



SPI-4.2 Interoperability with PMC-Sierra XENON Family in Stratix GX Devices

May 2003, ver. 1.0

Application Note 228

Introduction

The system packet interface level 4–phase 2 (SPI-4.2) specification, defined by the Optical Internetworking Forum (OIF), is fast becoming the most common interface for packet transfers between physical (PHY) and link layer devices in multi-gigabit applications, including: asynchronous transfer mode (ATM), packet over SONET/SDH (STS-192/STM-64), 10 Gigabit Ethernet, and multi-channel Gigabit Ethernet.

The SPI-4.2 protocol specifies a source-synchronous interface with differential data rates of at least 622 megabits per second (Mbps). Devices implementing the SPI-4.2 protocol are typically specified as having rates of 700 to 800 Mbps, and in some cases up to 1 gigabit per second (Gbps). At these high data rates, it becomes challenging to manage the skew between clock and data signals. It is common for designers to use the SPI-4.2 protocol in their PHY card designs to connect link layer devices to the mid-plane. This decreases timing margins and reduces channel skew. The SPI-4.2 protocol specifies a training sequence that can be used by receivers to correct skew by ± 1 -bit period, this is commonly known as dynamic phase alignment (DPA).

The Stratix™ GX family of devices are the first FPGAs with embedded DPA circuitry. The SPI-4.2 compliant POS-PHY Level 4 MegaCore® function integrates this circuitry, and offers a broad range of configuration parameters that allow designers to customize the core to meet the specific requirements of their system.

This application note discusses the following topics:

- Overview of the PMC-Sierra XENON™ family of OC-192 and 10 Gigabit Ethernet physical layer devices
- Configuring the POS-PHY Level 4 MegaCore function to interface with the XENON family of devices
- Interoperability testing of the POS-PHY Level 4 MegaCore function implemented in a Stratix GX device with the XENON devices

PMC-Sierra XENON Family

The PMC-Sierra XENON family of devices provides the physical layers for a range of applications, including 10 Gigabit Ethernet (WAN and LAN), OC-192 ATM/GFP/POS, and multi-channel Gigabit Ethernet into the industry standard SPI-4.2 interface.

- S/UNI-10xGE (PM3388)—10-Port Gigabit Ethernet PHY device with embedded standard Gigabit Media Access Controller (GMAC), Physical Coding Sublayer (PCS), Serializer/Deserializer (SERDES), and SPI-4.2 interface.
- S/UNI-1x10GE (PM3392)—10 Gigabit Ethernet LAN PHY for XSBI-based optics with embedded 10 Gigabit Media Access Controller (10GMAC), Physical Coding Sublayer (PCS) and a SPI-4.2 interface.
- S/UNI-1x10GE-XP (PM3393)—10 Gigabit Ethernet LAN PHY for XAUI-based optics incorporating an IEEE standard 10GE MAC and dual XAUI interfaces along with the industry standard SPI4.2/POS-PHY Level 4 system interface.
- S/UNI-9953 (PM5390)—10 Gigabit Ethernet LAN & WAN PHY and OC-192 ATM/GFP/POS concatenated framer to a SPI-4.2 interface. The S/UNI-9953 supports the SFI-4.1 and XSBI 10 Gigabit standard optical module interfaces.
- S/UNI-9953-POS (PM5392)—OC-192 ATM/POS concatenated to a SPI-4.2 interface. The S/UNI-9953-POS provides both ATM and POS functionality at both single-channel OC-192c and quad-channel OC-48c port densities.

SPI-4.2 Interface

The XENON family of devices share a common SPI-4.2 interface with integrated DPA capable of running at up to 700 Mbps. The S/UNI-1x10GE, S/UNI-1x10GE-XP, S/UNI-9953, and S/UNI-9953-POS are single-PHY devices, and the S/UNI-10xGE is a 10-port multi-PHY device. When configured for quad-channel OC-48c port densities, the S/UNI-9953-POS provides a 4-port multi-PHY interface.

Table 1 lists the FIFO buffer sizes of the XENON family of devices.

<i>Table 1. XENON Family FIFO Buffer Size</i>		
Device	Ingress (Kbytes)	Egress (Kbytes)
S/UNI-10xGE	224	64
S/UNI-1x10GE	128	16
S/UNI-1x10GE-XP	64	16
S/UNI-9953	16	20
S/UNI-9953-POS	16	20

Configuring the POS-PHY Level 4 MegaCore

The POS-PHY Level 4 MegaCore function has been designed to support a broad range of applications, and provides designers with a great deal of flexibility to customize the implementation to meet the specific requirements of their system. A number of different configurations can interoperate with the XENON family of devices—each one has a different set of benefits and trade-offs. This section outlines the general issues to consider.



This section assumes that you have a basic knowledge of the SPI-4.2 protocol, you understand the system in which the POS-PHY Level 4 core will be used, and you have read the “Choosing an Architecture” chapter of the *POS-PHY Level 4 MegaCore Function User Guide*.

Device Family

The Stratix GX device family with high-performance LVDS signaling, integrated DPA, and significant embedded memory resources is recommended for SPI-4.2 solutions.

The Stratix device family offers larger devices which may be preferred for high-density designs. However, without the DPA circuitry the need for precise board design is increased, as is the possibility of data rate limitations.

Dynamic Phase Alignment

As high-speed interfaces with source-synchronous clocking schemes approach 700 Mbps and beyond, the margin for clock-to-channel and channel-to-channel skew contracts significantly. To stay within the permitted skew, designers must use precise printed circuit board (PCB) design techniques, otherwise the slightest mismatch in trace lengths or the use of connectors could result in erroneous data transfers. Additionally, skew inducing effects such as process, voltage, and temperature variations compound the problem, thereby making static phase alignment techniques ineffective.

Recognizing the challenges that engineers face when designing systems that transfer high-speed data, Altera® has incorporated dedicated DPA circuitry in Stratix GX devices to dramatically simplify PCB design, eliminating the signal alignment problems introduced by skew-inducing effects.

DPA is recommended for data rates exceeding 622 Mbps, and is considered essential for high-quality signaling at 800 Mbps or across connectors at 700 Mbps. DPA is only available in Stratix GX devices.

LVDS Data Rate

XENON devices operate at a data rate of up to 700 Mbps, providing 11.2 Gbps of bandwidth across the interface. The 1.2 Gbps of additional bandwidth allows for in-band control word management and training, which do not impact the 10 Gbps line rate.

Data Path Width

The POS-PHY Level 4 core supports performance of 700 Mbps with both a full-size (128-bit internal data path width) configuration and a half-size (64-bit internal data path width) configuration.

Data Flow Direction

The core functions either as a receiver or transmitter. Receiver and transmitter functions operate independently, and are configured separately. Applicable configuration choices, such as buffer mode, buffer size, and buffer thresholds do not have to be the same for full-duplex systems.

Number of Ports

The S/UNI-1x10GE, S/UNI-1x10GE-XP, S/UNI-9953, and S/UNI-9953-POS are single-PHY (1-port) devices, and the S/UNI-10xGE is a 10-port multi-PHY device. When configured for quad-channel OC-48c port densities, the S/UNI-9953-POS provides a 4-port multi-PHY interface.

The POS-PHY Level 4 core operates in single-PHY, or multi-PHY mode with up to 256 ports.

Buffer Mode

The POS-PHY Level 4 core provides a configurable FIFO buffer with a single Atlantic™ interface when configured for single-PHY.

Table 2 lists the resources and internal speeds for a selection of single-PHY implementations.

Table 2. Single-PHY Performance for Stratix GX										
Parameters (1)	Receiver (2)					Transmitter				
	LEs	Memory			f_{MAX}	LEs	Memory			f_{MAX}
		M512	M4K	M-RAM			M512	M4K	M-RAM	
Individual FIFO buffer, 64-bit data path	3,673	0	13	0	184	2,912	0	16	0	204
Individual FIFO buffer, 128-bit data path	9,044	0	8	0	117	3,241	0	8	0	142

Notes from Table 2:

(1) FIFO buffer of 4,096 bytes.

(2) DPA enabled.

For multi-PHY applications, three different buffering options are available as follows:

- Individual FIFO buffer per port
- Virtual FIFO segments per port with buffer management
- Shared FIFO buffer with embedded addressing

Below is an outline of the benefits and trade-offs of each buffer mode for applications using the XENON family of devices.

Individual FIFO per Port

The primary benefit of this buffer mode is its simplicity, providing an independent Atlantic interface for each port. In applications where the data from each port is processed independently with parallel processing engines, one design can simply be replicated for each port.

For transmitters, this option provides the additional benefit of avoiding head of line blocking.

The trade-offs with this buffer mode relate to the resulting size of the configuration. For both receivers and transmitters, the size of the core increases linearly with each port. The half-size (64-bit) core may be used for 700 Mbps applications. However, since the logic utilization for 10-port configurations is not minimal, half-size (64-bits) should be limited to transmitter configurations, especially in applications using the S/UNI-10xGE where head of line blocking cannot be tolerated in the system.

Virtual FIFO Segments per Port with Buffer Management

This mode for multi-PHY applications configures the core to use a single buffer logically segmented to provide separate virtual FIFO buffers per port, managed by a single control block.

The benefits of this buffering option are its smaller size for high port densities (8 ports or more), and transmitter implementations lead to nominal head of line blocking. The disadvantages of this buffering option are its complexity (interfacing to a single Atlantic interface for multiple ports) and the addition of a fixed amount of logic (the common control block), making it less applicable for lower port densities (7 ports or less).



This option is currently only available for full-size (128-bit) configurations.

Shared FIFO Buffer with Embedded Addressing

This mode for multi-PHY applications configures the core for multiple ports to share a single Atlantic FIFO buffer with a single Atlantic interface (with an address field). While this buffer mode yields the lowest logic consumption and the highest performance (f_{MAX}), it can lead to head of line blocking.

For receivers, this means that the Atlantic-side logic cannot selectively pick a port to access. Instead, packets from all ports are stored collectively in one physical FIFO buffer and the ordering of the data bursts is maintained in the order in which they were received across the SPI-4.2 interface. The receiver's FIFO status channel communicates the relative fill level of the shared FIFO buffer, for all ports as opposed to a per-port basis.

For transmitters, the order in which packets are transmitted is ultimately dictated by the Atlantic-side logic. User-defined logic is required to schedule the ports to which to transmit. The data burst transmitted across the SPI-4.2 interface is the next data burst in the FIFO buffer regardless of its port. To prevent overflows on the receiving end, the transmitter decodes the incoming FIFO status channel and makes its transmission decision based on the worst-case port. For example, if one port out of ten is satisfied, the transmitter does not transmit data, because the FIFO buffer may contain data bursts for that port. Thus, the FIFO status of one port can block transmission to other ports in the system.

The benefits of this buffering mode are its small size and high performance in terms of the data bit rate (f_{MAX}). It can also be used for port densities up to 256 ports. However, the higher data bit rate performance does not necessarily yield the highest system level performance measured by the data rate of valid data. In applications where head of line blocking cannot be tolerated, this buffering mode should be avoided for the transmitter.

General Comments

Determining the optimal buffering scheme is application-specific and, in particular, depends on the nature of the Atlantic-side processing. For example, where parallel processing is done on each port, the simplicity of the individual FIFO buffer per port mode should be considered. However, in applications where the data from each port is multiplexed into a single stream in the receive direction, and demultiplexed from a single stream in the transmit direction, the virtual FIFO segments per port with buffer management mode offers advantages.



Each buffering scheme offers a different set of benefits and trade-offs and selecting the optimal configuration is beyond the scope of this application note.

Table 3 lists the resources and internal speeds for a selection of 10-port multi-PHY implementations.

Parameters (1)	Receiver (2)					Transmitter				
	LEs	Memory			f_{MAX}	LEs	Memory			f_{MAX}
		M512	M4K	M-RAM			M512	M4K	M-RAM	
Shared FIFO with embedded addressing, 64-bit data path	4,151	4	14	0	185	3,027	0	16	0	167
Shared FIFO with Embedded Addressing, 128-bit data path	12,688	1	19	0	121	3,551	0	8	0	115
Individual FIFO buffer per port, 64-bit data path	9,428	0	86	0	177	12,898	0	60	0	167

Parameters (1)	Receiver (2)					Transmitter				
	LEs	Memory			f _{MAX}	LEs	Memory			f _{MAX}
		M512	M4K	M-RAM			M512	M4K	M-RAM	
Virtual FIFO buffer segments with buffer management, 128-bit data path	19,674	15	270	0	104	9,042	15	270	0	104

Notes from Table 3:

- (1) FIFO buffer of 2,048 bytes per port.
(2) DPA & multiple continues enabled.

FIFO Buffer Size per Port

The core offers buffer sizes ranging from 512 bytes to 32 Kbytes per port. However, the amount of buffering is effectively limited by the amount of available device memory resources. To ensure proper operation, the core's FIFO buffers per port should be configured to be the same size or larger than the FIFO buffers provided in the adjacent device. However, this might not necessarily result in the optimal use of memory resources. In the case of 10-port applications, where the S/UNI-10xGE provides 224 Kbytes of buffering on the SPI-4.2 transmitter, and 64 Kbytes of buffering on the SPI-4.2 receiver, this would imply that the POS-PHY Level 4 MegaCore receiver be configured to have 16,384 bytes of buffering per port, and that the transmitter be configured to have 4,096 bytes of buffering per port. Full-duplex implementations would consume 200 Kbytes of memory, or 48% of the embedded memory of the larger Stratix GX (EP1SGX40) devices. Smaller FIFO buffers can be implemented in many applications, as demonstrated by the successful testing with the S/UNI-10xGE with 8,096 bytes of buffering per port in both the receiver and the transmitter, discussed later in this application note.

Almost Empty & Almost Full

The almost empty (AE) and almost full (AF) watermarks are used to segment the receiver FIFO buffers into three parts:

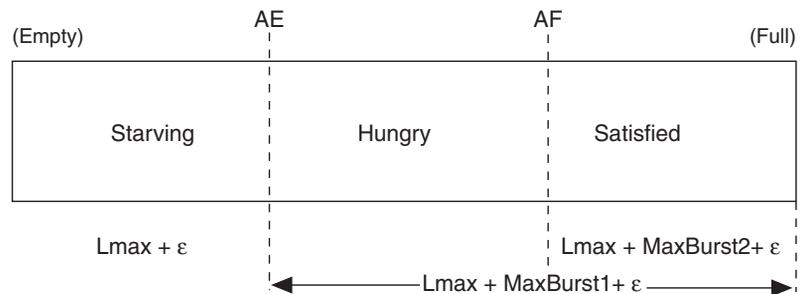
- **Starving** – Occurs when the FIFO buffer data is below the almost empty watermark, and data bursts up to the MaxBurst1 value may be transmitted from the adjacent device

- *Hungry* – Occurs when the FIFO buffer data is between the almost empty and almost full watermarks, and data bursts up to the MaxBurst2 value may be transmitted from the adjacent device
- *Satisfied* – Occurs when the FIFO buffer data is above the almost full watermark, and data bursts may not be transmitted from the adjacent device

The SPI-4.2 specification defines 2-bit values for starving, hungry, and satisfied. These 2-bit values are calculated based on the available space in the FIFO buffer, and on the user-defined parameters for almost empty and almost full for each FIFO buffer.

Figure 1 illustrates the relationship between the almost empty (AE) and almost full (AF) watermarks and the MaxBurst1 and MaxBurst2 values.

Figure 1. FIFO Buffer Watermarks



Notes:

- (1) L_{max} corresponds to the worst-case response time, from the delay in receiving a status update over the FIFO status channel, until observing the reaction to that update on the corresponding data path.
- (2) ϵ corresponds to the difference between the granted credit and the actual data transfer length. This difference arises from various protocol overheads.
- (3) The MaxBurst1 and MaxBurst2 values used to determine the AE and AF watermarks on the receiver FIFO are configured in the adjacent device's transmitter. Determining the optimal MaxBurst1 and MaxBurst2 values is application-specific, and requires an analysis of the data flows, beyond the scope of this application note.

MaxBurst1 & MaxBurst2

The MaxBurst1 and MaxBurst2 values are used in the transmitter to set the maximum data bursts (8 to 127 – 16 byte blocks) that may be transmitted when the adjacent receiver's FIFO buffer is in starving and hungry states, respectively. Determining the optimal MaxBurst1 and MaxBurst2 values is application-specific, and requires an analysis of the data flows, beyond the scope of this application note.

Atlantic Data Width

The core offers four Atlantic interface data width options: 32, 64, 128 or 256 bits. Determining the applicable Atlantic data width to be used when interfacing to a XENON device depends on the f_{MAX} target of the Atlantic-side logic. For 64-bit Atlantic data widths, user-logic needs to run at a minimum of 1/4 the LVDS data rate (e.g. minimum 175 MHz f_{MAX} for 700 Mbps LVDS data rate). Whereas, 128-bit Atlantic data widths require user-logic to run at a minimum of 1/8 the LVDS data rate (e.g. minimum 87.5 MHz f_{MAX} for 700 Mbps LVDS data rate). Internal f_{MAX} performance is dependent on the nature of the processing logic and the device speed grade.

Transmit Bandwidth Optimization

Enabling this optional parameter, applicable to transmitters only, results in the transmission of packets with minimal `IDLE` insertion, or in some cases no `IDLE` insertion (start-of-packet (`SOP`) and end-of-packet (`EOP`) markers share a common control word). Enabling the transmit bandwidth optimization parameter increases the effective bandwidth rate on the SPI-4.2 data bus, but significantly increases the LE consumption.

Disabling the transmit bandwidth optimization parameter reduces the LE consumption by approximately:

- 3,600 LEs for 128-bit configurations
- 1,400 LEs for 64-bit configurations

Disabling the transmit bandwidth optimization parameter allows `IDLES` to be inserted between `EOP` control words and the next `SOP` or payload control word. The maximum number of `IDLES` is a function of packet size, but is a maximum of six for full-size (128-bit) configurations, and two for half-size (64-bit) configurations. Although not using transmit bandwidth optimization lowers the effective bandwidth rate on the SPI-4.2 data bus, it greatly reduces the LE consumption. As such, this optional parameter should only be used where dictated by the system performance requirements. In particular, the need for bandwidth optimization reduces as the average packet size increases. In addition, the need for bandwidth optimization also reduces as the LVDS data rate increases. Operating the interface at the full 700 Mbps supported by the XENON device family would eliminate the need for bandwidth optimization in many applications.

Burst Mode

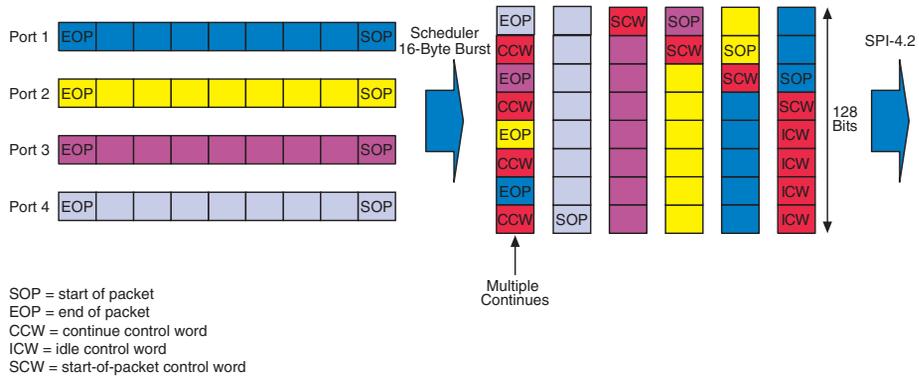
This optional parameter is applicable to certain network processors, with element based memories, and is not intended to be used when interfacing to XENON devices.

Multiple Continues

For ease of implementation, the SPI-4.2 standard specifies that start-of-packet (SOP) control words may not occur less than eight cycles apart. However, there are no specific restrictions for other control words. Accordingly, it is possible to encounter multiple continuation control words in less than eight cycles. A `continue` is a payload control word that is not immediately followed by a payload transfer corresponding to a start-of-packet (i.e. bit 15 = 1, indicates a payload transfer immediately follows the control word; bit 12 = 0, indicates the payload transfer does not correspond to a start-of-packet).

In worst case scenarios for standard SPI-4.2 implementations, the receiver may be required to process the remaining bytes of several packets, destined for various ports, occurring less than eight cycles apart. [Figure 2](#) shows an example where the multiple continue parameter is enabled.

Figure 2. Multiple Continue Example



The XENON family of devices use the standard (OIF compliant) implementation with zero idle insertions between data bursts, which can lead to multiple continuation control words in less than eight clock cycles in multi-PHY applications. The POS-PHY Level 4 MegaCore receiver must be configured with the multiple continue parameter enabled when interfacing to multi-PHY XENON devices.

Training Sequence Interval

The Maximum Training Sequence Interval (MaxT) parameter allows the user to select the interval at which the training sequence occurs—8 to 32,767—in units of 16-byte cycles. The training sequence is inserted after a payload burst, or during the transmission of idle control words. Once a training sequence is sent, the next training sequence is not initiated any earlier than MaxT 16-byte cycles. Extending the training sequence interval increases the effective bandwidth across the SPI-4.2 interface.

Training Pattern Repetitions

The training sequence includes one idle control word, plus ALPHA (α) \times 20 training words, where ALPHA (α) is a value in the range of 0 to 255.

The core require's approximately 1000 training pattern repetitions: 650 back-to-back for the core's byte aligner and channel aligner timing, and a few hundred for the Stratix GX DPA timing.

The XENON device's transmitter should be configured with an ALPHA (α) value of at least 1, however during initial training, the training pattern must be repeated continuously.

Interoperability Testing with S/UNI-10xGE

This section provides an overview of the interoperability testing of the POS-PHY Level 4 MegaCore function implemented in a Stratix GX device with the S/UNI-10xGE. The purpose of this testing is to confirm the correct SPI-4.2 protocol processing between the devices using typical Gigabit Ethernet traffic transmitted at 700 Mbps (LVDS data rate) across the SPI-4.2 interface. To exercise the embedded DPA circuitry, skew is introduced by intentionally mismatching trace lengths. This testing is in addition to the characterization of the Stratix GX device at various LVDS data rates—across process, voltage and temperature (PVT) ranges.

Test Equipment

The following components are used in the interoperability testing:

PM2381-KIT

The PMC-Sierra PM2381-KIT is a development board incorporating the S/UNI-10xGE 10-Port Gigabit Ethernet MAC, Small Form Pluggable (SFP) optical connectors, and the SPI-4.2 interface to a Molex 74057-1001 connector.

Altera Test Board

The Altera test board was developed for SPI-4.2 interoperability testing. The board features a Stratix GX (ES1SGX25) device with SPI-4.2 channels to a Tyco HM-Zd connector. To demonstrate DPA operation, no attempt was made to match trace lengths, and on the receive path, in particular, one channel-length was extended, introducing 0.3 UI of skew. Samtec connectors are inserted in the SPI-4.2 paths, providing logic analyzer access to both the receive and transmit signals.



The High-Speed Development Kit, Stratix GX Edition incorporates the same HM-Zd connector and pin-out used on the internal test board.

Interposer Card

An interposer card with reciprocal Molex 74057-1001 and Tyco HM-Zd connectors was developed to bridge the PMC-Sierra and Altera development boards. The interposer card is included in the High-Speed Development Kit, Stratix GX Edition.

Spirent SmartBits

The SmartBits 600 from Spirent is used for Gigabit Ethernet traffic stimulus, capture, and analysis.



For more information on the SmartBits 600 system, refer to Spirent's website at www.spirent.com.

Test Description

The Ethernet frames transmitted vary in size from the 64-byte minimum to the 5,000-byte maximum, applicable to Ethernet applications.

To exercise the embedded DPA circuitry in the Stratix GX FPGA, no attempt was made to match trace lengths on the Altera test board, and on the receive path, in particular, one channel-length was extended, introducing 0.3 UI of skew. Samtec connectors are inserted in the SPI-4.2 paths, providing logic analyzer access to both the receive and transmit signals.

[Figure 3 on page 14](#) shows a block diagram of the test set-up.

Figure 3. Test Set-up Block Diagram

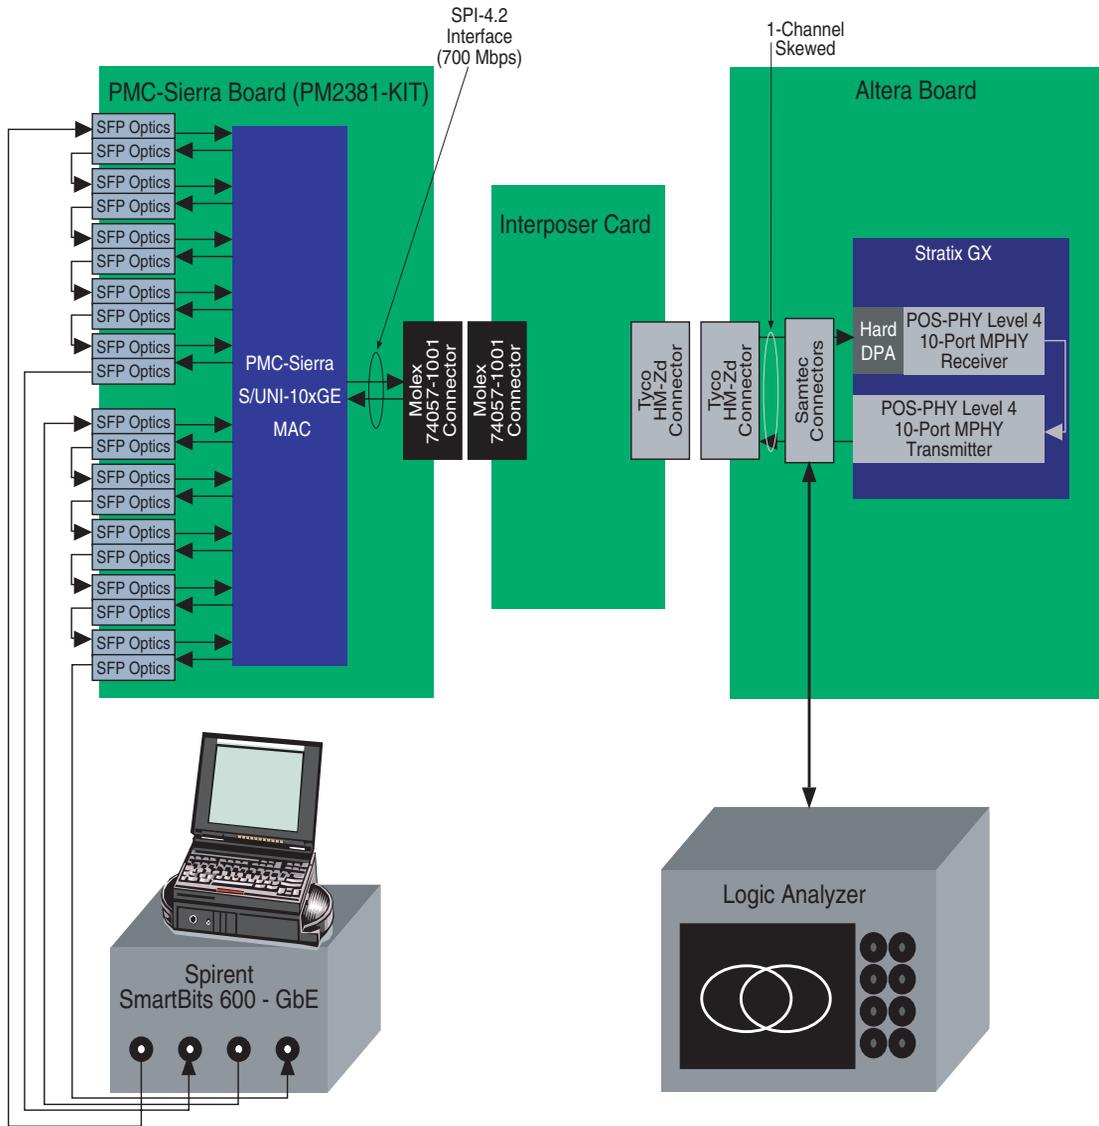


Figure 4 on page 15 shows a picture of the test set-up.

Figure 4. Test Set-Up

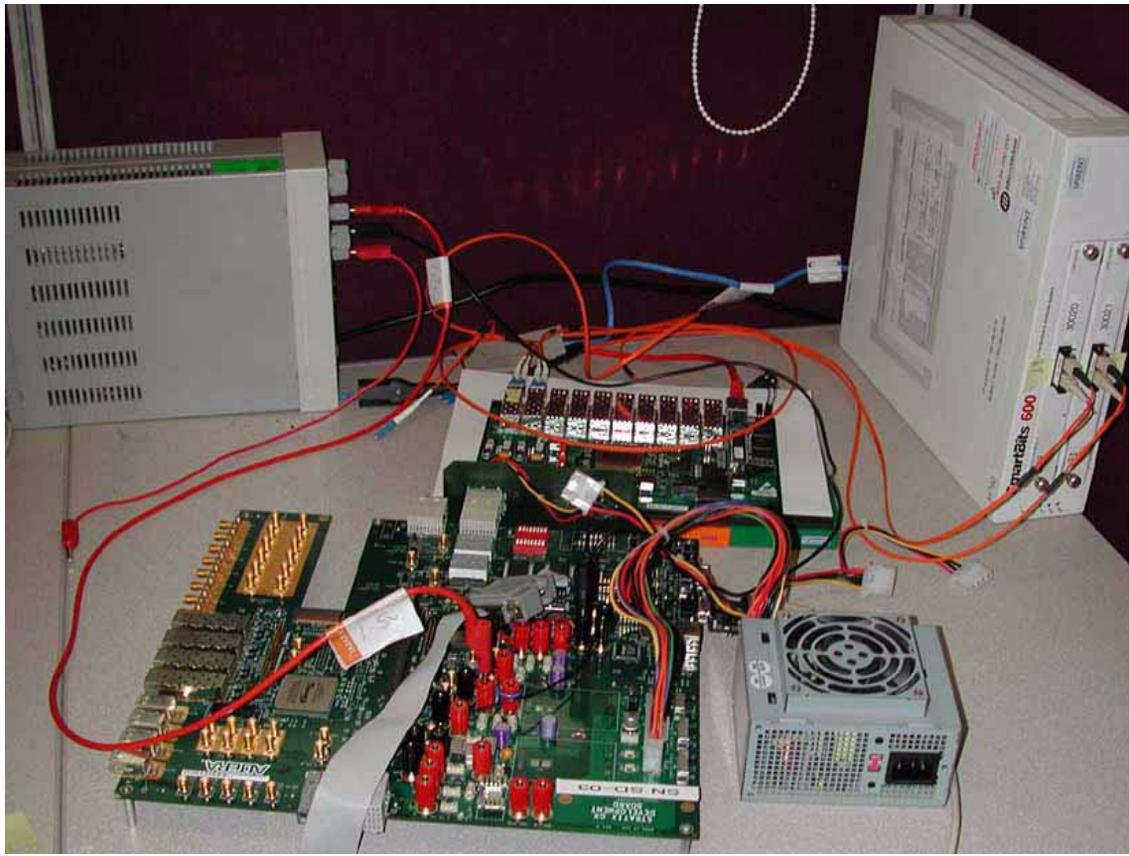


Figure 5 on page 16 shows a picture of the PMC-Sierra PM2381-KIT development board.

Figure 5. PMC-Sierra PM2381-KIT Development Board

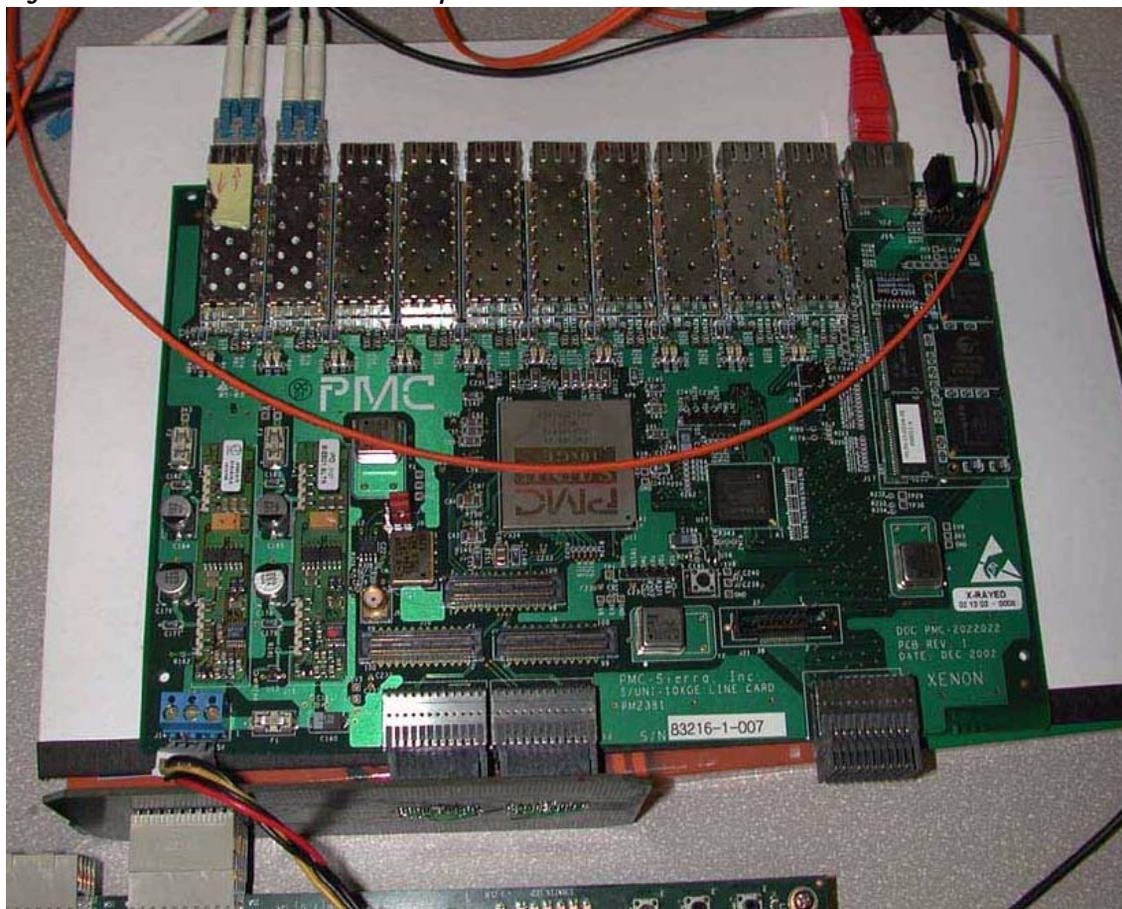
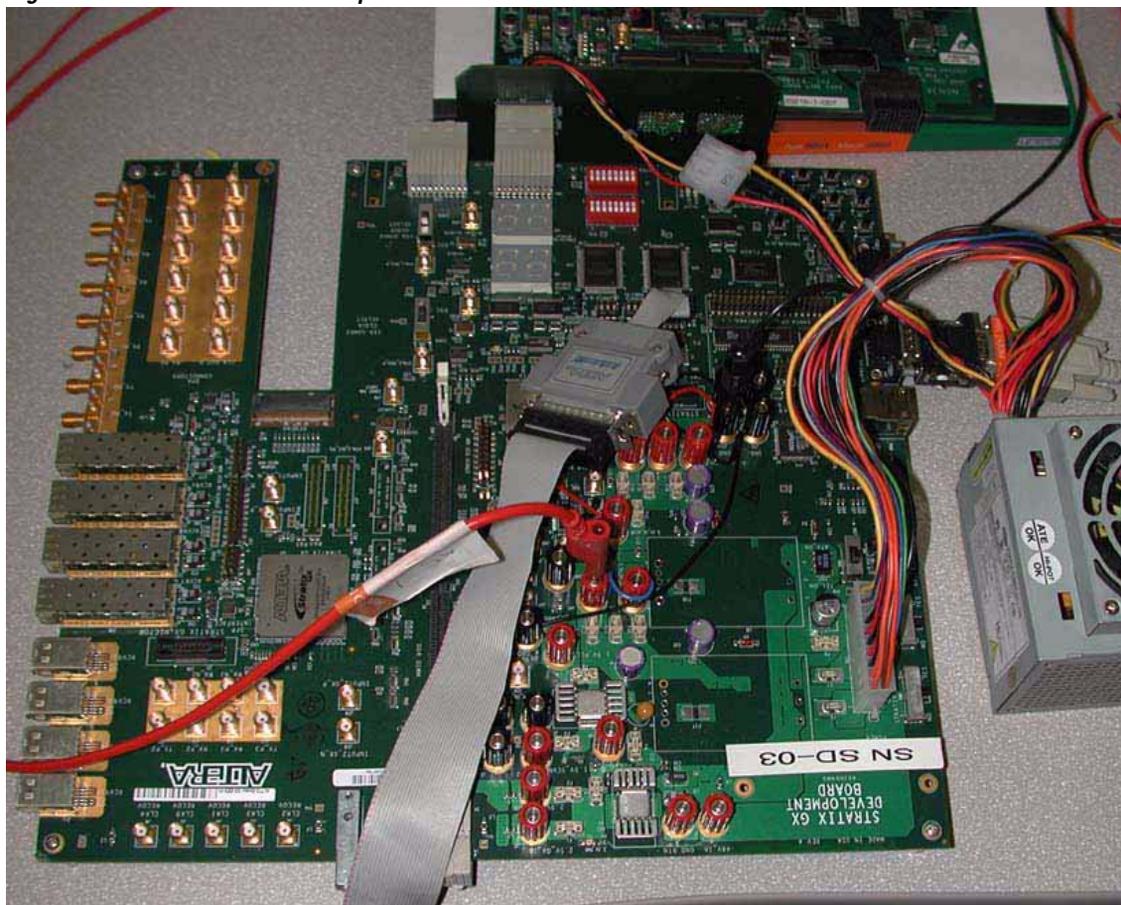


Figure 6 on page 17 shows a picture of the Altera Stratix GX development board.

Figure 6. Altera Stratix GX Development Board



Test Configuration

Table 2 lists the parameter values used in the loopback test system.

Table 4. POS-PHY Level 4 Parameter Values Used for S/UNI-10xGE Test (Part 1 of 2)

Parameter	Receiver	Transmitter
Device Family	Stratix GX	Stratix GX
Dynamic Phase Alignment	Yes	Not applicable
LVDS Data Rate	700 Mbps	700 Mbps
Data Path Width	128 bits	128 bits
Data Flow Direction	RX	TX
Number of Ports	10	10

Table 4. POS-PHY Level 4 Parameter Values Used for S/UNI-10xGE Test (Part 2 of 2)

Parameter	Receiver	Transmitter
Buffer Mode	Shared FIFO buffer with embedded addressing	Shared FIFO buffer with embedded addressing
FIFO Buffer Size	8,192 bytes	8,192 bytes
Almost Empty	512 bytes	Not applicable
Almost Full	1,024 bytes	Not applicable
MaxBurst1	Not applicable	1,024 bytes
MaxBurst2	Not applicable	512 bytes
Atlantic Data Width	128 bits	128 bits
FIFO Threshold Low	32 bytes	Not applicable
FIFO Threshold High	Not applicable	320 bytes
Transmit Bandwidth Optimization	Not applicable	No
Burst Mode	Not applicable	No
Multiple Continues	Yes	Not applicable
Maximum Training Sequence Interval	Not applicable	32,000
Training Pattern Repetitions	Not applicable	32
Status Channel Clock Edge	Negative	Negative
Calendar Multiplier	1	1

Interoperability Testing with S/UNI-9953

This section provides an overview of the interoperability testing of the POS-PHY Level 4 MegaCore function implemented in an APEX™ II device with the S/UNI-9953. The purpose of this testing is to confirm the correct SPI-4.2 protocol processing between the devices using typical packet traffic transmitted at 700 Mbps (LVDS data rate) across the SPI-4.2 interface.

Test Equipment

The following components were used in the interoperability testing:

APEX II – S/UNI-9953 Test Board

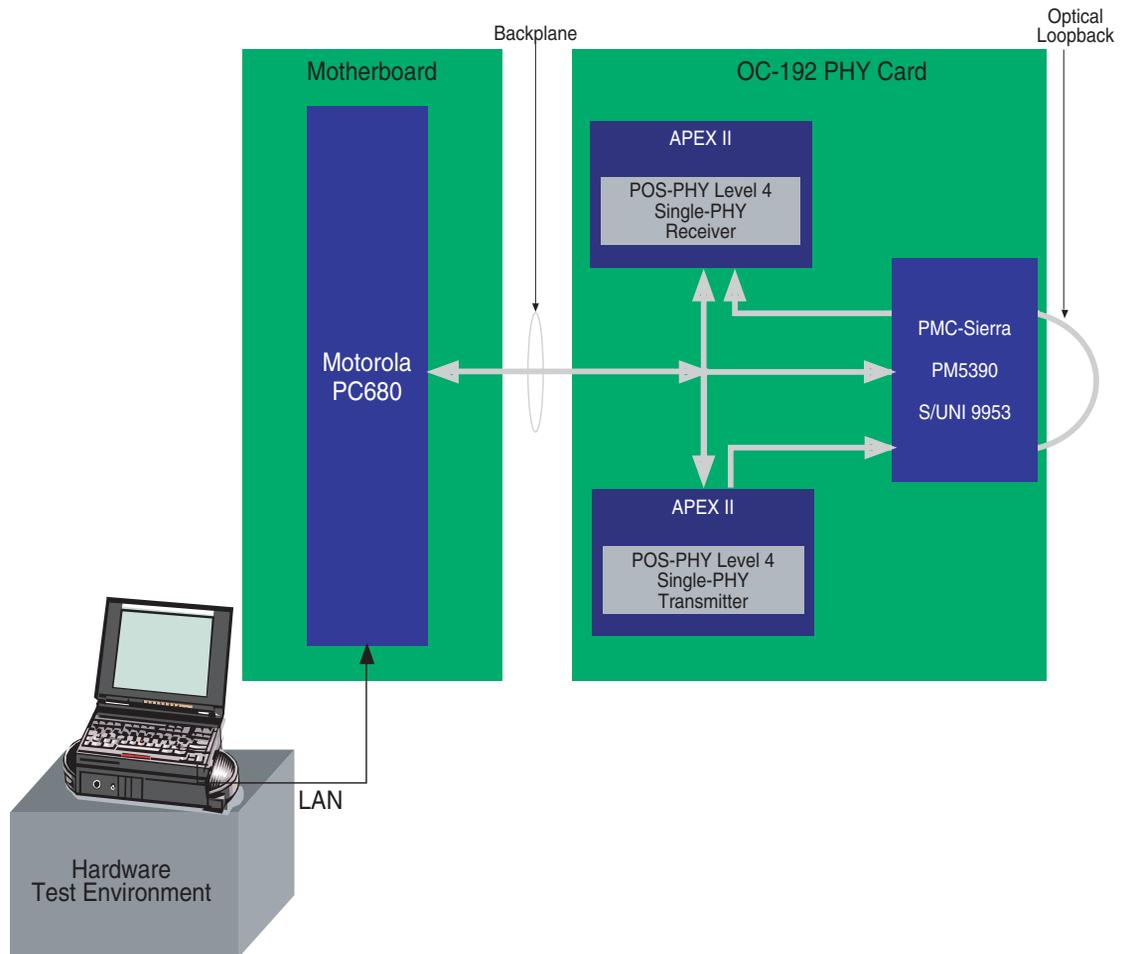
The APEX II – S/UNI-9953 test board features two APEX II (EP2A25) devices connected to the receive & transmit SPI-4.2 interfaces of the S/UNI-9953.

Test Description

Test traffic is generated and monitored using utilities implemented in the respective FPGAs with the S/UNI-9953 in loopback mode.

Figure 2 shows a block diagram of the test set-up.

Figure 7. Test Set-Up Block Diagram



The following test cases have been executed:

- Verified that packets with 1, 2, 3, 4, 5, 64, 256, and 1024 bytes were correctly transmitted and received
- Verified that data was correctly received with packets back-to-back on the Atlantic interface to transmitter core.
- Verified that data was correctly received with packets widely spaced on the Atlantic interface to transmitter core.
- Verified that no errors occurred during long transmit line idle
- Verified that no errors occurred during long receive line idle

- Verified that no errors occurred during long transmission of good packets (same size packets—16 bytes) (SOAK test)
- Verified that no errors occurred during long transmission of good packets (same size packets—15 bytes) (SOAK test)
- Verified that no errors occurred during long transmission of good packets (random size packets with same data payload) (SOAK test)
- Verified that no errors occurred during long transmission of good packets (random size packets with each packet having different payload) (SOAK test)
- Verified that no errors occurred during long duration of good receive packets (same size packets—16 bytes) (SOAK test)
- Verified that no errors occurred during long duration of good receive packets (same size packets—15 bytes) (SOAK test)
- Verified that no errors during long duration of good receive packets (random size packets with same data payload) (SOAK test)
- Verified that no errors occurred during long duration of good receive packets (random size packets with each packet having different payload) (SOAK test)
- Verified training pattern detection
- Verified DIP-4 generation
- Verified DIP-4 monitoring
- Verified framing, CALLEN generation and monitoring
- Verified DIP-2 generation
- Verified FIFO buffer overflow functionality on Atlantic transmit FIFO buffer interface
- Verified FIFO buffer underflow functionality on Atlantic transmit FIFO buffer interface
- Verified FIFO buffer overflow functionality on Atlantic receive FIFO buffer interface
- Verified FIFO buffer underflow functionality on Atlantic receive FIFO buffer interface
- Verified that packets marked with ERR on Atlantic transmit FIFO buffer interface are EOP-marked as ABORT.
- Verified that packets marked with EOP ABORT are received as ERR on Atlantic receive FIFO buffer interface.

Test Configuration

The S/UNI-9953 is configured for POS operation. [Table 5](#) lists the parameter values for the POS-PHY Level 4 MegaCore function.

<i>Table 5. POS-PHY Level 4 Parameter Values Used for S/UNI-9953 Test (Part 1 of 2)</i>		
Parameter	Receiver	Transmitter
Device Family	APEX II	APEX II
Dynamic Phase Alignment	No	Not applicable

Table 5. POS-PHY Level 4 Parameter Values Used for S/UNI-9953 Test (Part 2 of 2)

LVDS Data Rate	700 Mbps	700 Mbps
Data Path Width	128 bits	128 bits
Data Flow Direction	RX	TX
Number of Ports	1	1
Buffer Mode	Individual FIFO buffer	Individual FIFO buffer
FIFO Buffer Size	8,192 bytes	8,192 bytes
Almost Empty	25	Not applicable
Almost Full	166	Not applicable
MaxBurst1	Not applicable	50
MaxBurst2	Not applicable	50
Atlantic Data Width	128 bits	128 bits
FIFO Threshold Low	25	Not applicable
FIFO Threshold High	Not applicable	166
Transmit Bandwidth Optimization	Not applicable	No
Burst Mode	Not applicable	No
Multiple Continues	Not applicable	Not applicable
Maximum Training Sequence Interval	Not applicable	2,048
Training Pattern Repetitions	Not applicable	1
Status Channel Clock Edge	Negative	Negative
Calendar Multiplier	1	1



101 Innovation Drive
San Jose, CA 95134
(408) 544-7000
<http://www.altera.com>
Applications Hotline:
(800) 800-EPLD
Literature Services:
lit_req@altera.com

Copyright © 2003 Altera Corporation. All rights reserved. Altera, The Programmable Solutions Company, the stylized Altera logo, specific device designations, and all other words and logos that are identified as trademarks and/or service marks are, unless noted otherwise, the trademarks and service marks of Altera Corporation in the U.S. and other countries. All other product or service names are the property of their respective holders. Altera products are protected under numerous U.S. and foreign patents and pending applications, mask work rights, and copyrights. Altera warrants performance of its semiconductor products to current specifications in accordance with Altera's standard warranty, but reserves the right to make changes to any products and services at any time without notice. Altera assumes no responsibility or liability arising out of the application or use of any information, product, or service described herein except as expressly agreed to in writing by Altera Corporation. Altera customers are advised to obtain the latest version of device specifications before relying on any published information and before placing orders for products or services.

