

Introduction

This technical note describes a 1000BASE-X physical/MAC layer Gigabit Ethernet (GbE) interoperability test between a LatticeSC/M device and the Broadcom[®] BCM56800 network switch.

Specifically, this document discusses the following topics:

- Overview of LatticeSC[™] and LatticeSCM[™] devices and the Broadcom BCM56800 network switch
- 1000BASE-X physical/MAC layer interoperability setup and results

Two significant aspects of the interoperability test need to be highlighted:

- The BCM56800 uses a CX-4 HiGig[™] port, whereas the LatticeSC/M Communications Platform Evaluation Board provides an SMA connector. Therefore, a CX-4-to-SMA conversion board was used as a physical medium interface to create a physical link between both boards. The SMA side of the CX-4-to-SMA conversion board has four differential TX/RX channels (10-Gbps bandwidth total), but only one SMA channel (channel 0) was connected to the LatticeSC side.
- The 1000BASE-X Gigabit Ethernet standard specifies a 1-Gbps data rate (1.25-Gbps aggregated rate). Some Ethernet switch devices offer a 2.5-Gbps (3.125-Gbps aggregated) option for switched applications based on the 1000BASE-X standard. This applies to the LatticeSC/M device and the Broadcom BCM56800 network switch. The 1000BASE-X physical/MAC layer interoperability ran at a 2.5-Gbps data rate (3.125-Gbps aggregated rate). Both Lattice and Broadcom support this rate by running a faster speed reference clock that extends the bandwidth capability of 1-Gigabit Ethernet.

LatticeSC/M Overview

LatticeSC/M Features

The LatticeSC/M family, equipped with ASIC-based system level building blocks, was designed as a platform technology to facilitate the implementation of the many connectivity challenges that designers face today. This FPGA family includes features to meet the needs of today's communication network systems, including up to 7.8 Mbits of embedded block RAM, dedicated logic to support system level standards such as RapidIO, HyperTransport[™], SPI4.2, SFI-4, UTOPIA, XGMII and CSIX. Furthermore, the devices in this family feature clock multiply, divide and phase shift PLLs, numerous DLLs and a dynamic glitch-free clock MUX that are required in today's high-end system designs.

All LatticeSC/M devices also feature up to 32 channels of embedded SERDES with associated Physical Coding Sublayer (PCS) logic called flexiPCS[™]. flexiPCS logic can be configured to support numerous industry standard high-speed data transfer protocols. Each channel of flexiPCS logic contains dedicated transmit and receive SERDES.

flexiPCS in Gigabit Ethernet Mode

The Gigabit Ethernet mode of the flexiPCS (flexible Physical Coding Sublayer) block supports full 1000BASE-X compatibility, from the Serial I/O to the GMII interface of the IEEE 802.3-2002 Gigabit Ethernet standard. The flexiPCS also supports a 2.5-Gbps data rate (3.125-Gbps aggregated rate) by running a faster speed reference clock that extends the bandwidth capability of 1-Gigabit Ethernet.

The LatticeSC/M flexiPCS in 1-Gigabit Ethernet mode supports the following operations.

Transmit Path (from LatticeSC/M device to line):

- Cyclic Redundancy Check (CRC) generation and insertion into Gigabit Ethernet frame
- Transmit State Machine, which performs 8-bit data encapsulation and formatting, including the Auto-Negotiation code word insertion and outputting the correct 8-bit code/data word and k control characters according to the IEEE 802.3-2002 specification
- 8b10b Encoding

Receive Path (from line to LatticeSC/M device):

- Word Alignment based on IEEE 802.3-2002 defined alignment characters
- 8b10b Decoding
- Link State Machine functions compliant with the IEEE 802.3-2002 specification
- Clock Tolerance Compensation logic capable of accommodating clock domain differences
- Receive State Machine including Auto-Negotiation support compliant to the IEEE 802.3-2002 specification
- Cyclic Redundancy Code (CRC) checking

2.5-Gbps Ethernet Media Access Controller (2G5MAC) IP Core

The 2G5MAC soft core can support data rates of 1 Gbps or 2.5 Gbps in LatticeSC/M devices. The 2G5MAC transmits and receives data between a host processor and an Ethernet network. The main function of the Ethernet MAC is to ensure that the Media Access rules specified in the 802.3 IEEE standards are met while transmitting and receiving Ethernet frames.

Features:

- Compliant to IEEE 802.3z standard
- Generic 8-bit host interface
- 8-bit wide internal data path
- Generic transmit and receive FIFO interface
- Full duplex operation
- Transmit and receive statistics vector
- Programmable Interpacket Gap (IPG)
- Multicast address filtering
- Supports
 - Full duplex control using PAUSE frames
 - VLAN tagged frames
 - Automatic retransmission on collision
 - Automatic padding of short frames
 - Multicast and Broadcast frames
 - Optional FCS transmission and reception
 - Optional MII management interface module
- Supports jumbo frames up to 9600 bytes

Broadcom BCM56800 Overview

BCM56800 Features

The BCM56800 network switch is a high-density, 10-Gigabit Ethernet switching chip solution with 20 ports. Each of these flexible ports supports 10-GbE or 1-GbE. Additionally, the BCM56800 integrates all the SERDES required to interface to applicable copper and fiber physical interfaces. The integrated SERDES functionality includes 10-Gbps XAUI interfaces and 1-Gbps SGMII PHY interfaces. The integrated SERDES complies with the CX-4 standard and PICMG3.1 standard, which ensures interoperability with Ethernet line cards in an advanced TCA chassis.

BCM56800 GbE Ports

The BCM56800 has 20 10-GbE/1-GbE ports. The BCM56800 is based on the field-proven, robust StrataXGS™ architecture. This architecture offers integrated high-performance SERDES - integrated XAUI SERDES for all 20 10-GbE/HiGig ports, and it uses a single SERDES lane per port at GbE speeds. The device supports a 200-Gbps switching capacity at line rate.

In 10-Gbps modes (e.g., XAUI), each 10-GbE/HiGig port transmits and receives data on all four channels at 3.125 Gbps.

When configured in 2.5-GbE, only channel 0 of a 10-GbE/HiGig port is used.

Test Equipment

The equipment used in the interoperability tests is described below.

Broadcom BCM56800 Network Switch

Figure 1 shows the BCM56800 network switch.

Figure 1. Broadcom BCM56800 Network Switch



One can configure the Broadcom ports by connecting its serial port to a PC and starting a HyperTerminal session. Figure 1 shows a serial cable connected to the serial port at the back of the BCM56800.

Figure 1 also shows a CX-4 connector inserted into one of the four 10-GbE/HiGig ports available on the front right side of the BCM56800. This port, referred to as xe0, was selected for the interoperation with the LatticeSC/M device. It was configured in 2.5-GbE mode. In this mode, only channel 0 of the port was used.


LatticeSC/M Communications Platform Evaluation Board

The LatticeSC/M Communications Platform Evaluation Board provides a stable yet flexible platform designed to help the user quickly evaluate the performance of the LatticeSC/M FPGA or aid in development of custom designs. Each LatticeSC/M Communications Platform Evaluation Board contains among others:

- LFSC3GA25E-6F900C FPGA
- SMA test points for high-speed SERDES and Clock I/O
- Onboard power connections and power sources
- Onboard interchangeable clock oscillator
- Onboard reference clock management using Lattice ispClock™ devices
- Various high-speed layout structures
- On-board Flash configuration memory
- Various LEDs, switches, connectors, headers, SMA connections for external clocking, and on-board power control

Figure 2 shows the LatticeSC/M Communications Platform Evaluation Board.

Figure 2. LatticeSC/M Communications Platform Evaluation Board




ispVM System Software

ispVM System software is included with Lattice’s ispLEVER® design tool and is also available as a stand-alone device programming manager. ispVM System is a comprehensive design download package that provides an efficient method of programming ISPT™ devices using JEDEC and bitstream files generated by Lattice Semiconductor and other design tools. This complete device programming tool allows the user to quickly and easily download designs through an ispSTREAM to devices and includes features that facilitate ispATE™, ispTEST, and ispSVF programming as well as gang-programming with DLxConnect.

ispVM System is used in this interoperability test to download the LatticeSC/M bitstream, which configures the flexi-PCS in Gigabit Ethernet mode.

Figure 3 shows a screen shot of the ispVM System software.

Figure 3. ispVM System




ORCAstra System Software

The Lattice ORCAstra software is a PC-based graphical user interface that allows users to configure the operational mode of an FPGA by programming control bits in the on-chip registers. This helps users quickly explore configuration options without going through a lengthy recompile process or making changes to their board.

Configurations created in the GUI can be saved to memory and reloaded for later use. A macro capability is also available to support script-based configuration and testing. The GUI can also be used to display system status information in real time. Use of the ORCAstra software does not interfere with the programming of the FPGA portion of the LatticeSC/M FPGA.

Figure 4 is a screen shot of ORCAstra System.

Figure 4. ORCAstra System



Interoperability Testing

This section provides details on the 1000BASE-X 2.5 Gigabit Ethernet Physical/MAC layer interoperability between a LatticeSC/M device and the Broadcom BCM56800 network switch. This interoperability tests the correct processing of Gigabit Ethernet data from the BCM56800 network switch to the LatticeSC/M flexiPCS and 2.5-Gbps MAC IP core, and then back in the other direction. The test verifies the ability to transfer packets across the system in an asynchronous manner.

Test Setup

Figure 5 shows the LatticeSC/M and Broadcom board connections. Figure 6 is a block diagram of the test setup.

The setup includes:

- The Broadcom BCM56800 network switch.
- The LatticeSC/M Communications Platform Evaluation Board. In 2.5-GbE, the LatticeSC/M Communications Platform Evaluation Board uses an on-board 156.25-MHz differential oscillator to provide the reference clock to the LatticeSC/M flexiPCS. The flexiPCS multiplies the reference clock by 20 to achieve a 3.125-Gbps aggregated rate.
- A PC for software control/monitoring.
- A CX-4-to-SMA conversion board was used as a physical medium interface to create a physical link between both boards (see Figure 5). The SMA side of the CX-4-to-SMA conversion board has four differential TX/RX channels, or 16 SMA connectors for a total bandwidth of 10 Gbps (12.5-Gbps aggregated rate). However, for 2.5-Gigabit Ethernet, only one differential TX/RX channel (channel 0 with 4 SMA connectors) was used on the LatticeSC side (as shown in Figure 5).
- Cables
 - SMA for LatticeSC/M flexiPCS channel 0
 - SMA for Agilent clock generator
 - CX-4 for BCM56800 HiGig port xe0
 - ispVM JTAG cable for downloading LatticeSC bitstream and ORCAstra GUI access
 - Serial cable for BCM56800 HyperTerminal access

Figure 5. Board Connections





Figure 6. Test Setup Block Diagram



Test Description

This section describes how each interoperability partner is set up for GbE physical/MAC layer interoperability.

BCM56800

The BCM56800 switch generates and checks full protocol-compliant GbE packets. The BCM56800 is configured in 2.5-GbE.

Figure 7 illustrates the sequence of commands performed in a HyperTerminal from startup to configure HiGig port 0 (xe0) of BCM56800 in 2.5-GbE, while disabling auto-negotiation. To prevent port xe0 statistics counters from overflowing, a few lines are added to create Ethernet traffic on port xe1, and then redirect it to xe0.

As shown in Figure 7, xe1 generates Ethernet packets defined by the content of the bc_DA_1022 file and redirects the content to port xe0. Port xe0 then continuously outputs this data on channel 0 of its HiGig port. bc_DA_1022 is a VLAN tagged packet with 1022 bytes of data. The destination address in bc_DA_1022 is set to 00.00.00.02.

Figure 7. BCM56800 2.5-GbE Configuration Commands

```
'SETUP BC BOARD
'TRANSMIT PACKETS FOR 3600 seconds (1hr)
rc
'ENABLE VLAN PACKETS ON ALL PORTS
vlan remove 1 pbm=0x000000000001fffff
vlan add 1 pbm=0x000000000001fffff ubm=0x000000
vlan show
'SETUP XE0 AND XE1 IN 2.5 GBE MODE
'REDIRECT XE1 TRAFFIC TO XE0
fp qset clear
fp qset add InPorts
fp group create 1 1
fp entry create 1 1
fp qual 1 InPorts 0x1 0xfffff
fp action add 1 RedirectPbmp 0x2
fp entry install 1
fp entry create 1 2
fp qual 2 InPorts 0x2 0xfffff
fp action add 2 RedirectPbmp 0x1
fp entry install 2
port xe0 speed=2500 AN=OFF
port xe1 speed=2500 AN=OFF
port xe1 lb=phy
ps
clear counters
show counters
tx 4 pbm=xe1 file=bc_DA_1022
'RUN TRAFFIC FOR 1 HR, THEN SHOW COUNTERS
sleep 3600
port xe1 lb=none
sleep 1
show counters
```

LatticeSC/M Communications Platform Evaluation Board

The onboard 156.25-MHz clock oscillator sources the LatticeSC/M PCS reference clock for 2.5-GbE. The flexiPCS multiplies the reference clock by 20 to achieve a 2.5-Gbps data rate (3.125 Gbps aggregated rate).

In the RX direction, the LatticeSC/M SERDES recovers the 1000BASE-X packets from channel 0 of BCM56800 port xe0, and the flexiPCS converts them into GMII format. The data is presented to the MAC, which is configured in UNICAST mode with a local address of 00.00.00.02

The client loopback logic in the FPGA portion then loops the client data back into the TX direction. The LatticeSC/M device then transmits the 1000BASE-X packets back to the BCM56800 device.

Figure 8 shows the LatticeSC/M 2.5-GbE MAC Configuration script.

Figure 8. LatticeSC/M 2.5-GbE MAC Configuration Script

```
echo
echo Write MAC MODE REG to IDLE
load 00
write 00800

echo
echo Write 2g5 MAC address registers
load 00
write 0080a
load 00
write 0080b
load 00
write 0080c
load 00
write 0080d
load 02
write 0080e
load 00
write 0080f

echo
echo Write Test logic registers - no addr swap
load 06
write 00841

echo
echo Write MAC TX_RX_CTL --UNICAST MODE
load 9a 'normal mode
' load 9B ' promiscuous mode
write 00802

echo Write IPG REG
load 08
write 00808

echo
echo Write MAC MODE REG
load 0F
write 00800
```

Results

Figure 9 illustrates the section of the HyperTerminal output that resulted from running the last few commands in the BCM56800 2.5-GbE configuration sequence of Figure 7.

As shown in Figure 9, HiGig ports 0 (xe0) and 1 (xe1) of BCM56800 were configured for 2.5-GbE.

The last “show counter” command reports the status of BCM56800 port xe0 TX and RX packet (GTPKT.xe0 and GRPKT.xe0) and byte (GTBYT.xe0 and GRBYT.xe0) counters. Figure 9 does not show any error counters. This is an indication that the error counters have remained at a zero value during the test. Additionally, the Lattice ORCAstra System GUI (shown in Figure 4) was monitored for proper flexiPCS Gigabit Ethernet link state machine synchronization.

Figure 9. Output of BCM56800 Configuration Script

```

.
.

BCM.0> ps
      ena/  speed/  link  auto    STP      lrn  inter  max  loop
port  link  duplex  scan  neg?   state  pause  discrd  ops  face frame  back
xe0   up    2.5G FD    SW    No    Forward TX RX   None  FA  XGMII 16360
xe1   up    2.5G FD    SW    No    Forward TX RX   None  FA  XGMII 16360  PHY
xe2   down  -        SW    Yes   Forward TX RX   None  FA  XGMII 16360
xe3   down  -        SW    Yes   Forward TX RX   None  FA  XGMII 16360
xe4   down  -        SW    Yes   Forward TX RX   None  FA  XGMII 16360
xe5   down  -        SW    Yes   Forward TX RX   None  FA  XGMII 16360
xe6   down  -        SW    Yes   Forward TX RX   None  FA  XGMII 16360
xe7   down  -        SW    Yes   Forward TX RX   None  FA  XGMII 16360
xe8   down  -        SW    Yes   Forward TX RX   None  FA  XGMII 16360
xe9   down  -        SW    Yes   Forward TX RX   None  FA  XGMII 16360
xe10  down  -        SW    Yes   Forward TX RX   None  FA  XGMII 16360
xe11  down  -        SW    Yes   Forward TX RX   None  FA  XGMII 16360
xe12  down  -        SW    Yes   Forward TX RX   None  FA  XGMII 16360
xe13  down  -        SW    Yes   Forward TX RX   None  FA  XGMII 16360
xe14  down  -        SW    Yes   Forward TX RX   None  FA  XGMII 16360
xe15  down  -        SW    Yes   Forward TX RX   None  FA  XGMII 16360
xe16  down  -        SW    Yes   Forward TX RX   None  FA  XGMII 16360
xe17  down  -        SW    Yes   Forward TX RX   None  FA  XGMII 16360
xe18  down  -        SW    Yes   Forward TX RX   None  FA  XGMII 16360
xe19  down  -        SW    Yes   Forward TX RX   None  FA  XGMII 16360

BCM.0> clear counters
BCM.0> show counters
BCM.0> tx 4 pbm=xe1 file=bc_DA_1022
BCM.0> sleep 3600
Sleeping for 3600 seconds
BCM.0> port xe1 lb=none
BCM.0> sleep 1
Sleeping for 1 second
BCM.0> show counters
RUC.xe0      :          1,057,392,014      +1,057,392,014
GR1518.xe0   :          1,057,392,014      +1,057,392,014
GRPKT.xe0    :          1,057,392,014      +1,057,392,014
GRBYT.xe0    :      1,103,917,262,616  +1,103,917,262,616
GT1518.xe0   :          1,057,392,014      +1,057,392,014
GTPKT.xe0    :          1,057,392,014      +1,057,392,014
GTBYT.xe0    :      1,103,917,262,616  +1,103,917,262,616

```

The results show that over a billion Ethernet packets were successfully transmitted by port xe0 to the LatticeSC/M MAC and recovered at port xe0 error-free.

Summary

In conclusion, the LatticeSC/M FPGA family offers users a built-in 1000BASE-X 2.5-Gigabit Ethernet Physical/MAC layer support and is fully interoperable with Broadcom BCM56800 network switch.

Technical Support Assistance

Hotline: 1-800-LATTICE (North America)
+1-503-268-8001 (Outside North America)

e-mail: techsupport@latticesemi.com

Internet: www.latticesemi.com

Revision History

Date	Version	Change Summary
October 2007	01.0	Initial release.

HyperTransport is a licensed trademark of the HyperTransport Technology Consortium in the U.S. and other jurisdictions.