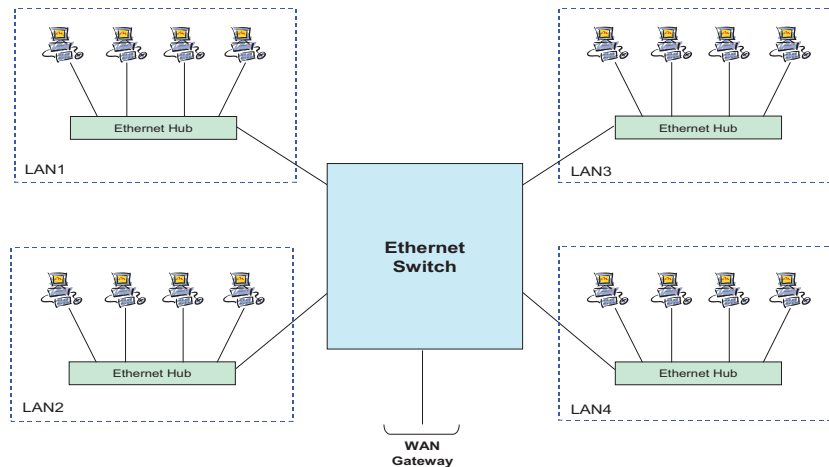


Increase Flexibility in Layer 2 Switches by Integrating Ethernet ASSP Functions Into FPGAs

Introduction

A Layer 2 Ethernet switch connects multiple Ethernet LAN segments. Because each port on the switch can be connected to a different LAN segment, this topology forms a larger Ethernet network. The switch stores the media access controller (MAC) address—observed in frames received through each port—to identify each network segment. Using the MAC address, the switch forwards frames from the source segment to the destination segment only, instead of forwarding the frame to all the connected ports, consequently reducing network traffic. Figure 1 shows the topology of a Layer 2 Ethernet network.

Figure 1. Switched Ethernet LANS

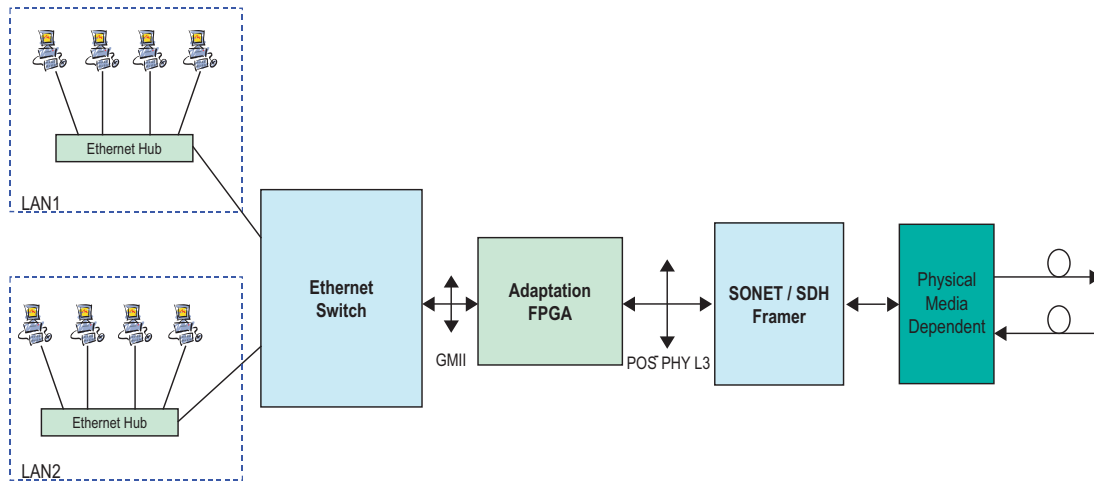


ASSPs for multi-port Ethernet switches are widely available from different vendors and can be used to meet the requirements of a number of applications. However, these ASSPs are designed for “typical” networking applications and do not provide a good solution for applications that require specific features, such as a configuration of an odd number of ports or a configuration of ports of varying speeds. This paper describes how, with a low-cost, programmable logic-based architecture, a Layer 2 Ethernet switch can be implemented to address these needs, increase integration, and enable highly customized solutions compared to an architecture based on ASSPs.

Limitations of Ethernet ASSPs for Layer 2 Switches

Ethernet switch ASSPs often implement physical Ethernet interfaces only, such as Medium Independent Interfaces (MIIs) for Fast Ethernet ports or Gigabit MIIs (GMIIs) for Gigabit Ethernet ports. These interfaces require a bridging device such as an FPGA, if, for example, one or more ports is connected to a PCI bus or a SONET/synchronous digital hierarchy (SDH) framer for Packet Over SONET (POS/SDH) or Ethernet Over SONET (EOS/SDH) applications, or to a voice processor for Voice-over-IP (VoIP) applications. Figure 2 shows an example of this implementation.

Figure 2. Discrete Application



In addition, Ethernet switch ASSPs implement a fixed number of ports (12, 16, 24, etc.) and a fixed distribution of ports, for example, eight 10/100 Ethernet ports plus one Gigabit Ethernet port. Standard Ethernet switches are a poor match for applications requiring differing numbers of ports or a different distribution of port speeds. Also, it is not possible to implement non-standard Ethernet ports such as a 2-Gbps port with ASSPs.

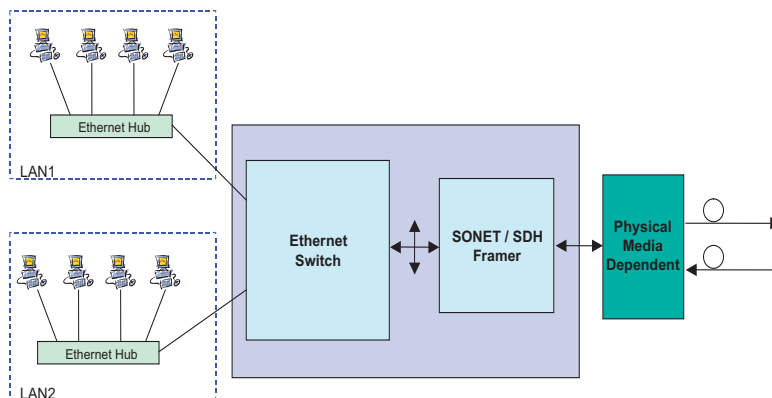
Programmable Logic-Based Ethernet Switch Enables Customization

To overcome the limitations of Ethernet switch ASSPs, MoreThanIP, an Altera® partner, developed a flexible Ethernet switching engine based on Altera FPGAs. The adaptable engine enables designers to develop single-chip solutions optimal for the following customization scenarios:

- Implements additional functions, such as PCI, POS-PHY/system packet interface (SPI), to complement the Ethernet switch
- Provides connectivity to a large number of standard parts such as SONET/SDH framers or VoIP processors
- Provides connectivity to systems such as proprietary backplanes or host computers via PCI/PCI-X

Figure 3 shows an example of this integrated application.

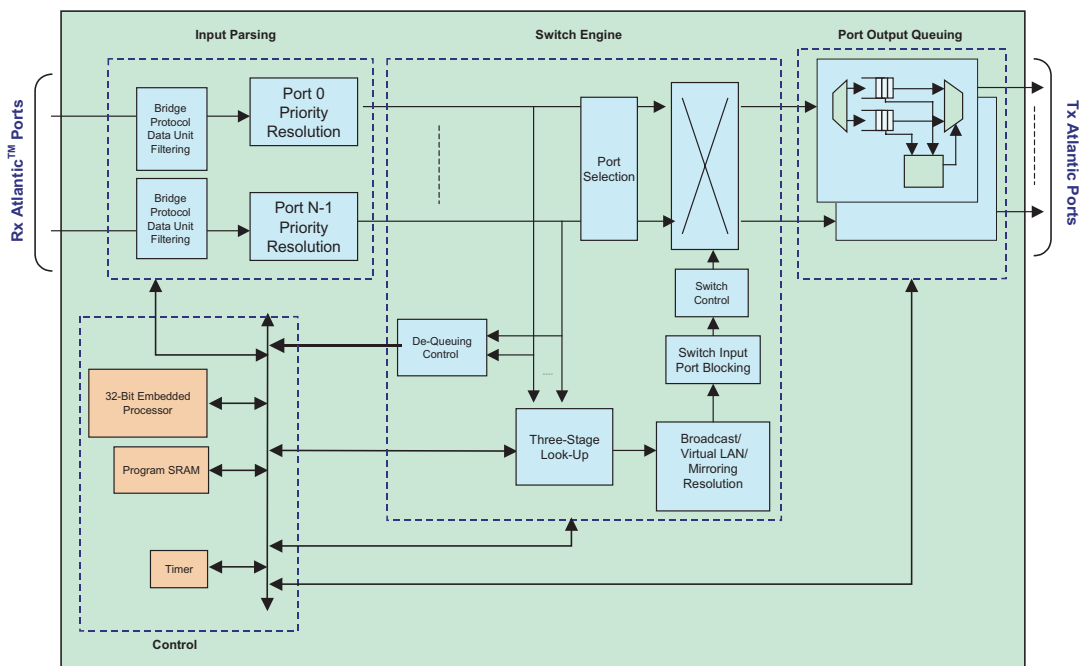
Figure 3. Integrated Application Using a Programmable Logic-Based Ethernet Switch



The Ethernet switch engine includes a hardware switch engine and a 32-bit soft-core embedded processor, which performs the table management tasks such as learning, aging, and migration. Using an embedded processor also enables designers to implement other high-level functions such as spanning tree algorithms, Ethernet termination, or other user-specific tasks. The switch implements a programmable number of ports, which are physically implemented with a simple master/slave FIFO-based interface on which any Layer 1 application or standard board-level interface can be connected.

The switch engine, shown in Figure 4, can be implemented in Altera's Stratix® II and Cyclone™ II devices and utilizes a Nios® II embedded processor for the table management tasks. When implemented in the latest high-density FPGAs, the Ethernet switching engine provides up to 8-Gbps switching capabilities. This bandwidth can be distributed flexibly among any number of ports to meet the specific needs of custom applications. For example, the switch can be configured to implement 4-Gbps Ethernet ports and 12 Fast Ethernet ports, or it can be configured to implement non-standard configurations such as a 2-Gbps port with multiple 200-Mbps ports. The switch also allows customization for adding higher quality-of-service (QoS) for certain classes of services.

Figure 4. Ethernet Switch Engine Block Diagram



Switch Engine Functional Overview

Frame switching is based on a two-stage hash code look-up associated with linear searching. This solution provides high performance and enough flexibility to extend the maximum number of MAC addresses supported by the standard implementation of the switch (2048 addresses). To provide maximum performance and a non-blocking operation, frame switching is performed at wire speed without any software or firmware overhead.

A 32-bit Nios II processor manages the switch look-up table (LUT). To control the switch, firmware was developed to perform the following tasks:

- MAC address learning
- LUT entry aging
- Port migration
- Hash code and table management

The switch also optionally implements a classification engine per input port and two prioritized queues per output port, which provide QoS to critical classes of services. The hardware and software parts of the switch have been designed for performance and are tightly coupled. In addition, the switch architecture is easily scalable to support higher throughputs such as a 10-Gigabit switch core.

QoS Support—Frame Classification

When a frame is received on an input port, several information fields are extracted from the frame header including the Ethernet MAC address, VLAN tag (IEEE 802.3q), the frame PDU, and Internet protocol (IPv4 and IPv6) headers. This information is used to determine the frame type, to classify the frames with up to eight levels of priorities, and to store the frames in output high and low priority queues. The classification is performed with a set of programmable tables, controlled by the embedded Nios II processor, and a programmable priority resolution tree.

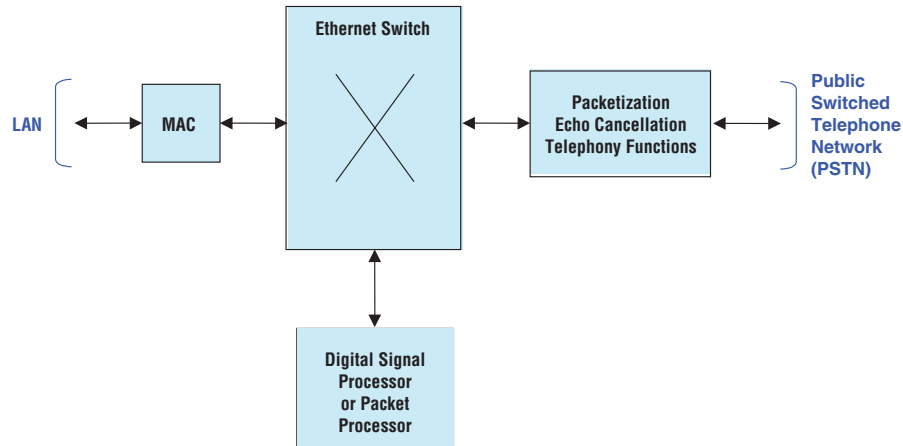
QoS Support—Scheduling

The queue arbitration implements a strict priority or a weighted round-robin algorithm to serve the output queues. In addition, the flexibility of the FPGA-based implementation allows the developer to use any custom proprietary arbitration scheme.

Application Example—VoIP Gateway

A VoIP gateway bridges a traditional time-division multiplexed (TDM) network to an IP network. To convert TDM traffic to IP packets and to perform echo cancellation and telephony functions, dedicated packet processor ASSPs or digital signal processors are typically used. The networking functions are implemented with an Ethernet switch ASSP and a MAC. A processor is also typically required for management and signaling functions. Figure 5 shows a typical VoIP gateway block diagram.

Figure 5. VoIP Gateway Block Diagram

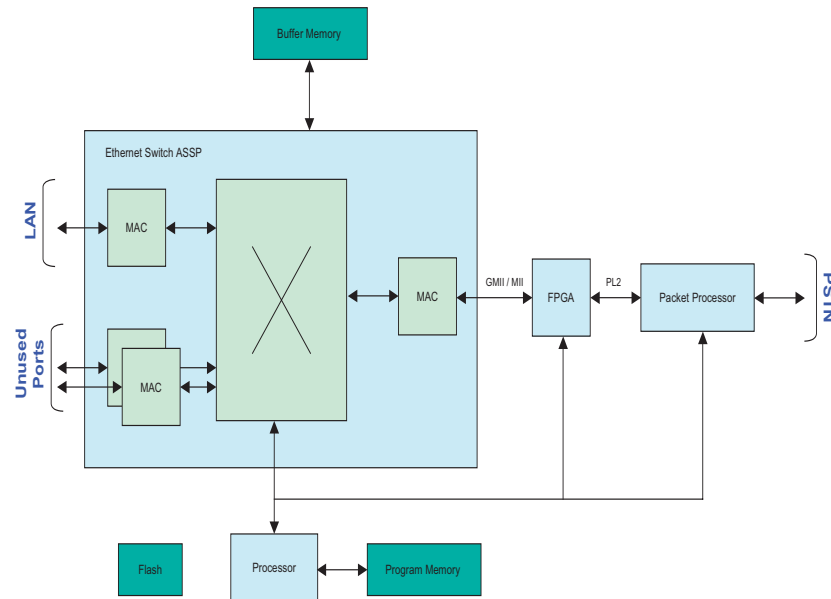


To implement a VoIP Gateway with ASSPs, the following devices are required:

- Digital signal processor or packet processor
- Ethernet switch
- Processor
- Memory
- System interface devices

These devices may have incompatible interfaces that need to be bridged and require additional devices, generally FPGAs, to do so. In addition, for a small number of voice channels, typical Ethernet switch ASSPs provide an over-sized solution consisting of a high number of non-scalable ports. Figure 6 shows an ASSP-based VoIP gateway.

Figure 6. VoIP Gateway Implementation—ASSPs



Ethernet Switch Engine-Based Solution

A highly integrated and cost-effective solution can be designed using an Altera FPGA and networking IP blocks from MoreThanIP. Figure 7 shows an example of a Stratix II FPGA-based implementation. The scalable switch IP block can be configured to implement the required number of ports (three are used in Figure 7) or more as required by the application, eliminating the waste of unused ports. The embedded Nios II processor can run signaling and management functions as well as quality and performance monitoring functions. The FPGA interface to the digital signal processor or packet processor can be implemented on the switch port without requiring an external bridging device.

Figure 7. VoIP Gateway Implementation—Stratix II FPGA

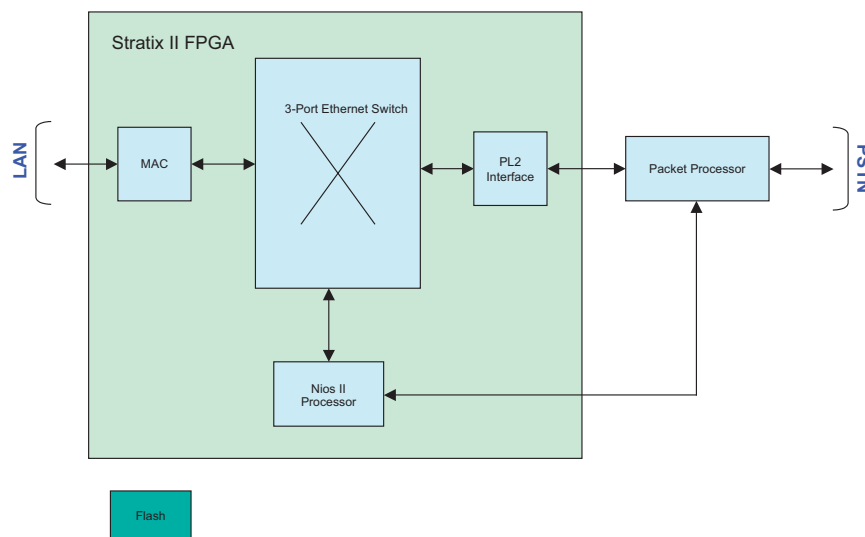


Table 1 summarizes the device utilization of the FPGA in terms of memory bits and logic elements (LEs)—the basic building blocks of Altera FPGAs. Table 2 shows target Stratix II speed grades and their switching capacities

Table 1. Gateway Stratix II FPGA Resource Utilization Summary

Function		LE Utilization	Memory Bits Utilization
10/100/1000 Ethernet MAC		3,500	30K
Ethernet switch		3,000	200K
Embedded control processor		4,000	16K
MPHY POS-PHY Layer 2		3,000	-
Optional QoS	Classification	500	-
	Scheduling	700	-
	Prioritized queues	-	120K
Total		14,700	366K

Table 2. Stratix II-Based Switch Performance

Device Speed Grade	Maximum Clock Frequency	Switching Capacity
-5	190 MHz	6 Gbps
-3	260 MHz	8 Gbps

Ease RoHS Transition With Altera Lead-Free Products

Altera maintains one of the most extensive lead-free product offerings in the industry, with over 1200 products in lead-free packages. As a preeminent supplier of environmentally friendly programmable logic solutions, Altera has shipped over 25 million lead-free products since 2002. Altera’s lead-free devices comply with the maximum concentration restrictions, as required in the EU Directive on the Restriction of Hazardous Substances (“RoHS Directive”) No.2002/95 with respect to lead (Pb), mercury, cadmium, hexavalent chromium, polybrominated biphenyls (PBB), and polybrominated diphenyl ethers (PBDE). Help ease your RoHS transition by integrating non-compliant ASSPs with Altera’s PLDs.

Conclusion

With the availability of Altera's high-performance, high-density programmable logic devices such as the Stratix II family, new integration and customization possibilities for Layer 2 switches are now available. When combined with external SDRAM memory, the Cyclone II-based version of this solution provides designers with a lower-cost implementation option for applications requiring less switching capacity. These programmable devices can achieve higher integration and enable more highly-optimized switch solutions than ASSP-based solutions, providing system developers the opportunity to cost-effectively differentiate their products.

Resources

For additional information, refer to the following resources on the Altera website.

- More System Integration Solutions
www.altera.com/technology/integration/int-index.html
- Customer Applications in Wireline Communications Products
www.altera.com/wirelinesuccess
- MoreThanIP Ethernet Layer 2 Switch IP Core
www.altera.com/products/ip/iup/ethernet/m-mtip.html
- AMPP Partner Profile on MoreThanIP
www.altera.com/products/ip/ampp/morethanip/morethanip.html



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

Copyright © 2006 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.