

Customizing Multi-Service Access Network Silicon

Introduction

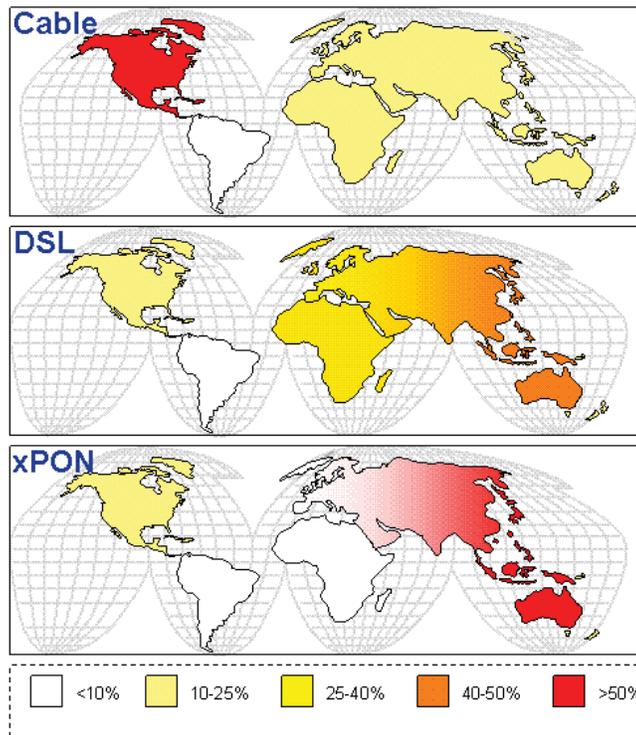
Network operators continue to invest in upgrading their access networks to improve delivery of high-bandwidth services. This investment has stimulated innovation throughout the supply chain to support the bandwidth, reliability, and scalability necessary for the continuing evolution of networks. Innovative markets, such as today's communications access market, tend to value technologies that can accelerate time to market while also adapting to changing requirements. FPGAs have been at the forefront of the rollout of multi-service access node (MSAN) equipment. This white paper discusses the evolution of access networks, the requirements for MSAN equipment, and the trend of access equipment manufacturers utilizing FPGAs to build custom solutions.

The Evolution of Access Networks

Broadband access networks are evolving from copper-based infrastructure to a mix of fiber and copper networks. Different technologies such as ADSL, VDSL, cable, Ethernet passive optical network (EPON), broadband PON (BPON), and Gigabit PON (GPON) are deployed in various regions of the world, and the development to support increasingly higher data rates has led to differing regional requirements.

As an example, Figure 1 highlights the regional use variations for broadband deployment. Infonetics forecasts that by 2010 over half of the cable broadband subscribers will be located in North America due to the high penetration of cable television service within the region (1). PON subscribers are expected to grow dramatically at a compound annual growth of 150% through 2010 in North America and Asia Pacific, with GPON gaining traction in China and North America, and EPON dominating Japan. European broadband subscribers will largely continue to use DSL, although each region has its own variations within it.

Figure 1. Percentage of Worldwide Broadband Subscribers by Geography by 2010



Based on these variations, standards organizations, such as the DSL Forum, target specifications for “agnostic access network architectures that deliver inherent quality, scalability, resiliency, and inter-working capabilities that enable services to be delivered via multiple business models.”⁽²⁾ The emerging network requires an access platform that can be configured uniquely for a specific deployment.

Birth of an MSAN

In 2004, British Telecom (BT) first announced an initiative to evolve their network, called the 21st Century Network (21CN). The stated objectives of this strategy include:

- To empower the customer with control, choice, and flexibility in a way that is simply not possible today
- To develop exciting new services for customers much faster than before
- To reduce costs and grow cash cost savings expected to amount to exceed \$2 billion per annum by 2008/2009 ⁽³⁾.

To deliver on these objectives, the 21CN architecture defined an access node with the capabilities to deliver voice, Internet protocol (IP), Ethernet, or fixed line services from a common platform. This MSAN enables BT to reduce their network elements while also supporting an evolution from a fixed network to a flexible IP network. As carriers around the world invest in infrastructure to support new services, this MSAN architecture will allow the evolution to support the necessary technology requirements.

Equipment Manufacturer Challenges

The market opportunity for access equipment manufacturers is considerable. Lightreading Insider forecasts a \$4.7 billion market for GPON equipment in 2011 ⁽⁴⁾. Infonetics forecasts the market for DSL and multi-service access platforms will soon surpass \$6 billion by 2009. The opportunity is significant, but so are the challenges.

Equipment manufacturers must handle both the technical and business challenges of the access market. These challenges include:

- Accelerating time to market
- Reducing development cost and protecting engineering investment
- Adapting to changing network demands
- Supporting tomorrow's scalability requirements

These challenges must be balanced with the seemingly unattainable goal of also reducing cost, power, and risk.

Accelerating Time to Market

In an emerging market, such as the GPON deployments within North America, the ability to be the first to deliver a working solution can be the difference between success and failure. The urgency of North American telco providers to combat the threat from cable operators has increased the need for accelerating time to market. After a high-volume market, such as the access market, has been supplied with a workable solution, the requirement to drive down the cost and power of the solution becomes critical.

Reducing Total Cost of Ownership

The past decade has marked a dramatic shift in the approach that OEMs take in developing products. For example, communications equipment manufacturers continue to consolidate, out source, and relocate engineering development to geographic areas where labor costs can be reduced. These efforts to reign in spending have led to a focus on reducing the total cost of ownership (TCO) in developing new technologies. TCO can be reduced for silicon investments when a common framework, methodology, and modules can be reused across multiple applications or platforms.

Protecting Engineering Investment

Knowledge of the business dynamics within the semiconductor market can help an OEM perform an appropriate risk analysis when selecting a vendor. Use of semiconductor devices, such as Ethernet multiplier-accumulator (MAC)

chipsets, memories, general-purpose processors, or FPGAs that benefit from broad usage across applications, can continue to justify investment in that technology over time. Devices that are targeted or optimized for a unique application will be more likely to exit the market as the cost to develop new products escalates in new process geometries, such as 65 and 45 nm.

For example, the network-processor silicon suppliers' landscape evolution shown in Table 1 demonstrates the risk associated with selection of application-specific components. Contrast this with stable growth with continuous access to leading-edge silicon with FPGAs, and it becomes easy to see that investment in intellectual property can be best protected by FPGAs.

Table 1. NPU Providers Continue to Exit the Market

High-End NPU Providers (2003)	High-End NPU Providers (2007)
Agere	Agere
AMCC	AMCC
Bay Micro	Bay Micro
EZ Chip	EZ Chip
Greenfield	Greenfield
IBM	IBM
Intel	Intel
Motorola	Motorola
Procket	Procket
Raza	Raza
Xelerated	Xelerated
	Cavium

Adapting to Changing Network Demands

Traditional access equipment provided a “dumb” interface to the home that was sufficient for delivering pure data services. This equipment has since evolved to support Ethernet and IP processing, which has led to a debate about how much intelligence should be embedded into access equipment.

Adding more intelligence into the access box may enable providers eventually to eliminate an additional piece of equipment such as a broadband remote access server (BRAS) box sitting next to the access box. Systems are being developed that support either centralized or distributed processing, as shown in Figure 2 and Figure 3. The distributed architecture pushes the packet processing functionality closer to the subscriber, which allows very quick response times for functions such as channel changes. The centralized architecture performs more advanced packet processing on a central switch card.

Figure 2. Typical Block Diagram of a Distributed GPON OLT Application

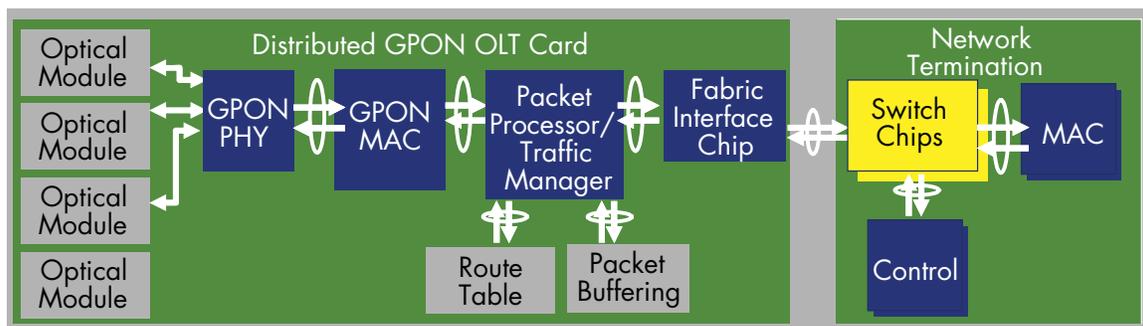
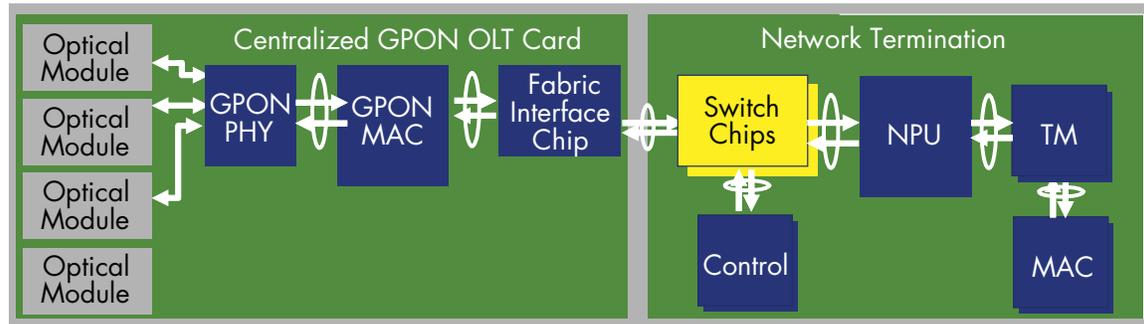


Figure 3. Typical Block Diagram of a Centralized GPON OLT Application



Supporting Tomorrow's Scalability Requirements

Video and peer-to-peer service adoption has created a demand for increasingly higher bandwidth access technologies, such as 10G PON and WDM PON. Additionally, equipment port density continues to increase, from 72 to 96 DSL ports and from 4 to 8 PON ports. These trends will affect silicon technology and system designers in a variety of ways. As data paths move past 10 Gbps, memory interfacing becomes a bottleneck that drives I/O counts, package sizes, and power consumption concerns. For system designers, evolving a design, such as a packet processing solution from 1 Gbps to 2.5 Gbps, to 5 Gbps, to 10 Gbps, to 40 Gbps and above, may require new technologies at every performance node. A limitation of off-the-shelf silicon technologies, such as network processor units (NPUs), is the inherently fixed architectures and resources. To efficiently manage TCO, a solution must be developed that can scale from 1 Gbps through 100 Gbps.

Customizing Optimal MSAN Solutions

What type of technology could enable such a wide range of scalability, while also being efficiently designed for each performance node? Looking in more depth at the packet-processing example, efficient architectures must have the optimal mix of software and hardware processing due to the increasing performance and protocol requirements. The benefits of software processing can be leveraged to eliminate hardware-coded state machine design, while hardware accelerators can be used to optimize a module for performance. Scalability can also be achieved through parallelism, which allows multiple tasks to operate simultaneously. Software parallelism can be realized with technologies such as multi-threaded processors, while hardware parallelism can be reached through multiple instantiations of modules. The optimal technology would also support a migration to future silicon process nodes without affecting the overall architecture of the design.

Architectural Trade-Offs Within Access Packet Processors

Packet processing solutions can be categorized into pipelined or non-pipelined, single threaded or multi-threaded, completely software programmable to purely hardware-based. Each of the fixed packet-processing architectures has limitations that prevent adequate support of a wide range of applications and requirements. For example, supporting channel bonding in DSL networks that deliver higher bandwidth could require custom silicon. Equipment manufacturers can add an additional device next to an off-the-shelf processor to overcome the shortcomings of the device, but having the ability to add custom hardware into an existing programmable device allows for greater flexibility in the future. A solution that is optimized for a particular application must be customized, but also must be reusable for future applications.

Reusable Design Methodologies

An emerging trend of MSAN equipment manufacturers is the building of custom solutions using reusable processor and hardware engines in programmable solutions, such as FPGAs. System architects can determine the appropriate mix of hardware and software components that will be optimal for a particular application. Building these components in an interchangeable and reprogrammable framework allows designs to be easily modified and

transferred, and thus allows designs to migrate through future process geometries and scale to the ever-increasing speed requirements at a lower cost.

Optimal Silicon Technology for MSAN

Programmable silicon technologies offer the benefits of customizing solutions while accelerating time to market. High-volume markets, such as the access market, may require aggressive cost reduction beyond the traditional FPGA cost structure. A path from an FPGA prototype to a structured ASIC, such as HardCopy®, allows OEMs to gain early adoption of their technology while offering lower cost when volumes ramp. Technical risk during this transition can be managed by working with a single silicon vendor from prototype to production.

Broadband Access Solutions From Altera

MSAN architectures contain a variety of functional requirements, including GPON MACs, packet processing, traffic management, and backplane interfaces. Designers must determine which blocks are most likely to change for the next product generation and which ones will allow product differentiation over the competition.

Altera offers scalable, cost-effective building blocks with flexible FPGA-based solutions. Building blocks for GPON and DSL line cards make the design process easier and more cost effective. These solutions are customizable for the unique needs of evolving applications such as triple-play IPTV networks. For example, if achieving 20-Gbps+ data rates and high-level integration is the design goal, the high-performance Altera® Stratix® series FPGAs with a migration path to low-cost HardCopy series structured ASICs provides a highly effective and efficient solution. For a low-cost packet-processing system, Altera's Cyclone® series FPGAs provide an attractive low-power solution. If an integrated serializer/deserializer (SERDES) for a Gigabit Ethernet backplane is needed, the transceiver-based Arria™ GX FPGA offers an optimized mix of cost and performance.

As it takes more than one device to complete a MSAN infrastructure application, it also takes more than one supplier to provide complete solution. Altera has formed a robust ecosystem with key technology partners to deliver an end-to-end solution for the MSAN market to support application development. From reference designs and IP cores to programmable devices, Altera and its partners offer a comprehensive set of solutions to support the development methodology of choice. Whether developing a MSAN application in-house or buying an off-the-shelf solution, Altera's programmable platform provides design scalability across applications, features, and data rates. New features, operator requirements, and densities can be incorporated into next-generation designs swiftly and easily.

- *In-house development:* Developing specialized IP using building blocks available from Altera or third-party IP vendors allows mixing and matching of flexible, modular functional blocks to meet exact design requirements. The final solution can be assembled easily while ownership and control is maintained, all with less design effort and time.
- *Off-the-shelf solutions:* Defining specific application specifications and turning to ecosystem partners for complete turnkey solutions achieves faster time-to-market and lowers overhead costs.

Designing Scalable, Flexible Triple-Play Equipment

Altera and its ecosystem partners provide programmable solutions for creating a single MSAN architecture that supports multiple access technologies:

- *Packet processing:* Altera offers an on-demand, customizable packet processing system implemented on a flexible building block, or a “shrink-wrapped” solution that is ready to plug into a design. Both options address cost, design ease, and technology longevity concerns. In addition, based on the FPGA used, these solutions meet increasing network traffic demands by scaling from 2 Gbps for ADSL to 20 Gbps for GPON line cards.
- *Industry-standard MAC solutions:* An Altera partner provides an off-the-shelf GPON MAC platform with 1.244- or 2.488-Gbps support, hardware validation, and full-service access network (FSAN) G.984 GPON encapsulation method (GEM). Altera provides a Gigabit Ethernet MAC solution.

- *Traffic management and channel bonding:* For developers of proprietary systems, Altera offers a development framework consisting of a hardware validation platform, high-speed packet memory controllers, and a multi-level scheduler reference design. For designers buying off-the-shelf technologies, Altera has complete traffic management IP with 10-Gbps traffic throughput, a packet buffer, and a scheduler to speed time-to-market.

Conclusion

The innovative access market demands flexible architectures that can support various technologies for delivering new services. Equipment manufacturers are challenged by the increasing demands to accelerate their time to market and adapt feature changes, while reducing costs and power. Selecting the right silicon technology that supports the current and future requirements for the access market is critical.

Further Information

1. Infonetics:
www.infonetics.com
2. DSL Forum:
www.dslforum.org
3. 21st Century Network (21CN):
www.btplc.com/21CN/Whatis21CN/BTsstrategyandkeyobjectives/index.htm
4. Lightreading Insider:
www.lightreading.com/insider

Acknowledgements

- Kevin Cackovic, Senior Strategic Marketing Manager, Communications Business Unit, Altera Corporation



101 Innovation Drive
San Jose, CA 95134
www.altera.com

Copyright © 2007 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.