

# Is CompactPCI a Replacement for VME?

*CompactPCI is an adaptation of the PCI electrical layer to Eurocard mechanicals. As such it offers the rugged mechanical and well disciplined thermal packaging which has historically been associated with VME. The question naturally arises as to whether CompactPCI will replace VME. This article analyzes the two architectures and the market forces at work in an effort to answer that question.*

Is CompactPCI a replacement for VME? Like most simple questions, this one has a simple answer. In fact it has two depending on who you ask. If the answer comes from a vendor of CompactPCI hardware with roots in the industrial PC market, the answer is likely to be an unqualified, "Yes". If the answer comes from a spokesman for the VME community, the answer is likely to be an emphatic, "No!"

The robust mechanical packaging afforded by CompactPCI is comparable with that of VME systems as shown in the following photograph of a Compaq product offering.




Figure 1. Compaq 233i-8 CompactPCI System

The similarity of CompactPCI and VME single board computers are evident from figure 2 and 3.

In the light of such obvious similarities, the objective observer would probably ask more complex questions, however, such as, "Now that you can get PCI processors and peripherals in Eurocard package, which components of the traditional VME system will migrate to PCI and will remain on the VME bus? How is VME system architecture likely to evolve? Was there any thought of evolution in the original architectural definition of VME or PCI?"

These are questions, which recognize the possibility of hybrid systems, incorporating both CompactPCI and VME components. The answers to these questions are less categorical, and hinge on a number of factors:

- the raw performance of the busses,
- the ways in which their architectures suit them for

particular system configurations,

- and how each leverages the prevailing trends in the market.

In this discussion we'll take a look at these factors in an effort to answer the more complex, and more interesting questions.

## PERFORMANCE COMPARISON

### VME

VME began as a unified memory/IO bus for the MC68000 (clearly inspired by Digital's Unibus) with a 16-bit data path and 24 bit addressing. It was an adaptation of Motorola's Versabus for the MC68000 microprocessor family (which was clearly inspired by the PDP-11), that made use of what was then the new DIN connector and Eurocard mechanicals. It was subsequently standardized in the US as IEEE 1014-1987 and in Europe as IEC 821.

As standardized by the IEEE in 1987, the VME data path was asynchronous, with separate 32-bit address and 32 bit data paths, and supported the following transactions:

- Single Transfer Read/Write, 8, 16, 32 bits
- Single Transfer Non-aligned 16, 24 bits
- Block Transfer Read/Write, (2-256 transfers of 8, 16, 32 bits each)
- Read-Modify-Write, 8, 16, 32 bits
- Interrupt Acknowledge (IACK) cycle

A claim of 40 Megabytes/second burst block transfer rate has been traditional for the basic VME bus, as described above, with a somewhat lower sustained rate due to bus arbitration and addressing overhead.

The most significant enhancement provided by VME64 was the introduction of a multiplexed 64 bit block data transfer protocol which uses the 32 A and 32 D lines together as a 64 bit address bus in the addressing phase, and as a 64 bit data path during the data phases, of a bus transaction. The result is a doubling of the burst data transfer rate to 80 Megabytes per second.

The VME64 Extensions standard (VME64x for short), recently recognized as ANSI/VITA 1.1, includes a new two-edged asynchronous block transfer signaling protocol, known as 2eVME, with nominal 160 Megabyte/sec performance

The 2eVME signaling protocol was originally proposed




Figure 2. Compaq 233I, a Pentium MMX based CompactPCI SBC

by Motorola to double the theoretical throughput of the D64 MBLT protocol by using both edges of the Data Strobe and the Data Transfer Acknowledge to transfer data. It supports a variety of transactions, including a slave-terminated protocol of particular interest to the high-energy physics community.

The most recent performance enhancements proposed for the VMEbus are embodied in a two edge source synchronous signaling protocol, 2eSST for short, which claims to support burst transfer rates up to 320 Megabytes/second. This protocol is an extension of the asynchronous 2eVME protocol defined in the VME64 Extensions.

At the same time a good deal of effort has been invested in proprietary backplane technologies capable of providing the electrical environment required to support these protocols.

### **PCI**

The PCI bus is a synchronous multiplexed structure reminiscent of Multibus II that may be operated 32 bits or 64 bits wide with full interoperability between devices supporting either width on bus segments of either width. The theoretical burst bandwidth is 133.33 Mbytes/sec for 32 bit transactions and 266.67 Mbytes/sec when operating 64 bits wide. These are theoretical maximum limits, which compare to the 40 and 80 Mbytes/sec figures often quoted respectively for the VME 32 bit BLT and 64 bit MBLT signaling protocols.

Because PCI does not have a source synchronous protocol, there's no way to make a fair comparison with 2eSST. The only provision made for performance scalability of the PCI bus is the specification of operation at 66 MHz. At 66 MHz, a 64-bit PCI bus operates at a theoretical burst rate of 533.33 Mbytes/sec, but this can only be supported on a bus segment of four electrical loads.

### **CompactPCI**

CompactPCI is a synthesis of the electrical interface definition and signaling protocols put forth by the PCI SIG in the PCI Local Bus specification, and traditional Eurocard packaging as modified in the most recent updates to the IEEE 1101 family of mechanical standards. But this technology enhances the functionality




Figure 3. AlphaVME 5/480, an Alpha 21164A VME SBC

of PCI in ways, which are far more important than mechanical robustness, convenient rackmount packaging, or improved ESD and EMC characteristics. While the mechanical enhancements are important, the real added value of CompactPCI is in increased slot count without bridging and in the potential for hot swap.

The CompactPCI specification has as its foundation an extensive simulation study, funded by a consortium of PICMG executive members, which established the basic design rules for CompactPCI boards and backplanes. The simulation study established the feasibility of eight slot backplanes without the necessity of bridging, which is almost twice the number which conventional PCI design rules allow using the edge connector identified in the base PCI specification.

These wider backplanes, with twice the number of primary slots, allow more IO devices to have low latency high throughput access to host memory while amortizing part of the incremental cost of the 2mm hard metric connector system through the reduced use of bridging.

While CompactPCI is no faster in theory than any other PCI implementation, it does afford greater access to the aggregate bandwidth of a bus segment by providing more connection points per bus segment.

### **Beyond PCI, the impact of NGIO and PCI-X**

Much has been written over the past year about hints from Intel that it is moving away from PCI as the IO port on its chip sets in favor of a serial interconnect. This so-called Next Generation IO strategy has led to comments from industry analysts that PCI is becoming obsolete, and to doubts about the longevity of CompactPCI.

At the same time, a PC industry group lead by Compaq has formulated a set of backward compatible architectural extensions to PCI under the name PCI-X. These extensions include a timing budget for register transfers on the PCI bus at 133 MHz clock rates, and a set of protocol enhancements to facilitate the use of block transfers on the bus. These extensions support data transfer rates in excess of 1 Gigabyte per second, which is another factor of two faster than 64 bit 66 MHz transfers.

With or without further enhancement, the PCI Local Bus as defined in Revision 2.1 of the PCI SIG specification, is architected for data path and clock rate scalability which have not yet been exploited. Claims that the PCI Bus is dead seem very premature for a technology, which has barely begun to deliver its full potential.

## Conclusions based on performance considerations

The survival of VME over the past seventeen years is a tribute to its ability to adapt to increasing bandwidth requirements. It is also a tribute to the ingenuity to those who have contributed to the evolution, for the VME bus was not crafted for evolution in its beginnings.

The PCI bus by contrast was architected for scalability from its beginnings, with a 64-bit data path and the potential for double speed clocking, to data rates of over a half gigabyte per second. As PCI becomes the dominant IO interconnect, displacing ISA and EISA, it is still rare to see a 64-bit option and far rarer still to see a 66 MHz bus segment. The scalability of the PCI bus is yet to exploit.

It must be acknowledged that PCI has a theoretical advantage over VME. Using 32 bit data transfers, PCI offers 133.3 Mbytes/sec burst throughput compared to the commonly accepted theoretical maximum of 40 for VME. Using 64 bit data transfers the throughputs of both technologies are theoretically doubled. PCI has a theoretical advantage of over three to one at both data path widths.

Even the 2eVME protocol, which has a performance target of 160 Mbytes/sec, reaches only about 60% of the theoretical throughput of 64 bit PCI. If 2eSST with its theoretical maximum of 320 Mbytes/sec is to be compared to PCI at all, it is most appropriately compared with 66 MHz PCI with a theoretical bandwidth of 533.3 Mbytes/sec.

The time seems to be at hand to recognize VME as a legacy IO bus, bearing the same kind of relationship to PCI as the ISA and EISA busses in PC desktop and server systems. There may well be a place for enhanced signaling protocols such as 2eVME and 2eSST in legacy systems, but such protocols will almost certainly be used to more fully exploit the full bandwidth of the PCI IO systems to which they are bridged.

## SYSTEM LEVEL ARCHITECTURAL FEATURES

### Processor Subsystem Architecture

When we consider the system configurations in which VME is likely to be deployed, we recognize that it will almost always be connected to host memory through a PCI bus! PCI was invented to be an architecturally neutral local bus spanning multiple generations of processor technology, and it has succeeded. Whether the central processor is a Pentium, or a Pentium Pro, or a PowerPC, or an Alpha the bridge to the IO subsystem is PCI. The functional block diagram in figure 4 for Compaq's AlphaVME 5/480 VME single board computer illustrates this point.

Unless the system designer is willing to invest in the design of a bridge direct from the PCI bus to the host memory interconnect, and to redesign it for each architecture at each new generation, no other IO bus is directly connected to memory. And unless an IO bridge is connected directly to host memory, its performance will never exceed that of the IO bus from which it is bridged.

So in order to realize the full potential of VME, it would have to be bridged to a 66 MHz PCI segment. But if a 66 MHz PCI segment is to be created, why not use it directly and fully exploit its performance?

These are the questions, which an objective system designer will ask in evaluating the competing architec-

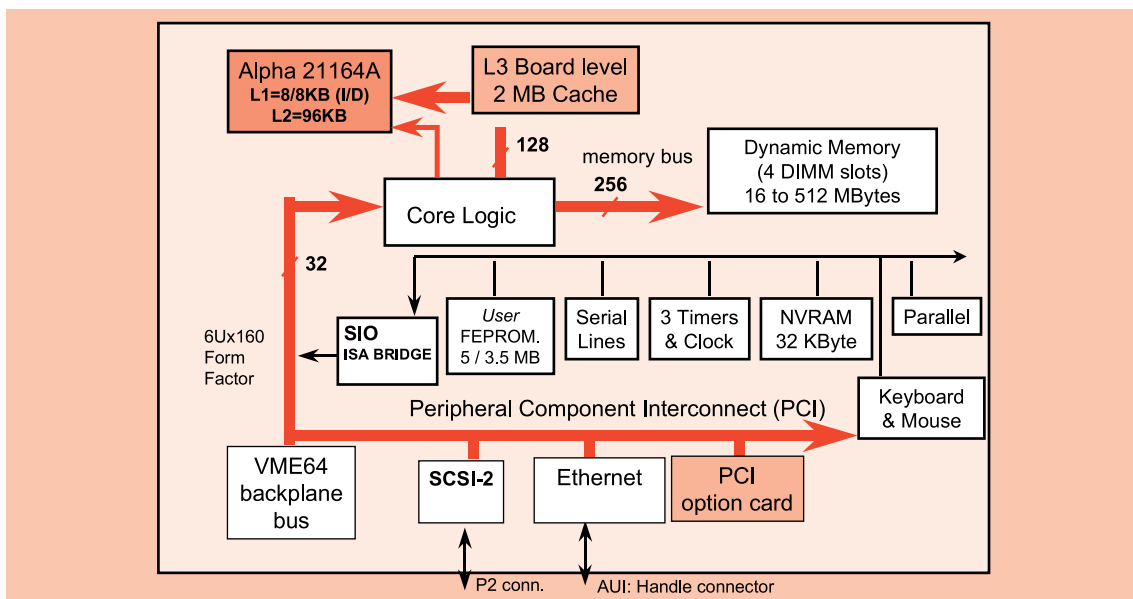


Figure 4. AlphaVME 5/480 VME SBC

## VME & COMPACTPCI

tures, and the answers seem obvious. Moreover PCI and VME are already complementary architectures, as evidenced by the chip level PCI to VME bridges which are now coming to market. CompactPCI and VME share a common mechanical specification, and the potential for creating a hybrid package bridging the two technologies has already been recognized in the form of a product offering from one of the traditional VME houses.

Conceptually a hybrid system might look something like figure 5 from the front:




Figure 5. Hybrid system from the front.

In this concept the CPU card is the double width module just left of center in the card cage. The seven slots to its left are CompactPCI while the twelve to the right are VME. The CPU module would include a PCI VME bridge electrically identical to the one shown in the block diagram. The VME bus could be pinned out to the CompactPCI J4 and J5 connector blocks as defined in a recently adopted PICMG specification, and routed to the appropriate pins of the VME connectors to the right of the CPU.

### Hot Swap and High Availability

The VME community has for some time recognized the need to improve the maintainability of platforms, and has responded with a number of schemes for supporting live insertion. A recommended practice for Board Level Live Insertion was adopted by VSO in 1995. Since then there has been continuing interest in using the pin sequencing feature of the five row VME64x connector to provide a lower cost approach to live insertion. Two draft documents, one attempting to define a High Availability VME architecture and the other focussed on the Live Insertion problem, are

available. Only the latter document is active.

In the CompactPCI community, by contrast, there has been a very aggressive effort to detail an approach to Hot Swap of modules. The 2mm hard metric connector system used for CompactPCI, which is already offered as standard product by at least three vendors, also provides sequenced pins. Pin sequencing is a necessary (though certainly far from sufficient) enabling technology for hot swap.

The CompactPCI Hot Swap specification was approved in July 1998, and Hot Swap ready backplanes are already beginning to appear on the market. Hot Swap PCI interfaces are also available, the first to market being the Intel (formerly Digital) 21554 "Drawbridge" chip.

It's interesting to contemplate why the development of Live Insertion capability has languished in the VME community, while the CompactPCI community (with many of the same members) has undertaken and completed its work without undue delay. One factor, which helped the CompactPCI effort, was the concurrent development of a Hot Plug Architecture for active motherboards by Compaq under the auspices of the PCI SIG. The Hot Plug Specification defines a hardware solution, which has very little applicability to passive backplane environments, but the Hot Plug software architecture is readily adaptable.

### Multicomputing and IO processing

One of the continuing criticisms of CompactPCI leveled by the VME community is its alleged lack of support for multicomputing configurations. VME single board computers are architected so that they can function as the system controller in slot 1, or in any other slot in the system with system controller functions disabled.

There are no "big" reasons why PCI can't support multiple processors. PCI processors can be both master and slave on the bus, just like VME processors. PCI processors can open up memory windows onto the bus just like VME processors.

PCI and VME system controllers are similar in the sense that both provide the bus arbiter for the bus segment, and the primary interrupt handler. So a VME processor needs to be configured differently as a co-processor with its arbiter disabled at a minimum. Interrupt handling is a bit more complicated, as we will discuss shortly.

The "small" differences between VME and PCI systems turn out to be substantial, however. Support of Plug and Play is one of them. Although the VME64 and VME64x standards include a configuration space, which could be used to support Plug and Play, it's optional and has not been widely implemented.

In PCI systems configuration space is mandatory on option cards, and the fact that most CPU core logic doesn't present a configuration space precludes use on an option card. A bridge is needed in order to present a configuration space on the bus, but traditional bridges are assumed to have more PCI devices on the other side. So all that can be found looking through the bridge on most processor boards are the local PCI

peripherals. Even with the arbitration and clock distribution circuits disabled, a bridged CPU card would look no different than a dumb multifunction peripheral to Plug and Play software when connected to an option slot.

The Intel (formerly Digital) 21554 Drawbridge solves this problem by presenting the devices behind it as a subsystem, operating in a separate clock domain and with a memory map independent of the rest of the system.

In CompactPCI systems the system slot pinout differs from the option slot. Rather than having bussed BR and BG lines on the same pins on each slot as in VME systems, it has radial REQ/GNT pairs from the system slot to each option slot. And unlike the asynchronous VME bus, the synchronous PCI has clocks running from the system slot to each option.

So it's pretty hard to build a CPU which has the bridging characteristics desired on a system slot (transparent with primary bus on board and secondary bus facing the backplane) and those of an option slot (non-transparent with primary side on the backplane and secondary facing the onboard peripherals) not to mention switchable clock distribution and arbitration pinouts.

And then there's the matter of interrupt processing. The good news is that the interrupt pins are same on the system slot as on the option slots, so it's feasible to reconfigure them. VME still has an advantage, however. In VME systems it's possible to hardwire the interrupts so that they're distributed among the coprocessors in the system.

### **Slot count**

One of VME's advantages over CompactPCI is slot count. The mechanical and electrical constraints on a VME backplane coincide, setting a limit of 21 slots. CompactPCI by contrast supports seven option slots per bus segment, constrained by the drive capability of PCI transceivers and the need to allow the reflected wave on a PCI segment to settle within 10 nanoseconds or so.

There are, of course, techniques for increasing CompactPCI slot count. The most common is to position the System Slot in the center of the enclosure and drive bridged CompactPCI segments both left and right from there.

Bridging on the CompactPCI backplane is less common, but is at least theoretically possible. Mechanical constraints make implementation a challenge, however.

### **Conclusions based on architectural considerations**

CompactPCI has one fundamental advantage over VME. When compared at the same data path width, and the same multiplexed block transfer protocol, PCI is faster. As such it is the ideal place to connect system peripherals, and bridges to slower IO busses such as VME. In a CompactPCI system, the CompactPCI slots are the place to configure storage, network, and display peripherals.

CompactPCI has another fundamental advantage over VME in terms of the connector used. CompactPCI uses only about 20% of the available connector pins to implement a 32 bit bus, and 40% to support 64 bits without counting the grounds provided by the outer shields. Conventional VME connectors use 66% of the available pins to implement 32 or 64 bit VME. In VME64x systems this number is reduced but the CompactPCI connector system still yields more user IO pinout.

Probably the most important attribute of the CompactPCI connector system, however, is its support for three levels of pin staging, which greatly facilitates implementation of Hot Swap. The VME64x five-row connector has only two pin lengths.

VME's advantages are currently in its support for multi-computing and distributed interrupt handling.

## **MARKET TRENDS**

The appeal of CompactPCI is in the marriage of the widely accepted Eurocard packaging standard with what has become the most widely deployed IO interconnect in the history of computing. The PCI Bus was invented by Intel and the other members of the PCI SIG to provide an architecturally neutral chip level peripheral interconnect for multiple generations of CPU and core logic technology. The success of the PCI "movement" is unquestioned, with chip level peripherals being offered by a growing number of manufacturers, and core logic chip sets for multiple CPU architectures over two and three generations of architecture. With the adoption of PCI the computing industry has finally achieved commodity economies of scale for extremely sophisticated graphics, storage, and communication options.

At this stage in the history of computing there are three things, which can almost be taken for granted:

- The next processor chip in any established family will have a PCI chip set.
- The next generation of any storage or network interconnect will connect to PCI.
- The next generation of any major operating system will have drivers for a wide variety of PCI devices.

With these realities virtually assured, the opportunity to add value is focussed on customization of packaging for particular industries and applications. This includes, to be sure, the engineering of IO devices and their drivers for particular niches, but most of all the engineering of the platform.

While there is certainly some diversity in the product offered by the CompactPCI community, it is still relatively minor compared to the proliferation of incompatible PCI packaging schemes:

- "standard" edge connector cards in a variety of lengths
- Cardbus
- PC-104+
- IEEE 1386.1 PCI Mezzanine
- Small PCI
- Mini PCI

- PC-MIP
- IPCI (European alternative to cPCI using Metral connectors)
- P2CI (VSO proposal for PCI on VME P2 connector) and probably more.

This proliferation of mechanical arrangements for PCI is enabled by the commonality of the electrical interface and programming model, which it provides. It is a sign that more users, in more industries, running more applications want to take advantage of the economies of scale which PCI chip sets, peripherals, and the software that runs them, make possible. The additional expense required to adapt packaging and options for these applications are minor compared to the development of suitable base platform silicon and its supporting software.

### **Conclusions based on market trends**

So the dynamic at work for PCI in general, and CompactPCI in particular, is quite different from that which has existed in the VME marketplace over the past 17 years. If VME systems have become more and more uniform in terms of their hardware attributes, it is probably because of the numbers of processors and operating systems that have had to be supported. This has forced the VME systems integrator to focus on the porting of applications and drivers from one processor to another and OS to another.

The proliferation of processor architectures in VME form factor actually has potential for stimulating the market for CompactPCI solutions. Most of the more recent entries into the VME market are really board level PCI based systems, which include a PCI to VME Bridge. The common mechanical properties shared by the VME and CompactPCI make it easy to repartition these designs into hybrid systems. Such systems would leverage the huge investments made by silicon vendors in chip sets and peripherals, and by base OS vendors in platform and PCI peripheral support, while providing a platform for the valuable legacy represented by specialty VME boards.

### **OVERALL CONCLUSIONS - COMPACTPCI VERSUS VME**

So, IS CompactPCI a replacement for VME?

Well, it could be eventually. PCI was architected from its beginnings to have the performance scalability, which VME has gradually developed on an ad hoc basis over its 17-year lifetime. Neither architecture has been exploited to its potential, though PCI with 64 bit 33 MHz chip sets and peripherals capable of 266 Mbytes/sec is clearly ahead. The 2eVME protocols, with the potential for 160 Mbytes/sec, have only just been standardized and have yet to be deployed.

The promise of further performance enhancements to VME, using the 2eSST protocol in conjunction with proprietary backplanes and advanced transceivers, is likely to be matched and exceeded by extensions such as PCI-X or a migration to the serial IO "hose" architecture exemplified by NGIO.

In the long run PCI can be expected to provide more

raw throughput to memory than VME, due to its theoretical advantages and to the practical consideration that the PCI bus will always be "closer" to memory in the hierarchy of interconnects than VME.

In the shorter term the system components which migrate to PCI are the essential system level components which need the greatest throughput and lowest latency access to memory, mass storage subsystems and devices, high performance networks, graphics, and multimedia.

But in the short term VME has architectural advantages in terms of slot count, interrupt handling, and multiprocessing support which will not be quickly overcome. If the application requires more processing power than the central processor can provide, the VME bus provides a better environment for a solution today than CompactPCI. If more than seven IO devices requiring peer to peer communication must be supported, the VME Bus is the answer today. If there is a requirement for distributed interrupt handling, VME is the current solution.

We can already foresee the direction to multicomputing support on the PCI bus, and we can project that the ability to configure intelligent IO subsystems that come with it will gradually obviate the need for high slot count and distributed interrupt handling. We can also foresee a gradual migration of the specialty IO devices, which have historically been VME peripherals to PCI. As these migrations occur, the need for VME slots will gradually diminish and "pure" CompactPCI systems will become increasingly viable as replacements for VME systems and hybrids.

In the meantime, a meantime of uncertain duration, CompactPCI and VME continue to be complementary technologies, and either of the simple answers to our original question continues to be wrong. ■

---

*Dick Somes is Technical Director of Compaq's, formerly Digital Equipment Corporation's, OEM Business Segment. He is also the Vice President of Technology for the PCI Industrial Computer Group, PICMG, and in that capacity has jurisdiction over development of all CompactPCI related specifications. Lest his objectivity be questioned, however, his is also Compaq's representative to the VMEbus International Trade Association and its ANSI accredited Standards Development Organization, the VSO. He was chairman of VSO for 1996. He has actively participated in the development of the recently approved VME64 Extensions standard. Previously he represented Digital in the ANSI X3T9.3 (now X3T11) technical committee which is responsible for HiPPI and Fibre Channel interconnects.*

*Dick has been with Compaq/Digital for 15 years in various capacities as engineering manager and individual contributor in the Laboratory Data Products, High Availability Systems, and Networks businesses prior to joining the OEM Business in 1994. He has a Bachelor's degree in Electrical Engineering from Tufts University and a Master's degree in Electrical Engineering from the University of Pennsylvania, and has over 32 years experience in the electronics business.*