



Performance Comparison of IDT Tsi381 and PLX PEX8111

80E2000_AN001_04

October 1, 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. Performance Comparison of IDT Tsi381 and PLX PEX8111

This report compares the IDT Tsi381 versus the PLX PEX 8111-BB and PEX 8112-AA. It highlights the performance advantages of using the Tsi381 over the two PLX devices.

This document discusses the following:

- “Throughput Measurements”
- “Latency Measurements”

Terms

- Upstream transaction – In the context of a PCIe-to-PCI bridge, this transaction flow starts on a PCI bus and ends on a PCIe link.
- Downstream – In the context of a PCIe-to-PCI bridge, this transaction flow starts on a PCIe link and ends on a PCI bus.
- Latency – The time required for a transaction to pass from one side of a bridge to another. The method of measurement depends on the type of transaction.

Revision History

80E2000_AN001_04, Formal, October 2009

This document was rebranded as IDT. It does not include any technical changes.

80E2000_AN001_03, Formal, April 2008

This document includes throughput measurements.

80E2000_AN001_02, Formal, November 2006

This document was updated to include 32- and 64-byte transaction test results.

80E2000_AN001_01, Formal, November 2006

This is the first version of this document.

1.1 Throughput Measurements

This section consists of a lab analysis comparison between the Tsi381 and the PLX 8112, as well as simulation throughput analysis of the Tsi381.

1.1.1 Lab Throughput Analysis

This section compares throughput measurements of the IDT Tsi381 and the PLX PEX 8112-AA. Please note that the throughput tests do not include bidirectional traffic. Each test is of a single direction of a particular size and type of transaction.

Unless noted, the default register settings for the devices were used.

1.1.1.1 Throughput Measurement Method

The test setup for the throughput measurements is described in “[Throughput and Latency Test Setup](#)”.

PCIe Performance Measurement Method

Throughput was monitored using the Agilent E2960 Protocol Tester Realtime Statistics display. This display measures card performance in megabytes per second (MBps).

The actual measurement was taken from a TCL script provided by Agilent called PerformancePrint. This script was a text list of sequential measurements reported every second for 10 seconds.



1.1.1.2 PCI Upstream Reads without Short-Term Caching

Table 1: PCI Upstream Reads without Short-Term Caching

| Transaction Type | Payload Size (Bytes) | PCI Bus Speed (MHz) | Tsi381 Throughput (MB/s) | PLX 8112 Throughput (MB/s) | IDT Performance Improvement | Footnote |
|-----------------------|----------------------|---------------------|--------------------------|----------------------------|-----------------------------|----------|
| PCI Mem Read Multiple | 4096 | 66 | 85.7 | 46.3 | 85.4% | a |
| | 2048 | 66 | 85.5 | 52.0 | 64.3% | a |
| | 1024 | 66 | 85.4 | 51.9 | 64.5% | a |
| | 512 | 66 | 80.3 | 50.3 | 59.7% | a |
| | 256 | 66 | 53.2 | 47.4 | 12.3% | |
| | 128 | 66 | 36.1 | 31.6 | 14.1% | |
| | 64 | 66 | 20.5 | 16.8 | 22.0% | |
| | 32 | 66 | 10.7 | 8.7 | 22.9% | |
| | 16 | 66 | 5.5 | 4.2 | 29.8% | |
| | 8 | 66 | 2.8 | 2.2 | 25.9% | |

Note a) The Tsi381 Registers were configured as follows during this test: Prefetch Control Register (offset BC) = 0x03FFFFFF.

Summary

The Tsi381 performance improvement over the 8112 increases according to the payload size.

1.1.1.3 PCI Upstream Reads with Short-term Caching

Table 2: PCI Upstream Reads with Short-Term Caching

| Transaction Type | Payload Size (Bytes) | PCI Bus Speed (MHz) | Tsi381 Throughput (MB/s) | PLX 8112 Throughput (MB/s) | IDT Performance Improvement | Footnote |
|-----------------------|----------------------|---------------------|--------------------------|----------------------------|-----------------------------|----------|
| PCI Mem Read Multiple | 4096 | 66 | 85.4 | 46.3 | 84.6% | a |
| | 2048 | 66 | 86.6 | 52.0 | 66.4% | a |
| | 1024 | 66 | 86.6 | 51.9 | 67.0% | a |
| | 512 | 66 | 86.7 | 50.3 | 72.4% | a |
| | 256 | 66 | 86.7 | 47.4 | 83.0% | a |
| | 128 | 66 | 86.8 | 31.6 | 174.7% | a |
| | 64 | 66 | 59.3 | 16.8 | 253.1% | a |
| | 32 | 66 | 32.9 | 8.7 | 278.2% | a |
| | 16 | 66 | 17.4 | 4.2 | 314.8% | a |
| 8 | 66 | 8.9 | 2.2 | 306.4% | a | |

Note a) The Tsi381 Registers were configured as follows during this test:

- Prefetch Control Register (offset 0x0BC) = 0x03FFFFFF.
- PCI Miscellaneous Control and Status Register (offset 0x044) = 0x7D9F_1900 (this sets the Short Term Caching Enable bit).

Summary

The Tsi381's short-term caching feature can cause noticeable performance improvements when sequential transfers of a small payload size are required.

1.2 Tsi381 Simulation Throughput Analysis

The Tsi381's simulation throughput was measured in the upstream and downstream directions for both read and write transactions. The results for both measurements are detailed in the following sections.

1.2.1 Simulation

1.2.1.1 Test Setup

The test results were derived from simulation. The Tsi381's simulation environment consists of the Tsi381 device with bus functional models (BFMs) on both the PCI and PCI-e interfaces. The BFMs were used to initiate read and write transactions through the Tsi381, as well as provide an ideal target response. The PCI bus frequency was equal to 66 MHz for all tests.

1.2.1.2 Test Results

Table 3: Upstream Writes

| PCI Burst Size (bytes) | Maximum Sustained Throughput (Mbytes/s) |
|------------------------|---|
| 32 | 147.1 |
| 64 | 183.3 |
| 128 | 208.9 |
| 256 | 206.2 |

Table 4: Upstream Reads

| PCI Burst Size (bytes) | Maximum Sustained Throughput (Mbytes/s) |
|------------------------|---|
| 32 | 128.7 |
| 64 | 159.3 |
| 128 | 162.2 |
| 256 | 151.4 |

Table 5: Downstream Writes

| PCI-e Data Payload Size (bytes) | Maximum Sustained Throughput (Mbytes/s) |
|---------------------------------|---|
| 32 | 144.8 |
| 64 | 179.9 |
| 128 | 197.7 |

Table 6: Downstream Reads

| PCI-e Data Payload Size (bytes) | Maximum Sustained Throughput (Mbytes/s) |
|---------------------------------|---|
| 32 | 110.9 |
| 64 | 159.4 |
| 128 | 193.9 |
| 256 | 197.9 |

1.3 Latency Measurements

This section compares latency measurements of the IDT Tsi381 and the PLX PEX 8111.

1.3.1 Latency Test Results

Table 7: Latency Test Results

| Transaction ^a | | | Device | | | IDT Performance Improvement |
|--------------------------|------------|-----------|------------|--------------|------------|-----------------------------|
| Type | Direction | Size | IDT Tsi381 | PLX PEX 8111 | Difference | |
| PCI posted write | Upstream | 128 bytes | 696 ns | 945 ns | 249 ns | 26% |
| | | 64 bytes | 456 ns | 725 ns | 269 ns | 37% |
| | | 32 bytes | 336 ns | 589 ns | 253 ns | 43% |
| | | 8 bytes | 248 ns | 485 ns | 237 ns | 49% |
| PCIe posted write | Downstream | 128 bytes | 904 ns | 934 ns | 30 ns | 3% |
| | | 64 bytes | 632 ns | 666 ns | 34 ns | 5% |
| | | 32 bytes | 499 ns | 552 ns | 53 ns | 10% |
| | | 8 bytes | 424 ns | 465 ns | 41 ns | 9% |
| PCI read request | Upstream | 128 bytes | 200 ns | 474 ns | 274 ns | 58% |
| | | 64 bytes | 200 ns | 472 ns | 272 ns | 58% |
| | | 32 bytes | 200 ns | 474 ns | 274 ns | 58% |
| | | 8 bytes | 200 ns | 461 ns | 261 ns | 57% |
| PCIe read request | Downstream | 128 bytes | 364 ns | 445 ns | 81 ns | 18% |
| | | 64 bytes | 364 ns | 436 ns | 72 ns | 17% |
| | | 32 bytes | 364 ns | 433 ns | 69 ns | 16% |
| | | 8 bytes | 364 ns | 438 ns | 74 ns | 17% |

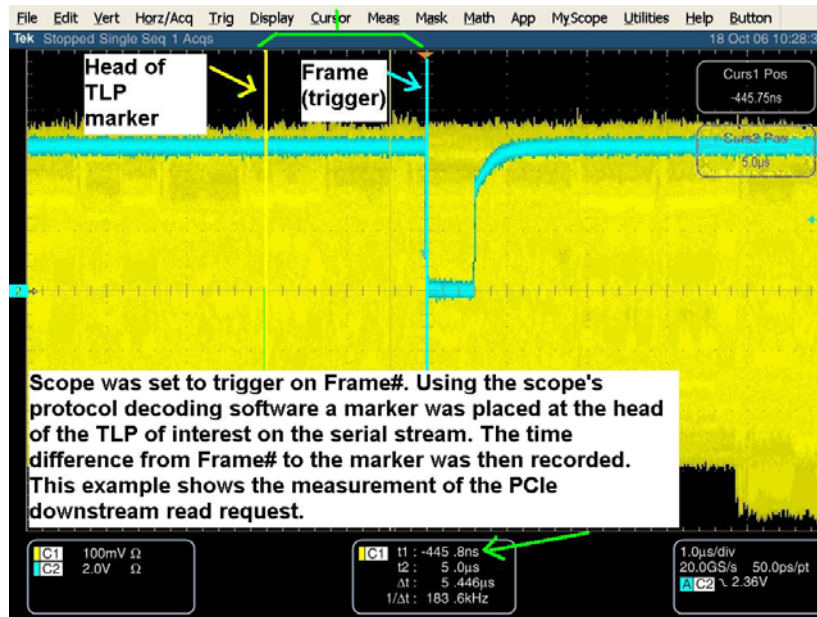
Table 7: Latency Test Results (Continued)

| Transaction ^a | | | Device | | | IDT Performance Improvement |
|--------------------------|-------------------------|-----------|------------|--------------|------------|-----------------------------|
| Type | Direction | Size | IDT Tsi381 | PLX PEX 8111 | Difference | |
| PCI read completion | Upstream ^b | 128 bytes | 743 ns | 915 ns | 172 ns | 19% |
| | | 64 bytes | 500 ns | 681 ns | 181 ns | 27% |
| | | 32 bytes | 380 ns | 564 ns | 184 ns | 33% |
| | | 8 bytes | 295 ns | 473 ns | 178 ns | 38% |
| PCIe read completion | Downstream ^c | 128 bytes | 795 ns | 967 ns | 172 ns | 18% |
| | | 64 bytes | 548 ns | 683 ns | 135 ns | 20% |
| | | 32 bytes | 417 ns | 488 ns | 71 ns | 15% |
| | | 8 bytes | 315 ns | 465 ns | 150 ns | 32% |

- All transactions were measured based on a x1 PCIe link and a 66-MHz PCI bus configuration.
- Upstream Completion with Data for Downstream non-posted request.
- Downstream Completion with Data for Upstream non-posted request.

1.3.2 Latency Test Cases

This section presents latency measurements for six test cases of the PEX 8111-BB. The following figure explains how the latency measurements were made for the test cases.



The measurements were made by capturing the serial PCIe stream and triggering on the FRAME# signal. The serial stream (upstream or downstream) is decoded by post processing the data using the serial protocol decoding software on board the scope. A marker is placed at the head of the TLP of interest, and the latency measurement is made by measuring the time from the marker to the trigger.

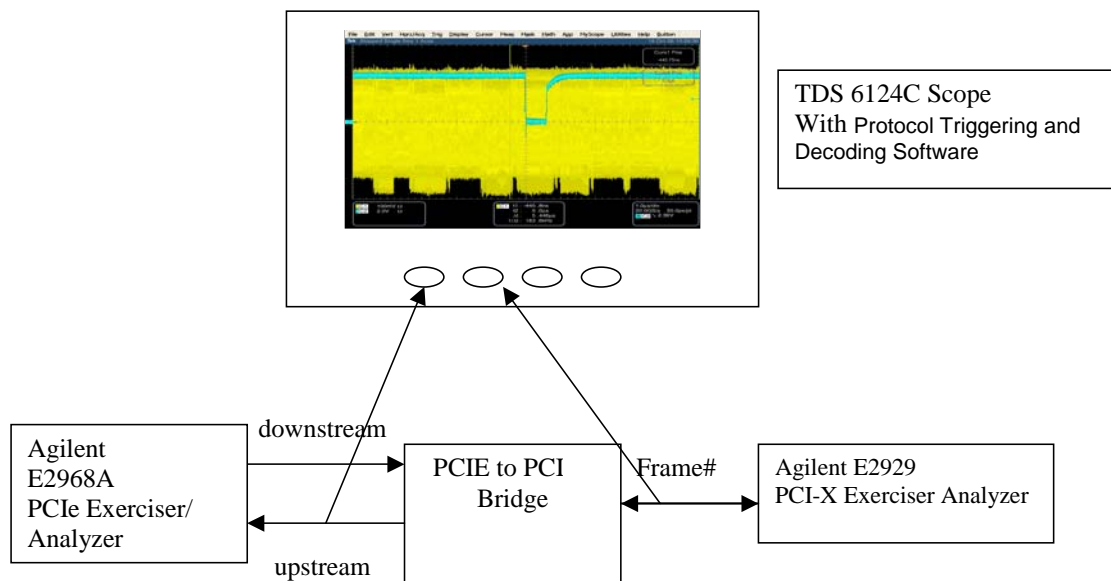
The trigger out from the Agilent Protocol Analyzer was not used for latency measurement due to its long trigger delay and its 64 ns uncertainty for the x1 lane width.



The Agilent E2928 PCI exerciser analyzer was used to measure transactions on the PCI bus.


1.3.2.1 PCI Posted Write (Upstream)

This measurement was made from the assertion of FRAME# on the PCI bus to the head of the upstream TLP.



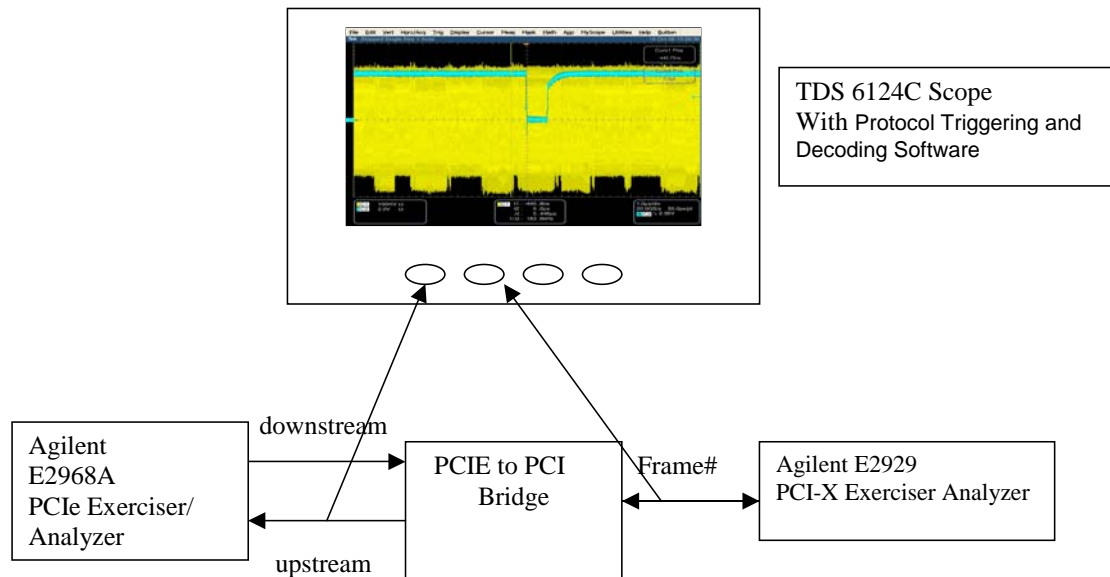
1.3.2.2 PCIe Posted Write (Downstream)

This measurement was made from the head of the downstream TLP to the assertion of FRAME# on the PCI bus.



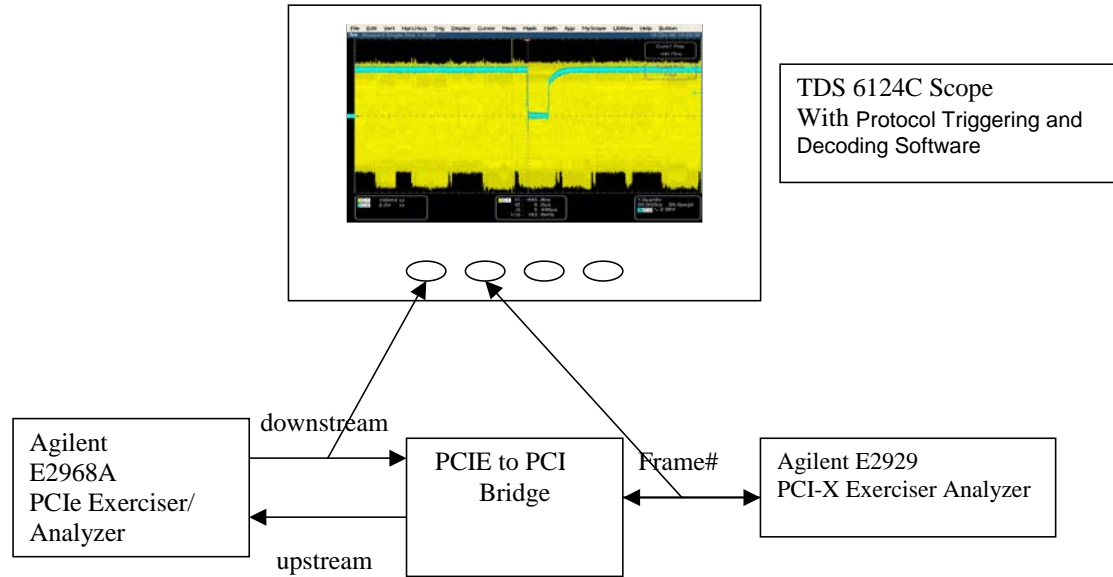
1.3.2.3 PCI Read Request (Upstream)

Initial Read request latency for PCI was measured from the assertion of FRAME# to the head of the upstream TLP.



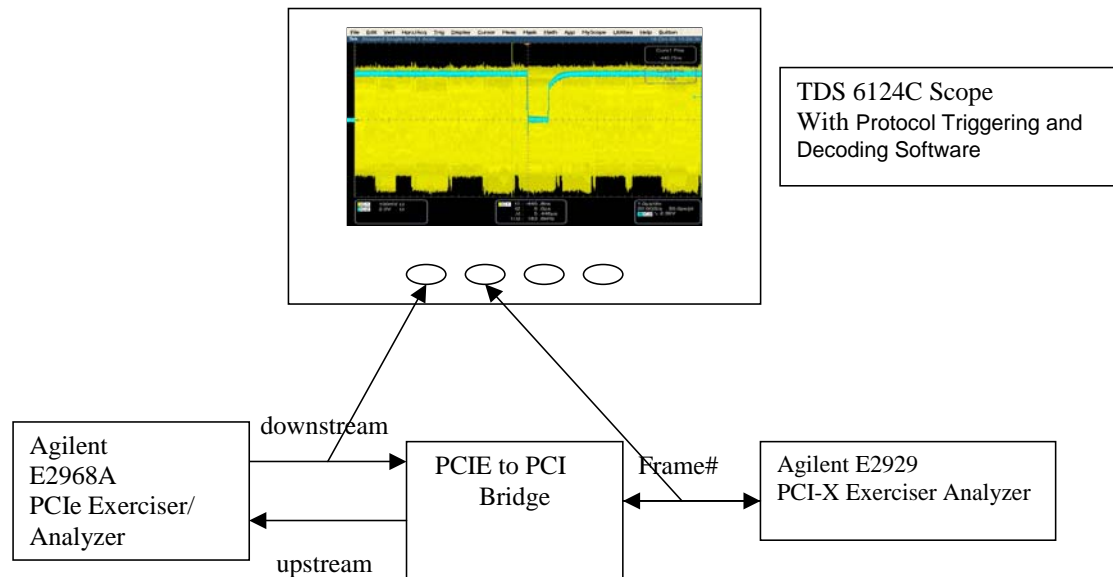
1.3.2.4 PCIe Read Request (Downstream)

Initial Read request latency for PCIe was measured from the head of the downstream TLP to the assertion of FRAME#.



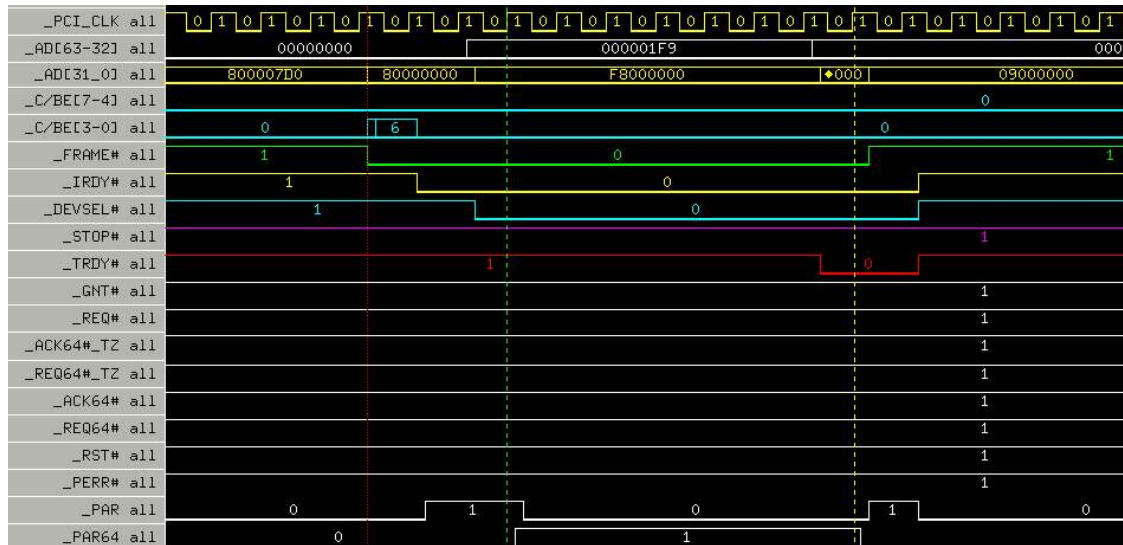
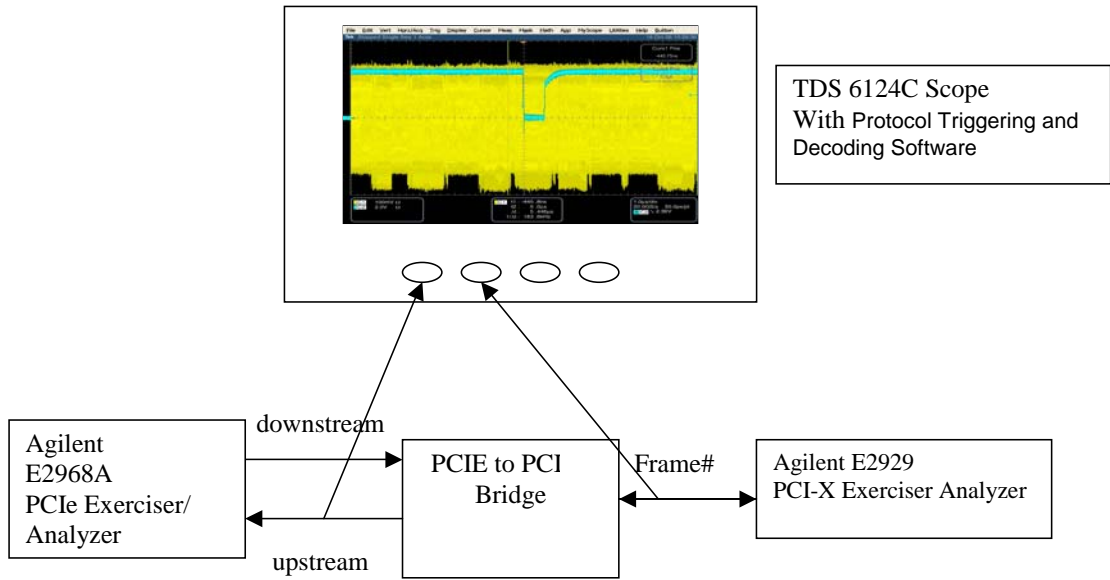
1.3.2.5 PCI Read Completion (Upstream)

This measurement was made from the starting character of the downstream TLP (data return) to the de-assertion of FRAME# on the PCI bus.



1.3.2.6 PCIe Read Completion (Downstream)

This measurement was made from the de-assertion of FRAME# on the PCI bus when data is being driven into the PEX 8111-BB until the starting character of the upstream TLP of data being driven out of the bridge. During the read, the Agilent PCI target was slow to respond on the initial read (this delay does not occur on subsequent reads but the measurement was taken on the first read). As seen in the second PCI waveform signal capture below, the 10 clocks of latency (delayed TRDY#) are present. This was removed from the measurement in order to not penalize the bridge under test. The target measurement is adjusted as if the target was responding as fast as the PCI protocol allows. As per the PCI spec, “The first data phase on a read transaction requires a turn around cycle (enforced by the target through TRDY#).”

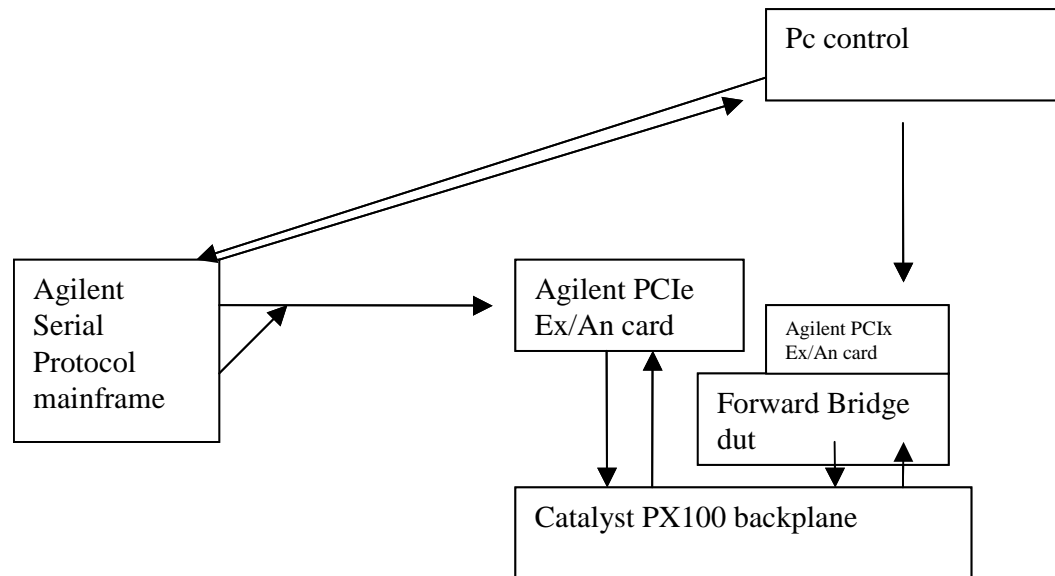


1.3.3 Throughout and Latency Test Setup

1.3.3.1 Test Environment for Throughput Measurements and PLX 8111 Latency Measurements

This section describes the hardware test environment.

System Setup



- A Catalyst two-slot backplane with model number PX100 was used.
- The appropriate (Tsi381, 8111, 8112) evaluation board was connected directly into slot 1 of the Catalyst backplane.
- The Agilent E2928 PCI Exerciser/Analyzer card was plugged into the top slot of the evaluation board.
- An ATX power supply provided power to the PCI bus through a connector on the evaluation board.
- The Agilent PCIe Exerciser/Analyzer card was connected to slot 2 of the Catalyst backplane.
- The Catalyst backplane supplied power and reset to both boards, as well as clocking to the evaluation board.
- The Agilent PCIe Exerciser/Analyzer clock source was internal.
- The Agilent serial protocol mainframe was connected to the Agilent PCIe Exerciser/Analyzer card through the E2942A single probe Y-cable. This cable allowed the simultaneous use of one active probe board for the exerciser and analyzer (using two I/O modules).
- The Agilent Exerciser/Analyzer software operated on a Control PC. The PC was connected to the Agilent PCIe hardware through Ethernet.


- The Agilent PCI exerciser analyzer was connected to the Control PC through its proprietary fast-bus interface. The Control PC operated the Exerciser/Analyzer control software.
- The Control PC provided full control of stimulus and response of the PCIe and PCI sides of the bridge.

Tektronix Scope



The latency measurements were made using a Tektronix TDS6124C oscilloscope. This scope has a 12-GHz bandwidth, 20-GSps sample rate, and 32 MB of storage on each of its four channels. It also has the Protocol Triggering and Decoding Software, which can easily decode the 8b/10b serial data streams, and set serial pattern markers after capture.

Complete Test Setup



1.3.3.2 Tsi381 Test Environment

All tests performed on the Tsi381 were based on Verilog simulation using PCIe and PCI bus functional models. The test bench setup consists of a PCIe bus functional model (BFM) on the primary side of the Tsi381, which acts as a Root Complex; and four PCI BFMs on the secondary side, which act as four different PCI devices. The PCIe BFM on the primary side of the Tsi381 generates TLPs on the link, while the PCI BFM generates all PCI transactions on the secondary side.

In the simulations, the downstream measurements were made from the STP (start of the TLP) symbol on the PIPE (PHY to PCIe) interface to the assertion of FRAME# on the PCI bus, while the upstream measurements were made from the assertion of FRAME# on the PCI bus to the STP symbol on the PIPE interface.

To these latency numbers, we added the latency numbers for the SerDes PHY Interface to obtain the overall Tsi381 latency number. This method of determining the Tsi381's latency was used because it required less time for data collection.



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: 80E2000_AN001_04

October 1, 2009