

Using Stratix GX in Switch Fabric Systems

Introduction

The switch fabric is one challenge designers face when creating a system that routes data from one of many inputs to any one of many outputs. The switch fabric is used in many types of applications, from high-speed telecommunications to networking to storage area networks.

Typical challenges designers face when implementing switch fabrics in their system include latency, bandwidth, queuing, backplane interfacing, scheduling algorithms, and traffic management interaction. To overcome these challenges, designers require a design platform that is flexible, customizable, and capable of real-time design and debug as well as in-field design updates/corrections. Stratix GX devices address all of these needs.

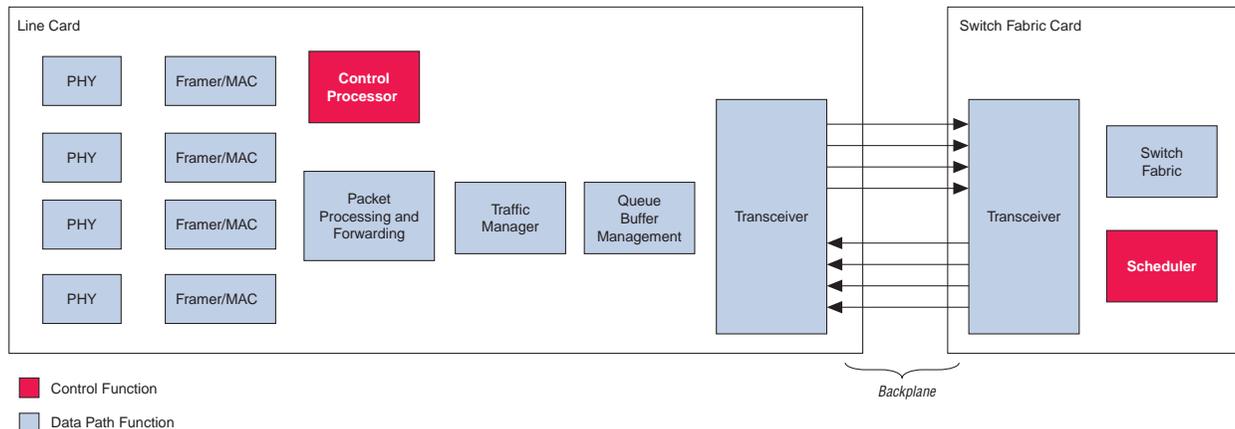
Built on Atera's feature-rich Stratix architecture, the Stratix GX device family allows hardware designers to transmit up to 20 channels of data at up to 3.125 gigabits per second (Gbps) each over a backplane and perform the necessary traffic management and switch fabric functions. The Stratix GX solution provides system designers with the flexibility, performance, integration, and design resources that are not available in any other solution.

System Overview

Figure 1 shows how a typical system uses switch fabrics. A typical system consists of multiple line cards, each containing a physical layer device (PHY), a MAC or framer, packet processing functions, a traffic manager, and a queuing system. The line cards are located in a chassis and electrically or optically connected by a backplane. This white paper provides information on using Stratix GX devices to implement traffic management, queue and buffer management, and backplane transceiver functions on a line card as well as the switch fabric, scheduler, and transceiver functions on a switch fabric card. A typical system uses a centralized architecture, where the switch fabric card connects all the line cards in the system. A distributed architecture does not use a centralized switch fabric card. Instead, the line cards complete the switching before sending data out over the backplane.

The PHYs and MACs/Framers process incoming network data. The MACs/Framers then pass the data on to the packet processor for forwarding, classification, prioritization and flow-control. Finally, the line card sends data to the switch fabric, which passes the data to other line cards. Switch fabric devices interconnect all ports on all the line cards in the system. The actual switch fabric devices are typically located on a separate card and are connected to the line cards using high-speed differential links utilizing CDR.

Figure 1. Data Flow Through One Line Card to the System's Switch Fabric Card



There are four main requirements for switch fabrics:

- Switch fabrics must provide a method to switch the packets from input ports to outputs ports.
- The switch fabric must arbitrate traffic when more than one packet arrives concurrently destined for the same output port.
- Switch fabrics must provide sufficient buffering to handle situations where the packet input rate is greater than the switch fabric's throughput capability. The two possible locations for buffering are at the input of the switch fabric (input queuing) or internally to the switch fabric (shared-memory).
- The switch fabric card must manage flow control on egress packets at the output of the switch fabric (output queuing).

Designing Stratix GX into a Switch Fabric System

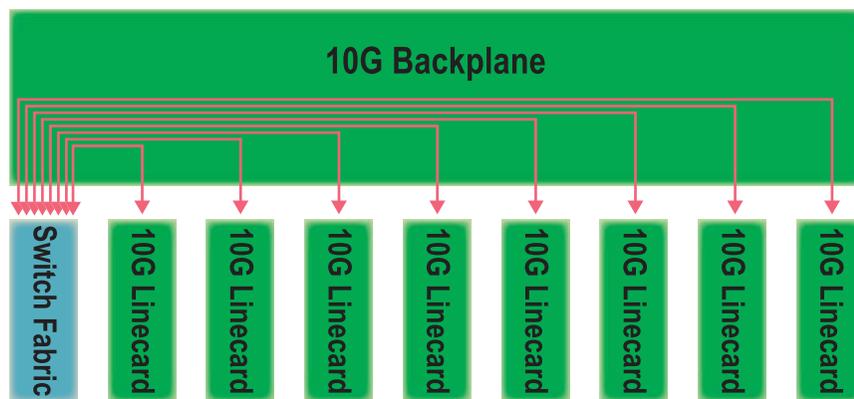
This white paper discusses the following two types of switch fabric architectures.

- Centralized (using the Crossbar function)
- Distributed

Centralized Switch Fabrics

A centralized switch fabric system contains a switch fabric card that is responsible for interconnecting all ports on all the line cards in the system. The switch fabric card routes the data and connects to the line cards using high-speed serial links across a backplane. Figure 2 shows a simplified eight line card system in which each line card connects to the central switch fabric card.

Figure 2. Centralized Switch Fabric System



Figures 3 and 4 show more detailed versions of a 4×3.125 -Gbps backplane interface on the line card and a switch fabric card with several sets of 4×3.125 -Gbps links. The designer can implement traffic management, buffer management, and the backplane transceiver in a Stratix GX device on the line card (see Figure 3). On the switch fabric card, designers can use Stratix GX devices as an integrated backplane transceiver/switch fabric solution. Figure 4 shows a Stratix GX device providing a total of twenty 3.125-Gbps serial link connections to five 10-Gbps line cards. A multi-Stratix GX device configuration can support systems with 16 or 32 line cards.

Figure 3. Backplane Interface on Line Card

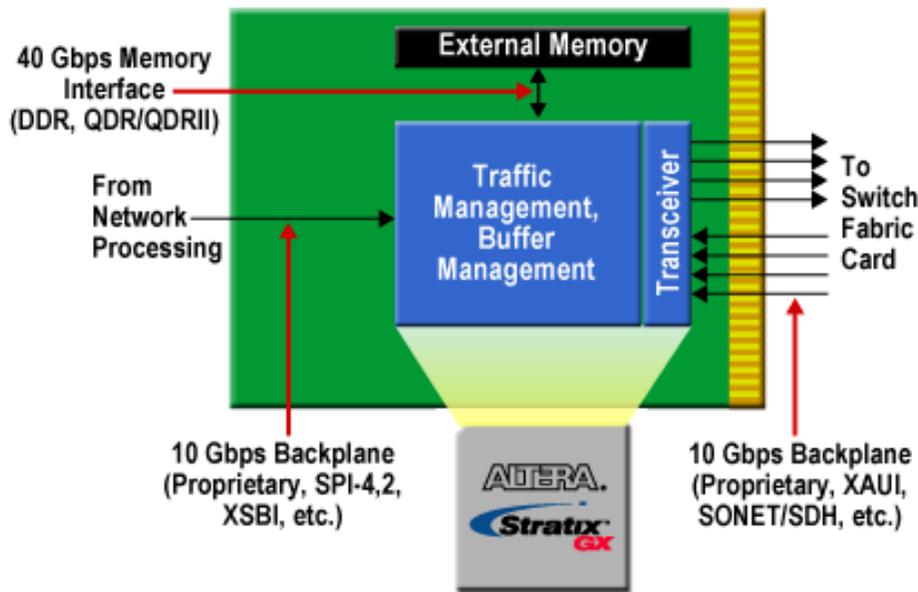
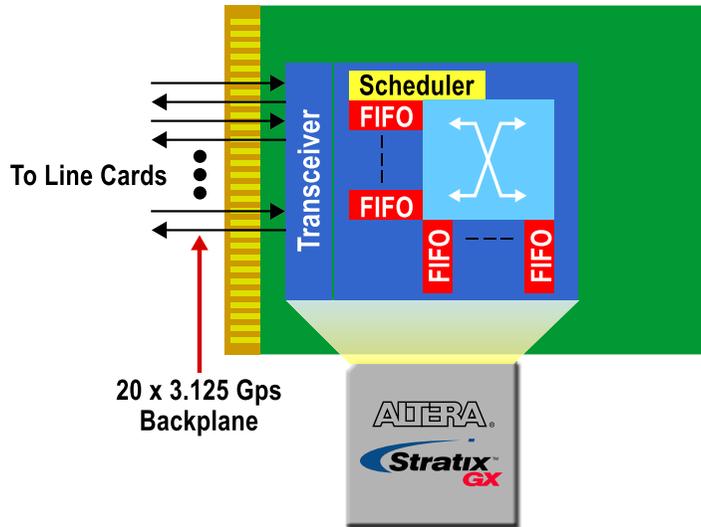


Figure 4. Backplane Interface on Switch Fabric Card



In addition to the backplane interface, Stratix GX offers critical features to ensure wire-speed data transmission. Stratix GX devices support high-speed chip-to-chip interfaces and offer 1-Gbps source-synchronous channels with dynamic phase alignment (DPA). The high-speed interconnect in the logic array can transmit packets at wire speed. Designers can implement packet buffering with virtual output queuing (VOQ) using the TriMatrix™ memory and support for high-speed external memory interfaces in Stratix GX devices. This advanced memory functionality allows the Stratix GX device to control the flow of varying input data rates.

The resources used for a sample implementation of a crossbar switch fabric in a Stratix GX device are shown in Table 1. The characteristics of the switch fabric design are as follows:

- 8×8 crossbar fabric
- Fixed cell size
- 16-bit data path
- Eight-cell VOQ depth
- 32-cell output queue depth
- Round robin scheduler

Table 1. Results of Crossbar Switch Fabric Design

Feature	Use
M4K blocks	96
Logic elements (LEs)	10,000
f_{MAX}	160 MHz
Transceiver channels	8

This design would fit in an EP1SGX25D device with over 15,000 LEs remaining.

Distributed Switch Fabrics

Stratix GX devices provide all of the functionality needed to effectively implement a distributed switch fabric solution. A distributed switch fabric implements the switching on the line cards before data is sent out onto the backplane. As shown in Figure 5, each line card connects to every other line card in the system using high-speed serial links, creating a high-speed mesh backplane. Each line card routes the packets before it sends them out to the appropriate line card. The Stratix GX device logic array provides ample programmable logic elements (LEs) to help process packets on the line card. Designers can implement traffic management, queue management, and the backplane interface in the Stratix GX device.

Figure 5. Distributed Switch Fabric System

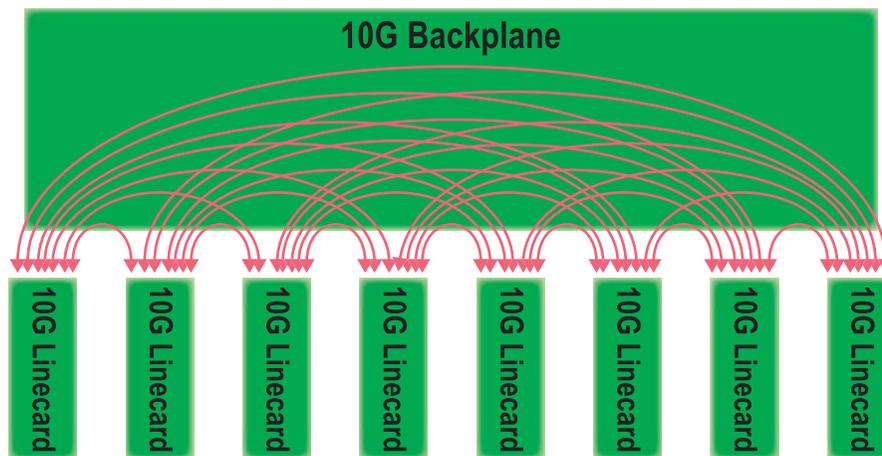
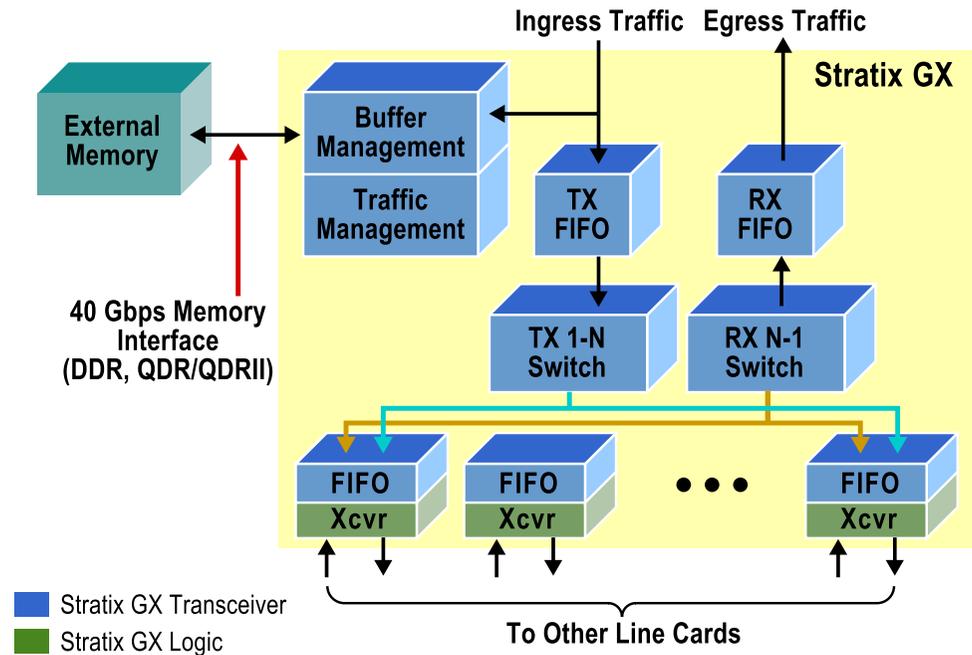


Figure 6 shows the details of an $N \times 3.125$ -Gbps backplane interface with an integrated distributed switch fabric system and traffic management and buffer management.

Figure 6. Integrated Backplane Transceiver with Traffic & Buffer Management



In addition to the backplane interface, Stratix GX devices offer critical features to ensure wire-speed data transmission. Stratix GX devices support high-speed chip-to-chip interfaces and offer 1-Gbps source-synchronous channels with DPA. The logic array high-speed interconnect transmits packets at wire speed. Stratix GX TriMatrix memory can implement packet buffering with virtual output queuing. Stratix GX devices can also interface to high-speed DDR, FCRAM, SDR, ZBT, and QDR/QDRII memory devices for external packet buffering. This advanced memory functionality allows the Stratix GX device to control the flow of varying input data rates.

The resources used for a sample implementation of a distributed switch fabric in a Stratix GX device are shown in Table 2. The characteristics of the switch fabric design are as follows:

- 16-line card system
- Fixed cell size
- 64-bit data path
- FIFO buffer depth (transmitter and receiver)
 - 128-cell line side depth
 - 32-cell backplane side depth
- Includes buffer and traffic manager with external memory interface

Table 2. Results of Crossbar Switch Fabric Design

Feature	Use
M-RAM blocks	1
M4K blocks	160
LEs	35,000
f_{MAX}	160 MHz
Transceiver channels	16

This design would fit in an EP1SGX40G device with over 5,000 LEs remaining.

Conclusion

Stratix GX devices provide programmable logic integrated with high-speed backplane interface capabilities and offer up to 20 integrated 3.125-Gbps transceivers. Additionally, Stratix GX devices support high-speed chip-to-chip interfaces and offer 1-Gbps source-synchronous channels with DPA. Stratix GX devices are also ideal for programmable digital functions such as scheduling, switch fabric, queue management, and traffic management. The superior external packet buffering is available through supported interfaces such as DDR, FCRAM, SDR, ZBT, QDR/QDRII. In addition, TriMatrix Memory offers over 400 Kbytes of internal buffering. Stratix GX devices offer the flexibility, time-to-market, and customizability for your switch fabric system.



101 Innovation Drive
San Jose, CA 95134
(408) 544-7000
<http://www.altera.com>

Copyright © 2002 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, maskwork 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.