



# SPI-4.2 Interoperability with the Intel IXF1110 in Stratix GX Devices

May 2003, ver. 1.0

Application Note 227

## 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  $\pm 1$ -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 Intel® IXF1110 10-port Gigabit Ethernet Media Access Controllers (MACs)
- Configuring the POS-PHY Level 4 MegaCore function to interface with the Intel IXF1110
- Interoperability testing of the POS-PHY Level 4 MegaCore function implemented in a Stratix GX device with the Intel IXF1110

## Intel IXF1110 10-Port Gigabit Ethernet Media Access Controllers

The Intel IXF1110 is a 10-port Gigabit Ethernet MAC that provides a 10 Gigabit, system-level interface to 10 individual 1000 Mbps full-duplex Ethernet MACs. An integrated clock data recovery (CDR) and serializer/deserializer (SERDES) interface provides direct connection to the optics.

### Features

- Ten IEEE 802.3 compliant 1000 BASE-T MAC ports with integrated CDR/SERDES
- 10-port SPI-4.2 interface with DPA from 640 to 800 Mbps
- A 17.0 Kbyte receive first-in first-out (FIFO) buffer and a 4.5 Kbyte transmit FIFO per port
- Compliance with IEEE 802.3x standard PAUSE command frames
- RMON statistics
- Broadcast, multicast, and unicast address filtering
- 32-bit microprocessor interface
- Independent enabling/disabling of any port

### SPI-4.2 Interface

- 640 to 800 Mbps source-synchronous data rate with LVDS signaling
- LVCMOS status channel at one-quarter the data rate, compatible with standard LVTTL signaling
- 10-port multi-PHY, or fewer if some of the Gigabit Ethernet MACs are disabled
- Integrated DPA circuitry on the receiver
- Programmable MaxBurst1 and MaxBurst2 values



Intel recommends an 800 Mbps data rate, where possible, providing 12.8 Gbps of bandwidth across the interface. The 2.8 Gbps of additional bandwidth allows for in-band control word management and training, which does not impact the 10 Gbps line rate.



For more information on the Intel IXF1110 10-Port Gigabit Ethernet MAC SPI-4.2 interface, refer to the Intel application note: *Interfacing with IXF1010 and IXF1110 10-Port Gigabit Ethernet Media Access Controllers*, available at [www.intel.com](http://www.intel.com).

## 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 Intel IXF1110—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

Intel recommends an 800 Mbps data rate, where possible, providing 12.8 Gbps of bandwidth across the interface. The 2.8 Gbps of additional bandwidth allows for in-band control word management and training, which do not impact the 10 Gbps line rate.

The IXF1110 supports data rates from 640 to 800 Mbps. The POS-PHY Level 4 core supports performance beyond 800 Mbps with a full-size (128-bit internal data path width) configuration, and performance up to 700 Mbps with a half-size (64-bit internal data path width) configuration. In systems where 700 Mbps performance provides sufficient additional bandwidth for control processing & training, using this lower data rate configuration could reduce the cost of the FPGA solution.

## Data Path Width

A full-size (128-bit) configuration is required for performance of 800 Mbps. A half-size (64-bit) configuration can be used in systems where 700 Mbps performance provides sufficient additional bandwidth for control processing & training.

## 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 IXF1110 supports up to 10 individual Gigabit Ethernet streams, and as such would typically require a 10-port configuration on the SPI-4.2 interface. However, since the IXF1110 supports the independent enabling or disabling of any port, when fewer Gigabit Ethernet streams are being supported in the system the number of ports on the SPI-4.2 interface is reduced accordingly.

## Buffer Mode

The POS-PHY Level 4 core offers three different buffering options for multi-PHY applications:

- 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 Intel IXF1110 in 10-port configurations.

### *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. In applications targeting 800 Mbps data rates, the full-size (128-bit) core is required and the resulting core consumes a large percentage of the FPGA. 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 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 1 lists the resources and internal speeds for a selection of MPHY implementations.

Parameters (1)	Receiver (2)					Transmitter				
	LEs	Memory			f <sub>MAX</sub>	LEs	Memory			f <sub>MAX</sub>
		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
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 1:**

- (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. Since the IXF1110 provides 17 Kbytes of buffering on the SPI-4.2 transmitter, and 4.5 Kbytes of buffering on the SPI-4.2 receiver, for 10-port applications the POS-PHY Level 4 MegaCore receiver should be configured to have 2,048 bytes of buffering per port, and the transmitter should be configured to have 512 bytes of buffering per port. Full-duplex implementations consume 25 Kbytes of memory, or approximately 6% of the embedded memory of the larger Stratix GX (EP1SGX40) devices.

## Almost Empty & Almost Full

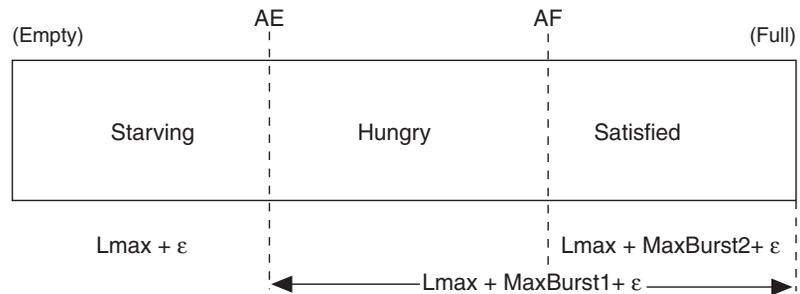
The almost empty (AE) and almost full (AF) watermarks are used to segment the receiver FIFO buffers into three parts:

- *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)  $\epsilon$  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 the IXF1110 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, as follows:

- 1,600 LEs for 10-port half-size (64-bit) configurations
- 3,800 LEs for 10-port full-size (128-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 800 Mbps would eliminate the need for bandwidth optimization in most applications with the IXF1110.

## Burst Mode

This optional parameter is applicable to certain network processors, such as the Intel IXP2800 Network Processor, and should not be used when interfacing to the IXF1110.

## 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).

The IXF1110 supports two modes relating to continuation control word transmission.

1. 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, and in the worst-case a repeating series of continuation control words followed by a single data word. When the IXF1110 is configured in this standard mode, the POS-PHY Level 4 MegaCore receiver configuration must be have the multiple continue parameter enabled.
2. The IXF1110 offers an optional mode on its transmitter which eases the design requirements of adjacent device SPI-4.2 receivers. This mode inserts up to four idle control words between data transfer bursts, ensuring that a half-size (64-bit) receiver has either a payload control word, data for a single port, or idle control words (i.e. a maximum of one continuation control word) to be processed in parallel. If the IXF1110 implementation uses this optional mode, and the half-size (64-bit) core is being used, the multiple continue parameter is not required. However, if the full-size (128-bit) core is being used, multiple continuation control words can appear on the 128-bit internal data path in parallel, and the multiple continue parameter must be enabled.

## 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 ( $\alpha$ )  $\times$  20 training words, where ALPHA ( $\alpha$ ) is a value in the range of 0 to 255.

The IXF1110 is capable of training reliably in an average of 32 training pattern repetitions. Accordingly, for transmitter configurations, the core's ALPHA ( $\alpha$ ) value can be set as low as 32.

The core requires 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 IXF1110 transmitter should be configured with an ALPHA ( $\alpha$ ) value of at least 1, however during initial training, the training pattern must be repeated continuously.

## Interoperability Testing

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 Intel IXF1110 10-Port Gigabit Ethernet MAC. 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 800 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:

### *Intel IXD1110*

The Intel IXD1110 is a development board incorporating the IXF1110 10-Port Gigabit Ethernet MAC, Small Form Pluggable (SFP) optical connectors, and the SPI-4.2 interface to an HM-Zd connector.

### *Altera Test Board*

The Altera test board was developed for SPI-4.2 interoperability testing with the IXF1110. The board features a Stratix GX (ES1SGX25) device with SPI-4.2 channels to the reciprocal Tyco HM-Zd connector for connection with the IXD1110 development board. 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, providing direct connection to the IXD1110 development board.

### *IneoQuest Singulus G1*

The Singulus G1 – Gigabit Ethernet system from IneoQuest Technologies, Inc. is used for Gigabit Ethernet traffic stimulus, capture, and analysis.



For more information on the Singulus G1 – Gigabit Ethernet system, refer to IneoQuest's website at [www.ineoquest.com](http://www.ineoquest.com).

### *Agilent 16702B*

The Agilent 16702B Logic Analysis System with SPI-4.2 software module and high-speed source-synchronous probes is used to decode the data traffic between the devices under test in both directions. In addition, the logic analyzer is used to generate eye-diagrams, illustrating the skew on the receive path and the subsequent alignment on the transmit path.

## Test Description

To fill the SPI-4.2 bandwidth using a single Gigabit Ethernet source, the Singulus G1 is used to generate ingress traffic to the optical receiver on Port 1, and through the loopback system. The optical output from Port 1 is connected to the optical input on Port 2, and similarly the optical output from Port 2 is connected to optical input on Port 3, Port 3 to Port 4, ... and finally Port 9 is connected to Port 10. The egress traffic from Port 10 is captured by the Singulus G1. See [Figure 2 on page 14](#). By using this scheme, the SPI-4.2 interface between the Intel IXF1110 10-Port Gigabit Ethernet MAC and the Stratix GX FPGA is filled with 10 Gbps of line-rate traffic.

The Ethernet frames transmitted vary in size from the 64-byte minimum to the 1,522-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 2 on page 14](#) shows a block diagram of the test set-up.

Figure 2. Test Set-up Block Diagram

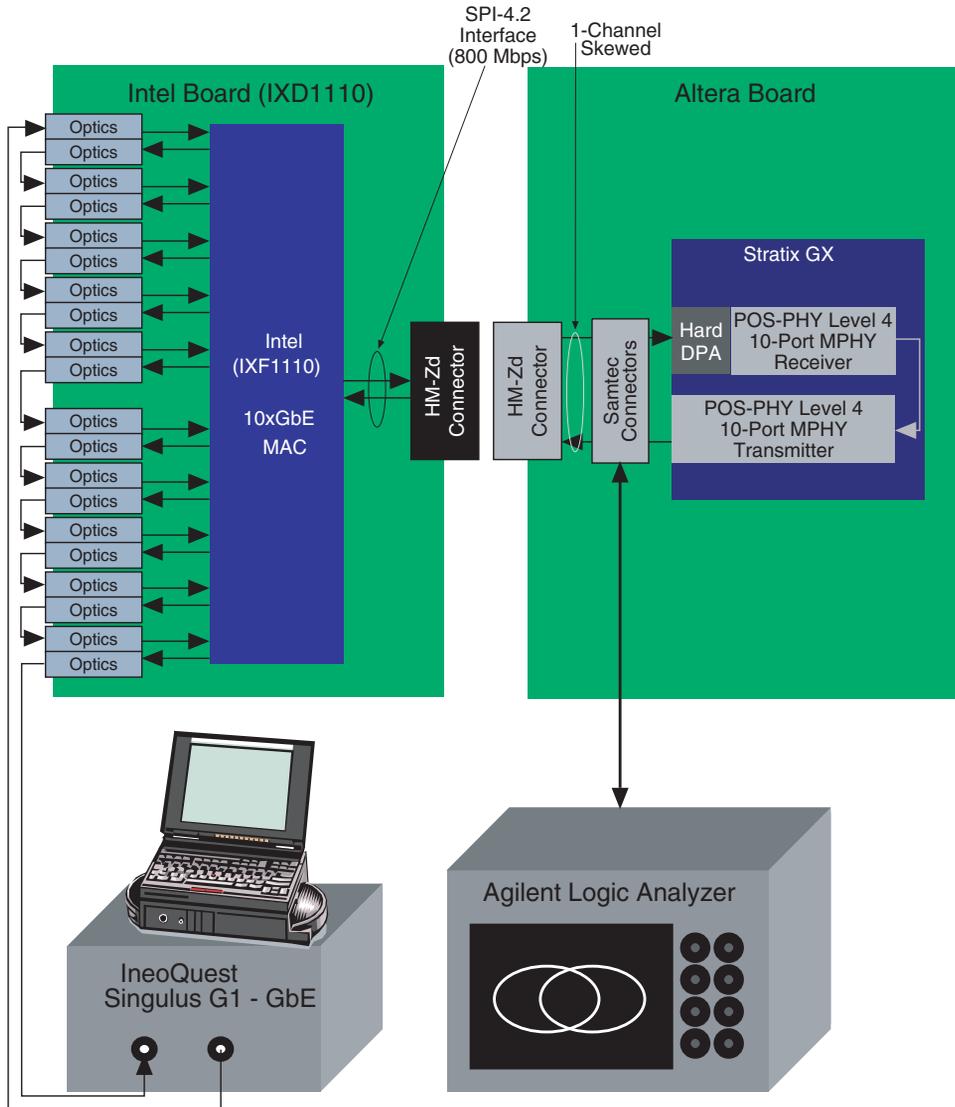


Figure 3 on page 15 shows a picture of the Intel IXD1110 development board.

**Figure 3. Intel IXD1110 Development Board**

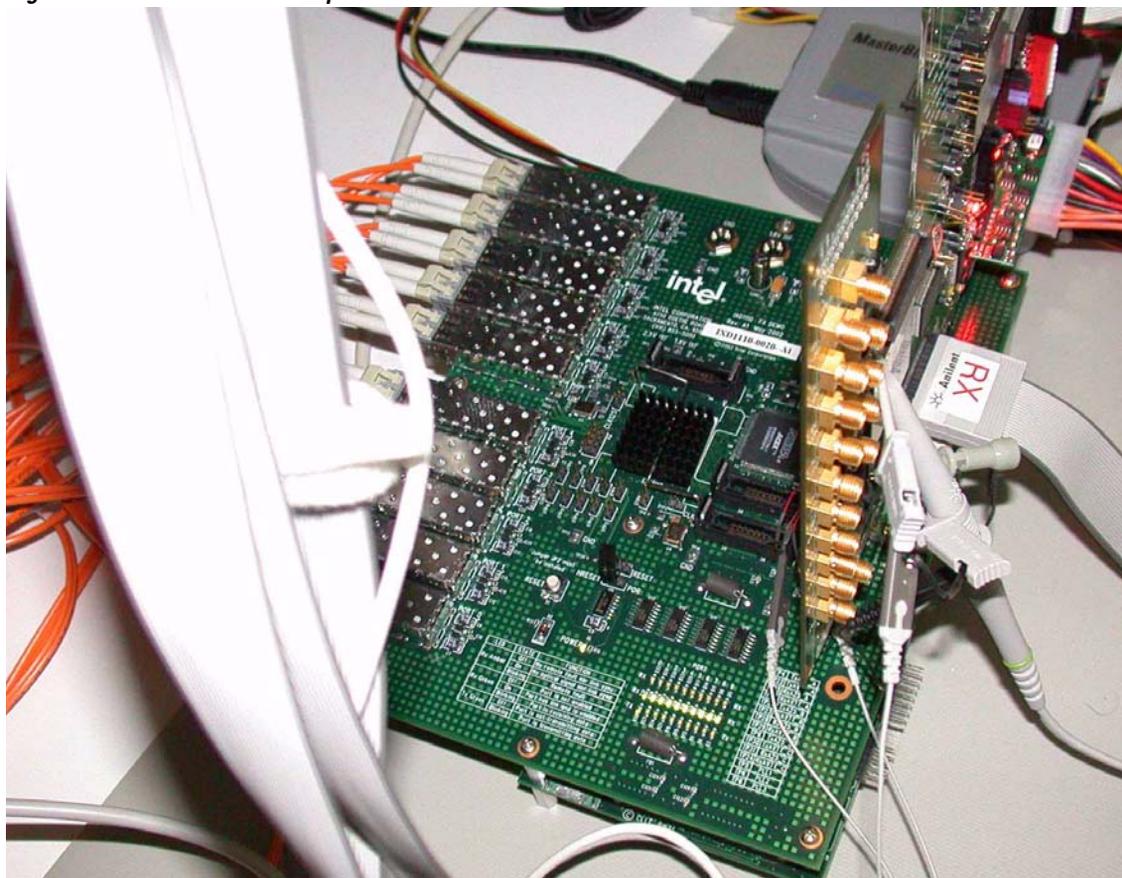
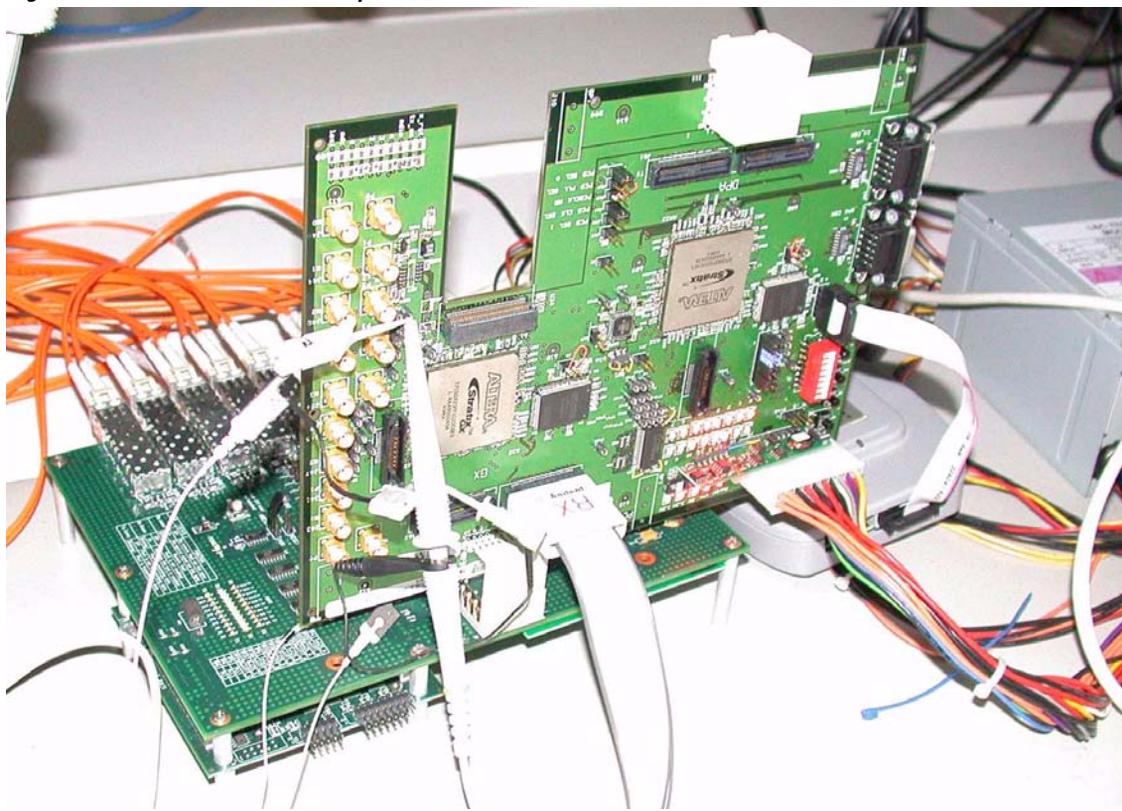


Figure 4 on page 16 shows a picture of the Altera Stratix GX development board mounted onto the Intel IXD1110 development board.

**Figure 4. Altera Stratix GX Development Board**



### Test Configuration

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

**Table 2. POS-PHY Level 4 Configuration Parameter Values Used for Testing (Part 1 of 2)**

Parameter	Receiver	Transmitter
Device Family	Stratix GX	Stratix GX
Dynamic Phase Alignment	Enabled	Not applicable
LVDS Data Rate	800 Mbps	800 Mbps
Data Path Width	128 bits	128 bits
Data Flow Direction	RX	TX
Number of Ports	10	10
Buffer Mode	Shared FIFO buffer with embedded addressing	Shared FIFO buffer with embedded addressing
FIFO Buffer Size	8,192 bytes	8,192 bytes

**Table 2. POS-PHY Level 4 Configuration Parameter Values Used for Testing (Part 2 of 2)**

<b>Parameter</b>	<b>Receiver</b>	<b>Transmitter</b>
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



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](mailto: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.

