
Traffic Management for Testing Triple-Play Services

Introduction

Triple-play services (voice, video, and data) are emerging as the key driver for telecom investment, with several factors driving this convergence. The Internet is maturing, not only for the transport of web information, but for e-commerce as well. Increasingly sophisticated protocols are emerging to manage and secure the network for business and consumer alike. With this acceptance comes increased scale, ubiquity, and cost effectiveness, far beyond regional telephone and cable networks. Additionally, the technology to digitally represent, store, and transport all types of data continues to accelerate. Managing these services on a common Internet protocol (IP) network leads to a dramatic reduction in operational expenses. Lastly, the deregulation of the communications industry has opened a Pandora's box of bundled services from incumbent suppliers and startups alike. The increased competition from the cable industry into data and voice services, coupled with the decreased profitability of traditional long-distance telephone, has lit a fire under traditional data carriers to begin serious investigation and deployment of triple-play capabilities via the Internet.

Triple-Play Requirements

Triple-play requirements are driven by both network technology and geographic influences. On the technology side, the transition of DSLAMs from Asynchronous Transfer Mode (ATM) to IP has been dramatic. Beginning in 2006, gigabit passive optical networking (GPON) and Ethernet-based access networks are also expected to replace ATM-based passive optical networks (PON). Although packet-based data transport networks such as Metro Ethernet are entering the mainstream, services such as IP television (IPTV), video on demand, and voice over IP (VoIP) are now being developed and deployed. Consequently, IP-based quality of service (QoS) needs to be enforced throughout the network to ensure critical, latency-sensitive traffic is appropriately prioritized. Additionally, the QoS requirements vary dramatically depending on the type of equipment being deployed, thus requiring very flexible traffic management implementations.

Regional differences among service providers further complicate network implementation with different legacy infrastructures and standards. For example, in North America, voice services are increasingly commoditized as consumers replace their local and long-distance phone service with comparable services from cable, VoIP, and wireless providers. In reaction, regional Bell operating companies (RBOCs) are responding by investing in high-speed access networks to deliver bundled voice, data, and IPTV services. In the greenfield environment of the Asia-Pacific region, where there are fewer legacy networks, government initiatives and innovative network operators have been instrumental in driving investment into packet-based networks such as EPON, Metro Ethernet, and RPR. European providers are also beginning the transition from ATM access networks. British Telecom's 21st Century Network is one example of many all-IP network investments, and several providers are beginning to roll out IPTV services.

With such a wide range of factors driving the testing requirements for triple play (such as emerging Internet protocols and QoS capabilities), no longer can all market requirements be met with fixed ASSP and ASIC solutions. Today's equipment must measure the device under test (DUT) response in terms of latency, throughput, missing packets, transaction rate, and mean opinion score (MOS). Further, equipment must be able to assign different QoS parameters (physical/logical ports, priorities, classes, distribution, etc.), allowing it to test a DUT's ability to correctly implement QoS policies on its traffic.

Triple-Play Testing

Test equipment for triple play must create or emulate many different types of Layer-4-7 traffic (HTTP, FTP, email, video streams, audio streams, and VoIP). Each traffic pattern needs to be tested under different QoS characteristics to effectively test a DUT's ability to correctly handle these different types. Measuring the performance for triple-play services means emulating and evaluating protocols in a step-by-step approach.

1. Create several baseline traffic types and measure the performance of a DUT when the traffic is run in isolation.
2. Combine all traffic types and re-assess the performance in terms of throughput, latency, and data loss.
3. Adjust QoS parameters on certain traffic flows and implement QoS policies on the DUT to measure its ability to properly prioritize certain streams within a triple-play environment.

Internet baseline traffic consists of web accesses, mail, ftp, P2P, and other forms of business traffic based on different port pairs (for instance, HTTP requests target TCP port 80, POP3 transactions target port 110, etc.). Video baseline traffic can be actual video streams or emulated through use of scripts. These streams simulate the behavior of video traffic through the DUT. Delay, jitter, and throughput must be measured. Both unicast for video on demand and multicast for broadcast services over UDP are required. VoIP traffic must test bi-directionally over multiple channels (48 VoIP pairs) with several different types of codec algorithms (G.711u, G.711a, G.723.1-ACELP, G.723.1-MPMLQ, G.729, and G.726). The quality of the voice calls is measured using the MOS of voice conversations and thus determines the effectiveness of the network for carrying voice traffic.

Once the baseline types are characterized, the traffic is combined and re-evaluated, usually with one type earmarked with QoS and the others without, until all combinations are measured. Then various QoS parameters are applied to quantify traffic throughput and response times under different QoS combinations. Generally, Internet baseline traffic has the lowest priority, since data services are not adversely affected by packet delays. Video traffic may have the next highest priority, since a few missing frames of video will not seriously degrade perceived appearance as long as the audio track (typically sent as a stream) is not broken. VoIP traffic typically has the highest priority, as voice services are very sensitive to latency and corrupted data problems.

Next-Generation System Architecture Challenges

Although internationally ubiquitous and highly cost efficient, the Internet and its transport language, TCP/IP, were not inherently designed to support time-sensitive traffic like voice and video. Thus, new emerging protocols and hardware providing QoS via traffic management are necessary to handle latency-sensitive traffic. These dramatic changes in communication equipment architectures create an enormous impact on the OEM suppliers. Successful OEMs must adapt to the triple-play environment with highly flexible architectures, enabled by FPGA-based network processing and traffic management functions.

A typical communications tester is a chassis-based heterogeneous architecture with line interface cards, network processing/traffic management cards, switch fabric cards, and a control CPU (smaller models may combine these functions together). Due to the variety of line interfaces and multiple protocols (classification, editing, and policing) to be processed, multiprotocol framer and network processor ASSPs have been developed. Many of these ASSPs require FPGA bridging solutions. Some OEMs have realized today's powerful FPGAs can now do the framing and network processing functions, too, providing significant flexibility and cost savings.

Previous data-only packet equipment had little need for traffic management for line interfaces (primarily for buffering under backpressure mechanisms) and only modest QoS requirements for the switch fabric. As the line interface types diversified and triple-play QoS requirements intensified, heavier traffic management became essential for both directions, each with different requirements for ingress and egress. FPGA-based traffic management is therefore an ideal solution, particularly for triple-play testing.

Altera Traffic Manger Overview

Altera has developed a 10-Gbps traffic manager (TM) solution targeting Stratix® II FPGAs. The inherent programmability of the FPGA allows the TM to be adapted to support changing test requirements and emerging services. This TM meets guaranteed line-rate performance for Ethernet, packet over SONET/SDH (POS), MPLS, and ATM traffic. The design was built in a modular fashion, enabling system designers to customize partial or complete sub-blocks of the TM.

Packet Flow: Ingress TM

The TM interfaces to a network processing unit (NPU) on one side and a fabric interface chip (FIC) on the other. In the ingress direction, traffic flows into the line card encapsulated in Ethernet, SONET/SDH, RPR, or OTN frames. The framer or MAC device performs the necessary processing and transmits the traffic to an NPU with a Layer 2 header attached. The NPU performs classification, modification, and forwarding operations as dictated by the protocols being processed. The NPU also adds additional headers to the packet for communicating information to downstream devices, including both ingress and egress TMs, as well as the egress NPU. The header for the ingress TM contains information such as class of service (CoS), multicast, and drop precedence, which are necessary for the device to appropriately prioritize the traffic.

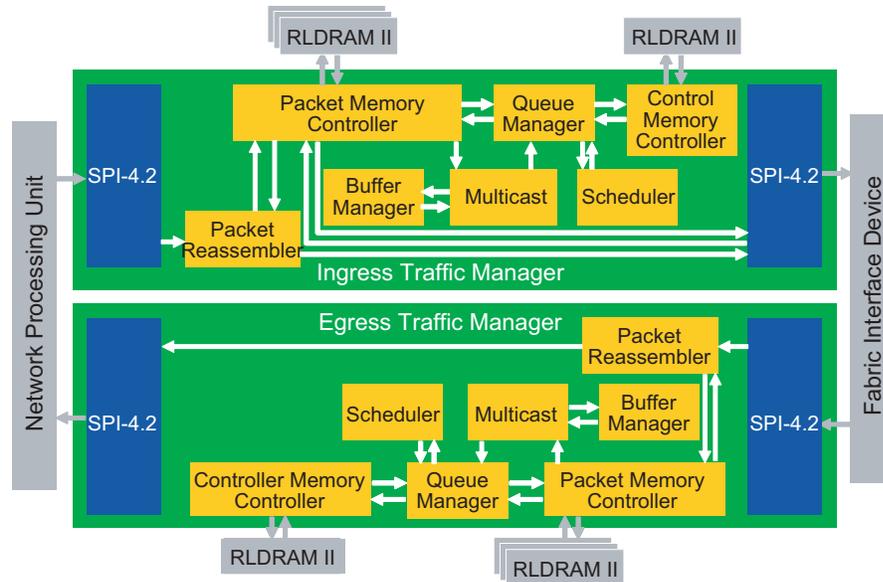


Figure 1. Altera's 10-Gbps TM

Inside the TM (Figure 1), the NPU interface utilizes the Altera® SPI-4.2 MegaCore® function. This core uses Altera's Atlantic™ flow-through packet interface, which enables designers to easily replace the SPI-4.2 interface with another (possibly proprietary) interface that also supports Atlantic.

The packet reassembler block receives data from the SPI-4.2 interface and parses the header for relevant parameters (such as CoS). The parsing can be customized to locate parameters within any header. This information is stored in the control memory, along with pointers to the packets. The reassembler block converts the received data into fixed-sized cells and communicates with the buffer manager block to ensure that the configured memory partitions are not violated by packet enqueues.

The buffer manager passes the enqueue request to the queue manager block, which maintains state and pointer information about each of the queues. If the enqueue is valid, the reassembler block selects an available pointer to external memory and passes the segmented packets to the packet-memory controller for storage in external memory. The list of available pointers is maintained both off-chip in the packet memory and in an on-chip cache using the FPGA's embedded memory. After the packet has been written to external memory, it is eligible to be transmitted by the scheduler.

The scheduler examines all ports that have data to send, then chooses the next port according to a hierarchical scheduling scheme. The scheduling algorithms are configurable and the entire scheduling block is customizable in order to support proprietary scheduling algorithms. When the scheduler has selected a cell from a packet for

transmission, it issues the request from the queue manager. The queue manager initiates the dequeue operation through the reassembler block.

Packets that have been scheduled are sent through the FIC interface. The current FIC interface also utilizes the SPI-4.2 MegaCore function. A header is added to each of the cells in order to enable the fabric to switch the cells to the appropriate port with the appropriate priority. The cells flow through the switch fabric, then the egress FIC converts the traffic to a SPI-4.2 interface for the egress TM.

Packet Flow: Egress TM

The egress TM's data flow is similar to that of the ingress TM. It also interfaces to the NPU through a SPI-4.2 interface, which determines the appropriate Layer 2 header to place on the packet for sending to the framer or MAC.

TM Functional Blocks

This section reviews the functionality of the individual blocks of the TM. Each function is designed to be a module that can be customized or replaced.

Scheduler

Traffic scheduling ensures that, during times of congestion, each port and each CoS gets its fair share of bandwidth. The scheduler interacts with the queue manager block, notifying it of scheduling events and receiving information about queue lengths. In addition to the scheduling of traffic, the scheduler is also responsible for performing the scheduling of RLDRAM refresh cycles for the external control memory.

Multicast Replication

As carriers add to their offerings by supporting streaming audio or video services, the need to replicate packets efficiently becomes essential. Altera's TM supports multicast replication in either the ingress, egress, or both directions, depending on the application requirements. The multicast replication block receives all packet enqueue messages (unicast and multicast) from the packet reassembler. Multicast packets are placed on a work queue (after passing a multicast partition check). When the multicast packet reaches the head of the work queue, the replication logic dequeues the packet and replicates.

Altera's TM provides both multicast and unicast support:

- Multicast replication supported on ingress and egress
- Dedicated multicast queue for each of eight classes
- 4K multicast groups supported
- 32K multicast leaves supported

Queue Manager

The queue manager block is responsible for initiating packet enqueue and dequeue requests based on notifications from the buffer manager (for enqueues) and the scheduler (for dequeues). The queue manager also maintains a prefetch FIFO buffer to cache packet descriptors and manages the head/tail pointers of each of the externally stored queues. The prefetch cache can be implemented using on-chip memory for applications with a low number of queues, or by using external QDR II SRAM for applications requiring a high number of queues.

Altera's TM provides a robust queue manager with:

- 10-Gbps line rate performance
- On-chip prefetch cache that reduces external memory

Memory Controllers

High-end traffic management applications require memory management that can achieve high bandwidths with reduced latency. Memory bandwidth is a major bottleneck in TM applications above 2.5 Gbps. Memory is required to

buffer packets, store pointers to the buffered packets, maintain control and state information for the queues, and collect statistics for billing and maintenance purposes. Development of efficient memory management architectures can be deceptively complex. Specialized external memories, such as RLDRAM II, are often required to meet the fast random access and response times.

To interface to these external memories, high-end traffic-management solutions require a large number of pins. In applications where pin counts drive the die size, standard product solutions do not have pricing advantages over FPGA implementations.

Altera's TM provides high-performance memory controllers and interfaces that:

- Use RLDRAM II
- Store up to 0.5M packets in 128 Mbytes (4.32-Mbyte RLDRAM II memories)
- Use RLDRAM II and QDR II for statistics and control memory

The Altera TM uses RLDRAM II memories for buffering packets and storing packet descriptors and statistics. The packet storage memory comprises four RLDRAM II devices in each direction, organized in two groups of two memory chips. The RLDRAM II interface is a bidirectional, double-data-rate memory interface running up to 300 MHz. The internal core clock is edge-aligned at one-half the frequency.

To achieve the high-bandwidth requirements for 10G traffic management, the Altera TM also uses a time division multiplex (TDM) strategy. This scheme dedicates specific time slices for the read direction and the write direction. It also dedicates no operation (NOP) and turn-around (TA) time slots to meet the required row-cycle times (tRC) and read and write latencies required by the RLDRAM II specification.

The flexibility of the FPGA implementation supports a variety of sequences for the TDM scheme, which allows system designers to optimize the memory management for a particular application.

A script was written to generate the memory bandwidth for each of the TDM structures at various frequencies. The results of the throughput analysis were then graphed to determine if the memory bandwidth was sufficient to sustain the packet throughput for each packet size. Figure 2 shows the results of the throughput, highlighting the ability to maintain line rate performance for 10-Gbyte SONET/SDH down to 40-byte packets, and Figure 3 shows the Altera TM demonstration board, developed to speed OEM development time.

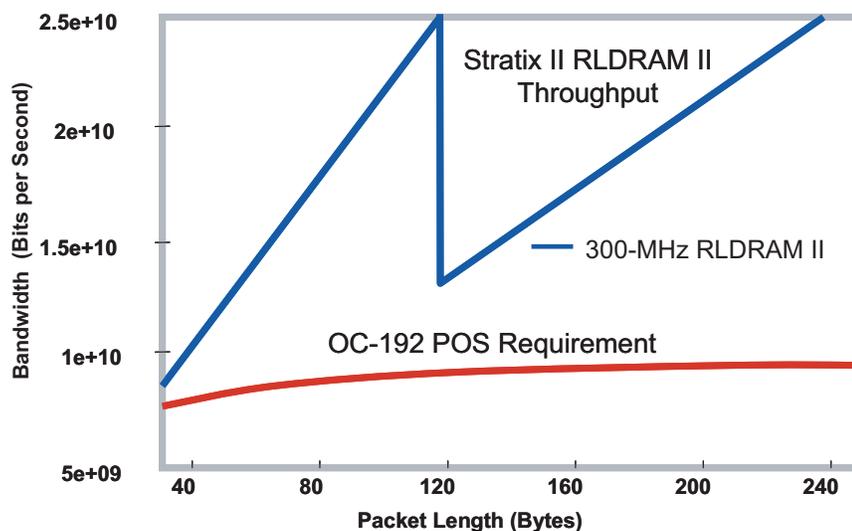


Figure 2. RLDRAM II Throughput for POS in Stratix II Devices

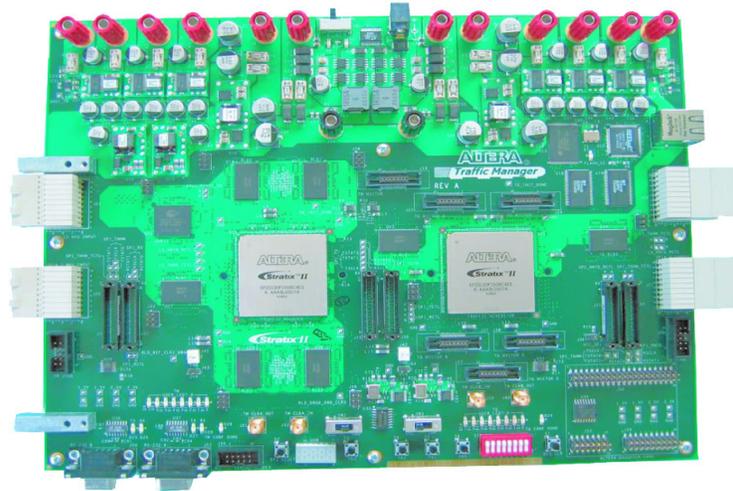


Figure 3. Altera Traffic Manager Development Board

Conclusion

By using FPGA-based silicon technology to enforce QoS, communications equipment OEMs can meet the changing needs of their customers by rapidly delivering customized test services to specific markets. Altera's TM provides OEMs the benefit of:

- Future-proofing their TM with the inherent flexibility of the FPGA
- Differentiating their solution using Altera's modular building blocks
- Reducing development costs and time-to-market by leveraging proven blocks
- Meeting 10-Gbps throughput for high-performance applications
- Reducing board space and cost by integrating external co-processing or bridging functions
- Scaling to adapt the solution across multiple platforms or cards

Altera's 10-Gbps TM solution meets the demands of next-generation networks by supporting high-speed throughput in a solution that can adapt to the changing market. The solution can be demonstrated in a hardware environment using Altera's demonstration board.

Further Information

- Other Altera communication IP and reference design offerings:
www.altera.com/technology/tc-index.html
- For more information on this solution, contact your Altera salesperson:
www.altera.com/corporate/contact/sales/na_reps/con-us_reps.html
- Triple Play Testing with Ixia IxChariot:
www.ixiacom.com/products/display?skey=ixchariot



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.