

MEDIA FOR

□
Solid-state memory takes over in niche military and aerospace applications

By J.R. Wilson

Solid-state memory grows in popularity for niche applications, including military and aerospace uses, yet magnetic and optical memory continue to dominate desktop commercial applications because of their low prices per megabyte.

Designers of solid-state memory systems have made significant strides in recent years, not only in increasing speed and density, but also in reducing costs. At the same time, however, those designing mechanical bulk memory have made equally significant advances in those same areas. As a result, the cost/performance computation continues to be a serious factor as systems integrators determine what kind of mass memory system they choose to implement in new programs, in upgrades, or in retrofits.

Yet lately, performance has begun to gain a new edge over cost. In an interesting side development, solid state is beginning to take over from older mechanical technologies in low-density applications where price is a deciding issue.



Radiation-hardened solid-state memory, such as that pictured above from Aeroflex, is starting to take over high-reliability niche applications of bulk memory from traditional magnetic and optical media.

Bulk memory applications represent "a very tough thing to go after with solid state," admits Steve Hudgens, chief technology officer for Ovonix Inc. in Santa Clara, Calif. "If you look at it just from a standpoint of cost, hard disks keep getting cheaper and cheaper. They don't have the read/write performance of solid state, but you can't beat the price. If you look at both cost and performance, there is an argument for solid state. People need archival storage, which would be a very tough market for solid state to address, unless there also is a performance requirement."

Rich Wawrzniak, director of non-volatile memory at Semico in Irvine, Calif., says the fast-dropping price of mechanical technology actually may lead solid state into the market that hard drives now dominate.

"Nobody should expect solid-state drives to replace hard-disk drives on a cost per megabyte basis," Wawrzniak says. "However, as hard-disk drive technology advances and they gain the capability of denser and denser drives, it is easier to do large drives than small drives. That means the cost for a 2-gigabyte hard drive will not be much less than a 60-gigabyte drive, so there are a lot of applications at the lower end of the spectrum that get squeezed out. And if you switch to a different format, you may find the footprint is different, the power requirements are different, and it just may not work. When you get to that point, the designers may just opt for solid state, but the cost of the end system has to be great enough to amortize the cost of the solid-state drive."

Cost differences

The cost difference per megabyte is substantial. For example, looking at available prices for consumer hard drives, a 36-gigabyte Hitachi hard drive costs \$476, compared to \$79,000 for a 32-gigabyte TMS RamSan-210 from Texas

Memory Systems in Houston —151 times as much. On the other hand, the RamSan-210 performs 200,000 input/output operations per second (IOPS), compared to only 300 for the hard drive —a 660-fold advantage.

The differences are not quite as great when making the comparison to a Fibre Channel-based Redundant Array of Inexpensive Disks (RAID) system, but are still noteworthy, says Woody Hutsell, marketing director at Texas Memory Systems. A 655-gigabyte Sun StorEdge T3 Array costs \$84,100, compared to more than 13 times as much (\$1,112,000) for a comparable 655-gigabyte RamSan-520 stack. However, he adds, the maximum IOPS they have seen for any RAID is 5,000, compared to 750,000 for the RamSan-520 —a 150X performance advantage for the solid-state disk.

"That wait time is largely due to latency in getting data from the RAID system to the process server, which can be 5 milliseconds, best case," he says. "With the RanSam, that is reduced to 20 microseconds in all circumstances. That is fundamentally because the hard disks inside the RAID are using mechanical components, while with a solid-state disk, the data is right there. Because so much of what is happening in high concurrent user applications is the server waiting on I/O, if you can eliminate that wait time, you will get better use of your servers."

But I/O and acquisition cost are not the only factors to consider, he adds.

"Historically, solid state has been used for direct access to accelerate a single server, but with the use of Fibre Channel you can use the one solid state disk and leverage it across multiple servers," Hutsell says. "That's particularly the case with our solid state disks, which are very high bandwidth. The 210 has a 700MB/s backplane and the 520 is 3 gigabytes per second. The 520 also has 15 Fibre Channel access ports (four on the 210), all of which could be used simultaneously with no performance degradation.

"So fast storage is a function of both bandwidth and I/Os per second. If you are doing database transactions with small, frequent, random blocks, you can't get the throughput out of a RAID you can out of a solid state disk."

Similar comparisons can be made for other formats, such as the low density (2GB or less) flash PC-cards and CompactFlash produced by SanDisk (Sunnyvale, Calif.), which currently retail at about 50-cents per megabyte, considerably more than modern hard drives.

"But if your application only uses 100 megabytes, the lowest-cost drive you'll find is probably \$70 because they have a minimum price they have to charge just to cover the parts —and you can probably get several gigabytes for that price," notes SanDisk marketing director Jeff Ellerbruch. "So flash is actually a better buy. But if you need a gigabyte or more, flash becomes more expensive than mechanical and you have to have a real need for it —such as improved environmental capabilities."

Solid-state memory

Solid-state bulk memory comes in a variety of formats, some still in development. In addition to the Texas Memory Systems solid-state disk and SanDisk's flash cards:

- Ovonyx is the sole proprietor of the basic intellectual property behind Ovonic Unified Memory (OUM);
- Semico and Ramtron International Corp. of Colorado Springs, Colo., are among several companies working on enhancing ferroelectric read only memory (FRAM), with the latter also working on enhanced Static Random Access Memory (SRAM);
- Aeroflex UTMC of Colorado Springs, Colo., produces one of the most common memory chip formats found in desktop computers, Synchronous Dynamic Random Access Memory (SDRAM);
- The Motorola Embedded Memories Center in Tempe, Ariz., is home to ongoing research into magneto-resistive read only memory (MRAM); and
- BiTMICRO Networks in Fremont, Calif., develops solid-state drive, flashdisk, network-attached storage (NAS), and storage area network (SAN) products based on its patented FlashBus technology.

"The new technologies have to prove they are truly viable, that they can be mass produced," Ellerbruch says. "Getting past those two hurdles is very difficult. A few years ago, some people expected to see holographic storage by the turn of the century. I'm not aware of any yet. Even looking at those new technologies and if they make their way into some of our segments, we still believe flash will have a place. But it's hard to say before we actually have something to compare ourselves against."

As for competing with the older technologies, he says that, too, is likely to remain a matter of niches and defining

requirements.

"If an application needs a lot of storage and can use a disk drive, then that drive will probably be cost-effective, but if that application requires rugged media or low power, the user may be willing to pay a premium for flash," Ellerbruch says. "And as we keep decreasing the price of flash, we increase the scope of where flash gets designed in. For example, if you have a battery environment, flash at 500 megabytes has advantages - flash will run considerably longer than a mechanical drive. There can be some advantages to flash as a kind of front end to some big applications. But you would have to do some things —that can be done —to get performance up on the flash. You would need a fast, beefed up controller on the front end, for example. And we have a partner who does that —Adtron Corp. [Phoenix, Ariz.].

Flash memory

"But I have a hard time believing flash will replace disk drives at the high-capacity level, cost-effectively, in the next 20 years, at least. It's hard to see it happening. We survive with disk drives in their market and us in ours and I don't see anything that is likely to change that in the foreseeable future."

Ellerbruch says SanDisk will continue to work on increasing flash densities, improving power, and lowering cost, but he does not expect the company to venture into any of the other technologies. There is no need, because he sees the market for flash expanding dramatically in the next decade, even to the point of replacing such storage technologies as videotape, DVD, cassette tapes, and CDs.

"Today our biggest application for flash is digital photography," he says. "That increases our production level and decreases our costs, which we can then pass on to other applications. As those other commercial applications evolve, the same cost and production factors come into play for our other customers, such as aerospace and defense."

Like flash, SDRAM is an established format that is increasing in performance and decreasing in price. Experts at Aeroflex UTMC, for example, are qualifying a 1-to-2-gigabyte SDRAM stack for a satellite mass memory system, company officials say.

"Once completed, we will take it to the open market as a 2-gigabyte SDRAM module for use as a key component in a mass memory system," says standard products director Tony Jordan. "It provides high-bit density —a high volume of bits in a small volume —which is key in a satellite system. It takes eight components to get to 2 gigabytes, but you could stack 80 of them and go to 20 gigabytes.

Saied Tehrani, manager of MRAM efforts at Motorola, says the highest capacity Motorola experts have achieved to date is 256 kilobytes of memory, but says that will increase dramatically soon.

Motorola MRAM

The Motorola process integrates a thin film of magnetic material with CMOS circuitry. The system stores information as magnetic polarization in the magnetic material; silicon circuitry reads and programs the bit. The basic cell architecture is based on one transistor and one magnetic tunnel junction. The transistor, during the read mode, turns on and detects the resistance of the magnetic tunnel junction; the two states of the memory are two different resistances in the junction. Programming happens when two currents flow in perpendicular lines —one setting on top of the magnetic tunnel junction and the other underneath it, with the bit at the cross-junction of the two lines.

"We have been able to demonstrate 256 kilobytes of memory with this architecture," Tehrani says. In that demonstration, we have shown we can integrate the magnetic materials with CMOS and at the same time achieve access time of 35 nanoseconds, both read and write. In microcircuitry, we are increasing the density rapidly by reducing the lithography geometry —the minimum size we can build in the processing. Initially, MRAM is somewhat behind the state of the art in density in solid state, but we will catch up."

He also cites factors in the technology that offer significant advantages over its competitors:


- Unlike DRAM, it is non-volatile and it is composed of radiation-hardened materials, making it ideal for military and aerospace applications;
- there are no cycle limitations on the number of times users can read/write to MRAM, unlike flash or FRAM;
- the bit can be programmed faster than flash (35 nanoseconds); and
- programming does not require high voltage.

"The technology is being driven into smaller geometries and higher densities," Tehrani says. "We expect to do

sampling in 2003 and volume production in 2004, initially in the 16 megabytes to 32 megabytes density. We anticipate the cost to be very competitive with existing solid-state memories. Our basic architecture minimizes the real estate consumed for storing a single bit, which should make it very competitive with DRAM and considerably smaller than SRAM, which requires six transistors per bit. MRAM will drive toward the pricing of flash, although it would take some time to reach the pricing of flash, which is a much more mature technology."

Motorola officials say they are not concerned about price or density comparisons with mechanical drives because they are going after an entirely different market segment —mainly embedded applications —where MRAM would be integrated with processors and other components on a chip.

"What we see MRAM being able to do is mainly replace multiple memory," Tehrani says. "Most systems today require multiple memory —they need flash for non-volatile storage, but because of the slow programming inherent in flash, they also must use SRAM. If you need to put information into memory and grab it back out, using flash alone would take a long time and cost a lot of performance. By putting SRAM in, as well, you can get the information in and out more quickly, but if you lose power, you lose all the data. So you are continuously going back and forth between the two memories. Because MRAM is more of a universal memory, it should be able to replace those multiple configurations and enhance performance at a cost savings."



Pictured above is a photomicrograph of a Motorola 256 kilobyte magnetoresistive random access memory, otherwise known as an MRAM.

The development program began with Defense Advanced Research Programs Agency (DARPA) funding in 1996, with an emphasis on its rad-hard capabilities for military and aerospace applications. Today Motorola officials are paying for the effort entirely with the company's own money —and are looking for expanded markets.

"Once we drive the technology to commercial applications, it obviously will mean improved performance and reduced costs for all users, including military and aerospace," Tehrani says. "We have a significant interest in this technology for wireless and automotive applications, which probably will be the first markets."

OUM memory

Ovonyx officials see similar capabilities —and potential —for OUM, which also is non-volatile, radiation hardened, requires little power, and is faster than flash; it performs read/writes in tens of nanoseconds and uses the same phase change process exploited for years in rewritable optical disks. But OUM, too, is still primarily a laboratory technology.

The Ovonyx approach uses a pulse of energy to switch the material between two structural states —crystalline or amorphous —going back and forth depending on the amount of energy applied and the time in which it is applied. When the energy input is from a laser, it is a rewritable optical disk; if the input is electric current, it is OUM.

An optical disk reads out stored information by measuring changes in reflectivity, depending on whether the material is amorphous or crystalline. OUM measures the electrical resistance, which changes by three or four magnitude when the structure changes. And unlike FRAM, the read-out is non-destructive; there is no need to "flip" the state when it is read, so there is no cycle wear-out on multiple reads.

That approach also

is fast. Ovonyx claims OUM can read in only a few nanoseconds and write in only tens of nanoseconds. By comparison, flash performs a block erase before it writes, which slows the process considerably. Flash also has a typical lifetime of 100,000 to 1 million rewrites; Ovonyx claims they have seen cycle life in excess of 1 quadrillion rewrites (10 to the 12th).

"Our technology is basically a back-end process," Hudgens says. "We apply the material layers patterning to build cells as a two- or three-mask process, which is easily integrated with the existing logic process without disrupting it,

so it is easily embedded. Our process simplifies computer-on-a-chip efforts because you just put it on top of the logic flow.

"Because it is so fast and nicely scalable, we think over time it may come to be seen as a replacement not only for flash but also for system memory —DRAM," Hudgens continues. "If we could extend the cycle life out to 10 to the 14th or 15th, there would be no need to distinguish between volatile and non-volatile. It is compact as DRAM and potentially as fast, so it could become a replacement. That has some very powerful consequences from a system standpoint. For example, if archival and system memory are the same thing, the computer could start instantly, without having to load an OS."

OUM is not in commercial production yet, but Hudgens says Ovonyx officials still can project costs that will be lower than flash, with a higher density —much lower than FRAM —and competitive with DRAM.

[Click here to enlarge image](#)

	DRAM	SRAM (BT)	FLASH	OUM	MUT-RAM	FERAM
Cell size (F ²)	6-12	50-80	7-11	5-8	???	Large
Volatile/non-volatile	Volatile	Volatile	Non-volatile	Non-volatile	Non-volatile	Non-volatile
Endurance write/read	= / =	= / =	1E6 / =	1E12 / =	1E12 / =	1E12 / 1E12
Read	Destructive	Partly-destructive	Non-destructive	Non-destructive	Non-destructive	Destructive
Direct Over-write	Yes	Yes	No	Yes	Yes	Yes
EE/Style (Write/Erase)	Yes	Yes	Block	Yes	Yes	Yes
Read (dynamic range (margin))	100-200mV	100-200mV	Delta Current	100-100K R	20-40% R	150-200mV
Programming energy	Medium	High	High	Low	Medium	Low
Write/Erase/read time	30ns/50ns/30ns	8ns/5ns/8ns	Sub/1-100ns/block/10ms	10ns/50ns/20ns	30ns/30ns/30ns	80ns/80ns/30ns
CMOS Logic compatibility	Bad	Good	OK, but Hi-V req'd	Good	???	OK, but Hi-V req'd
New Materials	Yes	No	No	Yes	Yes	Yes
Scalability Limits	Capacitor	ST	Tunnel Oxide/HV	Litho	current density	Possible Capacitor
Multi-bit storage	No	No	Yes	Yes	No	No
Relative Cost per Bit	Low	High	Medium	Low	???	High

FRAM memory

Ramtron and Semico officials say they see things from a slightly different perspective —one more favorable to FRAM, the oldest of the "new" technologies, with a start point dating back to 1984. Most of the advances in capacity are fairly recent, however, with more expected in the next few years.

FRAM stores information using ferroelectric crystals, which have the unusual property of possessing a partially free electron in the center that can occupy one of two stable energy states. The crystals then can use those two energy states to represent digital ones and zeros without an outside power source to maintain memory. While a wear-out mechanism has been identified, companies working with the current technology say it is measured in trillions of cycles before cell failure.

"We've been shipping production products at the 4 kilobit level since 1992 and it reached 256 kilobytes this year," says Ramtron marketing vice president David Bondurant. "We hit a couple of key milestones this year in that it is now shipping with a 3-volt power supply and achieving unlimited endurance. Projecting into the future, we recently announced a 0.35-micron processing technology we jointly developed with Fujitsu in Japan, which will support densities of 1 megabyte to 4 megabytes. Our roadmap expects FRAM to reach 1 megabyte next year and 4 megabytes to 16 megabytes in 2003. We also expect it to reach 64 megabytes by 2004."

Speed is not a problem —FRAM reads and writes at the same speed, with a current cycle time of about 130 nanoseconds. By comparison, E2PROM (electrically erasable programmable read-only memory) has a write speed measured in milliseconds and flash takes milliseconds to write after the time it spends in block erase.

FRAM's limited capacity forces it into a niche of its own, although Bondurant says he does not believe that niche is quite as restrictive as some might think.

"Cell phones want higher densities, but the leading supplier of cell phone chips is Texas Instruments and we are jointly developing a FRAM for their products to meet requirements in the 2004 timeframe," he says.

Today FRAM is approaching the cost of flash as it approaches 64 megabytes, "but I don't expect it to become a lower cost technology through that generation," Bondurant says. "FRAM is rapidly overtaking SRAM in density and will pass it in both cost and density in a couple of years. As the densities get higher, the solid-state disk gets interesting, although the first application in that arena probably will be memory card applications. An FRAM card would have all the benefits of a battery-backed SRAM card, but without the battery. Military computers today are being built with SRAM and DRAM, but FRAM densities are getting to the point where you could consider having a totally solid-state, non-volatile main memory. That is enabled by the fast read and write capability of the FRAM —something flash just can't do."

Another possible application involves the radio frequency identification (RFID) tags now being developed for airline

checked baggage and cargo. Cost is a significant consideration there because most of the tags would be single use/disposable.

"We are currently selling 4-kilobyte FRAMs for 41-cents," Bondurant says. "For an RFID tag, we probably could build one that would meet the airlines' price requirements. The FRAM technology is the only one that can read quickly as bags enter and exit the read field as it moves down the conveyor.

"The other thing FRAM is good for is non-volatile storage in any embedded application, especially compared to E2PROM or battery-backed SRAM. The majority of our current products also are shipping for the industrial temperature range, which makes them appropriate for most COTS applications in the military."

Wawrzniak, however, says Semico leaders have no hesitation in reciting the limitations of FRAM, despite his company's involvement in the technology.

"The density is too low for most applications; it still has a reputation of being hard to manufacture, so you have limited manufacturing sources; because cell phones are transitioning from 16 megabytes to 32 megabytes and in two or three years will be up to 64 megabytes, the question is whether FRAM will be able to keep up," he says. "There are other technologies on the horizon better suited to solid-state drives than FRAM. None is out yet —MRAM, for example, and Intel's colemetric ferroelectric and OUM probably hold more promise long-term.

Military & Aerospace Electronics December, 2001

Author(s) : J.R. Wilson

Interested in a subscription to *Military & Aerospace Electronics Magazine*?

[Click here](#) to subscribe!

Links referenced within this article

Click here to enlarge image

[javascript:OpenLargeWindow\(80024,650,257\);](javascript:OpenLargeWindow(80024,650,257);)

Click here

<http://www.ameda.com/cgi-win/mae.cgi?add&p=webedit>

Find this article at:

http://mae.pennnet.com/Articles/Article_Display.cfm?Section=Archives&Subsection=Display&ARTICLE_ID=129777

Uncheck the box to remove the list of links referenced in the article.

MEDIA FOR