

# Slot Card Constraints Complicate Memory Choices

Designers of board-level computers are chafing against a number of space and technology challenges as they craft new memory subsystems.

by Phu Hoang  
Virtium Technology

**S**ystem engineers designing for the real-time market are currently facing some of their greatest design challenges. Not only must they understand the industrial computing issues facing their customers in such growing fields as telecommunications, medical, military, servers and high-end workstations—they must answer them with systems offering speed, rugged endurance and mission-critical reliability.

The real-time system designer must unite myriad technical components—hard drives, microprocessors, system I/O and a significant range of memory technologies and densities—to support intense performance demands under adverse computing conditions. Memory subsystem design looms as a highly complex piece of this puzzle, complicated by technology and market issues that must be considered and understood to perfect the ideal system at the ideal cost.

The real-time user demands powerful, ruggedized memory solutions, such as higher density memory suited for extended temperature ranges. Real-time applications demand single slot VME or CompactPCI computers, in turn triggering demand for high-density, small form-factor memory

subsystems. Memory products that meet these demands are not commercially available and must be custom designed and built for individual real-time environments.

While real-time system engineers can remain thankful this is the role of the memory designer, it's crucial they understand the complexities, priorities and shifting market issues facing developments in memory subsystems for real-time computing.

## The Density/Space Conflict

Real-time applications—required to react in a timely, predictable and reliable manner under extreme load conditions—can include industrial automation, military command and avionics, transit systems, medical imaging or testing systems, network routers, or high-speed printing systems. These and other types of rigorous computing environments are frequently supported by VME or CompactPCI systems recognized for their deep roots in military applications and their powerful, rugged and streamlined performance.


VME systems offer a finite amount of space, which becomes extremely valuable when that space must support many critical functions. Power supply, hard drive, data

collection or maybe a tank control system must not only meet mission-critical performance requirements, but must do so within very limited physical space.

Commercial memory architectures, such as those used in standard PCs, represent unacceptable space requirements for these real-time computing environments. For example, a standup DIMM (Dual Inline Memory Module) is so tall, it requires two VME slots—space that simply cannot be sacrificed for telecommunications, military or similarly stringent applications. As a result, designers are questing after single slot solutions.

While moving to a smaller form-factor does free up precious VME slots, the memory designer must now address critical density issues, adding enough memory to maintain top real-time performance.

This highlights the ultimate challenge of real-time memory subsystem design—density vs. space creating a complicated “one step forward, two steps back” design environment. The challenge can only be answered by memory subsystems that incorporate small form-factor (Figure 1), high-density memory rugged enough to handle the extreme demands and conditions of real-time computing.



**Figure 1** DIMM packaging offers an easy fit into a standard DIMM socket but results in a height problem. The module can be placed into a 90° angled socket, however that only addresses the height issue by significantly increasing the amount of horizontal space required. SODIMM packaging answers the issue of space, but doesn't offer all the critical memory functions needed for real-time computing. Custom solutions such as Virtium's CompactRAM provides an alternative that meets reduced space requirements.

Memory Module Options	Advantages	Disadvantages
DIMM (standup socket)	<ul style="list-style-type: none"> <li>• Standard</li> </ul>	<ul style="list-style-type: none"> <li>• Requires two PCI or VME slots</li> </ul>
DIMM (90 degree angled socket)	<ul style="list-style-type: none"> <li>• Still standard, but down to one rack slot</li> </ul>	<ul style="list-style-type: none"> <li>• Requires more planar space</li> <li>• Limits to one or two DIMMs maximum</li> </ul>
SODIMM	<ul style="list-style-type: none"> <li>• Uses less space than DIMM</li> <li>• Still offers the standard</li> </ul>	<ul style="list-style-type: none"> <li>• Delivers half the density</li> <li>• ECC capability not commonly available</li> <li>• Non-standard registered capability</li> </ul>
Custom solutions such as Compact RAM	<ul style="list-style-type: none"> <li>• Highest density economically feasible</li> </ul>	<ul style="list-style-type: none"> <li>• Proprietary/custom solution</li> </ul>

**Table 1** Memory subsystem designers can choose from a number of module technologies, each with its own purpose and limitations. Commercial packaging standards such as DIMM and SODIMM form a basis for familiar and cost-effective memory subsystems. In slot card racks DIMM packaging presents a height problem. The module can be placed into a 90° angled socket, however that significantly increases the amount of horizontal space required.

## Standards and Practices

Issues of density and space are further complicated by the lack of standards for memory sockets in real-time applications. Rugged and powerful real-time memory solutions are simply not available as commercial off-the-shelf products, and must be developed from the ground up.

Which memory chip technology to select looms as the first and perhaps not-so-basic design issue. Considering perfor-

mance requirements of the memory subsystem leads the memory designer to any number of choices—each founded in commercial PC markets and each subject to change or becoming obsolete very quickly.

“DRAMs are distinguished by differences in their operating modes—EDO, synchronous DRAM, DDR and Rambus—each requiring a different controller design. When yesterday’s DRAMs become unavailable and new DRAM operating modes

are incompatible with existing controllers, memory engineers are forced into expensive controller redesigns or reliance on older DRAM types which have become relatively expensive niche products,” said Steve Cullen, director and principal analyst, semiconductor research, InStat/MDR.


Combined with issues of memory chip selection, memory designers must consider form-factor or packaging as a fundamental design element. Asking which mechanical functions the memory subsystem must execute leads the memory designer to a number of module technologies, each with its own purpose and limitations (Table 1). The memory designer must also consider whether custom products are the best choice, and whether to develop solutions internally or outsource for the best end memory product.

As a starting point for real-time memory designs, memory engineers frequently begin with longtime PC components such as DIMM and SODIMM (Small Outline Dual Inline Memory Module). These commercial packaging standards form a basis for familiar and cost-effective memory subsystems, however significant limitations must be addressed in order to adapt these packaging technologies for real-time demands.

For example, DIMM packaging offers an easy fit into a standard DIMM socket but results in a height problem. The



**Figure 2** Virtium's CompactRAM memory module design offers staged upgrades, required functions such as error correction and registered memory module technology, significantly reduced space requirements and two to three times the memory density in a single VME or CompactPCI slot. Adding stacked components can further increase density delivering 2 to 3 Gbytes within a single frame.



**Figure 3** In order to target blade servers supporting high-level communications and military applications, designers at Kontron expanded its product base from dual slot configurations into single slot CompactPCI. Single slot designs like Kontron's DT64 CompactPCI board faced extreme physical space limitations. That drove Kontron engineers to work closely with a custom memory module designer, and ultimately develop a custom form-factor memory module that delivers 2 Gbytes of registered SDRAM in a single slot.

module can be placed into a 90 degree angled socket, however that only addresses the height issue by significantly increasing the amount of horizontal space required. SODIMM packaging answers the issue of space, but doesn't offer all the critical memory functions needed for real-time computing. Ultimately, neither option offers a valid solution with the right amount of memory and function within the

right amount of space.

## Combining Art and Science

New and creative packaging technologies and designs are answering this need for high-density memory in space-constrained environments. Custom solutions such as CompactRAM (Figure 2) offer staged upgrades, required functions such as error correction and registered memory module technology, significantly reduced space requirements and two to three times the memory density in a single VME or CompactPCI slot. Adaptations such as adding stacked components can further increase density—for example delivering 2 to 3 Gbytes within a single frame.

Whichever packaging technology is used, translating commercial components into ruggedized, small form-factor memory solutions requires a creative approach. Innovation in the early stages of memory subsystem design is vital, and system engineers now routinely consult with memory manufacturers and subsystem designers at step one.

"We expanded our product base from dual slot configurations into single slot CompactPCI—mainly for blade servers supporting high-level communications and military applications. When designing memory subsystems within these types of extreme physical space limitations, there is just no way to rely on standard memory module designs without sacrificing density and capacity," said Jan-Philippe DeBroeck, product manager, Kontron. Kontron is a vendor for board- and system-level computers.

"Complex high-performance systems require acute understanding of memory design—and there is a great advantage provided by specific design experience," said DeBroeck. "We opted to work closely with a custom memory module designer, and ultimately developed a custom form-factor memory module that delivers 2 Gbytes of registered SDRAM in a single slot (Figure 3)."

## Fundamental Questions of Design

Beyond chip technology, density and form-factor, memory design and engineering decisions must reflect a solid understanding of the nature of the market and the factors affecting price volatility.

Forecasts cause supply and demand to fluctuate greatly, and general market confusion sometimes results from new technologies being announced or predicted. A good example is the market flux that took place when DDR was announced as an alternative to Rambus—memory market players weren't sure which to manufacture or use in subsystem designs, and for a time there was a shortage on both technologies.

"PC requirements continue to be the primary driver for DRAM development and volume sales, but PCs have a relatively short market life and their changing needs result in rapid obsolescence of DRAM technologies," said InStat's Cullen. "This presents a dilemma for real-time memory designers with business needs at the opposite end of the spectrum. Their primary focus is sustained availability of compatible DRAMs—and they recognize the value of sacrificing early cost breaks to assure long-term availability."

Even random speculation through the broker channel can cause shipping to hold and influence short-term prices. Memory designers—ultimately in need of a six-month pricing commitment—typically must work within 24-hour price guarantees and so must closely track these market and technology influencers.

Real-time memory subsystems are typically engineered for a five-year lifecycle, meaning flexibility and upgradeability is also a critical design concern. To assure selection of the right memory suppliers and future compatibility with die revisions, memory designers must leverage and expand business relationships that can help nurture a deeper understanding of this diverse market.

## Space, Density and Function

Memory subsystem designers must be keenly aware of engineering issues relating to the speed and performance of the memory bus—designing around critical issues such as impedance control, spacing, signal integrity and interconnectivity in memory designs.

"There are so many technical nuances—reflections from signal lines, correct series termination, proper trace routing and much more that directly affects overall signal integrity. Although we design an extensive array of high-speed boards, including full systems for factory

It's a given that memory subsystems will continue their path toward faster, better, higher density performance that requires less space—however it's the constantly shifting design dynamics that will fuel significant advances in memory packaging, density standards and design technology for real-time markets.

automation, data acquisition and even space station controls, sometimes there is a real business case for working in parallel with custom memory module manufacturers," said John Wyatt, senior design engineer, SBS Technologies Government Group-Raleigh. SBS is a developer of open architecture embedded computer products for COTS military, communications and commercial/rugged industrial applications.

"Long-term, consistent performance requires near perfect engineering—even an understanding of basic microwave transmission theories—and an almost microscopic design approach," said Wyatt. "Today's real-time memory designers are frequently tapping the module manufacturer's familiarity with specific chipsets and their access to high-performance simulation software."

Memory designs must also take into account the demanding performance requirements common in real-time applications. Features such as ECC (error correction code)—which tests the accuracy of data written to and read from memory—may be forfeited for certain commercial applications, but are a design necessity for real-time memory subsystems.

Real-time computing also needs registered memory module technology in place, meaning a single system clock is regenerated into multiple copies to synchronize all chips effectively. The quality of the duplicated clocks is extremely critical for reliable operation. It's also very difficult to achieve, as they are designed with complex feedback loops to ensure no propagation delay.

Reliable systems must integrate both ECC and registered memory module technology—and these are both standard with DIMMs. However when added to SODIMMs, designers are faced with a configuration that is not readily available

as a commercial product—a good example of ongoing real-time memory design issues.

### Commercial Components, Industrial Reliability

Before commercial components can be deemed suitable for real-time environments, reliability testing assures they can withstand harsh conditions and still maintain signal integrity. This is becoming a specialized field because most commercial memory manufacturers don't offer (and don't need to offer) this service.

Memory subsystem designers turn to specialized industrial manufacturers to push each commercial component to its limit through tests such as the HAST, or Highly Accelerated Stress Test. Testing screens non-industrialized components designed for commercial temperature ranges. Performance is measured under adverse conditions common in mission-critical and other real-time applications, such as high pressure, vibration or increased temperature.

### Big Picture Memory Design

Achieving super high-density memory subsystems using the smallest form-factor possible—and doing it better than anyone else—takes an overall approach that looks well beyond architectures or commercial components. Market and economic issues, mergers, supplier relationships, standards, and of course memory technology and function work together to create a "design dynamic" that must be interpreted and leveraged by memory designers for the most powerful and compact memory product.

It's a given that memory subsystems will continue their path toward faster, better, higher density performance that requires less space—however it's the constantly shifting design dynamics that will

fuel significant advances in memory packaging, density standards and design technology for real-time markets.

As a result, relationships are increasingly important to the memory subsystem design process. By working closely with custom memory manufacturers, system designers have access to more reliable supplier relationships, longer term design plans and veteran engineering teams that can customize and share insights on memory designs—positively affecting the development of the system itself.

System engineers should recognize memory manufacturers as a significant resource and asset—providing market information, access to DRAM in volume and the unique ability to understand, design and stress test commercial technology to the standards of a five-year design lifecycle. Development of trusted partnerships should be the top priority for today's system designers—seeking out the right memory manufacturer for their particular needs, getting them involved early and making the most of their hands-on design experience. ■

Virtium Technology  
Irvine, CA.  
(949) 460-0020.  
[www.virtium.com].