. A small current can read the field without flipping it. MRAM is like a hard disk with no moving parts. For reading and writing, MRAM is about as fast as SRAM. Unlike flash cells, MRAM cells do not wear out after millions of read or write cycles. Two types of cells are MRAM candidates.

One type of MRAM cell uses the "giant magnetoresistive effect" (GMR). The GMR effect has been studied and developed for thin-film heads in hard disks. One weakness of the GMR-based MRAM cell is a small, and therefore difficult to detect, difference between the read signal for a zero and for a one. Union Semiconductor demonstrated a 1-Mb GMR-based MRAM almost two years ago, but is not yet sampling these chips.

Another type of MRAM cell uses a "magnetic tunnel junction" (MTJ). The magnetic tunnel junction cell sandwiches an insulating layer between two ferromagnetic layers. The MRAM cell is like a DRAM cell where the ferromagnetic sandwich replaces the DRAM's capacitor. The magnetic alignment of one of the two ferromagnetic layers is "pinned" in one direction. The other magnetic layer is aligned in one direction or its opposite by passing a current through contacts on the top and bottom of the sandwich. If the magnetic fields are aligned in the same direction, the sandwich has lower resistance than if the fields are aligned in opposite directions.

For reading, a small electrical current perpendicular to the sandwich detects high or low resistance. For writing, a large current perpendicular to the sandwich flips the magnetic alignment of the free ferromagnetic layer. IBM has demonstrated MRAM write times of 2.3 nanoseconds, which makes it *hundreds* of times faster than flash and makes it the fastest of the new candidates. IBM also demonstrated read currents down to 2 milliamps, making MRAM 40 times more power efficient than DRAM. Motorola produced a 256-kb MRAM test chip in 2001 and this year IBM is testing a 1-Mb chip.

More than two years ago, Motorola demonstrated "full integration of MTJ with standard low-cost CMOS." That says it's possible to integrate MRAM and CMOS, but that doesn't say that it's easy. One of the challenges in commercializing MRAM is integrating it with CMOS. In late 2000, IBM and Infineon announced joint research with a goal of commercial MRAM chips by 2004.

Dynamic Silicon companies Cypress Semiconductor and TSMC are backing MRAM. Cypress Semiconductor recently invested $6 million in MRAM developer NVE Corporation. TSMC is working with the Taiwanese government's Electronics Research and Service Organization on the integration of MRAM and CMOS. Other MRAM backers include Bosch, IBM, Honeywell, Infineon, Intel, Motorola, NEC, Toshiba, and Union Semiconductor.

FRAM. Ferroelectric random-access memory is built like DRAM, but replaces DRAM's capacitor with a ferroelectric crystalline material sandwiched between two plates. Atoms literally move between two positions in the crystal, in response to an externally applied electric field. The material exhibits "hysteresis," meaning that the atoms stay in their last position until influenced by an electric field strong enough to coerce a change. Ramtron brought the first commercial product to market in 1988. This year, Ramtron offers 256-kb FRAMs, based on a design that uses one transistor and one ferroelectric capacitor per cell, for less than four dollars. (Earlier versions were based on two transistors and two ferroelectric capacitors per cell.)

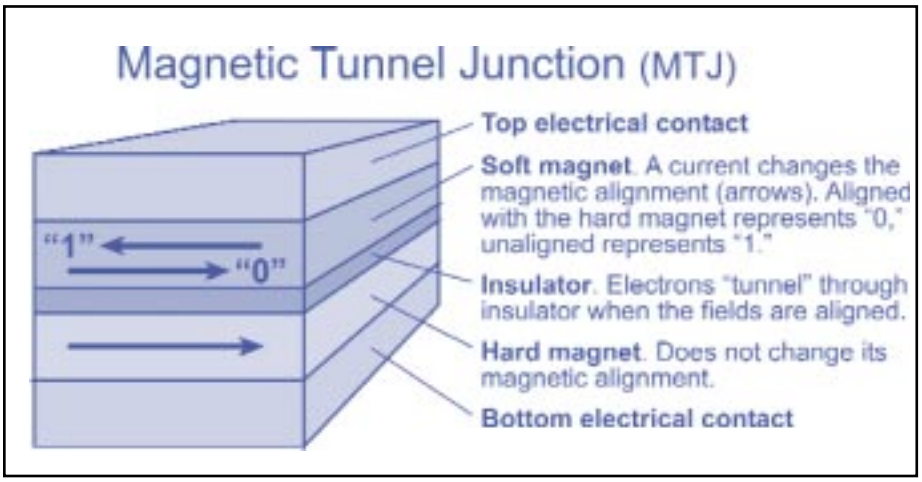

Toshiba and Infineon are Ramtron licensees. Toshiba built an 8-Mb FRAM prototype with one transistor and one ferroelectric capacitor per cell. In late 2000, Toshiba and Infineon announced plans to ship the 8-Mb FRAM and to introduce 32-Mb FRAM chips by 2003. Neither company seems to be selling

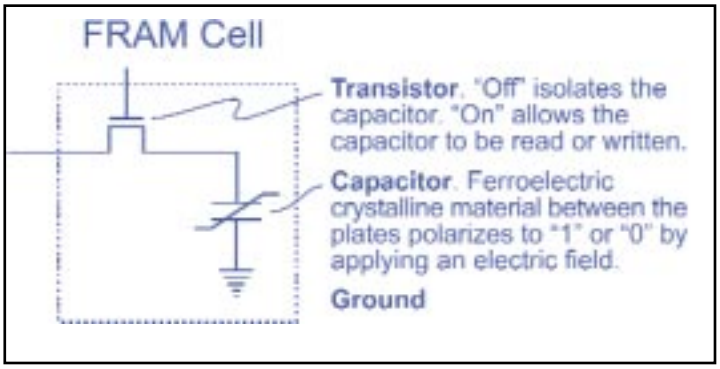

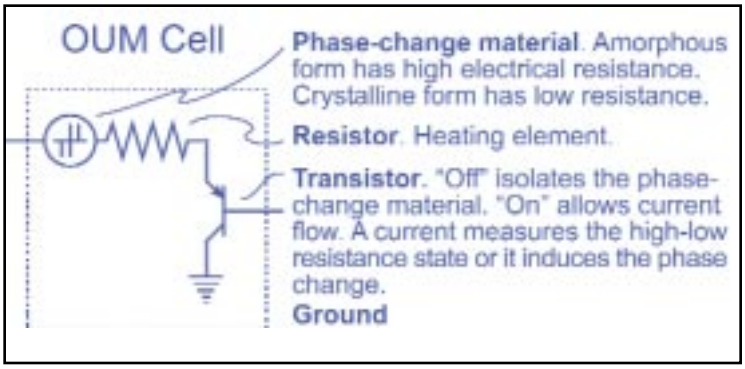

FRAM chips yet, making Ramtron's 256-kb FRAM the largest currently on the market. Product delays hint that the one transistor and one ferroelectric capacitor cell may be difficult to manufacture. The ferroelectric material must have uniform composition and thickness across the array's capacitors.

FRAM cells wear out just as flash cells do, but it takes *billions* of write cycles rather than the hundreds of thousands for flash memory. Like DRAM, FRAM cells have a destructive read cycle. The destructive read cycle means that reading takes about the same energy as writing and that reads have an associated write that contributes to wearing out the cell. The FRAM's destructive read cycle requires six to seven times the energy of the non-destructive read cycle of a comparable flash memory cell. So far, FRAMs are about half the speed of DRAM for reading and writing, which makes them faster than flash, but much slower than SRAM. Also, FRAM cells are large relative to DRAM cells, so FRAM makers will struggle to reach DRAM densities. It sounds bad for FRAM, but it is the only one of the leading candidates to have products in the market. So we can build it, but it may not be fast enough to replace SRAM.

FRAM integrates well with CMOS logic, using only two additional mask layers as opposed to the seven additional masks required for integrating flash memory. The ferroelectric capacitors are added after the CMOS circuits are fabricated and before the metal layers are added. FRAM has performance, endurance, and integration advantages over flash for system-on-chip designs. FRAM also has huge performance and power advantages over electrically erasable, programmable read-only memories (EEPROMs) in on-chip designs.

Fujitsu sells four chips with embedded FRAM. Two chips, with 32 kb of embedded FRAM, are 8-bit microcontrollers for smart card applications. The third chip, which also contains 32 kb of embedded FRAM, is a microcontroller for security applications. The fourth is a contactless RFID (radio-frequency identification) chip with a 5,000-gate logic circuit and 16 kb of embedded FRAM. Write cycles for the EEPROM version are thousands of times longer than the FRAM's write cycles, causing the EEPROM version to burn *hundreds* of times the power.

FRAM's backers include Celis Semiconductor, Fujitsu, Hynix Semiconductor, IBM, Infineon, Matsushita, NEC, Ramtron, Samsung, and Toshiba.

OUM. Ovonic unified memory is my choice, because of its simplicity. My explanation of OUM is shorter because of that simplicity. OUM uses a phase-change material and it is similar to that found in rewritable CDs and DVDs. By phase-change, I mean the material is either in a crystalline state, where its electrical resistance is low, or it is in an amorphous state,

Memory Scorecard

	Non-volatility	Speed	CMOS logic compatibility	Density	Standby power	Switching power	Cost	Lifetime	Future scalability	Backers
DRAM	👎	👉	👎	👍	👉	👉	👍	👍	👉	👍
SRAM	👎	👍	👍	👎	👎	👎	👎	👍	👍	👍
Flash	👍	👎	👉	👍	👍	👎	👉	👎	👍	👍
MRAM	👍	👍	👎	👉	👍	👉	👉	👉	👉	👍
FRAM	👍	👉	👍	👉	👍	👉	👉	👉	👎	👍
OUM	👍	👉	👍	👍	👍	👍	👍	👉	👍	👍

where its resistance is high. For writing, a current passing through the material and through an attached resistive heating element melts the material to allow a phase change. While it does wear out, ovonic memory can be read and written hundreds of billions of times. For reading, a small current measures whether the material is in the high or low resistance state. Reading is non-destructive (the cell value doesn't have to be written after reading). For both reading and writing, ovonic memory is slower than SRAM and is faster than DRAM.

The difference between the electrical resistance in the crystalline state and the resistance in the amorphous state is large enough to represent multiple bits. The material doesn't just have to be at the extremes of amorphous or crystalline; the range of resistance can be subdivided into four or even sixteen segments so that the cell can represent two or four bits, respectively. Two or four bits per cell doubles or quadruples the memory density.

Azalea Microelectronics has built a 4-Mb ovonic memory chip under contract from Intel. The cell size is potentially smaller than DRAM's cell size. Ovonic memory integrates well with standard CMOS. As with FRAM, the phase-change material is deposited after the CMOS circuits are built and before the metal layers are added.

Backers of ovonic unified memory include Azalea, BAE Systems, Intel, Ovonyx, and STMicroelectronics.

Other candidates

A $28-billion-dollar market is at stake. It's encouraging plenty of candidates and it's broad enough for niche technologies.

If you need the speed of SRAM and the non-volatility of flash, but can't wait for the new candidates' developers, there are at least two types of nvSRAM. Non-volatile SRAM sounds like an oxymoron. One version, from Apta, embeds a lithium battery in the package with a standard SRAM. The lithium battery holds the SRAM's contents when power is lost. In normal temperatures, the lithium battery should last 10 to 30 years (in high-temperatures, it may last only 3 years). The second type of nvSRAM, from Simtek (www.simtek.com), employs a standard SRAM chip that is shadowed by an electrically erasable programmable read-only memory (EEPROM) and some control logic that monitors the chip's power. When the control circuit detects unstable power, it backs up the SRAM data to the EEPROM. When power is applied, the controller copies the EEPROM's contents into the SRAM.

There are also candidates with carbon nanotubes, holographic memories, atomic resolution memories, organic compounds, quantum dots, and polymers. All have attractive attributes and major shortcomings. It will be many years before any of these is a commercial threat to DRAMs.

And the winner is…

All of the candidates are non-volatile. Each has powerful and impressive backers. All have been around for a long time—the *newest* is more than ten years old. Gordon Moore of Intel wrote about amorphous and crystalline non-volatile memories—ancestors of today's OUM—in 1970. Ramtron invented FRAM in 1984 and it brought the first chips to market in 1988. MRAM has been around since 1992. There are pluses and minuses to this situation. On the plus side, these aren't new and untested ideas brought to market by undercapitalized startups.

On the minus side, the semiconductor giants have been working to get products to market for years and the success so far has been unimpressive. MRAMs aren't on the market yet. The largest MRAM announced so far is IBM's 1-Mb test chip. Ramtron is shipping 256-kb FRAM chips. Fujitsu is shipping microcontrollers with as much as 32 kb of FRAM on chip. Toshiba has built an 8-Mb FRAM prototype. Ovonic memories aren't on the market yet. Azalea has built a 4-Mb OUM prototype for Intel. Meanwhile, *256-Mb* DRAMs have been shipping for three years. The largest-shipping candidate is five generations behind DRAM density.

MRAM is the fastest, it has the most backers, and it gets the most press. It also seems hard to build—it doesn't integrate well with CMOS. MRAM's cells could prove difficult to scale at a pace that matches the improvements in the CMOS process. FRAM is the only candidate that has real products in the market. Fujitsu's microcontrollers show that FRAM does integrate with

DynamicSilicon

Editors Nick Tredennick, Brion Shimamoto
Publisher Lauryn Franzoni
Designer Julie Ward
Subscription Services Melissa McNally
Editorial Director Richard Vigilante
Chairman George Gilder

Dynamic Silicon is published monthly by Gilder Publishing, LLC. Editorial and business address: 291A Main Street, Great Barrington, MA 01230. Editorial inquiries can be sent to: bozo@gilder.com.

CMOS. But FRAM's relatively large cell size will make it difficult to scale to compete with DRAM's density.

I like ovonic memory for the simplicity of its cell and for the straightforward nature of its manufacturing process. OUM integrates well with CMOS. It has low-power reads, it operates at low voltages, and its control circuits are simple. As Goldilocks would say, "It's just right." Unlike MRAM and FRAM, for which scaling is a challenge, OUM's performance should improve as its cell shrinks. OUM is lithography-limited; no breakthroughs are required to maintain scaling that tracks CMOS progress. It is potentially as dense and as cheap as DRAM. Since OUM doesn't store its 0s and 1s in electric or magnetic fields that limit proximity to other cells, there's potential to build 3D chips for higher storage density. Too bad there are no OUM chips on the market that prove its potential. Just as it was with MEMS-based storage (*Dynamic Silicon*, Vol. I, No. 5), there's lots of potential for better memories, but it's too early to name a winner. Today, it looks as if OUM will be just right.

Nick Tredennick and Brion Shimamoto
May 16, 2002

Manufacturers Scorecard

	Flash	MRAM	FRAM	OUM	Other
AMD	■				■
AZALEA MICROELECTRONICS	■			■	
BOSCH		■			
CELIS SEMICONDUCTOR			■		
CYPRESS SEMICONDUCTOR		■			
FUJITSU	■	■	■		
HITACHI	■	■	■		
HONEYWELL		■			
HP		■			
HYNIX SEMICONDUCTOR	■		■		
IBM	■	■	■		■
INFINEON		■	■		
INTEL	■	■		■	■
MATSUSHITA		■	■		
MICROMEM		■			
MICRON	■		■		
MITSUBISHI	■	■			
MOTOROLA	■	■	■		
NEC		■	■		
NVE CORP.		■			
OVONYX				■	
RAMTRON			■		
SAMSUNG	■	■	■		■
SANDISK	■				
SHARP	■				
STMICROELECTRONICS	■			■	
TOSHIBA	■	■	■		
TSMC		■			
UNION SEMICONDUCTOR		■			

Dynamic Silicon Companies

The world will split into the tethered fibersphere (computing, access ports, data transport, and storage) and the mobile devices that collect and consume data. Dynamic logic and MEMS will emerge as important application enablers to mobile devices and to devices plugged into the power grid. We add to this list those companies whose products best position them for growth in the environment of our projections. We do not consider the financial position of the company in the market. Since dynamic logic and MEMS are just emerging, some companies on this list are startups.

Company (Symbol)	Technology Leadership	Reference Date	Reference Price	5/2/02 Price	52-Week Range	Market Cap.
Altera (ALTR)	General Programmable Logic Devices (PLDs)	12/29/00	26.31	20.75	14.66 - 33.60	8.1B
Analog Devices (ADI)	RF Analog Devices, MEMS, DSPs	12/29/00	51.19	36.80	29.00 - 53.30	13.4B
ARC Cores (ARK**)	Configurable Microprocessors	12/29/00	£3.34	£0.27	£0.25 - £1.06	£115.6M
ARM Limited (ARMHY***)	Microprocessor and System-On-A-Chip Cores	11/26/01	16.59	10.15	8.39 - 19.20	3.5B
Calient (none*)	Photonic Switches	3/31/01				
Celoxica (none*)	DKI Development Suite	5/31/01				
Cepheid, Inc. (CPHD)	MEMS and Microfluidic Technology	12/17/01	4.73	3.40	1.48 - 11.48	92.0M
Chartered Semiconductor (CHRT)	CMOS Semiconductor Foundry	7/31/01	26.55	25.42	16.06 - 34.00	7.1B
Coventor (none*)	MEMS IP and Development Systems	7/31/01				
Cypress (CY)	MEMS Foundry, Dynamic Logic	12/29/00	19.69	22.20	14.00 - 28.95	2.7B
Cyrano Sciences, Inc. (none*)	MEMS Sensors	12/17/01				
QuickSilver Technology, Inc. (none*)	Dynamic Logic for Mobile Devices	12/29/00				
SiRF (none*)	Silicon for Wireless RF, GPS	12/29/00				
Taiwan Semiconductor (TSM†)	CMOS Semiconductor Foundry	5/31/01	14.18††	18.20	8.39 - 20.99	61.3B
Tensilica (none*)	Design Environment Licensing for Configurable Soft Core Processors	5/31/01				
Transmeta (TMTA)	Microprocessor Instruction Sets	12/29/00	23.50	2.26	1.17 - 18.75	287.0M
Triscend (none*)	Configurable Microcontrollers (Peripherals)	2/28/01				
United Microelectronics (UMC†)	CMOS Semiconductor Foundry	5/31/01	10.16	10.40	4.25 - 11.52	27.6B
Wind River Systems (WIND)	Embedded Operating Systems	7/31/01	14.32	10.00	9.71 - 29.25	781.8M
Xilinx (XLNX)	General Programmable Logic Devices (PLDs)	2/28/01	38.88	37.70	19.52 - 50.98	12.6B

† Also listed on the Taiwan Stock Exchange

†† TSM reported a stock split on 6/29/01. The Reference Price has been adjusted for the split.

* Pre-IPO startup companies.

** ARK is currently traded on the London Stock Exchange

*** ARM is traded on the London Stock Exchange (ARM) and on NASDAQ (ARMHY)

NOTE: This list of Dynamic Silicon companies is not a model portfolio. It is a list of technologies in the Dynamic Silicon paradigm and of companies that lead in their application. Companies appear on this list only for their technology leadership, without consideration of their current share price or the appropriate timing of an investment decision. The presence of a company on the list is not a recommendation to buy shares at the current price. Reference Price is the company's closing share price on the Reference Date, the day the company was added to the table, typically the last trading day of the month prior to publication. The authors and other Gilder Publishing, LLC staff may hold positions in some or all of the companies listed or discussed in the issue.