

ASPEN: the Serial PCI - Bus Solution

For already more than 10 years PCI is the most used parallel bus system in today's Personal Computers. This low-cost bus has in recent years also been adopted in various industrial applications, e.g. compact PCI, and telecom. The following article describes a serial PCI solution enabling connection of different PCI busses over distance up to 15m, for several new upcoming applications highlighted at the end.

DESIGN RULES

To develop a PCI bus system, designers have to take care of different rules where the following list is intended to be a brief summary.

Simplified PCI Design Rules:

- maximum of 10 loads: This rule is basically the limitation of the bus-length and should keep the overall line impedance within acceptable limits
- device-I/O on motherboard is 1 load, The 10pF I/O-capacitance is 1 load
- device-I/O on add-in-card is 2 loads, 10pF I/O-capacitance plus 10pF capacitance of the stub-line and the connector.
- max. length of stub-line on add-in-card is 0.38 cm (1.5 inch) for the 32-bit connector 5.1 cm (2 inch) for the 64-bit extension connector.

With this rule the stub-line and I/O-buffer on a add-in-card act like a lumped capacitor on the bus line. If the stub-line is longer, it has to be treated like a transmission line. This would be much more complicated.

The timing of the PCI-specification requests a maximum propagation delay time on the bus-trace of 10ns for 33MHz systems and 5ns for 66MHz systems. If we assume a signal speed of 25ns/m the maximum bus-length is 20cm for 33MHz and 10cm for 66MHz systems. This restriction does not allow building huge bus systems with one PCI bus.

PCI DESIGN EXAMPLES

Standard PC designs consist of CPU, level1 and level2 cache, the North-Bridge, a PCI-to-ISA bridge, a VGA controller and several PCI slots. According to the PCI rule of max. 10 loads, we have the following options:

VGA on motherboard:

- 1 load - North-Bridge
- 1 load - PCI-to-ISA bridge
- 1 load - VGA controller
- 6 loads - 3 PCI slots

VGA on add-in-card:

- 1 load - North-Bridge
- 1 load - PCI-to-ISA bridge
- 8 loads - 4 PCI slots

Therefore a PCI-PCI bridge is needed to expand the bus to more than 10 loads as well as to allow transmitting PCI Bus data to a secondary PCI-Bus system. As mentioned the max cable/wire lengths of a PCI bus is limited to a 10/20cm depending on bus speed. This leads to the conclusion that PCI is in general not suited to bridge a distance.

The following serialized PCI chipset from Texas Instruments will be looked at in more detail.

The PCI6050 and PCI6060 featuring a PCI-PCI bridge and a high speed Gigabit transceiver enabling connection to a secondary bus as well as bridge a distance up to 15m. It is the industry's first to support the theoretical maximum PCI bus speed of 132 Megabytes per second (MB/s) over inexpensive twisted-pair copper cabling. These devices enable easy expansion of embedded systems using Compact PCI (cPCI), and new applications for personal computers with serialized PCI.

The PCI6050 and PCI6060 are deployed at each end of the cable connecting two PCI buses. The entire configuration behaves as a single system even when the cable connects two or more distinct chassis or boxes. Because the number of serialized PCI signals being




Figure 1

Feature	PCI Dock	Serial CB	USB 2.0	1394.b	Serial PCI
% PCI BW supported	100%	< 76%	< 38%	< 73%	+100%
High-performance video	NO	NO	NO	NO	YES
Special Drivers Required	NO	YES	NO	YES	NO
PCI Expansion/legacy Supported	YES	YES	NO	NO	YES
Number of cable wires	> 70	4-6	2-4	4-6	4-6
Max. Cable Length	NA	+15m	5 m	4.5 m	15 m

Table 1. Comparison: Serialized PCI performance vs. other bus solutions.

transmitted over the connecting cable is reduced, the number of pins in the cable connectors is also reduced, allowing for very small connectors.

The Texas Instruments PCI6050 is a symmetric serialized PCI-PCI bridge node that provides a high-performance solution for connecting two PCI buses via a high-speed serial cable medium. The PCI6050 has two major interfaces: a PCI interface, and a transceiver (PCI6060) interface. The PCI6050 communicates with a PCI bus and through a generic parallel interface to the PCI6060 transceiver. Two separate PCI6050 nodes are required, each with its own PCI6060 serial link, for a complete system. This four-chip solution provides a single, transparent PCI-to-PCI bridge, as illustrated in figure 1.

THE PCI6050 AND PCI-PCI BEHAVIOUR

As a symmetrical device, the PCI6050 can be implemented as either a primary or secondary serial PCI-to-

PCI bridge node. The mode is determined by the level of the PRI/SEC# input terminal. At a high level, a primary PCI6050 node claims downstream transactions that appear on the primary PCI bus which target either the PCI6050-based bridge or any devices located on a target bus that is subordinate to the bridge.

These transactions are forwarded to the secondary PCI6050 node via the serial link so that they may be initiated and completed on the secondary PCI bus. In addition, a primary node will initiate upstream transactions that are forwarded from the secondary PCI bus by the secondary node in a similar manner.

The PCI6050 is a single-function PCI device, and provides an internal PCI bus arbiter supporting up to nine secondary PCI bus masters when configured as a secondary node. The PCI6050 also provides ten CLKOUT outputs that may be supplied to PCI devices by dividing the SYS_CLK input clock by 4 or through an external clock source provided by the system. The PCI6050 supports PCI clock frequencies up to a maximum of 33 MHz.

The PCI6050 serialized PCI-PCI Bridge implementation provides parallel PCI interrupt routing from secondary devices to the primary system. In addition, the PCI6050 supports Compact PCI Hot Swap on the primary interface by passing the Compact PCI Hot-Swap HS_ENUM# signal from the secondary bus to the primary interface, and supports four general purpose inputs and outputs which operate and are controlled independently for each node. A two-wire serial interface is provided for register pre-loading and subsystem identification.

The Texas Instruments PCI6050 symmetric serialized PCI-to-PCI bridge provides a high-performance connection path between two peripheral component interconnect (PCI) buses. Transactions occur between masters on one PCI bus and targets on another PCI bus, and the PCI6050 allows bridged transactions to occur concurrently on both buses.

The PCI6050 pipeline architecture supports burst-mode transfers to maximize data throughput, and the two bus traffic paths through the bridge act independently.

Furthermore a PCI6050-based PCI-to-PCI bridge is compliant with the PCI Local Bus Specification and can be used to overcome the trace length limitations by allowing for PCI-to-PCI implementations to span across separate system boards, as well as extending

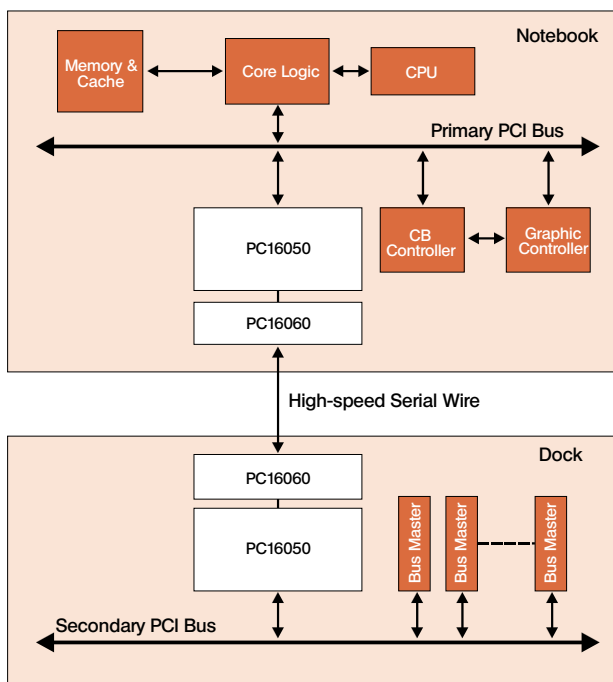


Figure 2. PCI6050 System Diagram.

the electrical loading limits of 10 devices per PCI bus and one PCI device per expansion slot by creating hierarchical buses. The PCI6050 provides two-tier internal arbitration for up to nine secondary bus masters and may be implemented with an external secondary PCI bus arbiter.

THE PCI6050 AND CPCI HOT SWAP

The PCI 6050 integrates the compact PCI hot-swap capability to address the industrial market, which makes this device an ideal solution for multifunction compact PCI cards, remote chassis-to-chassis interconnects, and the adaptation of single function cards for hot-swap compliance.

In addition to the hot-swap friendly behavior, the PCI6050 contains support for software control, and integrates circuitry required by the CPCI Hot-Swap Specification. To be hot-swap capable, the PCI6050 supports the following:

- Compliance with PCI Local Bus Specification
- Tolerance of V_{CC} from early power
- Asynchronous reset
- Tolerance of precharge voltage
- I/O buffers must meet modified V/I requirements
- Limited I/O pin voltage at precharge voltage
- Hot-swap control and status programming via extended PCI capabilities linked list
- Hot-swap terminals: HS_ENUM, HS_SWITCH, and HS_LED

CPCI hot-swap defines a process for installing and removing PCI boards without adversely affecting a running system. The PCI6050 provides this functionality such that it can be implemented on a board that can be removed and inserted in a hot-swap system.

A primary PCI6050 node provides three terminals to support Hot-Swap: HS_ENUM# (output), HS_SWITCH (input), and HS_LED (output). The HS_ENUM# output indicates to the system that an insertion event occurred or that a removal event is about to occur. The HS_SWITCH input indicates that state of a board ejector handle, and the HS_LED output illuminates a blue LED to signal insertion and removal ready status. The primary PCI6050 node Hot Swap functionality is controlled via the Primary CPCI Hot Swap Control and Status Register (offset E6h). The PCI6050 allows for a 2 ms debounce period on HS_SWITCH input to ensure that the state of the input has been stable for about 2 ms following the insertion and removal of board implementing PCI6050, before the state change is reported to the Hot Swap status bits.

When ENUM signal is received at the secondary PCI6050 node HS_ENUM# input pin, the hardware will forward the Hot Swap event to the primary system by way of the serial link.

THE PCI6050 AND GP I/Os

The PCI6050 implements additionally a four-pin general purpose I/O interface that controlled through software for both the primary and secondary nodes. Each node's GPIO terminals are controlled independently

through two separate locations with the PCI6050 configuration space. The primary GPIO control registers begin at PCI offset 65h. The secondary GPIO control registers begin at PCI offset A9h. Setting the GPIOTOI-IC bit in the Diagnostic Control register (Bit 3 of offset 41h) maps the control bits for GPIO1 and GPIO2 to SDA and SCL on the serial ROM interface, respectively. This allows for the system designer to implement in-circuit serial Rom programming through software.

In addition, The PCI6050 will stop forwarding I/O and Memory transactions if bit 5 of the chip control register (offset 40h) is set to 1 and GPIO3 is driven high. The bridge will complete all queued posted writes and delayed requests but delayed completions will not be returned until GPIO3 is driven low and transaction forwarding is resumed. The bridge will continue to accept configuration cycles in this mode.

THE PCI6050 AND POWER MANAGEMENT

The PCI Power Management Interface Specification establishes the infrastructure required to let the operating system control the power of PCI functions. This is done by defining a standard PCI interface and operations to manage the power of PCI functions on the bus. The PCI bus and the PCI functions can be assigned to one of four software visible power management states, which result in varying levels of power savings.

The four power management states of PCI functions are D0 "Fully on" state, D1,D2 "intermediate states" and D3 "Off" state. Similarly, bus power states are B0-B3. The bus power state B0-B3 are derived from the device power state of the originating device. The power state of the secondary bus is derived from the power state of the PCI6050.

For the operating system to power manage the device power states on the PCI bus, PCI functions support four power management operations:

- Capabilities reporting
- Power status reporting
- Setting the power state
- System wake-up

The operating system identifies the capabilities of the PCI function by traversing the new capabilities list. The presence of the new capabilities list is indicated by the PCI status register which provides access to the capabilities list.

The PCI6050 supports D0, D1, D2 and D3Hot power states. The PCI6050 is fully functional only in the D0 state. In the lower power states, the bridge does not accept any I/O or memory transactions. These transactions are master aborted. The bridge accepts type 0 configuration cycles in all power states except D3Cold. The bridge also accepts type 1 configuration cycles but does not pass these cycles to the secondary bus in any of the low power states. Type 1 configuration writes are discarded and type 1 configuration reads return all 1's. All error reporting is done in low power states. When in D2 and D3Hot states, the bridge turns off all secondary clocks for further power savings.

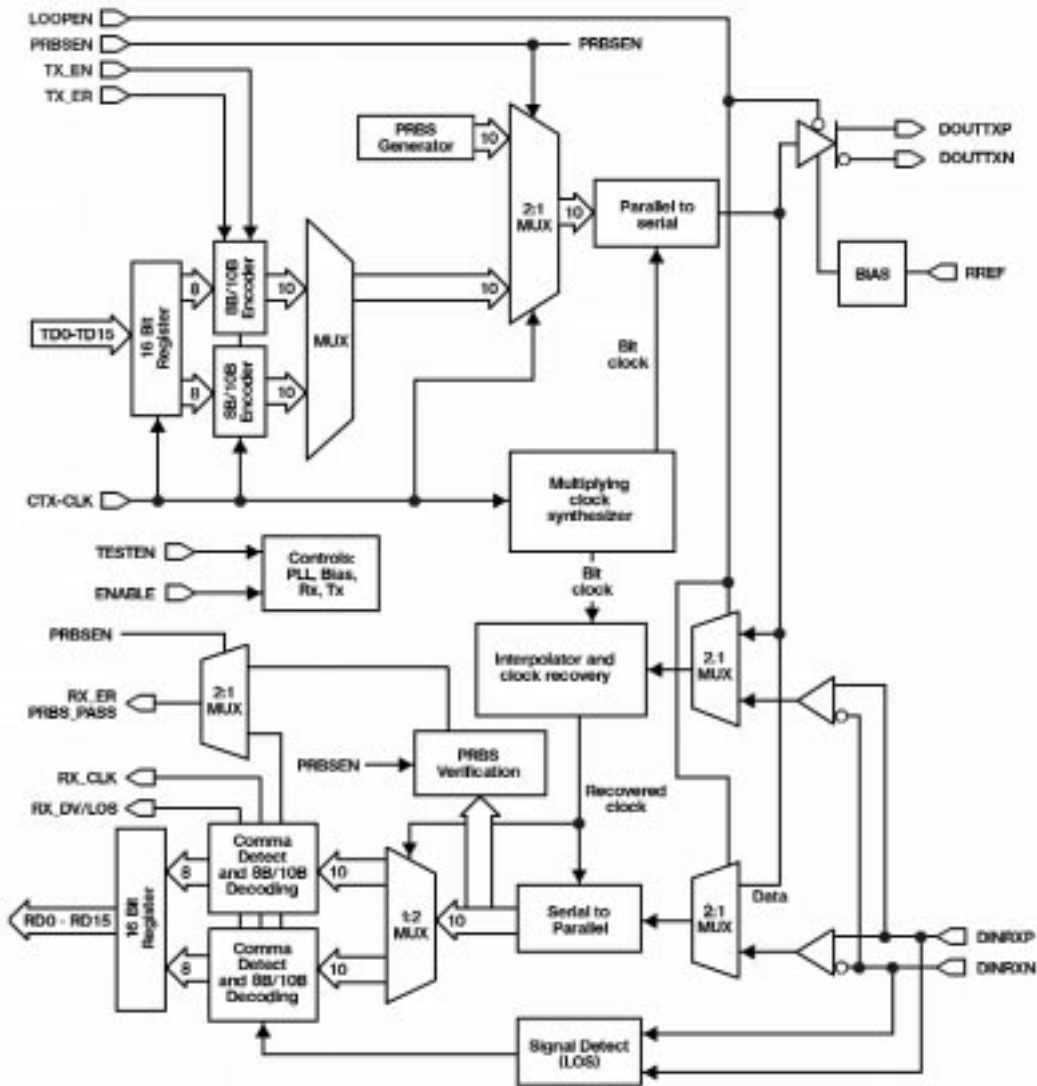


Figure 3. PCI6060 block diagram.

When going from D3Hot to D0, an internal reset is generated. This reset initializes all PCI configuration registers to their default values. All TI extension registers (40h-FFh) are not reset. The power management registers (beginning at offset E0h) are also not reset.

THE PCI6060 SERIALIZER/DESERIALIZER

Function

To enable to PCI bus to transmit data over distance the PCI6050 is converting the PCI data to a 16 bit wide interface format and send them to an Serializer/Deserializer device. As multi-gigabit transceiver the PCI6060 from Texas Instruments is used providing ultra high-speed input/output (I/O) data channels for point-to-point baseband data transmission over a variety of media types. The PCI6060 supports a serial interface speed of 1.6 Gbps to 2.5 Gbps, providing up to 2.0 Gbps of actual data bandwidth in each direction. The transceiver also can be used to

replace parallel data transmission architectures by providing a reduction in the number of traces, connector pins and transmit/receive pins. 16 bit parallel data loaded into the transmitter is delivered to the receiver over a serial channel, which can be a coaxial copper cable, a controlled impedance backplane or an optical link.

The data is then reconstructed into its original parallel format. The parallel interface supports the GMII (Gigabit Media Independent Interface) defined in IEEE P802.3z with extensions for data rates up to 2 Gbps. The interface provides a simple, inexpensive and easy way to implement interconnection between the media access control (MAC) sublayer and physical layer (PHY). The IEEE P802.3z GMII interface provides independent 8-bit wide transmit and receive data paths for full-duplex operation. The PCI6060 extends the standard GMII transmit and receive data paths to 16 bits, thus allowing for 2-Gbps operation at the same parallel-side clock speed.

Transmitter/Receiver detailed

The transmitter latches 16-bit parallel data at a rate based on the supplied reference clock (GTX_CLK). The 16-bit parallel data is internally encoded into 20 bits using an 8-bit/10-bit (8B/10B) encoding format. The resulting 20-bit word is then transmitted differentially at 20 times the reference clock (GTX_CLK) rate. The receiver section performs the serial-to-parallel conversion on the input data, synchronizing the resulting 20-bit wide parallel data to the extracted reference clock (RX_CLK). It then decodes the 20 bit wide data using 8-bit/10-bit decoding format resulting in 16 bits of parallel data at the receive data terminals (RXD0-15). The outcome is an effective data payload of 1.28 Gbps to 2.0 Gbps (16 bits data x the GTX_CLK frequency).

EXAMPLES

Several new applications will be enabled by implementing serialized PCI into different platforms. For example:

- Notebook docking
- Small-form-factor PC expansion
- Industrial/Telecom ■

Dipl.-Ing. Peter Peisker studied electronic engineering at the Rheinischen FH university of Cologne, and is responsible for PCI and IEEE1394 strategic product marketing in Europe. Prior to joining TI, he worked in the area of sales and project management of sensors systems and industrial controls for numerous years.

Furthermore Mr. Peisker worked for a couple of years as technical leader responsible for development of electronic end equipment and industrial controls.

SOURCES

- Texas Instruments, PCI6050/6060, data manual and ASPEN specification 2000

ADVERTISEMENT INDEX

QNX SOFTWARE SYSTEMS	2
ESMERTEC	11
EMBEDDED SYSTEMS 2001	15
PHARLAP	27
RTS 2001 - PARIS	31
POLYHEDRA	43
VMIC	47
IMAGE MEDIA	51
ENEA OSE SYSTEMS AB	55
LINEO	63
FORCE COMPUTERS	67
NATIONAL INSTRUMENTS	71
EONIC	83
ACT EUROPE	99
VMETRO	100

For more information on advertisement opportunities, call +32-2-520.55.77 or send an e-mail to info@dedicated-systems.com

EDITORIAL CALENDAR

1Q 01 Development methodologies & tools
(UML, lower and higher CASE tools, New compiler technologies, etc.)

2Q 01 Chips Issue
(Soc, DSP, Processors, Controllers, Debugging Tools,...)

Special Edition
Applications & Dedicated Software

3Q 01 RTOS Update
(Freeware, standardisation issues, safety-critical RTOS,...)

4Q 01 System Architectures
(SoC, Software Intensive Systems, Complexity issues, Networking...)