

VME in 2003: A Glimpse into the Future of Embedded Computing

Since its introduction in 1981, the VME standard has maintained its leadership position in medium to high-end embedded computing by undergoing a process of continual revision. Yet, it is possible that new technologies could successfully challenge this leadership position. Applications that are appropriate for VME today may be better served by emerging technologies in the future. The best way to analyze the situation is to ask tough questions such as: "Does VME still solve my problem? Will markets turn to other solutions such as CompactPCI? Will VME adapt to meet my needs in the future?" Nobody has a crystal ball, but a series of events that occurred during the past year have laid the groundwork for VME to improve its competitive position as we move towards the next millennium. The roadmap that is taking shape now indicates that when we reach the year 2003, VME will have undergone significant improvements in performance and will incorporate new features that maintain its value relative to competing technologies.

In the early 1980s, when the bus wars were raging between VME and Multibus, VME supporters felt it was time to unite and combine their efforts. As a result, the VMEbus International Trade Association (VITA) was formed. VITA soon created a group called the VITA Standards Organization (VSO) to drive those standards that would ensure the continuing viability of VME over time. Figure 1 shows some of the key standards currently working their way through the approval process. Let us consider how current VSO activity will benefit users as we move towards the year 2003.

VME64x introduces a number of new features including the 5 Row/160 pin DIN connector, with its additional user I/O pins, as well as the IEEE1101.10 front panel, with its updated handle that supports both ejection...and... insertion of boards. The new handles are somewhat larger than the traditional ones yet still allow extraction in most existing chassis. Their real advantage is evident in VME64x equipped chassis, where these handles ease the insertion of boards, especially those equipped with the new higher insertion force 5 row DIN connectors or with multi-slot CPU cards - no more scarred knuckles!

One of the key features introduced within VME64x is the long awaited 2eVME protocol. This interface standard doubles peak burst transfer rates to 160 MB/sec using existing VME backplanes. However, other ideas in the pipeline have been demonstrated which promise even greater potential to address embedded computing's insatiable appetite for increased performance.

ENTER 2ESST

At the Real Time Computer show in Santa Clara, CA, Drew Berding of Arizona Digital surprised the VME community by demonstrating a VME system transferring data over a backplane at 320 MB/sec. The new VME320 backplane contained several innovations. Instead of running bus signal traces from slot 1 to 2 to 3, etc. all the way to 21, the new design runs signal traces from each slot to slot 11 (at the center) forming what can be described as a star topology. The resulting circuit, as viewed from each slot, is a small series

inductance connected to a capacitive load. In other words, a low pass filter separates any one slot from each of the other 20 slots. This type of filter minimizes high frequency reflections and produces signal transitions that are smoother and essentially monotonic. More importantly, since each signal line is not capacitively loaded by connectors and boards distributed across 21 slots, signals stabilize more rapidly.

On a fully loaded VME64 backplane, a signal can take up to 8 nsec to propagate from one end to the other. This forces bus receivers to assume worst case propagation delays while waiting for the backplane signals to settle down. On even a fully loaded VME320 backplane stable signals reach each slot with minimal signal skew. To achieve an optimal source synchronous implementation, backplane signals require this minimal signal skew as the effects of propagation delay become less critical. As a source synchronous protocol, VME320 also allows both strobe and data to be simultaneously asserted while at the same time eliminating the data acknowledge cycles. This new VMEbus protocol has been titled 2eSST. A practical goal of the standard is to ease VME-to-PCI bridge design by supporting both 2eVME and 2eSST. Approval of the standard is expected to occur by the first half of 1999 and product introductions incorporating the new protocol are likely within the same timeframe. The initial draft of the 2eSST standard is available on the VITA web site at www.vita.com.

Mr. Berding also reports that improvements to the VME320 prototype system are possible, which would allow 2eSST transfers to take place at over 500 MB/sec. Beyond that, ideas are being bounced around - some of which appear very plausible - to move the data transfer rate to even higher levels. Although nothing has been formalized, it is fully possible that improvements in VME backplane performance will continue at the historic rate allowing VME to break the Gigabyte per second barrier by the year 2003.

Local bus faces increasing demands

These VME bandwidth improvements along with increased application performance requirements will

LI AND HA VME STANDARDS

Both the LI and HA VME standards are still in draft form. The following is a brief description that summarizes a few of the salient proposals. Changes are possible and likely, especially in HA VME which is less well defined.

The Live Insertion standard proposes a number of electrical, mechanical and software extensions to insure that field service personnel can exchange cards in a live system without disrupting its operation. New bus transceivers (specified by the separate ETL standard) are required to support a pre-charged, high impedance mode that prevents bus glitches during hot plug/hot pull. In support of pre-charging, the VME64X standard created some longer mate-first/break-last pins on the 5 row P1 and P2 DIN connectors. Sensing switches in the injector/ejector handles allow a board to physically detect when it is being inserted or removed so that software may be informed. Power ramp up and ramp down can then be controlled to tolerable levels. ESD contacts are added to discharge static from a board before contact with the backplane. Also, front-panel LED's indicate whether a board is ready for insertion or extraction.

The bus grant and interrupt daisy chains can be maintained by adding removable modules that plug into the backplane near each card slot. These modules use the newly defined I/O signal on P1 to determine whether to pass a daisy chain signal through or not. The HA VME standard specifies a centralized mechanism for managing the daisy chains.

continued on page 30

demand corresponding improvements in the bandwidth of the local bus used on most VME cards. While 33 Mhz PCI may meet the needs of the desktop for the next few years, its 264 MB/sec transfer rate (64-bit data bus) pales somewhat next to 2eSST, especially since the VME-to-PCI bridge chip on today's boards usually shares the PCI bus with a network interface, SCSI bus interface, and other devices as well. Moving to 66 MHz PCI doubles the bus throughput to over 500 MB/sec but still depends upon the delays associated with reflective wave switching. Until some innovative solution is proposed, 66 MHz designs significantly limit the number of PCI busloads. This makes 66 MHz PCI impractical for use in a backplane with many slots, or as an embedded bus with a large number of devices. There clearly is enough industry momentum behind PCI to ensure that this problem is resolved, but as yet no specific standard has emerged.

33 MHz PCI serves well as a VME board backbone when connected to today's interfaces such as Ultra SCSI, 100BaseT Ethernet and USB. In the future 500+ MB/sec backplanes will have adequate throughput to support emerging I/O technologies such as Gigabit Ethernet, ATM, IEEE 1394 (FireWire), and Fibre Channel. If VME steps up to the plate at these levels of performance, and in the unlikely event that PCI does not also substantially improve in speed, we could see a resurgent interest in smart VME I/O cards. That aside, the widespread use of multiprocessing architectures in VME applications will continue to place increasing demands on the primary system bus, and VME will evolve to meet these demands.

MEZZANINE POPULARITY GROWS

Mezzanine cards are an easy way to add support for that additional I/O interface your board is lacking. Emerging standards like Gigabit Ethernet are likely to appear first on these mezzanines. PCI Mezzanine

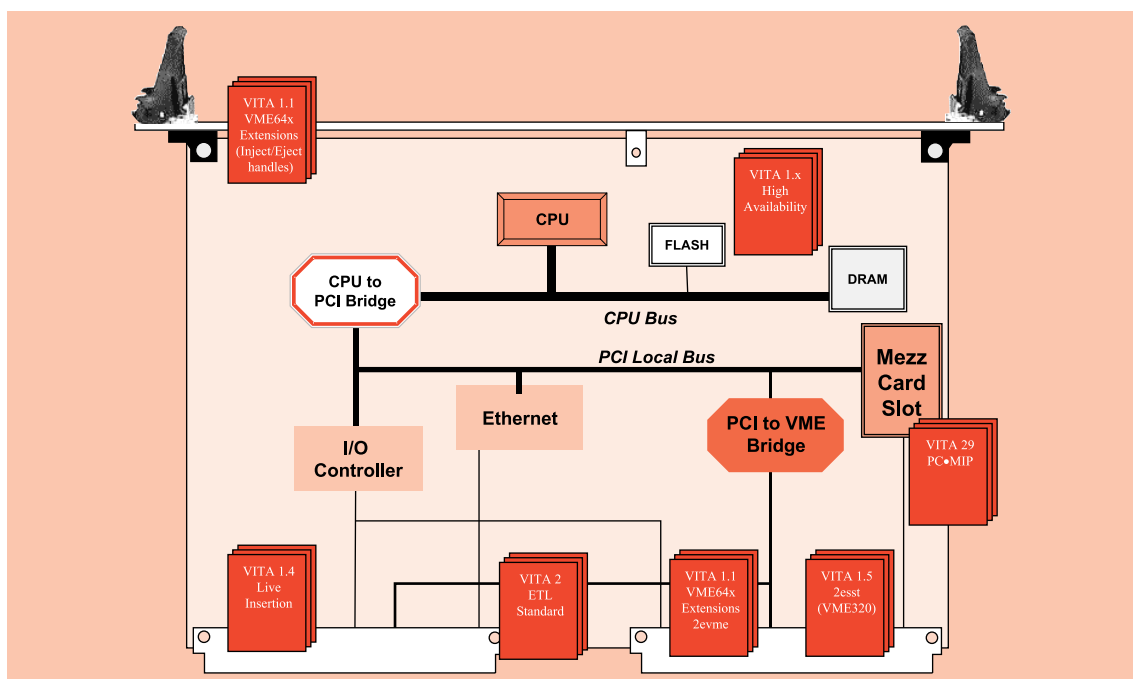


Figure 1. Summary of Key VSO Standards

Cards (PMCs) are the defacto standard on VME modules built around a local PCI bus; virtually every VME vendor offers such carriers and often a variety of PMCs as well. They are even rapidly appearing in other bus-based environments such as CompactPCI. However, the newly proposed PCoMIP mezzanine standard is sparking lively discussion with its PCI interface on a smaller than PMC form factor. This proposal might be considered a combination of the best of PMCs and IndustryPacks. It is difficult to predict how user's will respond to the advantages (and disadvantages) of this new mezzanine standard, but it is certain that PCI based mezzanines will continue well into the next millennium. The initial draft of the PCoMIP mezzanine standard is available on the VITA web site at www.vita.com.

VME MOVES TOWARDS LIVE INSERTION

One possible criticism of VME is its lack of support for both Live Insertion (hot plug/hot pull of cards) and High Availability (elimination of single points of failure, with software hooks for fault detection and recovery). Users requiring these capabilities are currently using standard VME64 backplanes with interposer cards (a board that isolates a VME card from the backplane) to implement board-level hot swap, or alternatively are using split backplanes (electrically isolated groups of slots) to provide redundancy at the chassis level. VSO and its numerous volunteer members are busy rolling these features into standards. Even pessimists are predicting Live Insertion (LI) standard approval in 1999 with High Availability (HA VME) following two years later. Beyond that, Nortel has proposed TEMPE (Telecom Enterprise Multimedia Platform and Environment), a standard to extend High Availability VME with enhancements targeted at the telecommunications industry. Computer telephony could also use the P2 and P0 connector for H100 applications. The P0 connector, located on VME boards between P1 and P2, is defined in the VME64x standard.

With the changes described above, VME should remain a viable technology into and beyond the year 2003. One thing is certain, the VME system in the year 2003 will have advanced a long way beyond the VME64 systems in common use today. The key to VME's survival has been and will continue to be based in its ability to evolve. Competing technologies continue to force developers to improve VME so that it retains its relevance and vitality. The level of standards activity in the VME space supports the view that VME will remain a premier platform in high-end embedded computing. ■

Eran Strod serves as MCG's product and technology consultant for companies using embedded computers in the greater Boston area. He also publishes technical articles and speaks at seminars on topics related to embedded technology. Before joining the Computer Group in 1996, Eran was an Applications Engineering Manager at Alacron, Inc. He also spent several years as a software engineer at Data general, CSPI and Shemer Computers. Eran holds

Ad Embedded Systems 1

continued from page 28

The 2eVME and 2eSST protocols have minimal software implications at the application level other than the fact that bus transfers proceed more rapidly. Board Support Packages for operating systems must talk to a new bus interface device, but this is not uncommon when porting an operating system to a new board. Software vendors and application developers are quite familiar with this process.

Live Insertion and HA VME, on the other hand, represent a new paradigm for software development with implications that radically affect application development. For example, the LI draft standard recognizes that software may, by accident, address a device that has been hot pulled. To manage this inevitability, the LI draft introduces a Bus Error (catch) Module that detects bus errors and intervenes to complete the cycle. In a LI system, the Host (System Controller or module fulfilling that function) is interrupted and provided with address information describing the access that caused the bus error. This address information allows software to identify the device in question and takes an appropriate action. A system crash is no longer an acceptable consequence of an errant bus cycle. Software must also be able to restore the integrity of the system by releasing resources that were associated with the board that caused the problem.

When a board is inserted or removed, an interrupt is generated to the Host. In addition, a software utility polls a memory region where control registers are mapped to identify devices that have been added or removed. Software is expected to adjust application execution accordingly - no trivial task. In the case of an insertion, the Host identifies the new board, dynamically downloads the appropriate executable code to it, and configures software drivers to communicate with the new device. A high priority process informs device drivers if a board is no longer responding. Software device drivers must be robust enough to recover from accessing devices no longer present or responding. When self-diagnosis indicates that a hardware failure has occurred, the user is informed so that appropriate action may be taken.

The public version of the HA VME standard (Draft 0.2) is still under development, drawing upon high availability systems that have been available in the computer industry for some time. Some items in the HA standard build upon topics addressed in LI. To achieve high availability, certain aspects of system operation are likely to be addressed:

- redundant system controllers
- the concept of a "Monarch" or control board that performs HA administration
- a utility communications channel that gives the Monarch access to boards which are isolated from the bus via ETL high impedance mode.
- extensive board level diagnostics, which run on demand, to provide meaningful and well defined error codes
- an additional mechanism to get around the bus grant and interrupt daisy chains
- pre-defined fault characterizations
- methods to monitor the system for faults
- guidelines for reporting faults to the Monarch
- methods to isolate faulty components from the bus
- system alarms

The software implications are still being defined, but parallels can be drawn from existing High Availability solutions. Software will need to spawn redundant processes with the ability to re-map any logical resource to a different physical device at any time. It will also need to periodically check the system for failures. Upon detecting an anomaly, processes that use the resources in question will pause, transfer state to another resource and continue service. Unlike fault tolerant systems costing much more, some data loss is tolerated in the High Availability scenario. The requirement to handle bus errors mandates that operating systems provide MMU support. This in itself is a paradigm change for some operating system vendors. The need for robust pre-emption points in each operation will fundamentally transform the way in which developers conceive and implement applications.

Overall, the roadmap to both higher bandwidth on the VMEbus backplane and to Live Insertion is pretty well defined. High Availability, on the other hand, still requires a great deal of work. VME has the potential to lead other technologies in terms of data throughput. Development also continues on the standardization of High Availability for VME. By 2003, VME may indeed incorporate High Availability functionality into the standard, but a significant enhancement in the performance roadmap for CompactPCI would even the score. Who will win? Clearly the users will. Healthy competition from emerging technologies will motivate VME vendors to deliver higher performance products with an optimum mix of features.

Bachelor Degrees in both Computer Science and Psychology from the University of California at Santa Cruz.

Greg Novak is responsible for definition, program management and marketing of Motorola's VME and CompactPCI CPU boards. He has published a number of papers and articles on embedded architectures and markets. Greg joined Motorola in 1977 with the Microsystems Division, the team that even-

tually evolved into the Computer Group. He has held various positions in engineering and marketing. His industry experience includes hardware and firmware development at Honeywell Information Systems. Greg holds a Bachelor's Degree in Electrical Engineering and Computer Science, a Master's Degree in Electrical Engineering and also a Master's Degree in Business Administration.