
IDT[®]

Overview of PCI/X Bridging Techniques

80B6000_AN002_05

September 19, 2009

6024 Silver Creek Valley Road San Jose, California 95138

Telephone: (408) 284-8200 • FAX: (408) 284-3572

Printed in U.S.A.

©2009 Integrated Device Technology, Inc.

GENERAL DISCLAIMER

Integrated Device Technology, Inc. ("IDT") reserves the right to make changes to its products or specifications at any time, without notice, in order to improve design or performance. IDT does not assume responsibility for use of any circuitry described herein other than the circuitry embodied in an IDT product. Disclosure of the information herein does not convey a license or any other right, by implication or otherwise, in any patent, trademark, or other intellectual property right of IDT. IDT products may contain errata which can affect product performance to a minor or immaterial degree. Current characterized errata will be made available upon request. Items identified herein as "reserved" or "undefined" are reserved for future definition. IDT does not assume responsibility for conflicts or incompatibilities arising from the future definition of such items. IDT products have not been designed, tested, or manufactured for use in, and thus are not warranted for, applications where the failure, malfunction, or any inaccuracy in the application carries a risk of death, serious bodily injury, or damage to tangible property. Code examples provided herein by IDT are for illustrative purposes only and should not be relied upon for developing applications. Any use of such code examples shall be at the user's sole risk.

Copyright © 2009 Integrated Device Technology, Inc.
All Rights Reserved.

The IDT logo is registered to Integrated Device Technology, Inc. IDT is a trademark of Integrated Device Technology, Inc.

1. Overview of PCI/X Bridging Techniques: Transparent, Non-transparent, and Opaque

PCI-X is the next-generation, PCI-based bus protocol for the embedded systems industry. It builds on PCI, the proven multi-drop bus protocol that has been used in PC and server applications for over ten years. This document explains how PCI-X and the IDT Tsi310 create new bridging techniques for systems designers. It discusses the following topics:

- “Brief History of Bus Bridging: PCI to PCI-X” on page 3
- “PCI/X Bridging Techniques” on page 4
 - “Transparent Bridge” on page 4
 - “Non-Transparent Bridge” on page 5
 - “Opaque Bridge” on page 6
- “Typical Applications” on page 7
 - “Host Board and Embedded Board” on page 7
 - “Adapter Board” on page 9
 - “Intelligent Adapter Board” on page 10
- “Overview of the Tsi310” on page 11

1.1 Brief History of Bus Bridging: PCI to PCI-X

The PCI Special Interest Group (PCI-SIG) is the governing standard body for managing the PCI protocol. In 1992, the first PCI specification was created. The PCI standard allowed multiple semiconductor vendors to converge on a single interface that provided device interoperability. It was designed primarily for the desktop PC market; however, it proliferated throughout other market segments soon after its launch as developers began leveraging this new high-speed, multi-drop bus to interconnect their devices.

The multi-drop bus architecture resulted in the emerging requirement for PCI bridges. In order to ensure interoperability in the desktop market, a PCI-to-PCI bridge specification was created. This specification became the basis of today’s transparent PCI-to-PCI bridge, which defines a contiguous memory map across both primary and secondary PCI segments.

As the PCI interface became ubiquitous for I/O systems, it found its way into the embedded communications market. The embedded market had greater need for scalability and density than the desktop platform. The PCI bus in an embedded backplane could support a large number of peripheral cards. This posed a problem, however, since several devices that required large amounts of addressable memory could easily consume the PCI memory space. This problem was addressed with the creation of the non-transparent PCI bridge, which allowed for independent processor sub-systems within a PCI host centric system.

In 1999, the PCI-X specification was released by the PCI-SIG. PCI-X is a faster and more efficient multi-drop architecture that is backward compatible with PCI. PCI-X increased the bus frequency to 133 MHz and added transaction types to the protocol, of which the split transaction is probably the best known and most important. This increase in bandwidth addressed the increasing I/O speed found on the PCI-X endpoints. These PCI-X I/O devices were primarily used with either a transparent bridge as previously defined, or with an opaque bridge. The opaque bridge provided PCI-X I/O device isolation, similar to the non-transparent bridge, but still contained transparent bridge features.

1.2 PCI/X Bridging Techniques

As previously explained, the PCI standard has evolved to offer several bridging options for designers of multi-drop systems. The following sections explain three of the most common PCI/X bridging techniques: transparent, non-transparent, and opaque.



PCI/X refers to both PCI and PCI-X.

1.2.1 Transparent Bridge

A transparent bridge is commonly used on system controller cards ([Figure 1 on page 8](#)) or embedded boards ([Figure 2 on page 9](#)) where a contiguous memory map across both primary and secondary PCI/X segments is sufficient. It is also used in I/O adapter boards that do not contain a local CPU and require the host to initialize the I/O devices behind the bridge. The original function of the bridge was to provide greater fan-out of I/O across PCI. However, the multiple revisions of the PCI specification and various link speeds have resulted in this bridge often being used as a mechanism to bridge two different PCI/X bus speeds in an application.

The main technical features of a transparent bridge include the following:

- Address memory decodes use base and limit address registers to decode downstream traffic from the primary PCI/X interface to the secondary PCI/X interface. A second set of base and limit registers is used to move data from the secondary PCI/X interface to the primary PCI/X interface.

- Primary PCI/X bus initialization occurs from the BIOS bus scan of the PCI/X interface using Type 0 configuration cycles. Devices on the immediate bus segment respond to Type 0 configuration commands when a device's IDSEL pin is asserted.
- Secondary PCI/X bus initialization occurs from the BIOS bus scan of the PCI/X interface using Type 1 configuration bus cycles. These cycles traverse the bridge to the secondary bus segments where they appear as Type 0 configuration cycles, initializing all devices on that segment when the specified device's IDSEL pin is asserted. The BIOS bus scan will continue scanning all IDSEL lines until every device is identified.

1.2.2 Non-Transparent Bridge

A non-transparent bridge is commonly used on intelligent I/O adapter boards where isolation of the local PCI/X processor or PCI/X chipset and the local PCI/X I/O devices are required (see [Figure 3 on page 10](#)). System architectures designed with a non-transparent bridge result in completely independent address memory spaces on the primary and secondary sides of the bridge. Independent address memory maps result in the need for greater register memory space, beyond the 256 bytes of configuration memory space assigned by the PCI Bridge Specification.

The main technical features of a non-transparent bridge include the following:

- Address memory decodes use base address registers on both the primary and secondary bus segments, similar to any PCI/X endpoint, to decode transactions passed through the bridge. The target address on the destined bus is determined through device specific translation registers.
- Primary PCI/X bus initialization occurs from the BIOS bus scan of the PCI/X interface using Type 0 Configuration cycles. Devices on the immediate bus segment respond to Type 0 configuration commands when a device's IDSEL pin is asserted. Initialization of registers beyond the 256 bytes of configuration space is accessed through a memory address command.
- Secondary PCI/X bus initialization occurs from either the local CPU on the intelligent adapter card, or by the host itself using a custom OS software driver providing access to a memory space register within the non-transparent bridge that can generate configuration cycles on the secondary PCI/X bus segment.

1.2.3 Opaque Bridge

An opaque bridge is the latest bridging technique for PCI/X designers. This type of bridge contains attributes from both transparent bridges (for example, address decoding using base and limit address registers) and non-transparent bridges (for example, device initialization on primary and secondary buses).



The Tsi310 was the first bridge to support opaque mode and PCI-X. It was developed by IBM, later acquired by Tundra Semiconductor, and subsequently Tundra was acquired by IDT.

Opaque bridging was created to support specialized functionality contained in PCI/X chipsets or native PCI/X processors not typically found in PCI/X-to-PCI/X bridges. It provides the host CPU with a transparent view into the functionality of PCI chipsets or native PCI/X processors on the secondary segment, while hiding from host specific PCI I/O devices on the same secondary segment.

An opaque bridge is commonly used on intelligent I/O adapter boards where host CPU visibility of the local PCI processor or PCI chipset, and the isolation of local PCI I/O devices, are required (see [Figure 4 on page 11](#)).

The main technical features of an opaque bridge include the following:

- Address memory decodes use base and limit registers to move data from the secondary PCI/X interface to the primary PCI/X interface, with the exception of the defined opaque region. The opaque region on the secondary PCI/X interface is determined through a base and limit register. Any memory address in that region is not forwarded upstream and is decoded by the local PCI/X I/O devices on the intelligent adapter.
- Address memory decodes use base and limit address registers to decode downstream traffic from the primary PCI/X interface to the secondary PCI/X interface.
- Primary PCI/X bus initialization occurs from the BIOS bus scan of the PCI/X interface using Type 0 configuration cycles. Devices on the immediate bus segment respond to Type 0 configuration commands when the device's IDSEL pin is asserted.
- Secondary PCI/X bus initialization occurs from the BIOS bus scan of the PCI/X interface using Type 1 configuration bus cycles that traverse the bridge to the secondary bus segments as Type 0 configuration cycles. This, in turn, initializes all devices on that segment when the appropriate IDSEL pin is asserted. IDSEL pins of devices within the opaque region are masked to prevent identification and initialization by the host CPU. Instead, the local CPU on the secondary segment is responsible for initialization of these devices (much like a non-transparent bridge). The BIOS bus scan will continue scanning the IDSEL lines until every device that is accessible is identified.

1.3 Typical Applications

This section describes typical applications for transparent, non-transparent, and opaque PCI/X bridges.



The Tsi310 can support two to three loads operating at PCI-X 133 MHz in a system that has an optimized layout and well designed buffers. This capability exceeds the recommendation of the *PCI-X Addendum to PCI Local Bus Specification (Revision 1.0a)*, which states a PCI-X bus is capable of supporting only one 133-MHz load.


1.3.1 Host Board and Embedded Board

A host board is typically the principle card in a system that contains a CPU; it performs initialization and control for the system. This board can take the form of a motherboard or, in the case of a backplane application, a host system controller board. A motherboard is typically found in PC and server applications, while a host system controller is found in compactPCI or rack-mount systems.

An embedded board is a specialized board that performs a single function. A PCI/X bridge is often used in an embedded application to either increase the fan out, or to separate high-performance ASICs from legacy PCI in order to prevent the de-rating of the PCI-X bus speed.

Host and embedded boards use transparent PCI/X bridges since a flat addressing model is required for the CPU to perform full initialization for the system (see [Figure 1](#)).


Integrated Device Technology
www.idt.com



1.3.2 Adapter Board

An adapter/expansion board is used to extend the I/O capability of the host board or add I/O functionality not contained on the host board. Common form factors for adapter cards include standard PCI/X cards, PCI Mezzanine cards (PMC), and compactPCI cards. Since devices on the secondary interface of an adapter card are non-intelligent I/O devices, they require initialization by the host CPU. The simplest and most common way to do this is to use a PCI/X bridge that supports transparent addressing, such as the Tsi310 (see [Figure 2](#)). The transparent bridge results in a flat address between the I/O devices on the adapter card and the host CPU.

Figure 2: PCI/X Adapter with Transparent Bridge — Flat Address Map



1.3.3 Intelligent Adapter Board

An intelligent adapter board is used to extend the I/O capability of the host board and off load the processing tasks required for the I/O devices (see [Figures 3 and 4](#)).

In many cases, intelligent adapter boards must be initialized prior to the loading of the OS by the host. A RAID adapter is a good example of this. In order to perform this initialization, the BIOS uses an Expansion ROM, which is ROM located on the adapter card itself. In the case of a transparent or opaque bridge, the Expansion ROM can reside on the secondary PCI bus segment as part of the local processor's memory space in the transparent region so that the host has access to it. In the case of a non-transparent bridge, however, the Expansion ROM is often attached to the bridge itself, resulting in the need for additional pins on the bridge.

Figure 3: PCI/X Intelligent Adapter with Non-transparent Bridge — Independent Address Map





Figure 4: PCI/X Intelligent Adapter with Opaque Bridge — Shared and Independent Memory Map



1.4 Overview of the Tsi310

The IDT Tsi310 is the most mature PCI-X bridge on the market and has undergone rigorous and lengthy qualification testing on numerous customer designs. It is a 64-bit PCI-X bus bridge that operates at speeds up to 133 MHz, and supports transfer rates up to 1 Gbps. The PCI-X protocol is backward compatible with the PCI 2.2 bus standard ensuring that legacy PCI-based systems are portable to the faster PCI-X environment.

The Tsi310 connects two electrically separate PCI-X bus domains, allowing concurrent operations on both buses. This results in optimal use of the buses in various system configurations and enables hierarchical expansion of I/O bus structures. The device supports configurations of PCI or PCI-X modes on either bus, and in any combination.

The Tsi310 is used in numerous storage applications, and is the *de facto* standard for transparent and opaque mode PCI-X bridging. It is the only PCI-X bridge on the PCI-SIG Integrators' List, and the only bridge that can claim PCI-SIG compliance to the PCI-X 1.0b standard.



CORPORATE HEADQUARTERS
6024 Silver Creek Valley Road
San Jose, CA 95138

for SALES:
800-345-7015 or 408-284-8200
fax: 408-284-2775
www.idt.com

for Tech Support:
email: ssdhelp@idt.com
phone: 408-284-8208
document: 80B6000_AN002_05