

Extending PCI Performance Beyond the Desktop



PCI buses are not only prevalent in desktop systems, they are also being reconfigured for industrial applications, embedded systems, laptops, other mobile devices, and graphics applications.

Shlomo Weiss
Tel Aviv
University

Ehud Finkelstein
Seagull
Semiconductor
Ltd.

In just a few years since its introduction, the PCI (peripheral component interconnect) bus¹ has become an industry standard, implemented in most PC systems as well as in some workstations. Designed for the desktop PC market, the PCI bus cannot be used directly in most embedded systems or in industrial applications due to packaging, reliability, or maintenance considerations.

As a result of its widespread use, various groups have developed several PCI-related standards. These standards bring PCI performance to other segments of the computer industry and enable the use of PCI silicon, software, and firmware in nondesktop systems. While the proliferation of PCI-related standards certainly attests to PCI's success, some confusion arises because it is not always clear why there are so many standards, what the differences are between them, and which standards apply to a specific application.

PCI BUS BASICS

The I/O expansion bus connects most I/O peripherals into the system. It consists of connector slots into which I/O controller cards can be inserted. A bus standard provides an interface between a computer system and I/O peripherals; it specifies control and transfer protocols that are supported, and mechanical and electrical requirements that compatible I/O cards must meet. By far the most popular of the expansion buses is the Industry Standard Architecture (ISA) bus, based on the original IBM PC/AT bus. At only 5 Mbytes/s, the ISA bus is inadequate for many of today's high-performance peripherals, however, and higher performance buses have been developed.

The PCI bus is designed to overcome many of the limitations in previous buses. It can transfer data at a peak rate of 132 Mbytes/s using the current 32-bit data path and a 33-MHz clock speed. High-end implementations featuring a 64-bit data path and 66-MHz clock may transfer data as fast as 524 Mbytes/s. (For more information see the "PCI Bus Extensions and

Options" sidebar.) PCI is a local bus, sometimes also called an intermediate local bus, to distinguish it from the CPU bus. The concept of the local bus solves the downward compatibility problem in an elegant way. The system may incorporate other buses, such as ISA, and adapters compatible with these buses. On the other hand, high-performance adapters, such as graphics or network cards, may plug directly into the PCI, as shown in Figure 1.

The PCI bus can be expanded to a large number of slots using PCI-to-PCI bridges, which connect separate, small PCI buses to form a single, unified, hierarchical PCI bus. When traffic is local to each bus, more than one bus may be active at a time. This allows load balancing, while still allowing any PCI master on any bus to access any PCI target on any other PCI bus.

All PCI-compliant cards are automatically configured for PCI bus transactions. There is no need to set up jumpers to set the card's I/O address, interrupt request number, or direct memory access (DMA) channel number. By incorporating the optional OpenBoot standard,² any device with OpenBoot firmware can boot systems containing any microprocessor, without any operating-system- or CPU-specific software support. Even without OpenBoot, it is common to see drivers for many video cards and SCSI controllers that cover multiple CPU architectures and operating systems.

The PCI bus supports advanced features such as bus snooping—to enforce cache coherency even with multiple bus masters—and a locking mechanism to support semaphores. At the same time, it also supports older standards. For example, the PCI interoperates with existing ISA cards by supporting subtractive decoding of addresses (allowing addresses not decoded by PCI cards to go to a non-PCI bus such as ISA). It supports standards such as those for VGA (video graphics array) and IDE (integrated drive electronics) controllers. To do so, the PCI allows these devices to be locked at fixed I/O addresses, which the operating system requires to boot successfully.

PCI Bus Extensions and Options

PCI 2.1 specifies several optional features intended to provide room for future growth and accommodate multiple platforms and architectures. We cover the most important options and extensions.

64-Bit Extension

A PCI device can optionally support 64-bit addressing or data transfer or both. The 64-bit address extension supports address space sizes beyond 4 Gbytes. The 64-bit data transfer extension lets a 64-bit master transfer eight bytes in each data phase to a 64-bit target. The transaction defaults to 32-bit data transfers if the selected target does not support the 64-bit extension. A 64-bit master can fit into a 32-bit card slot. Lacking the signals that

support the 64-bit data transfer capability, the targets will not sense the 64-bit transfer request, forcing the bus master to use the lower half of the bus.

66-MHz Extension

The PCI 2.1 specification includes a 66-MHz extension with the intent of providing high bandwidth for advanced graphics and video applications. It appears that the 66-MHz PCI and AGP interfaces aim at solving the same problem. However, use of the 66-MHz PCI bus is limited because 66-MHz operation requires all devices on the bus to support the 66-MHz extension. If any 33-MHz PCI device connects to a 66-MHz PCI bus, the bus defaults to 33-MHz operation. Thus, using a 66-MHz, shared PCI bus unnecessarily increases the cost of

PCI devices that do not use the extra bandwidth. This is why Intel defined AGP as a point-to-point connection, rather than a shared bus.

3.3 V, 5 V, and Universal Cards

The core of a PCI card is a PCI chipset that uses either 5 V or 3.3 V logic. Systems must support the older 5 V PCI cards. Mobile systems must operate at 3.3 V to save power, and 3.3 V is also required by all 66-MHz implementations. So systems will continue to use one or the other voltage, and manufacturers must key 5 V and 3.3 V cards to prevent insertion into the wrong slot.

Some vendors provide universal cards with logic that can work at either 5 V or 3.3 V. They key such cards to plug into either 5 V or 3.3 V slots.




Figure 1. PCI bus architecture.

INDUSTRIAL APPLICATIONS

Price is a significant design consideration in desktop PC systems for consumers. In contrast, designs for critical embedded computing systems—like those used in telecommunications, industrial automation, and military applications—must also account for two other primary factors, reliability and maintenance. To meet these needs, industry has typically turned to the

VME bus, but is increasingly considering different forms of PCI.

Reliability

In desktop PCs, ISA and PCI adapters are inserted into card edge connectors located on a motherboard. A PCI card is typically fastened only at one point of the edge opposite to the connector. This mechanical

Table 1. VME card characteristics.

Card	No. of 96-pin connectors	No. of address bits	No. of data bits
3U	1	24	16
6U	2	32	32

arrangement cannot withstand shock and vibration very well, and the card edge connectors are subject to shifting or even disconnection. This poor situation is made worse by the use of an *active backplane*, a motherboard on which both ICs and connectors are mounted. In an industrial or military environment, the mechanical stresses involved in supporting cards may lead to failures in the motherboard.

To work around these problems, critical embedded systems use pin-and-socket connectors, front-panel retainers that lock the card to the frame, and card guides. These devices mechanically support every card on all four edges. A *passive backplane* also enhances reliability; its primary function is to provide electrical connections and shielding. Mechanical stresses on such a backplane are low because of the front-panel retainers and card guides.

Maintenance

Low mean time to repair is a key requirement in critical embedded systems. Failures in a passive backplane are uncommon due to robust mechanical characteris-

tics and the lack of active components. The use of a passive backplane considerably simplifies maintenance because ICs and other active components are on cards plugged into the backplane that can be speedily removed. This simple replacement procedure applies to all the cards in the system, such as processor cards or peripherals. The system controller is simply one of the cards plugged into the backplane. In fact, in telecommunications products, it is totally unacceptable for the backplane to be a field-replaceable assembly.

VME bus

To meet their reliability and maintenance requirements, the telecommunications, industrial automation, military, and medical markets have made wide use of the VME bus.³ VME began its life as a 16-bit data/24-bit address bus called VERSAbus; it worked with Motorola 680×0-based machines. VERSAbus was later modified to use the EuroCard form factor, after which it was renamed VME (Versa Module Eurocard). Today, VME supports 32-bit data/32-bit address cards. VME64 supports 64-bit cards. The VME system now supports all popular microprocessors, including the Motorola 680×0, Sun Sparc, Compaq Alpha, and Intel ×86.

VME systems are usually based on a passive backplane with one or more CPU cards, and one or more I/O cards. VME backplanes may have as many as 21 I/O cards. VME cards come in two form factors: 3U and 6U, which are defined by the VMEbus International Trade Association (VITA).³ Table 1 lists how these two cards vary.

Glossary

AGP (Accelerated Graphics Port): Standard that specifies a high-bandwidth connection between the graphics controller and main memory.

CardBus: Laptop version of the PCI bus.

CompactPCI: Passive backplane bus that attempts to replace VME. Cards are the same size as VME cards.

HiRelPCI: Standard currently being developed by IEEE. It supports the scalable coherent interface (SCI).

IRQ: Interrupt request.

IDE (intelligent drive electronics): Standard specifying how to connect secondary storage devices that have an embedded controller with the ISA bus.

PC/104: The embedded system version of the ISA bus. Also the name of the consortium that serves as a custodian of the standard.

PC/104-Plus: Standard that specifies PC/104 size cards with both ISA and PCI bus.

PC Card: New name for the 1994 release of the PCMCIA standard.

PCI (peripheral component interconnect): Local-bus standard.

PCI SIG (PCI Special Interest Group): Association established to monitor and enhance the development of the PCI bus.

PCMCIA (Personal Computer Memory Card International Association): Organization that developed the PCMCIA standard, the laptop version of an ISA compatible I/O expansion bus.

PICMG (PCI Industrial Computer Manufacturers Group): Consortium of computer product vendors established to extend the PCI specification for use in industrial computing applications. Developed the PCI-ISA passive backplane and CompactPCI standards.

PMC (PCI Mezzanine Card): Defined as IEEE Std P1386.1, PMC cards use PCI chips and may be mounted on VME cards.

SmallPCI: Expansion card with form factor identical to PC Card and CardBus. Primarily intended for original-equipment manufacturer products.

VGA (video graphics array): Originally a PC video card made by IBM, VGA today serves as a hardware compatibility standard for PC video cards.

VITA (VMEbus International Trade Association): Organization of VME bus manufacturers that supports the PCI Mezzanine Card standard.

VME (Versa Module Eurocard): Passive backplane bus.

The second connector in 6U cards provides additional address and data bits, and 32 user pins. The user pins allow extra user cables to connect to the card through the rear of the backplane. This arrangement allows users to rapidly switch cards without removing the cables first. The 3U and 6U VMEbus cards can be adapted to 9U racks if they use the first two connectors to define a VMEbus. Today, most standard cards are 32-bit-wide 6U cards.

PCI-based buses

Although analysts expect the VME architecture to continue dominating the industrial system market (at least in the near future), PCI's success has led to three new standards that apply PCI technology to industrial systems:

- PCI-ISA passive backplane standard,
- PMC (PCI Mezzanine Card), and
- CompactPCI.

A fourth standard, HiRelPCI, is currently under development by an IEEE working group.

PCI-ISA passive backplane. This standard specifies a passive backplane with two buses, ISA and PCI, and a CPU card with dimensions identical to the dimensions of the longer PCI cards. The CPU card contains all the components normally located on the motherboard (including the PCI chipset that implements the PCI-ISA bridge), and two connectors, one each for the ISA and PCI buses. Hence the "active" motherboard function is replaced with a plug-in CPU card. All field-replaceable expansion cards plug into a passive backplane that has only connectors on it. As we mentioned earlier, passive backplanes improve reliability, simplify maintenance, and are required in many applications.

The PCI-ISA standard supports up to 15 ISA slots and four PCI slots. As in desktop PCs, these multiple slots allow the use of off-the-shelf ISA and PCI adapters, which use standard ISA and PCI edge connectors. These connectors are less reliable than the pin-and-socket connectors used in VME and CompactPCI systems.

PCI Mezzanine Card. The PCI Mezzanine Card standard⁴ uses existing PCI silicon and packages it in a different form factor suitable for mounting on existing CPU cards. These cards are used most in VME-based systems. In this way, manufacturers of VME-based systems enjoy VME's expandability with a wide selection of PCI-based solutions. PMC is defined as IEEE P1386.1 and backed by VITA, the VMEbus International Trade Association.

PMC is basically an I/O expansion card for single-board computers. In desktop PCs, PCI cards are perpendicular to the motherboard, a space-consuming




Figure 2. 3U and 6U form factor Eurocards. A slot may contain one double-height 6U card or two single-height 3U cards. The card guides are not shown.

installation. In contrast, PMCs provide a low-profile configuration for systems that cannot directly use PCI technology due to space limitations. Two single PMCs or one double PMC can be attached in parallel to a VME board. PMC does not replace the VME bus.

CompactPCI. The CompactPCI standard^{5,6} attempts to replace VME directly by using PCI chips on 3U or 6U form factor cards, called Eurocards. Figure 2 shows the size of these cards and how they insert into slots. The Eurocard mechanical form factor is common in embedded systems due to the VME bus's popularity, and the use of card guides and pin-and-socket connectors results in a rugged and reliable package. This is the main reason for adopting the Eurocard mechanics in CompactPCI. An added benefit is that because VME and CompactPCI cards have the same form factor, you can mix them in the same frame. The two buses can interconnect via a CompactPCI-VME bridge.

CompactPCI uses high-density 2-mm connectors designed for telecom products. 3U cards use two connectors, as shown in Figure 2. The lower connector contains 32-bit PCI signals; the upper connector is for 64-bit cards. 6U cards use three additional connectors for user-defined I/O. Note that the lower two connectors are identically defined in 3U and 6U cards. Hence, as far as PCI bus signals are concerned, 3U and 6U cards are electrically interchangeable. Unlike normal PCI, CompactPCI allows up to eight cards on a backplane without a bridge chip. This improvement is mostly due to the improved noise margins obtained from the use of high-quality connectors and large numbers of ground pins.

In comparison with VME, CompactPCI has the following benefits:




Figure 3. Views from (a) above a single PC/104-Plus card and (b) the side of cards stacked one on top of another using space-through connectors. Spacers are not shown.

- CompactPCI systems use the same PCI chips and the same software as desktop PC systems. Operating systems, drivers, and applications that run on desktop PCs also run on CompactPCI systems. Hence several operating systems and applications work with CompactPCI platforms.
- PCI chips are plentiful. As a result, these chips are inexpensive and can be used in CompactPCI systems. Due to high integration levels, a PCI chipset consists of only a few chips, which considerably simplifies system design.
- The CompactPCI peak bandwidth is 132 Mbytes/s for the 33-MHz 32-bit implementation, and 264 Mbytes/s for the 33-MHz 64-bit implementation. This bandwidth is identical to the PCI bus peak bandwidth.

HiRelPCI. The IEEE 1996 working group is now developing HiRelPCI.⁷ The group intends this bus to be a high-reliability, high-availability system for transportation control and telecommunications. Basic bus operation follows the PCI specification,¹ and the current draft proposal specifies a connector similar to that of CompactPCI. It also specifies two metric board formats,⁸ 6su and 12su, which are both larger than the corresponding VME and CompactPCI inch-based 3U and 6U Eurocard formats shown in Figure 2.

HiRelPCI adds features to support a large number of nodes. For example, large systems often involve long delays, but PCI was not designed to handle long latency devices. PCI-2.0-compliant target devices had to hold the bus until they could deliver the data, no matter how long it took. The bus couldn't be used during this latency time, nor was the latency time limited. To solve these problems, the PCI Special Interest Group—an industry consortium—modified the PCI

standard in a few subtle ways and included the following changes in revision 2.1:

- Target latency is limited to a specified number of cycles. Any target requiring longer latency must terminate the transaction.
- A PCI 2.1 master must reissue a request if the target disconnects without delivering data. The master must retry the request until the target is ready.
- After disconnecting, the target must refuse any request other than the previous one. This preserves the read order.

This solution lets other masters compete for the bus between the request and the target reply, but it has two problems:

- The master has to poll the target until it receives the data, wasting bus bandwidth that other masters could use.
- Because polling occurs at fixed time intervals, the master will get the data when it polls the target after the data is ready, rather than getting the data as soon as the target is ready. This makes a long-latency transaction even longer.

HiRelPCI addresses the latency problem by extending PCI to include a packet mode. This mode provides split response operation, which improves system performance by freeing the bus during the time the target needs to fetch or process data. Packet operations are transparent to normal PCI transactions. A HiRelPCI bus can operate in packet or PCI mode or both.

EMBEDDED SYSTEMS: SINGLE-BOARD COMPUTER EXPANSION

PC/104 is the embedded system version of the ISA

bus. PC/104-based embedded systems incorporate 16- and 32-bit x86 processors running at 16- to 33-MHz clock rates. Many of these systems are upgraded versions of earlier implementations that used 8-bit micro-controllers, and the 5-Mbytes/s ISA (PC/104) bus meets their bandwidth requirements. Introducing the Pentium CPU into embedded systems, however, leads to a need for an extended PC/104 bus. The result is the PC/104-Plus standard,⁹ which defines a new form factor and a pass-through connector for PCI but remains compatible with existing PC/104 cards. PC/104-Plus offers

- compact 3.6-inch by 3.8-inch cards,
- self-stacked cards (which provide expansion without backplanes or card cages), and
- the use of pin-and-socket connectors and four-corner mounting (which have reliable electric connections and good shock and vibration characteristics).

As illustrated in Figure 3b, a typical system may consist of PC/104 (ISA) cards and PC/104-Plus (PCI) cards stacked using pass-through connectors; stacking lets you use multiple cards on a single bus.

As in other ISA systems, the 104-pin ISA bus is split into two connectors, shown on the right sides of Figures 3a and 3b. The 120-pin PCI bus in PC/104-Plus modules uses a high-density (2-mm) connector, which consumes less space than ISA connectors. On the other hand, PCI connector pins (to the left in Figure 3b) are thinner and more vulnerable than ISA connector pins (on the right). Because PC/104 systems do not use card guides, a connector pin “shroud” (which is not shown in Figure 3) serves as a guide for the PCI connector and protects the connector pins.

LAPTOPS

As its name indicates, CardBus, release 1.0 of the PCMCIA (Personal Computer Memory Card International Association) standard, was designed for memory cards that are primarily used in laptops. These Type I cards are about the size of a credit card and 3.3 mm thick. The next version, release 2.0, enabled the use of network, modem, and other I/O expansion cards by introducing a higher capacity, 5-mm thick, Type II form factor. Type II maintains backward compatibility (Type I cards will work in Type II slots). Release 2.0 also defines a Type III (10.5-mm thickness) primarily for use in PCMCIA hard disk drives. Two vertical Type II slots may contain either two Type II cards or a single Type III device.

The version of the standard released in 1994 (and updated a few times since) was renamed PC Card¹⁰ instead of PCMCIA release 3.0. PC Cards use the same form factor and Type I, Type II, or Type III




Figure 4. PC Card and CardBus need socket services software to deal with different hardware interfaces. Card services software manages system resources among PC Cards.

thickness as PCMCIA release 2.0.

The PC Card standard also defines several enhancements; the enhancement pertinent to this discussion is a 32-bit, 33-MHz interface called CardBus. CardBus is very similar to PCI, but addresses additional issues such as hot swap and power down, which are essential in laptop implementations. CardBus cannot use PCI chips directly, but it is very easy to modify existing PCI chips to be CardBus compatible. CardBus uses Type I, Type II, or Type III thickness, and the same form factor and connector as PC Card. A PC Card can use a CardBus slot; a connector key, however, prevents a user from inserting a CardBus card into a PC Card slot.


Each PC Card (or CardBus) slot connects to an adapter that can control multiple slots (called sockets in PC Card literature). Because adapters from different vendors often have different hardware interfaces, a software layer called socket services provides a generic interface to PC Card adapter hardware. Hence, socket services hides inconsistencies in the hardware interface. The next software layer, card services, enables the use of the same slot for different PC Cards. Card services manages a pool of resources (such as memory maps, I/O addresses, and interrupt request lines) and assigns resources to PC Cards plugged into the system.

Figure 4 shows where these layers of support software are in the system implementation.

MOBILE SYSTEMS: SMALLPCI

SmallPCI¹¹ is a small form factor of PCI endorsed by the PCI Special Interest Group; it has dimensions and connectors very close to those of CardBus. CardBus and SmallPCI target different market segments, however.

Figure 5. In (a) an implementation without AGP, the graphics controller must access the system memory through the PCI bus; (b) with AGP, the graphics controller has direct access to the system memory through the bridge. For clarity, I/O devices connected to the PCI bus and the connection of the CPU bus to the bridge are not shown.



CardBus developers intend it to be a standard for end-user cards and, as such, it supports both hot swapping and PCMCIA backward compatibility. As we have discussed, however, CardBus requires special software support (socket services) for its operation.

On the other hand, the developers of SmallPCI intend it for OEMs who want to install small functional modules based on PCI chips inside their products. Potential applications include set-top boxes, PDAs, traffic controllers, elevator controllers, navigational systems, and other products that require small-form-factor packages.

SmallPCI supports all the features of a standard PCI card (without a 64-bit extension), but adds a CLKRUN signal. CLKRUN controls the PCI clock frequency to support a reduced power mode. A SmallPCI card has the same form factor as PC Card and CardBus, and inserts into a connector mounted on and parallel to the system board. The connector is keyed to prevent insertion of PC Card or CardBus devices.

GRAPHICS APPLICATIONS

Intel developed the Accelerated Graphics Port standard¹² to give PCs a high-bandwidth graphics expansion slot. Because of its high speed, they designed AGP for point-to-point operation. Accordingly, it supports a single graphics expansion card. As shown in Figure 5, AGP provides a high-bandwidth path between the graphics controller and the system memory, which allows the graphics engine to use both local and main memory efficiently. As a result, some 3D data structures may shift from the graphics card's local memory into main memory.

Texture data, for example, is read only and resides in main memory without causing data consistency problems. Data loaded from secondary storage into the graphics card's local memory must pass through main memory anyway; leaving it in main memory saves the transfer to the graphics card.

Clearly, moving data to main memory reduces the cost of the graphics local memory, but requires allocation of additional main memory. This alternative is attractive for two reasons. First, main memory uses less-expensive memory chips. Second, the extra main memory purchased to support graphics applications is also available to other applications.

As Figure 5 shows, AGP does not replace the PCI bus. Essentially, AGP is an additional connection point into the main memory, through the bridge chipset. This path is both logically and physically independent of the PCI bus. Most I/O devices use the PCI bus and PCI transactions, leaving AGP for the exclusive use of graphics devices. AGP maintains compatibility with the standard PCI protocol, but uses additional signals to add new AGP data transfer modes, called AGP transactions.

There is a distinction between a bus and its transactions. For example, a PCI bus can only run PCI transactions. Because the AGP standard is an extension of the PCI standard, AGP buses can run both PCI and AGP transactions. So AGP bus traffic can interleave AGP and PCI transactions.

Compared with the PCI standard, AGP offers the following benefits:

- *Data transfer rates of 133 MHz.* With a sustained throughput of more than 500 Mbytes/s, AGP

throughput is substantially higher than that of current 33-MHz PCI buses.

- *Demultiplexed address and data transfers.* By adding an optional, 8-bit address bus, multiple read requests can be queued up through this bus, while the main address/data bus is used for data transfer. Request queuing helps achieve 100 percent actual bus usage (rather than an asymptotic performance level on infinitely long bursts, as in the PCI bus).
- *Pipelined memory operations.* By decoupling read data transfers from read requests, a system can initiate multiple requests without waiting for the first request to finish. This feature helps hide memory access latency.

Pipelined AGP transactions provide split response operation. The graphics controller initiates a pipelined transaction by placing a memory read request. The bridge responds by providing the data later. The graphics controller can issue additional requests, which may be also pipelined without waiting for pending transactions to complete. The number of outstanding transactions is limited by the amount of buffering in the graphics and bridge chips.

PCI's success gave birth to multiple PCI-related standards. These standards build upon PCI technology and leverage it for the benefit of other areas for which normal PCI cards are unsuitable. Because of its huge volume, the desktop PC market enjoys ongoing investments of billions of dollars in hardware and software innovation. PCI-related standards enable the use of these innovations in telecommunications, instrumentation, industrial control, military applications, and other segments of the computer industry.

The future of PCI lies in 64-bit-wide and 66-MHz implementations. Further into the future, implementations of PCI-X, a specification currently under development, will support throughputs of up to one gigabyte per second. ♦

Acknowledgment

We thank the referees for numerous comments and suggestions that have improved this article's contents, structure, and presentation.

References

1. *PCI Local Bus Specification, Revision 2.1*, PCI Special Interest Group, Hillsboro, Ore., 1995; <http://www.pcisig.com>.

2. *IEEE Std. 1275, IEEE Standard for Boot (Initialization Configuration) Firmware: Core Requirements and Practices*, IEEE, Piscataway, N.J., 1994; <http://standards.ieee.org/catalog/bus.html>.
3. *The VMEbus Handbook*, VMEbus Int'l Trade Assoc., Scottsdale, Ariz., 1989; <http://www.vita.com>.
4. *IEEE Draft Std. P1386.1 Physical/Environmental Layers for PCI Mezzanine Card, PMC*, IEEE, Piscataway, N.J., 1995; <http://standards.ieee.org/catalog/drafts.html>.
5. *CompactPCI Specification, Revision 1.0*, PCI Industrial and Computer Manufacturers Group (PICMG), Wakefield, Mass., 1995; <http://www.picmg.com>.
6. *CompactPCI vs. Industrial PCs*, Force Computers, San Jose, Calif., 1997; <http://www.forcecomputers.com/technology.htm>.
7. *IEEE P1996 HiRelPCI* (unapproved draft), Summit Computer Systems Inc., San Jose, Calif., 1996; <http://www.scsi.com/IEEE.htm>.
8. *IEEE Std. 1301.2, IEEE Recommended Practices for the Implementation of a Metric Equipment Practice*, IEEE, Piscataway, N.J., 1993; <http://standards.ieee.org/catalog/bus.html>.
9. *PC/104-Plus Standard*, Ampro Computers Inc., San Jose, Calif., 1997; <http://www.ampro.com/forum/specs/pc104p10.pdf>.
10. *PC Card Standard*, Personal Computer Memory Card Int'l Assoc. (PCMCIA), San Jose, Calif., 1997; <http://www.pcmcia.org>.
11. *SmallPCI Specification, Revision 1.0*, PCI Special Interest Group, Hillsboro, Ore., 1996; <http://www.pcisig.com>.
12. *Accelerated Graphics Port Interface Specification, Revision 1.0*, Intel Corp., Santa Clara, Calif., 1996.

Shlomo Weiss is a faculty member of the department of electrical engineering/systems at Tel Aviv University. His research interests include branch prediction, cache memory, and embedded systems. Weiss has a PhD in computer science from the University of Wisconsin at Madison. He is a member of the IEEE and the IEEE Computer Society. Contact Weiss at shlomow@msil.sps.mot.com.

Ehud Finkelstein is a VLSI design engineer at Seagull Semiconductor Ltd. His research interests include CPU architectures, CAD tools, and programmable logic. He has an MSc in electrical engineering from Tel Aviv University.