FAIRCHILD SEMICONDUCTOR™

February 2002
Revised April 2002

Section 2 - Backplane Protocols Backplane Designer's Guide

Backplane bus protocols determine - electrically and mechanically - how signals are dropped off or picked up at every device attached to the line. Protocols determine how cards pass information, the speed of data transfer, the maximum number and type of peripheral devices allowed, and the physical size and layout of the system. In general, the choice of the correct protocol for a backplane design is a complex decision that involves many design compromises. No one solution will give the best performance in each design category.

This section discusses the readily available Peripheral Component Interconnect (PCI)-based and VERSAmodule Eurocard (VME)-based backplane protocols. The goal of this section is to provide an understanding of what is necessary to implement a backplane protocol and how to choose the protocol that best fits the design needs.

Section Reference

This section presents the following information:

- A summary comparison of PCI-based and VME-based protocols
- A comparison, by feature, of the strengths and weaknesses of each protocol
- Information about PCI, CompactPCI®, PCI-X®, and VME bus.

Section	Section Title	Contents
1	Introduction	Application demands, basic backplane considerations, and how to use this guide.
2	**Backplane Protocols**	**Descriptions of different backplane bus protocols, including PCI- and VME-based protocols.**
3	Backplane Architecture	Topics relevant to backplane configuration, including parallel versus serial configuration and different configuration topologies and timing architectures
4	Backplane Design Considerations	Issues relevant to backplane layout, including distributed capacitance, transmission line effect, stub length, termination, and throughput.
5	Backplane Signal Driving and Conditioning	Signal driving and conditioning, including power consumption, rise/fall time, propagation delay, flight time, device drive, pin conditioning, live insertion, and incident wave switching.
6	Noise, Cross-talk, Jitter, Skew and EMI	A review of the enemies of signal integrity and high frequency.
7	Transceiver Technologies	Detailed information about the following technologies: TTL-based (ABT, FCT, and LVT); ECL; and GTLP.
8	Mechanical Considerations	Information about mechanical considerations such as backplane chassis/cages and connectors.
9	Layout Considerations	Physical layout of the receiver and driver cards plugged into the backplane, primarily focusing on construction of the physical layer and the configuration of the devices that comprise the cards.

Feature Comparison

Choosing the best platform for an application requires design compromises due to various constraints such as cost, size, and power requirements. To illustrate these compromises, this section summarizes the strengths and weaknesses found in PCI-based and VME-based backplane bus protocols. These strengths and weaknesses include the maximum number of cards on a trace for parallel backplanes, data path widths, maximum data rates, hot swap capabilities, and costs. The table below is a summary of the PCI- and VME-based protocols highlighting their individual features.

Parallel Protocol Summary

Attribute	PCI-Based Solution	VME-Based Solution
Address Space	32-Bit Address Space	Up to 64-Bit addressing
Data Path	Up to 64-Bit Wide	Up to 64-Bit wide
Bus Clock Speed	33, 66, or 133 MHz Bus Clock Speeds	Asynchronous transfers do not use a standard clock
Data Rate	Up to 1.06 Gbyte/s Transfer Rate	Up to 512 Mbyte/s transfer rate using proprietary Arizona Digital design
Hot-Swap Capabilities	Yes has been specified	Yes has been specified
Operating Voltages	3.3 or 5 Volts	3.3 or 5 Volts
Cost	Reduces Cost	Can be high: licensing costs as well as lower volumes and historically higher margins compared to PCI

- **Highest number of cards on a segment: VME64.** VME technology is the clear leader for the number of cards found in a segment (or continuous trace). VME is able to fit up to 21 cards on a standard 19-inch backplane. It is important to note that, although PCI can have only four cards (eight for CompactPCI and two for PCI-X) on a segment, the technology can support up to 256 segments. However, this does have drawbacks because segments are arranged in a hierarchical tree that requires more logic bridges and longer propagation delay times when adding additional segments to the bus.
- **Widest data path: VME64 (64-bit); PCI (64-bit).** The width of the data path has increased to 64 bits. This size is a standard option in VME64, VME64x, VME320, PCI, CompactPCI, and PCI-X.
- **Fastest data rate: PCI-X (133 MHz, 64-bit, 1.06 Gbyte/s).** The PCI protocol is, by far, the leader in speed. With the new PCI-X specification operating at 133 MHz and 64-bit data path, speeds of up to 1.06 Gbyte/s can be achieved. Current limitations allow only two cards on this high-speed segment and only one segment can run at 133 MHz, leaving the other segments to run at 66 MHz. The fastest competitor in the VME family is VME320, licensed by Arizona Digital, which has a bus bandwidth of over 320 Mbyte/s and a peak speed over 500 Mbyte/s.
- **Hot-swap capabilities: CompactPCI, PCI-X, VME64, VME64x**. To hot swap means to leave the bus interface component on, or in a known state while removing or inserting a card into the system. Hot swapping or live insertion is important for customers who require systems with zero downtime. All the major protocols are adopting specifications to support hot-swap capabilities. CompactPCI and VME64 both include special mechanical configurations in the connector to make and break the ground, pre-bias, control, and V_{CC} connections before the data pins touch. This allows a card to be precharged and reduces insertion transients associated with noise and voltage spikes.
- **Low Cost: PCI.** Devices manufactured for the PCI specification, including CompactPCI and PCI-X, are generally less expensive than VME technology. PCI devices are an inexpensive, mass-produced bus and board standard, with literally thousands of suppliers competing at low margins. If VME is used, it will have some PCI in the final product because compatibility issues force VME system designers to include PCI in their products. With over 800 companies supporting PCI and designing cards and accessories for the PCI open specification, VME customers demand that products they buy are compatible with their current and future PCI devices. The addition of a VME bus and a local PCI bus, of bridges, and of other logic increases complexity and cost. This provides PCI an advantage by being less costly and compatible with the widest range of devices currently on the market.

PCI

Peripheral Component Interconnect (PCI) architecture is the most widely used multifunction application platform available to meet Personal Computer (PC) computing needs. This architecture was designed to alleviate problems such as Interrupt Request (IRQ) and address conflicts associated with installing a new card in an ISA bus-based computer. The advantages of PCI include data rates up to 532 Mbyte/s, low cost, multimedia support, and the open-standard aspect.

The open-standard interface is an advantage for PCI device and application development because all compliant host systems are capable of hosting any PCI-compliant peripheral device. The PCI Special Interest Group comprises 800 member companies that develop products based on the PCI specification. More information about the group can be found at their Web site:
http://www.pcisig.com

PCI Electrical Characteristics

- 32- or 64-bit data path
- 33 MHz and 66 MHz bus clock speed
- Maximum data rate of 128 Mbyte/s for 33 MHz/32-bit bus, and 532 Mbyte/s for 66 MHz/64-bit bus
- CMOS drivers
- TTL voltage levels
- 5V, 3.3V interoperable operating voltage
- Reflected wave switching
  (Bus flight time is calculated into the timing window. This requires the bus to be short.)
- Device connections above 4 require PCI-to-PCI bridges
- Direct drive - no external buffers

PCI Features

- Multiple bus masters on the same bus
- Autoconfiguring: All components that plug into the PCI bus will be configured by the Binary Input Output System (BIOS), making jumpers unnecessary.
- IRQ sharing: The PCI bus is capable of sharing a single IRQ between cards.
- High bus bandwidth: The PCI bus runs at a standard rate of 33 MHz up to 66 MHz for PCI 2.1. Standard data width is 32 bits (4 bytes) per clock cycle and can be expanded up to 64 bits.
- Multiple functions on one card: The PCI specification allows for up to eight functions (video, sound, etc.) on a single card.

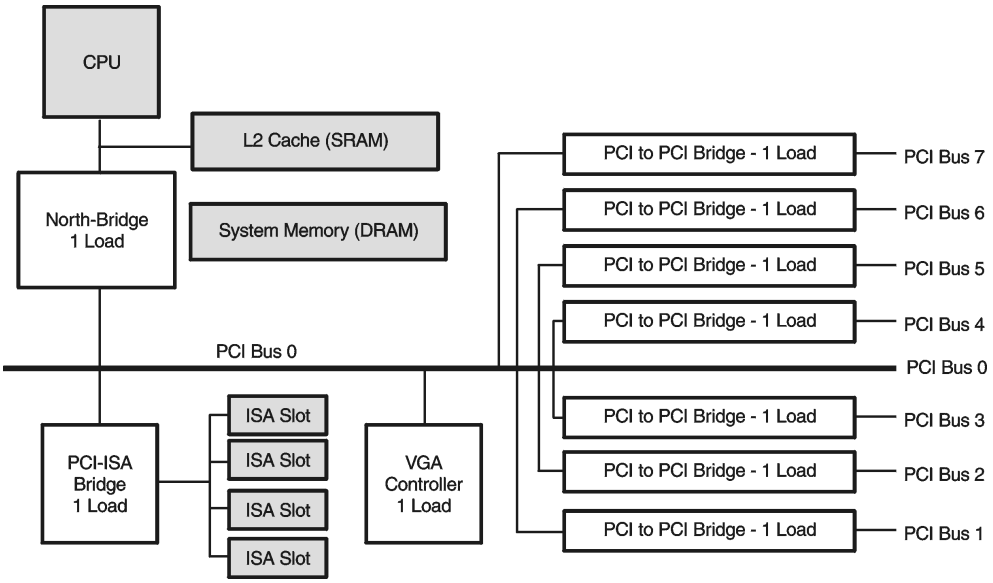

FIGURE 1. PCI Bus System with 8 PCI Busses

PCI Operation

PCI architecture is a local bus design. This means PCI devices have access directly to the CPU local bus. A PCI-HOST bridge accomplishes this access by connecting the PCI bus to the CPU and system memory. The bridge is a caching device that lets the CPU fetch information from the bridge while still allowing the cache controller expansion device access to the system memory. This is called "concurrent bus operation", a feature that previous architectures did not support.

Communication between peripheral devices is accomplished using burst transfers. To eliminate contention on the bus, only one bus master device can send information at a time. Other architectures use arbitration and streaming techniques that are typically not as efficient. The PCI method ensures that a burst communication will not be

PCI (Continued)

interrupted while providing a method to allow the next device to access the bus line.

Another advantage of the PCI design is that up to 256 devices can be attached to one PCI local bus, and up to 256 PCI busses can exist in one system. This allows for over 65,000 individual components in one system. In reality, a system designer cannot have 256 devices on one bus. The actual number is four devices. However, only four devices per bus permits over 1,000 devices in the system.

PCI is a plug-and-play architecture that is autoconfigured by system BIOS. The result is system resource assignments destined for expansion devices that are all automated. This eliminates difficulties with setting jumpers or dual in-line-package (DIP) switches on peripheral cards, creating a system that can dynamically handle additional cards, adjusting resources accordingly, without having to physically select and implement changes on the remaining devices in the system.

Disadvantages of the PCI Bus

Although PCI is widely accepted as the standard in PC applications, there are some drawbacks to this protocol. These include both physical and logical constraints that limit PCI's ability to scale to support large numbers of processors and peripherals. On the physical side, PCI has electrical and mechanical constraints to ensure proper signals, including limits on loading, trace lengths, and connectors. A real-world maximum on a single PCI-bus segment is a total of ten loads. Each loaded connector and device combination is considered two loads. Thus, a motherboard or passive backplane typically has a limit of four plug-in boards on a single PCI-bus segment. With 66 MHz PCI, this limit will be two plug-in boards on a motherboard. In the future, speeds greater than 66 MHz could mandate point-to-point connections. This would require a switching fabric interconnect between PCI agents.

PCI busses are card-limited because PCI operates using reflected wave switching. This means that a signal must pass the length of the segment and back before the system will recognize the change. At higher frequencies, the delay time between successive clocks is shortened causing the physical length of the bus to become the limiting factor.

PCI-to-PCI bridges are used to overcome the physical limitations of a single PCI-bus segment. Bridging allows connection of 256 PCI-bus segments. However, the bridge specification defines only the dual-port PCI-to-PCI bus bridge (P2P). Using P2P bridges results in the 256 bus segments connected as a hierarchical tree. This tree is an inefficient interconnection topology for larger systems because of significant contention, high latency, and poor bisection bandwidth characteristics.

CompactPCI

The CompactPCI protocol is an adaptation of PCI for industrial and embedded applications. The design has a more robust mechanical form factor than desktop PCI in order to compete with other protocols such as VME. CompactPCI uses industry-standard mechanical components and high-performance connector technologies to provide an optimized system that is intended for rugged applications. Compact PCI provides a system that is electrically compatible with the PCI specification and that allows low-cost PCI components to in a form factor suited for rugged environments.

As with PCI, CompactPCI is an open specification supported by the PCI Industrial Computer Manufacturers Group (PICMG), a consortium of companies that use PCI for embedded applications. More information on this specification can be found at: http://www.picmg.org

Compact PCI Form Factors

One of the most important features of the CompactPCI specification is the design of the physical board. The form factor defined for CompactPCI boards is based on the Eurocard industry standard. Two sizes are specified: 3U, which is 100 mm by 160 mm, and 6U, which is 233.35 mm by 160 mm.

The 3U form factor is the minimum for CompactPCI because it accommodates the full 64-bit CompactPCI bus. The 6U extension is defined for boards that need extra space on the card or more user-defined connections to the backplane. Rear connectors are numbered J1 through J5 starting at the bottom connector. J1 is used for 32-bit signaling, and the use of both J1 and J2 allows 64-bit PCI bus signaling.

CompactPCI Features

- Pin-and-socket interconnect mechanism
- Multivendor support
- Coding mechanism providing positive keying (prevents 3.3V card insertion into a 5 V slot)
- Staggered make-break pin connections for hot-swap capability
- Rear pin option for through- the- backplane and "sequencing" I/O applications
- High-density PCI capability
- Expandability for end-user applications

The connector is the second most important physical aspect of CompactPCI. Unlike PCI-specified components that use the edge of the printed circuit board as the means to connect to the backplane, CompactPCI uses a gas-tight, high-density pin-and socket connector that meets the International Electro Technical Commission (IEC) 1076 international standard. The connector's low inductance and controlled impedance make it ideal for PCI signaling. This 2 millimeter "hard metric" connector has 47 rows of five pins per row, with a total of 220 pins (15 pins are lost to the keying area). An additional external metal shield is also used. The large number of ground pins ensures adequate shielding and grounding for low ground bounce and reliable operation in noisy environments. This connector's controlled impedance minimizes unwanted signal reflections and enables CompactPCI systems to have eight slots, as compared to the desktop PC's four on a single bus segment.

PCI-X

Peripheral Component Interconnect Extended (PCI-X) is a bus technology that increases the maximum speed that data can move within a system from 66 MHz to 133 MHz. As with CompactPCI, the PCI-X Specification is an open-industry standard, available from the PCI Special Interest Group (SIG) as an addendum to the PCI Local Bus 2.2 Specification.

PCI-X was developed jointly by IBM, Hewlett Packard, and Compaq to be compatible with standard PCI interfaces. PCI-X is important because of the wide acceptance and use of PCI. PCI-X is backward compatible with the existing PCI bus at the adapter, device driver, and system level. A PCI-X system automatically switches between conventional PCI and PCI-X mode, depending on the type of adapters installed on the bus segment. PCI-X is designed to meet the crucial I/O demands for high-end computing requirements. More bandwidth is available to the system by using features like a register-to-register design and a more efficient use of the bus.

PCI-X Features

PCI-X doubles the speed and amount of data exchanged between the computer processor and peripherals. With current PCI design, one 64-bit bus runs at 66 MHz and additional busses move 32-bits at 66 MHz or 64-bits at 33 MHz. Using the standard PCI design, the maximum amount of data exchanged between the processor and peripherals is 532 Mbyte/s. With PCI-X, one 64-bit bus runs at 133 MHz with the rest running at 66 MHz, allowing a data exchange of 1.06 Gbyte/s. This is very important with the rise of technologies like gigabit Ethernet, fiber channel, ultra3 SCSI, and other bandwidth-intensive peripherals.

Compared to the conventional PCI protocol that uses an immediate method of sending information, PCI-X uses a buffer that works using the following sequence:

1. Operation begins on the rising clock edge. When detected, the device switches the output signal to a HIGH or LOW state on the PCI-X bus.
2. The signal propagates across the bus and is then sent to a register or flip-flop, which holds the signal state until the next clock cycle.
3. The receiving device has a full clock cycle to decode the signal and determine the appropriate response.
4. The receiving device responds two clock cycles after the sending device first sent the signal.

With the register-to-register protocol, PCI-X transactions require one clock cycle more for each transaction. For example, a write transaction that completes in nine clock cycles for conventional PCI will complete in ten clock cycles for PCI-X, adding an extra cycle. In conventional PCI specifications, each clock cycle has to give the peripheral enough time to decode an incoming signal and process the instruction, all in one clock cycle. This makes increasing the clock speed of the bus very difficult. The individual cycle for processing allows PCI-X to reduce total transaction time. A Compaq technology brief states: “A transaction that takes nine cycles at 33 MHz will finish in 270ns, while a PCI-X transaction that takes ten cycles at 133 MHz will finish in 75ns, improving the transaction time.”

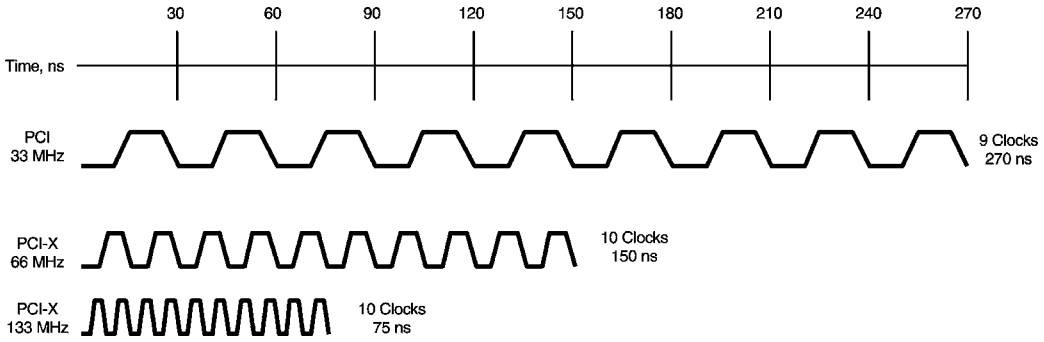

FIGURE 2. Time Difference Between PCI and PCI-X Protocols for Write Transaction

These speeds are possible because the extra clock cycle slices outsaves time that is otherwise wasted. The general objective of PCI-X is to identify inefficiencies in the standard PCI specification and eliminate them. Speeding up the transaction time by adding a dedicated clock cycle for processing is just one example of increased efficiency.

Enhancements for Bus Efficiency

PCI-X's second important throughput benefit is better efficiency when dealing with the bus. This efficiency is accomplished with four enhancements: the attribute phase, split transaction support, optimized wait states, and a standard block size movement.

ATTRIBUTE PHASE

The attribute phase uses a 36-bit attribute field to describe bus transactions in more detail than allowed by conventional PCI specifications. The use of the more flexible attribute phase allows four areas of improvement:

1. Relaxed ordering: With the standard PCI protocol, bus transaction requests are handled in the order in which they are received because the bridge cannot identify which transaction came from which device. The attribute phase includes information to allow higher priority requests to be serviced before lower priority ones. This benefits real-time applications such as streaming audio and video.

PCI-X (Continued)

2. Non-cache-coherent transactions: A common problem in multiprocessor systems is a consistent view of memory during a transaction. Because data is shared among multiple caches on the processor bus and main memory, old or inconsistent data may reside in different caches throughout the system. The standard PCI specification includes a processor “snoop function” whereby the CPU checks for inconsistent data. This snoop activity adds traffic to the bus and hurts performance. PCI-X provides more control over whether the CPU should snoop, based on the transaction.
3. Transaction byte count: With conventional PCI specifications, the PCI bridge fetches a default number of cache lines for every data request. Because there is no way to know how much data will be requested, the PCI bridge uses a default number of one or two lines. With PCI-X, the bridge knows exactly how much data to fetch because the byte count is included. This enhancement allows for better and more efficient buffer management.
4. Sequence number: This enhancement increases efficiency in buffer management. The sequence number uniquely identifies transactions that are a part of the same sequence.

SPLIT TRANSACTION SUPPORT

Split transaction support increases bus efficiency. The split transaction operation is performed when the device requesting the data sends a signal to the target. The target device informs the requester that it has accepted the request, and the requester is then free to process other information until the target device initiates a new transaction and sends the data.

OPTIMIZED WAIT STATES

In traditional PCI protocol, devices often add extra clock cycles, or wait states, into their transactions to hold the bus if the target PCI device is not ready to proceed with the transaction. PCI-X eliminates use of wait states, except for initial target latency. When there is nothing to transfer, the device removes itself from the bus to allow other peripherals a chance to send information. This reduces wasted clock cycles.

STANDARD BLOCK SIZE MOVEMENT

The PCI-X standard includes a transaction size boundary of 128-bytes, allowing for longer bursts of transmission. This standard block size improves the efficiency of the processor bus and main memory.

VMEbus

In 1980, Motorola, Mostek and Signetics designed VERSAmodel Eurocard (VME) as a computer bus protocol. The most recent version of VMEbus is the American National Standards Institute (ANSI) and VMEbus International Trade Association (VITA) 1-1994 Specification, also known as MVE64. VME64x and VME320 are other enhancements to the VME64 standard, facilitating faster operation, similar to that of PCI-X. Unlike PCI, which is an open specification, VME is a licenses technology and this fact can increase overall design costs.

VMEbus closely follows the IEEE Standard for Mechanical Core Specifications for Microcomputers (IEEE 1101). Single-height boards are referred to as 3U and double-height boards are referred to as 6U. “U” is a unit of measure for the front panel, where each U is equal to 1.75 inches. The use of two standard card sizes is a plus for VMEbus: single-height 3U boards are available for tight space requirements, and the larger, more common 6U card can be used when space is not restricted.

VME Electrical Characteristics

- 16-, 24-, 32-, 40-, or 64-bit addressing range picked dynamically
- 8-, 16-, 24-, 32-, or 64-bit data path picked dynamically
- Asynchronous transfer with multiplexed and nonmultiplexed bus cycles
- Primary/secondary architecture
- 0 to 500+ Mbyte/s data transfer rate
- 1 to 21 processors
- A maximum of 21 card slots in the backplane

VME64

VME64 is the new standard in the VMEbus family. Key improvements include an upgrade to a larger 64-bit data path and 64-bit address range for the larger 6U boards and double the bandwidth up to 80 Mbyte/s.

The VME64 architecture is configured as primary-secondary. Functional modules that are called “primary” transfer data to and from modules called “secondary”. Before a primary can transfer data, it must first gain control of the bus using a central arbiter. This arbiter is part of a system controller and its function is to determine which primary will receive access to the bus. All activity takes place on five sub-busses: the data transfer bus, the data transfer arbitration bus, the priority interrupt bus, the utility bus, and the serial bus.

VME64 is an asynchronous bus, meaning no clock cycles are used. Data is passed between modules using interlocked handshaking signals. The slowest module in the transfer sets the cycle speed of each data transfer.

VME64x

In 1997, the VITA Standards Organization adopted a superset to the VME64 standard. To add function and speed, the new capabilities in this standard include:

- New 160-pin connector family
- 3.3V power supply pins
- Higher bandwidth bus cycles (up to 160 Mbyte/s)
- Live insertion (hot-swap) capability
- VME64x modules are compatible with the older generation VME64 architecture. Older VME64 modules can also be plugged into newer VME64x backplanes, allowing for greater flexibility and reduced cost in system upgrades

VME320

In 1997, Arizona Digital released a modified VMEbus architecture called VME320. VME320 increased the bus bandwidth over 320 Mbyte/s, with peak bandwidths of over 500 Mbyte/s. Unlike the original VME64, VME320 is a proprietary architecture. However, VMEbus modules can be designed, built, and sold without license.

To achieve these fast speeds, the VME320 backplane uses a “star” interconnection. This means all connections on the backplane are routed to the middle slot, and the leading edges of signals pass through only one slot on their way to their destination. The result is a shorter signal travel distance when compared to other architectures. This allows for tighter skew delays and increased speed.

VMEbus Backplanes

VMEbus backplane cards are located on 0.8-inch centers (i.e., 0.8-inch slot pitch). This slot pitch allows from one to 21 cards to be supported on the backplane. This number is derived from the fact that a standard rack accommodates a backplane no larger than 19 inches, allowing a maximum of 21 cards on 0.8-inch centers.

Three general types of backplanes are available: standard, 3-row connector, VME64x; 5-row connector with 160-pin enhanced; and VME320. Many different options are available, including:

- Active termination: This reduces power consumption to less than a few milliamps. Most backplanes are supplied with passive termination that is generally simpler and less expensive, but passive termination consumes power, typically about 1.3 Amps.
- Automatic chain: Backplanes are available that automatically close the interrupt acknowledge and bus grant daisy chains.
- Rear I/O connector shrouds: Standard backplanes do not necessarily come with connector shrouds or specific connector hardware on the rear of the backplane.
- Power supply connection: No standard way of connecting power supply leads to the backplane.
- J2 high-speed backplane: Offers extra rows of ground pins.

Summary

This discussion of backplane protocols shows performance compromises are necessary when deciding which protocol to use. All backplanes contain some performance compromises regardless of whether they are an industry-specified standard architecture or a custom design. These compromises are due to electrical and mechanical constraints, and no solution can provide the highest performance in every category. All design choices are results of system requirements and cost; and, ultimately, the backplane design chosen will define overall system performance.