

Bridging PCI to PCI: Filling in the missing elements

CompactPCI board designers need to use PCI-to-PCI bridge chips to interface their onboard PCI bus to the CompactPCI backplane. There are two basic types of bridge chips:

- Transparent bridge chips
- Embedded bridge chips

There are several approaches to using these bridge chips in a CompactPCI system:

- The board designer can use a transparent bridge chip on the system slot board, and use an embedded bridge chip for each of the peripheral slot boards. (See Figure 1) However, this does not allow a user to use the same board in any slot of the CompactPCI system. The only way to provide that kind of versatility is to put

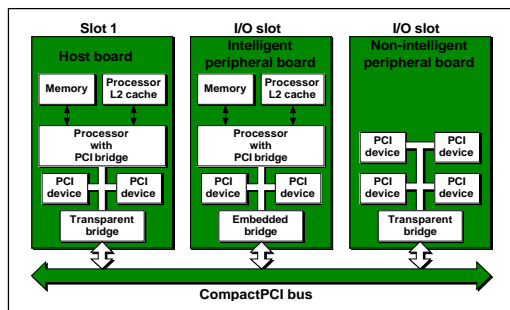


Figure 1

both types of bridge chips on the board.

- The board designer can use an embedded bridge on the system slot board along with some additional glue logic. However, because this is a non-standard hardware configuration, it might result in incompatibilities with existing software which has been designed to work with standard bridge chips.

Even when one of these approaches is employed, it does nothing to solve problems that engineers face when designing embedded scalable, mission-critical CompactPCI systems. For example, when designing an embedded telecom or data-com system, engineers need a solution that provides:

- Standardized support for asymmetrical multiprocessing
- Advanced interrupt support, to ensure deterministic response time
- I/O capable of high bandwidth throughput
- High availability, which is provided by high mean time between failure (MTBF) and low mean time to repair (MTTR) times

Asymmetrical multiprocessing

When performing asymmetrical multiprocessing (multiple processors, possibly running different operating systems, inside a single CompactPCI system, as shown in Figure 2) in an embedded CompactPCI application, intelligent boards must be able to communicate with each other.




Figure 2

One method of communicating is through the CompactPCI backplane. Today, CompactPCI backplane communication is based on proprietary hardware and software technologies. These hardware/software solutions are not standardized, nor are they supported by competitors. This lack of standardization prevents the integration of board-level components from multiple vendors.

Clearly, a standard approach is needed for the construction of asymmetrical multiprocessing systems – an approach that will guarantee that a board from Vendor X



will communicate with a board from Vendor Y. It is very important that this interprocessor communication method provide for high speed data transfers. Many embedded systems require very high data transfer rates – especially scalable systems.

Advanced interrupt support

Another problem that must be solved in designing a universal PCI-to-PCI bridge for CompactPCI embedded systems is the limiting interrupt structure on the CompactPCI backplane.

The CompactPCI backplane has only four interrupt lines. If there are more than four I/O boards plugged into the backplane, then some of these boards must share an interrupt line. For systems with many I/O boards, this might result in very long (or even non-deterministic) response times, when interrupts are requested. Mission critical, hard real-time applications might not be able to tolerate these slow or unpredictable response times.

Efficient I/O performance for high-bandwidth throughput

Many CompactPCI embedded applications require that the I/O support high-bandwidth throughput. Theoretically, a 32-bit, 33 MHz CompactPCI backplane is capable of moving data at a rate of 132 Mbytes/sec. However, in the real world, the data transfer rates through a standard PCI-to-PCI bridge are far less than this.

The main problem is that FIFO buffers in the PCI-to-PCI bridge chip are not deep enough to buffer an adequately-sized block of data at the theoretical maximum data transfer rate. When the buffer fills, the data transfer must pause, and this slows down the overall data transfer rate.

High Availability, better mean time between failures (MTBF) and mean time to repairs (MTTR)

In order to provide High Availability, a CompactPCI system must provide:

- a high mean time between failure (MTBF)
- a low mean time to repair (MTTR)

In general, CompactPCI systems have an excellent MTBF and MTTR, due to:

- good cooling
- robust connectors
- easy accessibility
- minimal cabling

However, to fully support high-availability the system's hot swapping mechanism should support:

- hot swapping of system slot boards
- hot swapping of non-intelligent I/O boards
- hot swapping of intelligent I/O boards

Unfortunately, the transparent bridges (which are used on non-intelligent I/O boards and on system slot boards) do not support hot-swap with special built-in registers. Currently, "hot-swap friendly" silicon is available only for intelligent I/O boards, which employ embedded bridge chips.

What's needed to fill in the missing pieces?

Embedded systems engineers need a universal PCI-to-PCI bridge chip that provides these features, in order to offer a robust, complete CompactPCI interface solution (See Figure 3).

Asymmetrical multiprocessing support

To provide a standard architecture for asymmetrical multiprocessing, a universal PCI-to-PCI bridge chip should include I₂O support hardware.

Interprocessor communication over the CompactPCI backplane in asymmetrical multiprocessing systems has traditionally been supported by *mailbox registers*. When sending data to an intelligent board with a mailbox register, the sending board first writes the data into the receiving board's on-board memory, and then writes to the receiving board's mailbox register. The write operation to the mailbox register generates an interrupt request to the receiver's onboard processor. This causes that processor to read the data from the onboard memory.

However, a PCI-to-PCI bridge chip only has one mailbox register to receive messages from all of the other CompactPCI boards in the system. Because two or more boards might write to that mailbox register at almost the same time, the receiving processor must check the message buffers for all of the other boards each time it detects an interrupt from the mailbox register.

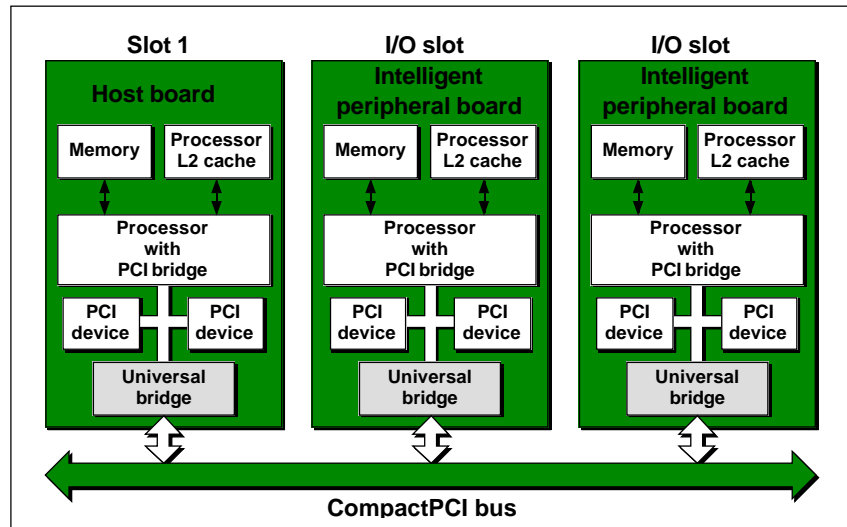


Figure 3

Ideally, this universal bridge chip should:

- support asymmetrical multiprocessing architectures (with intelligent I/O) using a standard such as I₂O
- include FIFOs and control logic
- include advanced interrupt support that bypasses the CompactPCI backplane interrupt lines
- include deep read/write FIFO buffers, to support high throughput
- provide enhanced availability through the hot-swapping of non-intelligent I/O boards

A better approach is to replace the destination mailbox register with a hardware-controlled FIFO. Ideally, this hardware controller would support I₂O protocols, and allow bi-directional data transfers between intelligent I/O boards.

The advantage of this FIFO method (over a mailbox register method) is that every incoming data message gets stored in the FIFO, and the associated interrupt requests are never overwritten due to near-simultaneous arrival. This eliminates the need for a time-consuming scanning loop each time a new message arrives, while ensuring that



the destination processor does not overlook the arrival of any new data.

This method provides another advantage over the mailbox method. Because it is implemented entirely in hardware, it is both operating system and processor independent. The hardware implementation also allows a board to receive data from multiple boards without putting an added burden on the board's processor.

Advanced interrupt support

Another problem facing engineers who are implementing embedded CompactPCI systems is the need for advanced interrupt support. In a CompactPCI-based system there are only four interrupt request lines. That might be adequate for a desktop PC. However, for embedded systems (especially CompactPCI systems) this is inadequate.

A CompactPCI system might have as many as seven I/O boards – each of which can generate interrupts. Most CompactPCI I/O boards drive only one CompactPCI interrupt request line. However, a multiple-function I/O board might drive as many as four interrupt request lines.

Real-time systems require deterministic response times. Sharing of interrupt lines can delay interrupt response, and might even make interrupt handling non-deterministic – a major drawback for mission-critical applications.

For example, consider the CompactPCI system in Figure 4. There are seven I/O boards and a slot 1 controller board running a real-time application. To keep things simple, let's assume that each of these I/O boards:

- contains a "single-function" PCI device
- drives only one interrupt request line to the CompactPCI backplane

With seven I/O boards, three of the four interrupt request lines will be *shared* by a pair of I/O boards. (In this case IntA, IntB, and IntC are shared.)

When the interrupt controller in slot 1 receives an interrupt request on one of the shared interrupt request lines, it will not know which of two boards generated the

request. To determine which board generated that request, the onboard processor will need to use the CompactPCI bus to access the interrupt registers on *both* boards. This delays identification of the requestor, especially if there is heavy traffic on the CompactPCI backplane.

A better solution is to avoid the use of shared interrupt request lines by employing *message signaled interrupts*, which are described in PCI specification 2.2. An I/O board sends a message signaled interrupt to the processor on the slot 1 board by accessing a specific register in the slot 1 PCI-to-PCI bridge chip. This access triggers an onboard interrupt request to the processor. Since the source of this interrupt request is known (based on which bridge register was accessed) it can be handled much faster than an interrupt request routed through a shared backplane signal (which requires polling).

High bandwidth

As previously mentioned, the maximum theoretical throughput that a CompactPCI system can achieve over a 33 MHz, 32-bit-wide backplane is 132 Mbytes/sec. However, to approach this data transfer rate, overhead cycles on the CompactPCI bus must be kept to a minimum. To do this a PCI-to-PCI bridge must support an efficient read/write strategy for handling data transfers.

One of the best ways to improve I/O performance in a CompactPCI system is to provide deep and advanced FIFO buffers. These buffers should be able to hold multiple transactions in each direction. This enables the PCI-to-PCI bridge to use an intelligent data transfer strategy (such as *flow through* of read transfers, and *posting* of write transfers) to come as close as possible to the maximum data transfer rate.

High Availability

High availability implies that mean time between failure (MTBF) has been maxi-

mized and the mean time to repair (MTTR) has been minimized. To maximize MTBF, a CompactPCI system should have good cooling and robust connectors. To minimize MTTR a CompactPCI system should have easy accessibility, hot-swapping capability, and minimal cabling.

As mentioned earlier, non-intelligent I/O boards that employ transparent PCI-to-PCI bridge chips cannot be swapped. To provide higher availability, hot-swapping capability should be extended to include non-intelligent I/O boards.

To provide truly comprehensive hot swapping support, there should be some mechanism for hot swapping the slot 1 board. As per the PICMG 2.13 hot-swapping specification (which is still in the standardization process) slot 1 hot-swapping will be supported in the future.

Force's Sentinel universal PCI-to-PCI bridge

Current PCI-to-PCI bridges cannot deliver all of the features that are essential for designing mission-critical and scalable systems. To solve this problem, Force Computers has introduced the Sentinel chip – an integrated circuit that provides a high-speed PCI-to-PCI bridge for a new generation of high-performance, processor-independent CompactPCI boards. This Sentinel chip:

- has been optimized for high-speed data transfers and multiprocessing
- provides message signaled interrupts
- offers improved I/O performance, with an advanced FIFO buffering architecture
- enhances system availability by extending hot-swap capability to non-intelligent I/O boards

CompactPCI boards that employ the universal Sentinel PCI-to-PCI bridge can be designed to meet the needs of a full spectrum of applications.

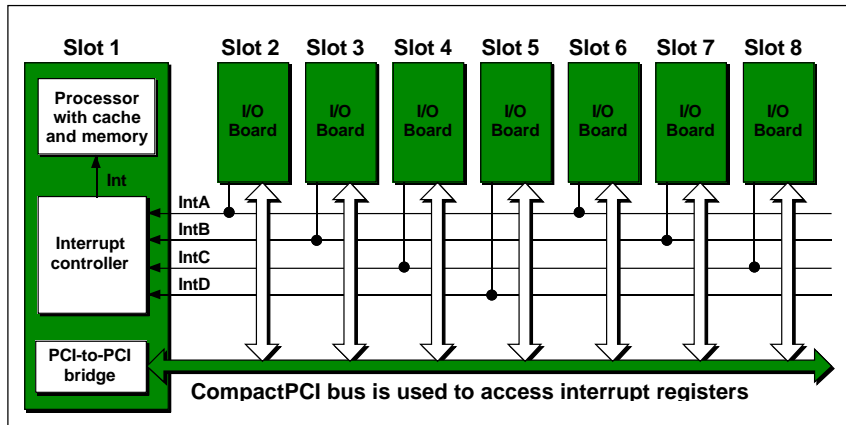


Figure 4

Support for asymmetrical multiprocessing

The major barrier to supporting asymmetrical multiprocessing in an embedded application has been the low data transfer rate through the CompactPCI backplane. However, the Force Sentinel chip provides a FIFO structure with sophisticated control logic that handles fast data transfers through the CompactPCI backplane. With the intelligent I/O capability built into the Sentinel chip, fast data transfers can be provided (in both directions) between intelligent I/O boards. Because all the protocols are handled by the chip hardware, this fast data communication is both operating system and processor independent.

Advanced interrupt support through message signaled interrupts

Embedded PCI systems and desktop PC systems have some similarities as well as some significant differences. The most significant difference is the ability to tolerate delayed (or non-deterministic) response times.

While PCs can tolerate compromised response time, embedded CompactPCI systems that are used for mission critical real-time applications cannot. A CompactPCI system might have as many as eight CompactPCI boards sharing just four interrupt request lines. The Sentinel chip solves this problem with message signaled interrupts, bypassing the backplane interrupt structure to allow deterministic response times.

The Sentinel chip resides on the slot 1 board (as shown in Figure 5) and can pass interrupt requests received from other CompactPCI boards directly to its onboard processor. Any CompactPCI board can send a message signaled interrupt to the Sentinel chip by simply accessing a register in the chip.

Deep read/write FIFO buffers support high bandwidth

The Sentinel PCI-to-PCI bridge offers a deep and advanced FIFO buffer structure, implemented as four data FIFOs, each with 128 entries (each 32 bits wide) to allow buffered reads and writes in both directions. These FIFOs can hold as many as eight transactions in each direction. Given a cache line size of eight data words, the Sentinel can even buffer up to 16 cache lines.

This enables the PCI-to-PCI bridge (together with an intelligent data transfer strategy, like *flow through* for read transfers and *posting* of write transfers) to approach the maximum data transfer rate of 132 Mbytes/sec.

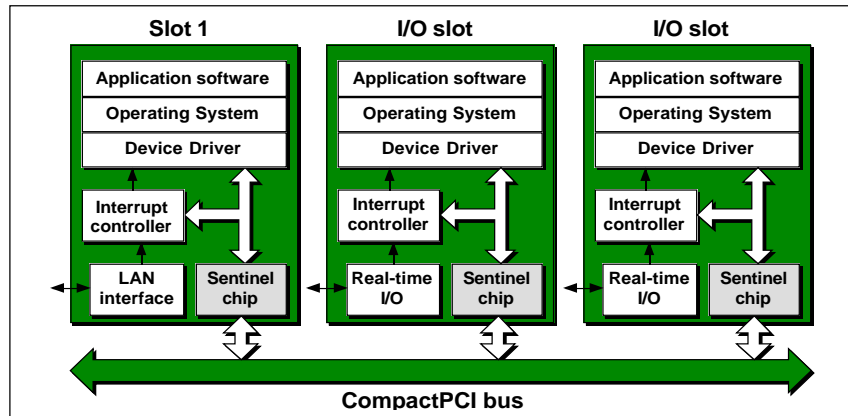


Figure 4

Hot-swappable silicon for enhanced availability

How swapping is vital when providing highly available CompactPCI systems. This includes:

- fan units
- power supplies
- CompactPCI boards

Transparent PCI-to-PCI bridges (which are used on slot 1 and non-intelligent I/O boards) do not support hot swapping. However, the Sentinel chip provides the necessary hot-swap friendly silicon for non-intelligent I/O boards, as well as for intelligent I/O boards.

More importantly, the Sentinel chip provides this hot-swap capability without the need for any additional glue logic or software changes, because all of the necessary logic and registers have been integrated into the chip.

A solution that fills the gaps

Until now, engineers had to make due with PCI-to-PCI bridge chips that only solved a portion of their embedded application design problems. However, Force's universal, single-chip PCI-to-PCI bridge chip combines the functionality of an embedded bridge chip and a transparent bridge chip into a single device. In addition, it offers the performance and the enhanced features that are needed to build a complete, robust CompactPCI system. The Sentinel chip offers:

- asymmetrical multiprocessing support
- advanced interrupt support
- increased backplane communication speed
- enhanced system availability

With all these features, the Sentinel chip eliminates the need for costly, time-consuming, custom interface designs in CompactPCI systems.



Christoph Adam is a strategic marketing manager at Force Computers. His current responsibilities include defining future directions and technologies for Force Computers products. Prior to that, he worked for two years in the technical marketing department, interfacing with customers and defining new VMEbus and CompactPCI products. He joined the company in 1993 as an ASIC Design Engineer, designing high-speed VMEbus Interfaces and PCI controllers. Adam holds a degree in Electrical Engineering/Information Technology.

If you have questions about this article, or if you would like to know more about Force's products, you can contact Force at:

Corporate Headquarters

Force Computers

5799 Fontanoso Way
San Jose, CA 95138-1015
Tel: 408-369-6000 • Fax: 408-371-3382

European Headquarters

Force Computers GmbH
Prof.-Messerschmitt-Str. 1
D-85579 Neubiberg/Munchen
Tel: (0 89) 608 14-0
Fax: (0 89) 609 77 93
Email: info@force.de
Web: www.forcecomputers.com

Trademark notices:

Sentinel is a trademark of Force Computers