VME320 Enables High Performance COTS

The Need:

Military systems have demanding needs. In additional to the traditional needs for ruggedness and ability to operate in hostile environments, the ever increasing importance of electronic warfare and electronic countermeasures requires staggering performance. Features such as multiple target acquistion and tracking, require faster speeds than legacy VME systems. Although CPUs and memory have increased dramatically in performance, the overall system performance has not increased proportionally because the backplane (which is the system backbone) has traditionally limited the system performance gains.

Figure 1 shows the performance of legacy VME backplanes up until fairly recent. This is not a very dramatic chart since performance stayed constant at 40 MB/s (peak theoretical) until 1991 when Performance Technologies announced their VME64 (a very clever way to utilize inactive address lines as additional data lines in order to double the data width and thus double the backplane bandwidth).

Today, the serial interconnect technologies are getting a lot of attention, particularly Gigabit Ethernet, Infiniband, and RapidIO. Infiniband has been hailed as an excellent solution for high-speed server clusters. But, will that technology, perhaps via hybrids, cross over to military applications? Well, current specifications don't show much scalability and it is more of a "forklift" upgrade---that is, remove the old system (hardware, software and architecture) and completely replace it with a new one. More importantly, Infiniband will not use backwards compatible components. Gigabit Ethernet, Infiniband, and RapidIO will all require the drivers and software to be re-written. Support for long life cycles and compatibility with legacy hardware and software are mandatory in military applications. There will likely be a vicious marketplace battle to be the next generation system. Gigabit Ethernet has the name recognition and the largest installed base, Infiniband has the heavyweights behind them, but it is still in question whether RapidIO will have the marketing muscle to gain the size and compatibility that the military demands.

Then there is StarFabric, PLX, and others. StarFabric looks to be an excellent development for telecom applications, particularly for scalable systems in bay stations, working up to central offices. With features such as high availability and quality of service, utilizing backwards compatible PCI technology, it has far-reaching potential. But StarFabric and most of these technologies are based on PCI, so they may not be able to handle real-time interrupt and tightly coupled multi-processing. These technologies would likely have to adapt their technology and create hybrids for the military market unless the real-time reliability issues with PCI-based systems are resolved. Indeterminantcy, for example, is one issue. With PCI, your can't predict with certainty the latency of the system. In target acquisition, that uncertainty can be the difference in making the kill or being killed. Furthermore, Intel with its incredible marketing muscle, is determined to eliminate PCI in the desktop which will jeopardize component availability for embedded systems.

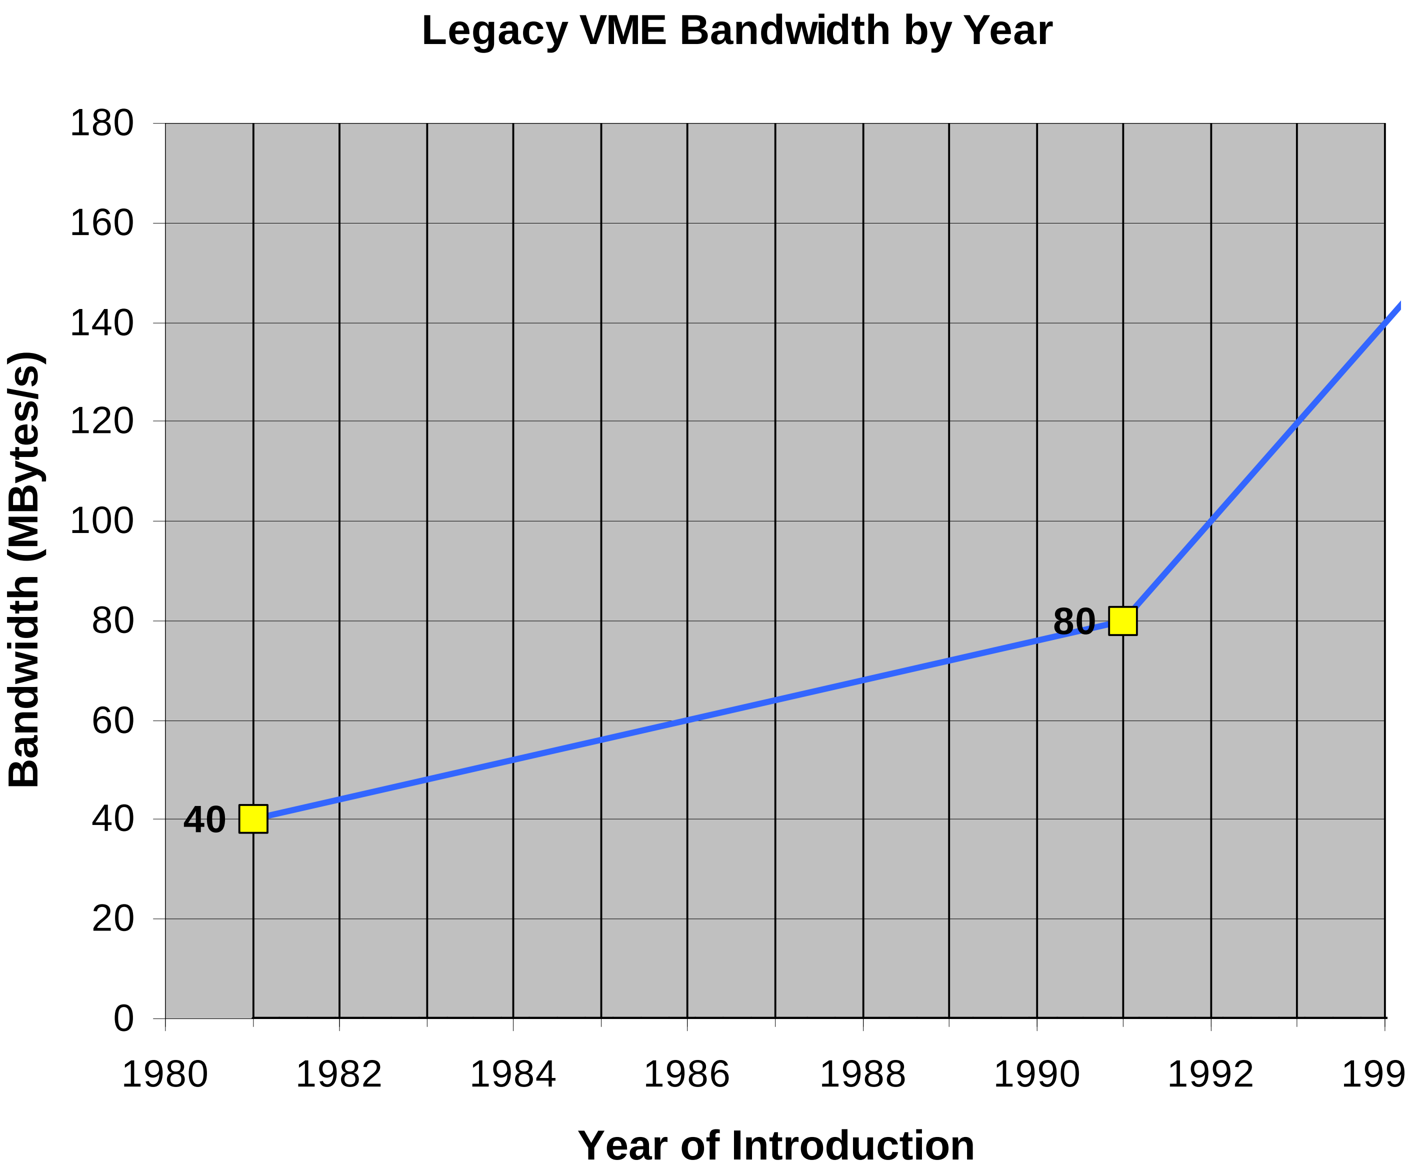

Most deployed systems are presently limited to 80 MB/s peak theoretical although much more performance is needed.

Additional Requirements:

Increasing system performance is not enough. Support for long life cycles and compatibility with legacy hardware and software are also mandatory. To further add to the difficulty of finding an adequate solution, there is a new emphasis on lower cost resulting from the COTS[1] initiative.

[1] COTS = "Commercial Off The Shelf" technology.

The Solution:

Recent VME Bandwidth Improvements by Year

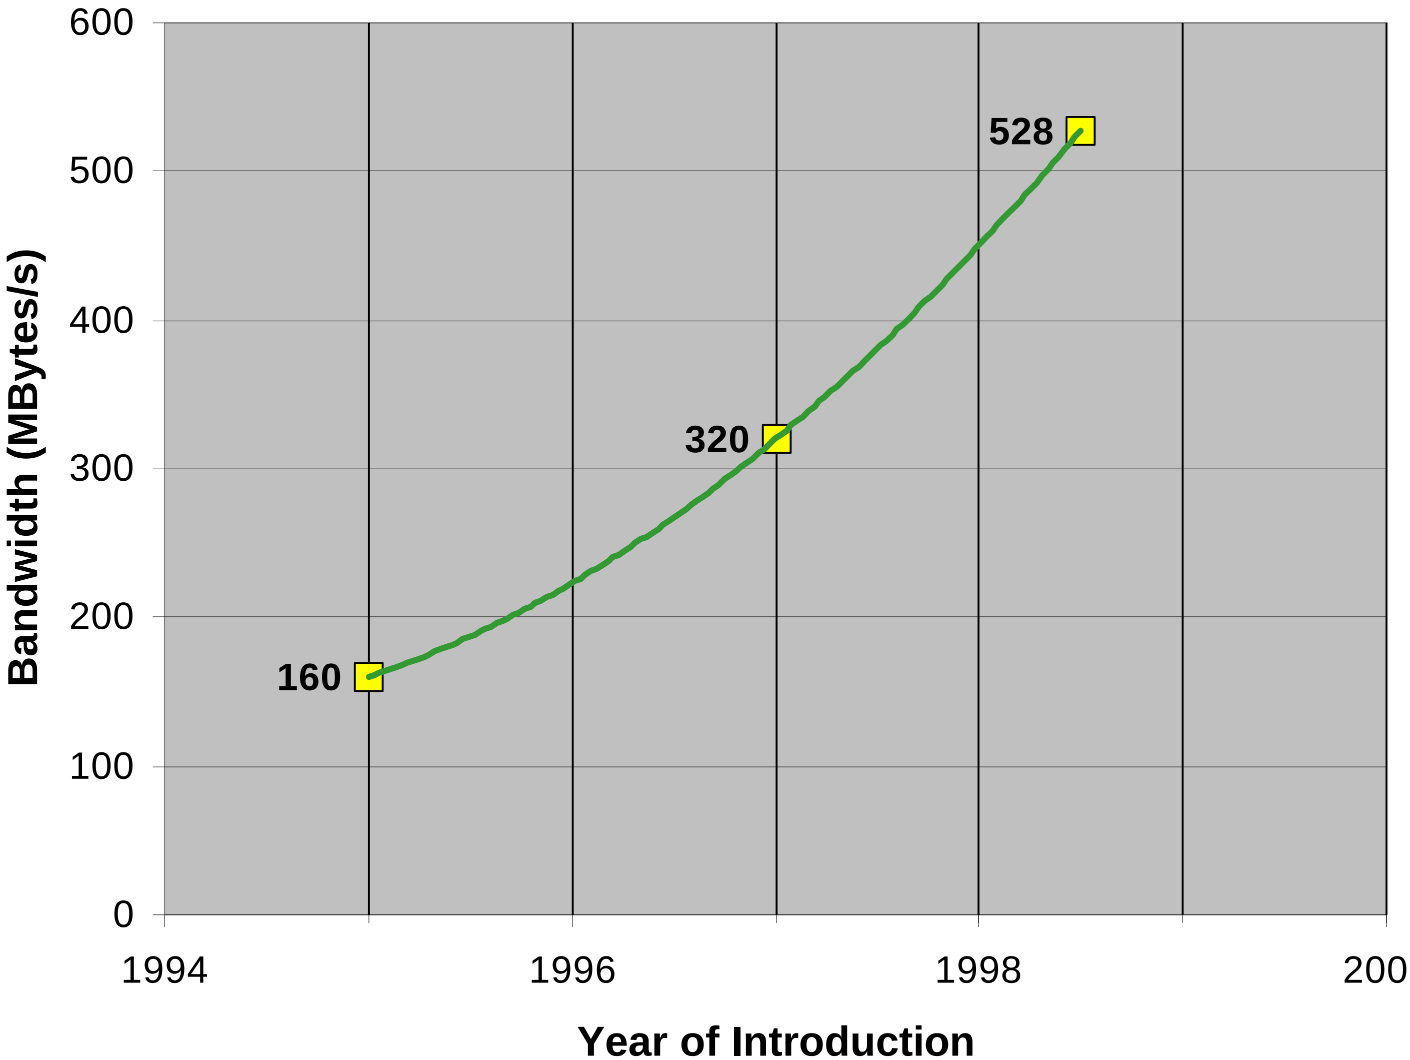

Fortunately, there is a solution that meets all of these needs---the VME320 technology from Arizona Digital in cooperation with Bustronic. It builds upon the solid VME systems that have long met military needs, it is 100% hardware and software compatible with every VME card ever shipped, it costs no more to implement and it offers spectacular performance. It is considered the logical growth path.

Chart 2 shows how VME320 has recently increased the performance of VME systems.

The first point on the chart actually refers to VME64x---an innovative idea from Force Computers that utilizes both edges of the Strobe signal to qualify data and increase the performance of VME64 systems. Prior to that time, only the negative going edge of Strobe was used to qualify data with the result that the bandwidth required by Strobe was double the bandwidth required by the data. This had the unfortunate consequence of limiting the data bandwidth to half of what would otherwise have been achievable. However, by using both edges of the Strobe, VME64x allowed the data bus to run at the same bandwidth as the Strobe which then doubled the theoretical peak data bandwidth. Unfortunately, although VME64x has seen some acceptance, by the time its specification was approved, VME320 had eclipsed it.

The Technology:

So how does the VME320 technology work? The higher VME performance has been achieved by using a different approach to backplane design.

As is well known, conventional backplanes have a major problem due to reflections and ringing because the backplane wiring acts as heavily loaded, slow transmission lines. Designers have attempted (unsuccessfully) to solve this problem by terminating the transmission lines or using more powerful drivers---both approaches are incompatible with the VME specifications and with legacy cards.

A much better philosophy when you have a problem, instead of trying to mitigate the problem, eliminate the problem. Instead of trying to solve the effects of the transmission lines, eliminate the transmission lines entirely. Therefore, instead of stitching the wiring from slot to slot as in a conventional backplane, the VME320 uses a novel star wiring scheme where all lines are connected radially to the common center slot[2]. The effect is that the backplane behaves as a lumped capacitance located at the center slot. Since each driver is driving a simple capacitance rather than a transmission line, this results in clean monotonic waveforms suitable for very high-speed operation using the 2eSST (Two Edge Source Synchronous Transfer) protocol. Furthermore, it maintains 100% compatibility with every VME card ever shipped.

How Fast Can It Go?

This original simple passive VME320 technology can be utilized up to about 640 MB/s where the lumped inductance L of the longest trace and the lumped capacitance C of the backplane begin to act like an LC low-pass filter. The resultant delay and limited rise-time prevent going faster unless something is done. However, recent improvements to the VME320 technology[3] have allowed even more impressive performance gains while still maintaining 100% hardware/software compatibility with legacy systems. These same improvements had the unexpected fortuitous benefit of reducing the number of wiring layers required---thus reducing the cost. Therefore, this outstanding performance can be achieved for similar costs to conventional backplanes.

These recent improvements limit the lumped inductance L of individual traces while having a negligible effect on the total backplane capacitance C. This resulted in a greatly reduced LC product which allows about 1.5 times improvement in risetime, delay and bandwidth. These improved backplanes show a clear roadmap to achieving VME bandwidth up to 1.066 GB/s for

---

[2] http://www.arizonadigital.com./P5696667.pdf
[3] http://www.arizonadigital.com./P5930119.pdf

21-slot backplanes while at the same time maintaining 100% backwards compatibility with legacy VME cards. Side benefits of this technology compared to conventional backplanes are greatly reduced crosstalk, inherent hot-swappability, reduced number of layers and reduced cost. The projected performance improvements by year are shown in chart 3.

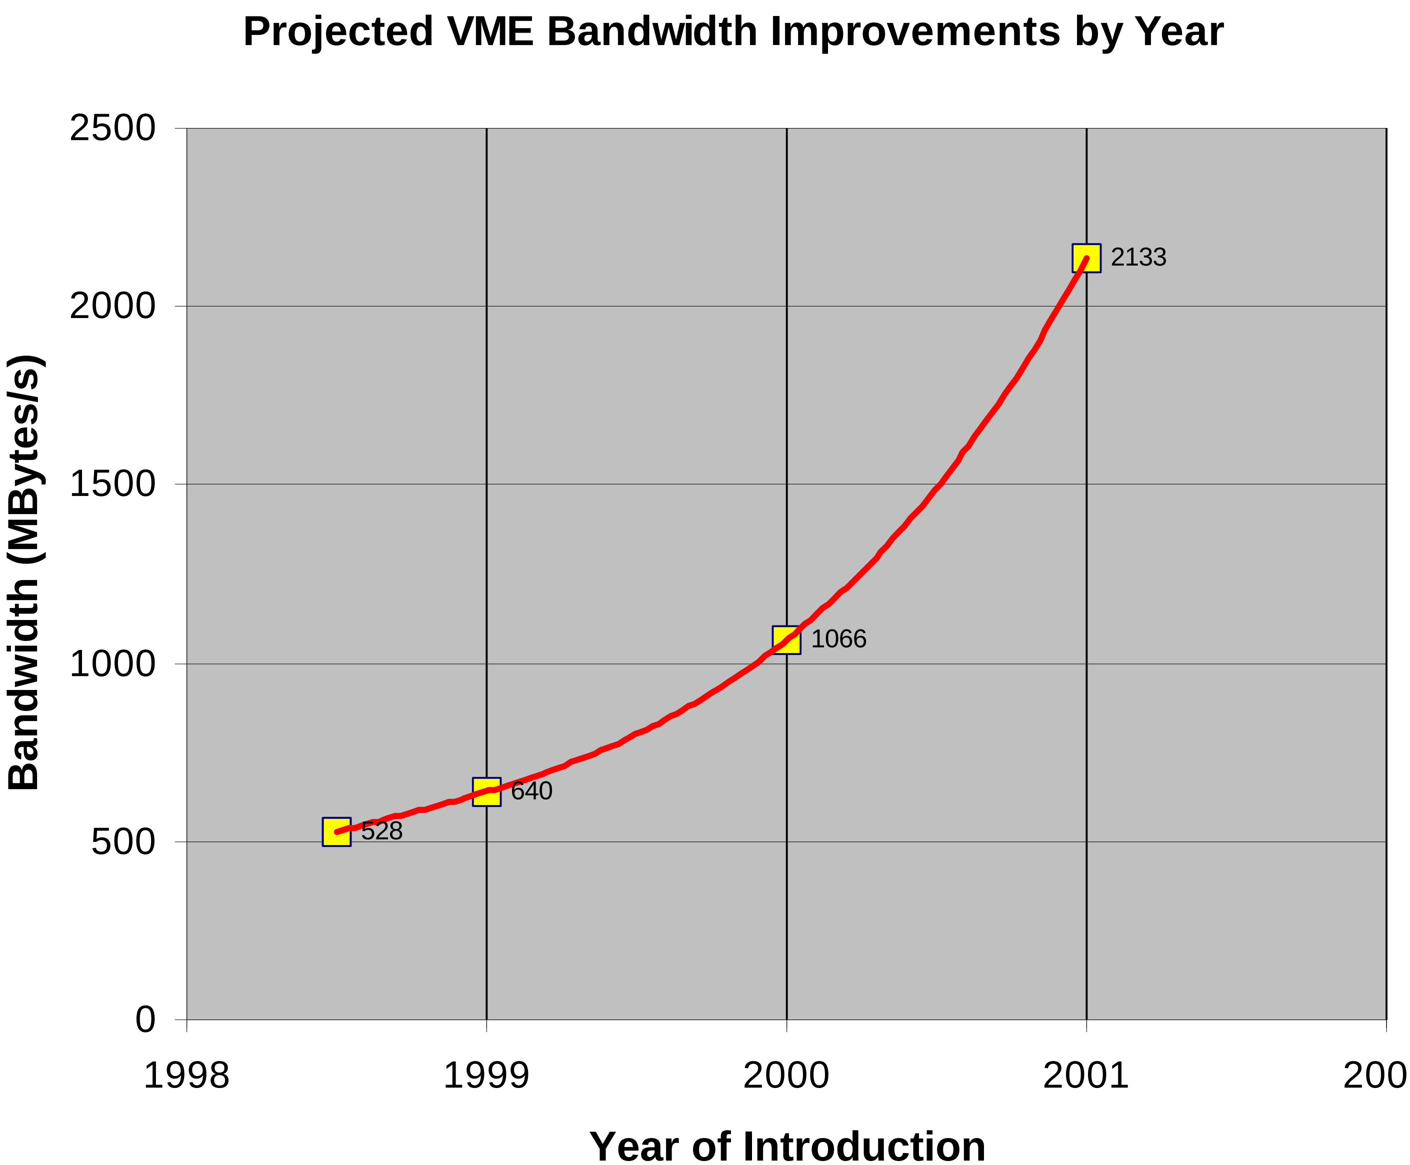

It is even more dramatic to plot all three of these past/present/future charts as a single chart showing VME Bandwidth versus year of introduction. It climbs so steeply that it even climbs if plotted on a logarithmic scale!

The Traditional VME protocol:

The traditional VME protocol is a very robust, full-handshake protocol but it is consequently slow. First of all, the Master broadcasts the Address along with an Address Modifier code (to indicate the kind of cycle desired---for example: a VME64 cycle involving the transfer of 8 data bytes in parallel).

For a Write cycle, the Master places the data on the bus, then asserts a Strobe to tell the Slave that the data is valid. Then the Slave asserts DTACK* to acknowledge that it received the data. As soon as the Master sees that the Slave has asserted DTACK*, it can remove the Strobe and as soon as the Slave sees that the Strobe has been removed, it can remove DTACK*.

A Read cycle is similar except that the Master asserts Strobe (DS1* and/or DS0*) to say it is ready to receive the data and the Slave asserts DTACK* to say that the data is valid. In a similar manner to the Write cycle, as soon as the Master sees that the Slave has asserted DTACK*, it can remove the Strobe and as soon as the Slave sees that the Strobe has been removed, it can remove DTACK*.

Thus, whether it is a Read or a Write, in order to do a single transfer of data between the Master and the Slave involves four driver delays, four receiver delays, four logic delays, four wire delays on the backplane and four settling times on the backplane. That is why legacy VME is limited to 10 MHz (best theoretical if all the silicon delays are zero!). Actual speeds in a well tuned system are around 6-8 MHz which translates to 48-64 MB/s data bandwidth for VME64.

The 2eSST Protocol:

In contrast, the 2eSST protocol is not limited by delay---whether it be wiring or silicon. As with a normal VME cycle, the Master broadcasts the Address and Address Modifier code. However, in this case, only boards that support the 2eSST protocol will respond since the Address Modifier code (20h) is a code that legacy boards do not recognize and therefore will ignore.

For a Write cycle, the Master places 8-bytes of data on the bus and sends a Strobe (DS0*) to indicate that the data is valid. However, the Master does not wait for an acknowledgement from the Slave that the data was received. The Master places the next 8-bytes of data on the bus and toggles the Strobe to the opposite state to indicate that new data is valid. And so on.

For a Read cycle, the Slave places the data on the bus and sends a Strobe (DTACK*) to indicate that the data is valid. However, it does not wait for an acknowledgement from the Master that the data was received. The Slave places the next data on the bus and toggles the Strobe to the opposite state to indicate that new data is valid. And so on.

Thus, all of the delays (data drivers, Strobe drivers, logic or wire) do not affect the bandwidth for the 2eSST protocol. The bandwidth is only affected by the variations in delays which contribute to timing skew or uncertainty. Since the data paths and the Strobe paths are similar (although not identical) most of the delays cancel out.

The 2eSST protocol includes cycle terminations with error indications in case of data overruns or other error conditions. Therefore the data is guaranteed to be delivered.

Difficulties:

Up to now, there have been two things that have delayed designs using VME320. These have now been eliminated:

1) The lack of an approved specification for the 2eSST protocol.

2) The related lack of support silicon (remember that Tundra decided not to upgrade their Universe II chip to support 64-bit/66-MHz PCI or the 2eSST protocol).

Concerning the first issue, the VITA 1.5 specification[4] was approved on December 1, 1999.

Concerning the second issue, there are several sources for support silicon coming on line:

a) General Micro Systems has announced an ASIC 64-bit/66-MHz support chip (the Omni chip) that will be commercially available.

a) IBM France has developed a PowerPC-to-VME chip for CETIA which presumably will be made commercially available.

b) Silicore is developing their CoreVME[TM] which is a 64-bit/66-MHz PCI-to-VME VHDL core which supports 2eSST.

c) Arizona Digital has also had extensive negotiations with Xilinx to have Xilinx provide a VME solution based on their Virtex XVC300 FPGA chip. This will use the Xilinx soft 64-bit/66-MHz PCI core along with Arizona Digital's VME interface.

Arizona Digital, with support from Bustronic, is providing the complete turn-key reference design for the buffered VME interface including schematics, BOMs, Gerber files, drivers, APIs, and documentation. This is a comprehensive solution that does the total interface from the on-board PCI local bus to the VME bus and totally relieves the burden on the designer.

[4] http://www.vita.com/vso/draftstd/2esstd21.pdf

Applications:

VME320 is absolutely positioned for applications in the COTS market with its rugged, passive backplanes and 100% backwards hardware/software compatibility with legacy cards. It can achieve 528 MB/s at 66 MHz with up to 21 slots (or over 1 GB/s at 133 MHz with 10 slots). This is 12 to 25 times the bandwidth of the original VME specification! Since most COTS systems are 10 slots or less, this gives a huge performance upgrade potential to any system using VME320 backplanes without incurring any additional cost.

DY-4's Application:

Recognizing this fact, DY-4 in Canada is using the VME320 for several products since that gives them an easy upgrade path and it doesn't cost them any more to do it[5]. Much of their business and a substantial part of their profit comes from upgrades. They are aggressively advertising 15-year designs. In fact, they are upgrading the Abrams M1 battle tank with these designs.

Themis's Application:

Using the same logic, Themis Computer is planning on supporting VME320 backplanes in the modular I/O bay of their Rugged Enterprise Server (RES). The Themis RES is a very powerful 14-way UltraSPARC Symmetrical Multiprocessing server with over 28 Gigabytes of main memory, 140 GigaBytes of high performance RAID-5 storage, I/O controllers, high performance graphics engines, and Digital Signal Processing (DSP) engines in one environmentally hardened chassis.

Themis initially plans to utilize only the VME64x features of the new backplanes in the RES I/O bay. However, the VME320 backplanes give Themis' customers a dramatic backplane performance upgrade path, allowing them to simply plug in new cards that support the 2eSST protocol in addition to the standard VME protocol. The existing cards and the new cards coexist together with no interoperability problems.

General Micro Systems:

General Micro Systems in Q1, 2001, will be shipping a Power PC 74xx Single Board Computer based on the OmniVME chip. This board, "Atlantis", is the first SBC to fully utilize the PCI 66Mhz / 64-Bit bandwidth by providing VME transfers of up 533MB/sec, using the VME-320 backplane. To allow OEM users to implement the GMS OmniVME, GMS will be shipping a design reference kit, which includes schematics, BOM, design layout consideration, and software drivers for VxWorks (Windows 2000 /NT in the future). The OmniVME device will be available both in Source level for OEM use or as a device from GMS.

Conclusion:

The VME320 is the acknowledged upgrade path for COTS systems because of ruggedness, performance, cost and compatibility with legacy hardware/software.

---

[5] It is a lot easier to put a high-performance VME320 backplane into a system initially than to try to rip out and replace a conventional one in the field.

Drew Berding
Arizona Digital, Inc.
480-443-3069
www.arizonadigital.com

Fred Hirsch
Bustronic
510-490-7388
www.bustronic.com

Benjamin Sharfi
General Micro Systems
909-980-4863
www.gms4vme.com