

MEDIA FOR



▪
Bulk memory designers eye new technologies

By J.R. Wilson

Military and aerospace systems designers need ever-increasing amounts of bulk data memory systems in a wide variety of technologies, and engineers in the data-recording and data-storage industries are crafting several new candidates.

The digitized battlefield is information intensive by its very nature, as hundreds of gigabytes of data flow in all directions and require storage, filtering, evaluation, ranking, and distribution with the highest possible reliability and speed. These qualities are placing new requirements on bulk memory capabilities as well as very high-speed networking to link the media where data reside.

New memory and data-transfer technologies under development will blend into this scenario as quickly as experts can validate them for military applications — assuming, of course, that costs are not prohibitive. Nevertheless, a major concern remains compatibility in areas such as protocols, interconnects, and power. New memory systems — particularly solid-state bulk memory systems — also must be able to sustain exposure to radiation and still operate reliably, not only in space but also on Earth.

"We've been pushed quite hard recently to develop denser and denser memories," says Joe Benedetto, manager for standard products technology at Aeroflex UTMC in Colorado Springs, Colo. "In our main line, we sell to satellites and avionics, both of which have concerns regarding single-event upsets (SEUs). It has become more of a challenge to build products needed for avionics and space as new, smaller geometry bits hold less and less charge, which can be upset by what are called 'terrestrial neutrons', which are in the atmosphere."

An SEU occurs when a charged particle travels through a transistor substrate and generates electrical signals within the transistor. This usually occurs in near-Earth orbit to spacecraft passing through the Van Allen radiation belts. Solid-state memory chips are particularly vulnerable to the effects of single-event upsets.

Memory chips

One approach to boosting the densities of solid-state memory chips without going to smaller geometries is called "stacked packaging," which evolved to meet the cellular telephone industry's constant need to shrink the size of products. Several companies, such as Intel of Santa Clara, Calif., and Tessera Technologies in San Jose, Calif., have developed this technology for commercial applications. Aeroflex UTMC, which has concentrated on military and aerospace markets, currently offers 1 gigabit of SDRAM in a cube measuring 0.3 inches high by 0.5 inches long by 0.2 inches wide.

Aeroflex UTMC experts say they hope to increase the memory capacity within that same volume by stripping the typical 500-micron bottom layer of plastic and reducing the size of the silicon itself.

"A typical silicon die is about 250 microns, but only about 10 microns are actually needed for the part to function; the remaining 240 are there for rigidity and structural integrity," Benedetto says. "We're looking to see how thin we can go; we hope to get as close to 10 microns as possible. But in space applications, the final product still has to be pretty rugged to survive shock and vibration, so we have to determine how much we can remove without causing any cracks in the silicon during the lifetime of the application.

"The ultimate goal would be to increase capacity by up to 25 times in the same volume, if the physics allow us to get

there," Benedetto continues. "We are looking to go to 8x (30 microns of silicon) in the next six months, which would double our current density, and we believe that will qualify without much trouble. The next doubling will probably be about a year after that. After that, we will evaluate the data to see if it makes sense to go any further. We have a much tougher challenge in qualifications because our market is aerospace."

To reach those goals, Aeroflex will leverage commercial developments, where a 1-gigabit die is being introduced. Providing it passes all the qualification tests, that would enable Aeroflex UTMC engineers to stack 16 layers of silicon by early 2004, and put 16 gigabits into the same space currently required for only one gigabit.

"Our customers are requesting these largely for solid-state recorders," Benedetto says. "The goal is to replace the existing spinning-media drives with solid state components that are much smaller, more reliable, and much, much faster, by a factor of about 1,000. With 16-gigabit cubes, it would take 25 cubes to replace a 50-gigabyte hard drive."

While speed, weight, and power will hold fairly constant as this density increases, heat dissipation will become a major challenge, experts say. As chip components are packed more tightly within the cube, wattage per cubic millimeter will increase dramatically.

Optical memory

Another version of stackable storage is optical holography, which stacks information throughout the thickness of a storage medium rather than only writing it to the surface, as is currently done on the face of a disk. This relies on a combination of lasers and a photosensitive material, such as a crystal, to save data. The laser beam splits in two; one half carries data and the other, called a "reference beam," holds the location of that data. The two beams intersect to create a pattern of light and dark bands. The process then engraves a replica of that pattern three-dimensionally into the crystal to create a hologram. Users can recover the stored data by shining the reference beam into the hologram, which refracts the light to replicate the data beam.

Using this technique, one 12-centimeter disk theoretically could hold one terabyte of data (roughly equivalent to 200 DVDs). Theory also says a holographic storage system could read and write data one million bits at a time instead of one-by-one as is the case with existing magnetic storage. This means that users could move or duplicate massive amounts of data in seconds.



The Series 4 DTS, pictured above, from the Targa Systems Division of L-3 Communications Canada provides a removable means for moving Targa SCSI Flash Disks between mission platforms and data retrieval systems.

Researchers are looking into the viability of several different volume holographic techniques. Scientists report the most promising to date are angle-multiplexed, wavelength-multiplexed, spectral, and phase-conjugate holography.

Angle- and wavelength-multiplexed methods differ only in the way they store and retrieve data, using either different angles of reference beam incidence, or different wavelengths. Spectral combines a photorefractive crystal with a time-sequencing scheme, partitioning holograms into individual subvolumes of the crystal, then using the collision of ultrashort laser pulses to differentiate between the image and the time-delayed reference beam. Phase-conjugate holography eliminates the need for some optical parts. It does this by replacing the reference beam that reads the hologram with a conjugate reference beam that moves oppositely to the recording beam. The hologram diffracts a

signal, and this signal moves back along the path from which it came and refocuses onto the detector.

One of the leading developers of holographic memory is InPhase Technologies of Longmont, Colo., a spin-off of Lucent, which has set a goal of bringing its first product — Tapestry, which is initially write-once and capable of recording 100 gigabytes of video — to market in a couple of years. InPhase officials say their technology breakthrough is based on Tapestry's two-chemistry polymer media, which they claim is dimensionally and thermally stable, has high photosensitivity, will be competitively priced, and will have an archival life of 30 years. It also will support high-speed interfaces, such as SCSI and Fibre Channel.

Scientists at the NASA Jet Propulsion Laboratory (JPL) in Pasadena, Calif., also have been working on a high-density, high-speed holographic system using a newly developed electro-optic (E/O) beam steering technology to achieve a design goal of storing as much as 250 gigabits in a cubic photorefractive crystal with a transfer rate as fast as 1 gigabit per second.

Insufficient storage

Experts say such capabilities are essential to NASA's future Earth Science missions, which will require massive and high-speed onboard data storage capability — as much as 40 terabits stored between downlinks by 2003. However, JPL scientists estimate that onboard storage capability next year will be only 4 terabits — 10 percent of the requirement — and by 2006, will only be able to support 1 percent of onboard storage requirements.

Because the current leading commercial technologies — Static Random Access Memory (SRAM), Dynamic Random Access Memory (DRAM), and Flash memories — suffer from problems ranging from high volatility to limited endurance and poor radiation-resistance, JPL scientists are trying to develop compact holographic memory using electro-optic beam steering that will satisfy all of NASA's mission needs for non-volatility, radiation hardness, long endurance, high density, high transfer rate, as well as low power, lightweight mass, and small volume. Officials say the scientists are developing and demonstrating holographic data recording and retrieval, with storage capacity of as much as 250 gigabits, with as much as two orders of magnitude increase envisioned.

These projects, however, still represent technologies in development, and are likely to encounter problems facing any new technology when it comes to military applications, points out David Huestis, principle investigator on smart time-domain optical memory at SRI International in Menlo Park, Calif.

"In the whole field, there are no devices in use in a computer or network situation based on holography, even though the theory originally was explained in the 1960s in a book by Joe Goodman, Introduction to Fourier Optics — roughly, how to use light to represent information in amplitude and phase," he says. "So the idea is 40 years old. How you might interface it is still being thought about. The computer's architecture was invented in the 1940s — the concept of a CPU and peripheral — so it is 60 years old. And that is one of the difficulties in incorporating optical components into computer components — they don't match."

SRI's own optical program began in 1986 with funding from a Japanese telephone company and has continued through the years with money from the U.S. Defense Advanced Research Projects Agency (DARPA) and the U.S. Air Force. SRI's program involved storing information based on the spin orientation of impurity atoms in a crystal — such as neodymium in an yttrium aluminum garnet (a YAG crystal), which readily absorbs available light. A sequence of light pulses sent into the crystal will be recorded, and then reproduced if the crystal is later re-excited.

"Roughly, we could store 100 gigabits per cubic centimeter, which in 1985 was an unbelievable amount; in 2002, it isn't," Huestis says. "But if I store 100 gigabits in a cubic centimeter, I can read the whole thing in 10 microseconds. You could easily get bit rates of 10 to the 15th bits per second. The problem is, you don't have anything you can put that data into other than another optical crystal. The data format and access procedures are very different from any other media we have."

"What we have done so far is a sequence of demonstrations of concept," Huestis continues. "Some progress has been made in how to adapt optical storage ideas into a digital system. In the follow-on proposal we wrote for the Air Force, we tried to pick out things of particular importance to the military that won't necessarily be automatically done by the commercial developers of computers and networks, such as very fast space/time processing. If we have a live video stream we want to analyze in real time, computers can do that with conventional memory, but intelligent screens may have much higher bandwidth. The data also would have bit-to-bit relationships in space and time."

This kind of optical system also could help control phased-array radar, he adds, by programming as many as 1,000 patterns of different intensity and time delay into a crystal, then pulling them out one at a time as needed to steer the radar's different antennas.

At the moment, however, funding to continue the effort is in doubt, in part because of the success of U.S. military technology.

"This is all still in the future," Huestis says. "The community of people involved is starting to drop away as funding diminishes and the need for this sophistication is falling with the falling sophistication of our opponents. It's not obvious how the military should be doing its investments as a result. The U.S. has always wanted to have the fanciest of everything — and that has usually done well for us — but now there is some contemplation about going back to older, proven systems.

"What the Defense Department should do for technology is now a very difficult subject. On the research side, the defense research organizations are much crippled by the Mansfield Amendment, which requires them to have a specific defense application for whatever they fund, which makes it difficult for them to fund things into the distant future. We have aging inventory but we still want the most advanced technologies. The new technology is still important, but maybe we haven't been doing the necessary life analysis of it as we invent it."

Molecular memory

Another new approach with strong potential military applications is solid-state molecular memory, being developed by Rolltronics Corp. in Menlo Park, Calif., using what is called a "roll-to-roll" process.

"The technology is similar to some other memory technologies being developed in that it utilizes a crossbar architecture, where the layer where the memory data is stored is between two layers of patterned electrodes, perpendicular to each other," explains Glenn Sanders, the Rolltronics vice president of business development. "At that point, the similarity ends. The others we are aware of are based on a chemical change of state or some other physical change. Ours stores data through a charge migration within the material itself.

"The material (zinc porphyrin) is an insulator, but the class of materials we will be using reacts to the presence of an electrical bias, where there is a greater potential on one side than the other, plus the presence of light," Sanders continues. "When both elements are there, there is a migration of electrons within the layer itself. The molecules in the layer arrange themselves in what resemble little stacks of quarters — billions of stacks of molecules next to each other. There seems to be little if any crosstalk between those stacks."

Removing either light or bias causes the material to revert to its natural state as an insulator, which locks the molecules. Sanders says his engineers have tested it for as long as 7,000 hours (nearly 10 months) with no measurable change in the charge. They also cycled it as many as 1.5 billion times with no measurable change in the amount of charge stored. That compares to a couple of hundred thousand cycles for flash before the reliability of stored data diminishes severely.

Simply put, the flexible plastic device — produced in sheets 1 meter wide and 10 kilometers long — comprises four main layers: Zinc porphyrin sandwiched between two layers of electrodes (any standard material that can produce a bias) and the light source. Sanders says engineers in the future may produce a low-luminance light source inexpensively in the roll process, which would make it easier to integrate with the memory array. The memory array is then mated with existing chips, circuits, and form factors.

"In principle, if you take a PCMCIA flash memory card, you can take out the memory there and replace it with ours, giving you a new memory device with higher density, longer life and lower cost than flash. In the first generation, the PC Card form factor could hold about 640 megabytes; of course, the first generation of any electronic product is always the simplest architecture," he says.

"By the third generation, we might reach 5 terabytes in the volume now used by a standard 3.5 inch hard disk. Every computer has a bay that can hold that form factor, so it could conceivably be plug-and-play, hooked up through USB or whatever the interconnect may be at that time. We're looking at first generation in about 18 months to two years and third generation in about five years."

Sanders claims the process will produce a rugged, entirely solid-state memory that will meet military environmental requirements better than most others now available, with a temperature range from -40 to 50 degrees Celsius, shock

tolerance of 1,000 Gs (on or off), power usage in the milliwatts per centimeter range, and weight measured in milligrams per megabyte. Because it is a roll-to-roll process, the material also must be flexible, which lends itself to substantial design flexibility. Reliability and durability also come into play, comparing an extremely tough, flexible plastic sheet to a rigid green circuit board or glass.

"For data density, it would be lighter than any other solution we know about," Sanders says. "In the beginning, we expect it to be much less expensive than flash, perhaps half as much per megabyte. One of the most compelling attractions to roll-to-roll is high volume, low cost. Compare creating a single, very expensive 12-inch round silicon wafer to a single sheet about 5 feet wide by 5 miles long (some 26,000 square feet); the thinner you get, the more you can fit into the same size roll, so there would be some potential for even longer sheets. With that process, to produce a comparable storage capacity might be a little as 10 percent the cost.

"Speed is dependent on read/write circuitry; the memory itself responds in a sub-nanosecond," Sanders says. "We won't know precisely what the first generation will be able to do, in terms of speed, until we actually build the first prototypes and begin interfacing them. We expect to have multi-megahertz speeds depending on the interconnect available at that time."

Sanders says the Rolltronics' Nanoscale Molecular Memory (NanoMem) device technology "is totally agnostic as far as interface goes" because they will not be making memory devices for sale directly to end-users. Instead, manufacturers will license the technology to build into their own systems, probably retaining existing products but replacing the current storage medium with NanoMem.

The long sheets of material produced in the roll-to-roll process can be cut and shaped to fit any form factor and layered on top of each other to create stacked storage options such as the 5 terabyte hard disk replacement.

Even a PC Card format — at 64 gigabytes — would provide a massive amount of storage in an extremely compact form, enough for:

The conversion of that to military information requirements would provide significant storage capacity for moving maps, photoreconnaissance, manuals, etc.

"Imagine being able to carry all of that in a PC Card — and at a cost of only a few hundred dollars," says Rolltronics CEO Michael Sauvante. "That is our goal — high-density memory at affordable prices. It will take several years, but our molecular memory technology will make it possible."

Chalcogenide memory

Another new technology, in which the Air Force has invested research funding, is chalcogenide super-high-density memory. The term "chalcogen" refers to the Group VI elements of the periodic table; "chalcogenide" refers to alloys containing at least one of those elements.

According to industry researchers, chalcogenide is super-dense for non-volatile memory applications and has intrinsic benefits for the space community because it is inherently rad-hard and extremely fast, changing states in the 50-nanosecond range. Most of the current applications are 256 kilobits or at best 1 megabit, but are expected to reach 16 megabits around 2005 and eventually 64 megabits, although when that happens is largely dependent on funding. While those are considered tiny by modern commercial standards, it would still constitute a considerable level of memory for a non-volatile space application.


BAE Systems and Ovonyx Inc., both in Manassas, Va., have integrated a chalcogenide memory element with BAE Systems radiation-hardened CMOS technology, and used another test chip to develop 64-kilobit arrays with full write-read circuitry suitable for environmental and radiation testing.

"Our goal is to have produced and characterized chalcogenide-based memory structures and to have gathered enough data to begin a product design targeting a 1-to-4-megabit C-RAM device that is latch-up and SEU immune and total dose hard to greater than 1 megarad, operating across the full temperature range commonly specified for space applications," says Laura Burcin, BAE Systems's program manager for chalcogenide RAM.

According to BAE Systems, current state-of-the-art suggests that radiation-hardened, non-volatile, 1-to-16-megabit devices with unlimited read cycle endurance and read and write access times less than 100 nanoseconds could be available by 2004.

Despite the shortcomings cited at JPL, Flash memory remains one of the leading current technologies for fast, portable transfer and storage of large quantities of data.

Designers at the L3 Communications Targa Systems Division in Ottawa, plan to introduce an 18-gigabyte flash disk storage system with a high-grade military connector in the first quarter of 2003. Targa's Series 4 Data Transfer System (DTS), consisting of a Data Transfer Unit (DTU) and a removable Data Transfer Device (DTD), provides a compact self-contained system to store and retrieve data from Targa's 2.5" SCSI Flash Disks and meets the most demanding MIL-STD-810 and RTCA DO-160 environments, says Joe Fronsee, Targa's vice president of marketing. Targa introduced a 9-gigabyte version in 2002, which is part of a map system upgrade on British RAF Tornado jet fighter-bombers.



Engineers at UTMC Aeroflex are stacking solid-state memory chips as thinly as possible, with a goal of increasing systems as much as 25 times

The DTS flash disk storage component integrates E-Disk FlashBus technology from BiTMICRO Networks in Fremont, Calif. The DTS features a locking access door in a ruggedized housing, SCSI-2 and SCSI-3, as well as USB 1.1 and USB 2.0 interfaces and supports 28-volt input power.

"The Series 4 DTS provides the high capacities [larger than 5 gigabytes] and high-speed data transfer [18 megabytes per second] that is required in today's demanding military and aerospace applications, such as map displays, imaging and mission loads." Fronsee says. "We chose the BiTMICRO technology because it helps us meet our customer's system requirements."

Officials of Conduant Corp. of Longmont, Colo., recently introduced an advanced version of their StreamStor high-speed digital recording technology capable of supporting as much as 5 terabytes of storage with sustained 200 megabytes-per-second record and playback capability. The Compact PCI-816 digital recorder also includes a new daughter board architecture providing industry standard high-speed interfaces.

"Our products are intended as alternatives to slower, larger, and more expensive tape drives," says Mark Walker, Targa's vice president of marketing. "And we are being adopted widely in the military because the database requirements for those applications are going through the roof. StreamStor is the only product that can sustain the needed data rates for hours and hours. Beyond the speed, which is significant, is our new product, which is an infinite recorder.

"We have replaced 20-year-old reel-to-reel cabinets with rack-mounted, 5-U high systems that can record seven times as fast and provide all the advantages of disks — random access, low maintenance, high speed. And they can go on infinitely because we have created a system that banks eight drives together in a module, two modules in a cabinet. The user can record to one bank until it is full, and then have the stream switch to the second module while the operator removes the first bank and installs a new one. That can be done over and over as long as necessary."

With its speed and capacity, Walker says, StreamStor can record the entire Encyclopedia Britannica in 5.12 seconds and the 20 million books contained in the Library of Congress — about 20 terabytes of text — in 29.13 hours.

Switched-fabric memory

Another relatively new technology that continues to evolve in the bulk storage and rapid data transfer arena is very

high-speed serial Switched Fabric interconnect, such as StarFabric from StarGen Inc. in Marlborough, Mass. StarFabric products are silicon devices that implement switched fabric protocol for embedded distributed processing applications.

"The initial bridge we have is a PCI bridge, a protocol-agnostic switched interconnect architecture," explains Greg Whelan product marketing director at StarGen. "This overcomes a number of problems inherent in a standard bridged architecture, such as scaling. In a rack, every blade would have the full bandwidth of the PCI bus. You can have redundant switched fabric with automatic failover, high priority provisioning for high priority traffic, etc. In January, we will have parts ready for TDM (time division multiplexing) voice switch, which will allow both voice and data traffic in a single, unified bus.

"You can start today with StarFabric at 2.5 gigabytes and migrate to 10 gigabytes advanced switching (AS) in five years and take most of your software with you," Whelan says. "Because the alignment between the two is so close, what we call StarFabric 2.0 will, in fact, be AS 1.0. StarFabric will continue to exist for applications requiring 2.5 gigabytes or less. Ten gigabytes requires more expensive cables and connectors, especially in the comparatively low unit demands of military systems."

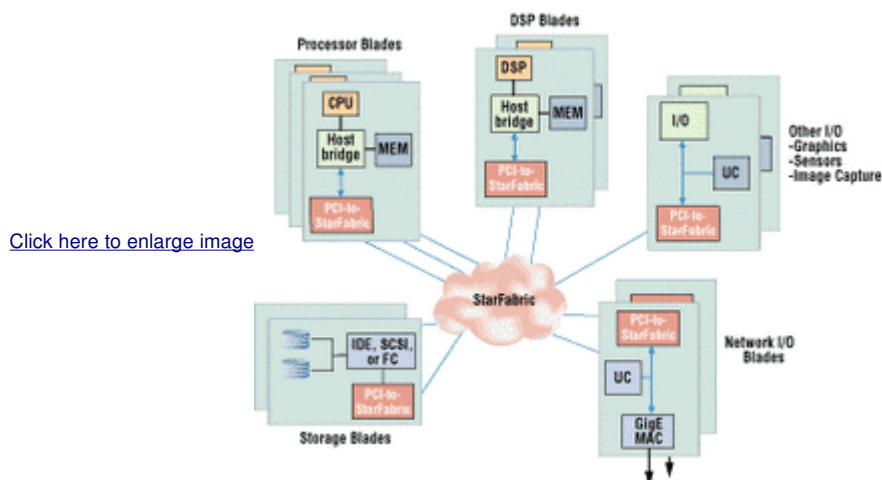
He says StarGen officials probably will begin producing silicon on version 2.0 at the end of 2004 or early 2005, with the first products using it appearing perhaps by late 2005.

"It too will be a very scalable technology. I think you will, over time, see up to 16x switches [using sixteen 2.5-gigabyte links] to create a 40 gigabytes per second interconnect, but that's some time down the road," Whelan says. "Initially, you will see 1x, 4x and maybe 8x [for graphics applications]. "There are a number of military applications, including a number already in place today. You won't find us in an Apache helicopter, but we are finding our way into a lot of communications equipment infrastructures, including radar applications.

"In the battlefield of the future, you have a massive amount of data coming in, a lot of it analog, and you need multiple processors to handle it, crunch that into digital, then crunch it again," he says. "That is massive amounts of data and you don't have time for a lot of different protocols — you want to send a lot of data through the system as quickly as possible. With our switched interconnect, we send data through a system in a series of memory read/writes, with the switching decisions based on memory address, not protocols; we can encapsulate any protocol."

A recent StarFabric application is PMC-StarLite, an ultra-high-speed switched fabric interconnect PMC module from Radstone Technology in Woodcliff Lake, N.J. PMC-StarLite provides a 6-port StarFabric switch, with each port capable of 5 gigabits per second peak communications — 2.5 gigabits per second peak transmit plus 2.5 gigabits per second peak receive duplex. The 6-port switch is connected to a StarFabric-to-PCI bridge (64 bit/66 MHz) to provide a seamless connection between multiple compute nodes on the host board.

"As signal processing problems become increasingly demanding of bandwidth, using many multiples of processors and boards, there needs to be a corresponding method of transferring large amounts of data between processors — both on-board and off-board," says Stuart Heptonstall, Radstone's DSP product marketing manager. "PMC-StarLite is designed to dramatically increase interboard communications bandwidth over existing shared bus technologies."



High-speed serial switched fabric interconnects such as StarFabric from StarGen, illustrated above, implement

switched-fabric protocol for embedded distributed processing and bulk memory applications.

A primary military problem with virtually all bulk memory systems is the high environmental demands of military and space applications compared to the relatively benign needs of the commercial world. COTS components are not tested to full military shock, vibration, radiation, or temperature extremes. The traditional MILSPEC temperature range, for example, is -55 to 125 C, compared to an official range of 0 to 45 C for the new 64-bit PCI System-On-Module Blazor architecture from PFU Systems in Santa Clara, Calif.

"Traditionally, we have defined temperature specs very conservatively," notes Kishan Jainandusing, the PFU vice president of marketing. "With our current 32-bit products, the military has been using them at much lower and higher temperatures than we have listed. Some data points have been able to drive all the way down to -30 C and as high as 70 C. The company traditionally has not been too forthcoming in terms of wide temperature use of these components. We are making some changes so we eventually will be screening for our customers and offer products in wider temperature ranges. Those will, of course, be more expensive because of the additional screening and lower yields.

"However, we have been very successful in the military so far, with our current components being used in a wide variety of temperature ranges," Jainandusing says. "Most of our customers essentially do what is necessary to design the enclosure properly. We have seen the requirement for robust electronics, but not a willingness to pay the high prices necessary to provide full mil specs. That prohibits a lot of military subcontractors from starting designs with those components."

Jainandusing says the industry has been in a steady movement away from mil-spec components and focusing on the container. Efforts, however, continue to determine what designers can do with commercial components in efforts to meet reliability requirements of the defense environment.

"We are looking at about 10-G shock in operational conditions and a quarter G in vibrations (for the new Blazor product). Essentially, that is a lot less than what we have on our RazorBlade products [shock 50 Gs and vibration 20 Gs]," he says. "But when you look at where Blazor is going to be deployed, it won't be moved around a lot, so we rely on the container approach a lot more than we did for the other products. But that still gives the designer a fair range of shock-and-vibration ruggedness to design against. "We are continuing to work with our customer base to define other products based on this architecture. If our customer needs it at a wider temperature range, we can build that from the architecture without having to redesign a final product. That enables us to bring things to market much more quickly."

The 64 bit bus on Blazor (which begins shipping in January 2003) can carry about four times as much data as a 32-bit bus and eight times when coupled with a 133 MHz SDR/DDR PCI-X bus, which PFU experts say they will be able to incorporate rapidly, probably by the second half of 2003.

"If you go to double-data-rate PCI-X, you are handling 16 times more in terms of bandwidth. Down the road, we are looking at PCI Express, which gives you even more bandwidth," Jainandusing says. "Most of the peripherals today are PCI/PCI-X based; PCI Express is probably two years down the road."

That capability becomes increasingly important with the military's increasing requirements for high-speed, high-quantity data movement on the digitized battlefield and in sophisticated built-in-training simulation systems on platforms such as tanks. All of this data must be processed and interpreted in the field, often at the platform level.

"Some of the interesting applications we've seen with respect to Blazor include centralized video service on an aircraft carrier, which is used for briefings, post-mission analysis, training, etc. As the massive amounts of data collected during and after a raid come back, they need to be analyzed on the aircraft carrier for immediate use," Jainandusing says. "Further down the line, we see other things, such as back-end signal processing for things like sonar and radar, going toward higher resolution and with more data to represent the results.

"All of this today can be done today essentially with CompactPCI, but that remains a bulky form factor. Even onboard an aircraft carrier, you still have to be sensitive to space; that is even more true aboard a tank. To integrate this kind of data capability, having the module come in as a component rather than even a single board computer gives them better options and freedom to design the form factor."

Military & Aerospace Electronics December, 2002

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\(111694,353,257\);](javascript:OpenLargeWindow(111694,353,257);)

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=163721

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

MEDIA FOR