
Custom NPUs for Broadband Access Line Cards

Introduction

Telecommunications (telecom) equipment makers are facing tough challenges in their Digital Subscriber Line Access Multiplexer (DSLAM) designs. These challenges translate into numerous specific requirements for the access processors (APs) resident on DSLAM line cards. To examine these AP requirements, here is a brief look at semiconductor industry trends, as well as the architectural features available in a specific FPGA-based access flow processor (AFP). The paper also highlights why custom network processing units (NPUs) are an ideal choice for DSLAM line cards.

DSLAM Challenges

The first of these challenges is the one that designers usually most enjoy: offering more! While there are multiple facets to this challenge of offering more, it primarily involves enabling the triple play. This much-talked-about service offering is wanted by telecom operators so they can offer their consumers more revenue-generating services. In order to compete, DSLAMs must be designed with significantly more bandwidth to support video, additional protocol handling, and more stringent traffic management to deliver high quality of service (QoS), particularly for voice and video traffic.

Accomplishing this, of course, presents the next major challenge: offering everything for less! Increasing global competition has put tremendous pricing pressure on DSLAMs. The dollar per port price that telecom operators are paying for DSLAMs has fallen dramatically in recent years. Based on these falling prices, the challenge is maintaining some margin while lowering the cost to build DSLAM hardware.

A third challenge faced by DSLAM designers is to create an adaptable platform. The true value of a DSLAM is its adaptability to handle new requirements, including new protocols and possibly higher throughput, without replacing the entire chassis or its line cards.

DSLAM Line Card Processing Requirements

There are a number of features that the ideal access processor must possess to address the challenges outlined above.

Throughput

It is critical that DSLAMs have the throughput necessary to enable new revenue-generating services and deliver the QoS contracted with all customers. Throughput must be predictable and cannot vary based on network traffic type. The Internet Protocol DSLAMs (IP-DSLAMs) must allow the bandwidth to be flexibly partitioned based on port, service, and direction, while providing low line card cost and low power. As the video services aspect of the triple play dominates, the need for throughput in the DSLAM line card increases perhaps by an order of magnitude. DSL aggregation bandwidth can no longer be so largely oversubscribed. The demands of video streams are too large and too constant to allow the level of over-subscription tolerated in pure data service DSL deployments.

A 72-port ADSL2+ line card would likely need to handle over 3.5 Gbps of throughput just for video. This assumes three unique high-definition (HD)-broadcast or video-on-demand channels per port at 6 Mbps per HD video stream. Another example, a 32-port VDSL line card offering 100M/100M service per port, requires 5.0 Gbps of throughput at 25 percent utilization.

Scalability

While sufficient throughput is needed for high-speed and high-port-count line cards, a solution that is downward scalable in price (including external memory), power, and board area is needed for other line cards. System vendors gain product lifecycle efficiencies by using the same technology, components, and suppliers across a wide range of

products. APs must have the scalability and flexibility to fit in products covering a range of performance and applications.

Service providers today add new features and scale the service offerings by doing a truck-roll that adds to the operating expenses (OPEX). Then these expenses usually are recovered over multi-year contracts or through increased billing. By having a programmable silicon and software, the service providers can provide flexible and scalable solutions by remote configuration, thus avoiding a truck roll and the associated OPEX.

Flexible Interfaces

A flexible IP-DSLAM solution has provisions to support the exact interfaces needed for the line card it is designed into. Since it is common for access equipment vendors to have a family of line cards for their platform(s), the ideal IP-DSLAM must be available as either a family of devices or as a device with flexible interfaces. If the IP-DSLAM is to meet the requirements of multiple interfaces, then the line card design will likely be burdened with additional board area, cost, and power in order to support unneeded functionality. Even worse would be a situation where the IP-DSLAM does not have a sufficient number of interfaces to meet the requirements. In this case, IP-DSLAM can not be used as is, and will require additional external devices or logic to compensate for this shortcoming. A low-port-count DSLAM may only require a single GbE interface for the upstream connection and a single Utopia 2 (extended) or POS interface, while a high-throughput, high-port-count may need two GbE interfaces and three Utopia/POS2 interfaces.

Flexible Network Traffic Processing

An AP should be able to support a full range of protocols as defined by TR-101 (i.e., PPPoA, IPoA, PPPoEoA, IpoEoA, EoA, Tagged Ethernet, SNAP, and Q-in-Q) and have the flexibility to handle the yet-to-be-defined protocols of the future, including PPP Termination, PBT, and others. In addition to these processing requirements, operators enable a multitude of services that draw constraints to support the service level agreements (SLAs) by requiring the IP-DSLAM to support differentiated services (through traffic scheduling attributes) and guarantee quality of experience (by integrating traffic policing and shaping attributes). This requires the support of advanced traffic management including multicast scheduling (SP, WFQ, WRR) and per-flow shaping and policing.

A DSLAM should be able to deliver deterministic performance if the same throughput is delivered of: 1) the upstream or downstream protocols of the network traffic, and 2) the simultaneous mix of protocols across the various logical ports. The throughput of a DSLAM solution lacking deterministic performance must be measured based on its worst-case performance, thus assuring that committed service rates are not impacted when handling “tortuous” protocols.

Cost and Time to Market

Service providers and operations are competing fiercely to gain a fair share of the average revenue per user (ARPU) by providing even more to the end user so that the user makes the switch from the current service subscription. In doing so, they are trying to deliver a flexible, scalable solution at a very competitive price structure. This competition forces the OEMs building DSLAMs to not only re-visit the architectures and features as discussed and deliver this complete package at a very low cost, but at the same time provide the ability to upgrade feature sets to enable providers to be first to market.

The IP-DSLAM solution cost includes the required external memory, external control plane processor (if needed), board area, and power and cooling.

Power

The silicon content and the associated memory must remain within the power budget of the system. Some designs operate without external power and total system power is restricted to that available from the twisted pairs. Line cards designed for existing platforms usually face a strict budget, which was set when the DSLAM platform was first developed, but line cards designed for new systems also face stringent power budgets. Conversely, a low-power AP provides potential benefits, including lower system cost, greater system reliability, and greater port densities.

Drawback of Current Approaches to Address Processing Requirements

OEM vendors like to have access to predictable and sustained roadmaps, provided by internal R&D teams, that address the processing challenges. However, to reduce the internal research and development and non-recurring engineering (NRE) costs and to expedite time-to-market requirements, OEMs are leaning toward an off-the-shelf solution that looks for ASSP vendors and processor makers to support them in the challenges they face.

However, currently available ASSP-based NPU development options come with a host of drawbacks. Multi-core processing solutions suffer from software complexity and fixed resources. General-purpose CPUs are unable to deliver the performance requirements desired by data-path processing elements. ASICs attempt to address some of these limitations; however, they usually end up being unable to scale, inflexible, and too costly in terms of time and money.

Enabling an FPGA Solution

A more fundamental concern for OEMs is the lack of vendors available for them to work with as well as a lack of viable roadmaps that meet their timelines and requirements. This is where a market-leading company with a processing ecosystem of partner companies can enable these OEMs. With packet processing being a primary focus as well as an area of innovation, the two most important factors that can enable these OEMs to meet their wireline access infrastructure market requirements are: 1) having a silicon platform based on leading-edge process technology for high performance and flexible processing and 2) having the option to build their own NPU or access to a pre-built processor that can easily be customized to a specific target application (Table 1).

Table 1. Demonstrated FPGA Advantages Over ASSPs

Packet-Processing Requirements	FPGA-Based Packet-Processing Platform	Traditional ASSP-Based NPUs
Development Speed (1)	Fast: <ul style="list-style-type: none"> • Three months for “shrink-wrapped” NPUs on FPGA • Nine months or less for custom FPGA-based NPUs 	Slow: <ul style="list-style-type: none"> • Six to nine months in software development for NPU applications on ASSP
Acquisition and Development Cost	Low: <ul style="list-style-type: none"> • Simple design methodology and tools • Ease of design change • Many choices 	High: <ul style="list-style-type: none"> • Complex design methodology and tools • Difficult to change/debug • Very few choices
Maintenance and Upgrade	Low: <ul style="list-style-type: none"> • Easy to debug and port 	High: <ul style="list-style-type: none"> • Difficult to debug and port
Flexibility in Adding New Features	Fully flexible: <ul style="list-style-type: none"> • Due to hardware and software reprogrammability 	Limited: <ul style="list-style-type: none"> • Design can only be reprogrammed at the software level
Silicon Roadmap/Sustainability (2)	High: <ul style="list-style-type: none"> • Total available market (TAM) for FPGAs over \$3 billion • Two dominant players 	Low: <ul style="list-style-type: none"> • Small TAM of less than \$200 million (3) • Many players

Notes:

- (1) Estimates based on customer experiences
- (2) Many NPU vendors have exited the market due to bankruptcy or acquisition
- (3) According to industry estimates

Whether OEMs are looking to develop a processing technology in-house using FPGAs or use pre-built NPU solutions on FPGAs, reducing time-to-market, increasing cost advantages, and lowering risk are all advantages gained. Programmability at the hardware and software layers provides the flexibility that is needed to adapt a specific solution.

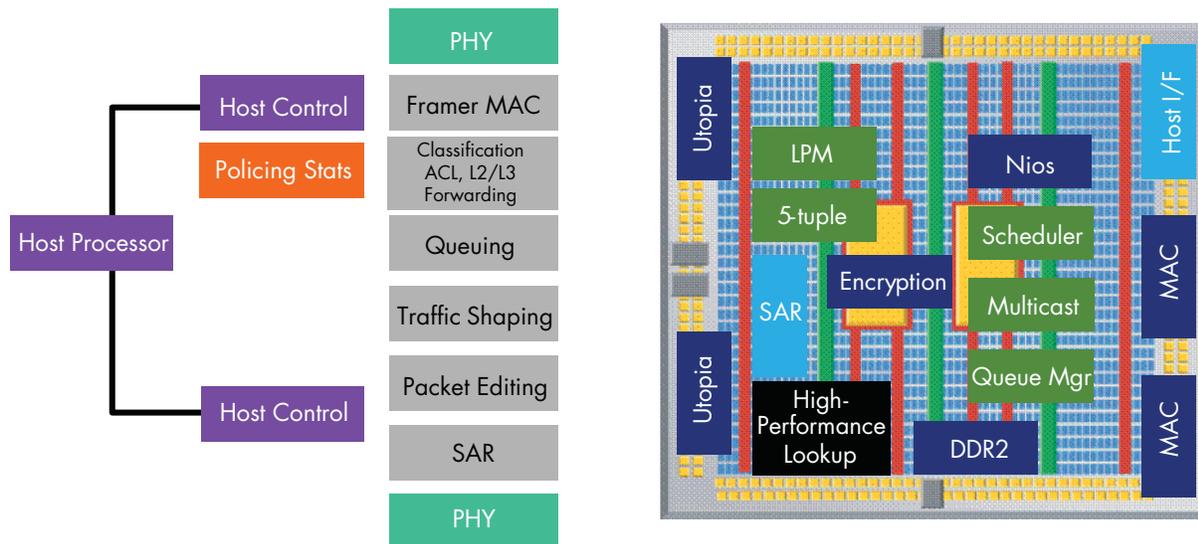
A processing ecosystem made up of a number of partner companies can provide intellectual property cores that will also reduce development time and costs for OEMs who want to “make” their own solution. Such an ecosystem also brings field-proven partnerships that can provide OEMs full access to a versatile processing technology that can be customized to meet the target application requirement.

Altera’s Flexible, Scalable FPGA-Based Platform

Altera® FPGAs, from the low-cost Cyclone® series to the high-performance Stratix® series devices, are equipped to provide a packet-processing platform that meets all of the challenges of quickly and cost-effectively differentiating wireline access infrastructure (i.e., new protocols, access methods, evolving system architectures as well as scaling system bandwidth).

As shown in Figure 1, Altera programmable solutions can be easily implemented in flexible, scalable building blocks or as part of a customized system. Whether developing a differentiated NPU in-house or buying a pre-built, customizable processor, working with Altera and its ecosystem partners enables the development of an ideal architecture quickly, cost effectively, and without component obsolescence concerns.

Figure 1. A Flexible Packet-Processing Architecture Supports Custom Requirements



For faster product availability, prototype in a high-performance Stratix series FPGA and migrate seamlessly to a low-cost HardCopy® series structured ASIC for volume production. For a low-cost packet-processing system, look to the Cyclone series FPGAs and, if an integrated serializer/deserializer (SERDES) for Gigabit Ethernet backplane is needed, consider the transceiver-based Arria® GX FPGAs.

Showcase: Ethernity Networks

Ethernity Networks’s ENET3000 FPGA AFP features a unique architecture that supports the use of low-cost FPGAs utilizing a combination of DDR2 memory to provide advanced functionality, flexibility, throughput, and low power, while competing with off-the-shelf packet processors and network processor ASSPs. The ENET Family (Figure 2 and Figure 3) is implemented in the Altera Cyclone family FPGA, supporting up to 5 Gbps, and in Stratix family FPGAs for 20-Gbps support. The ENET3000 leverages a paradigm shift in ASIC development that enables complex and high-volume designs to be manufactured with FPGAs cost-effectively.

Figure 2. ENET Family

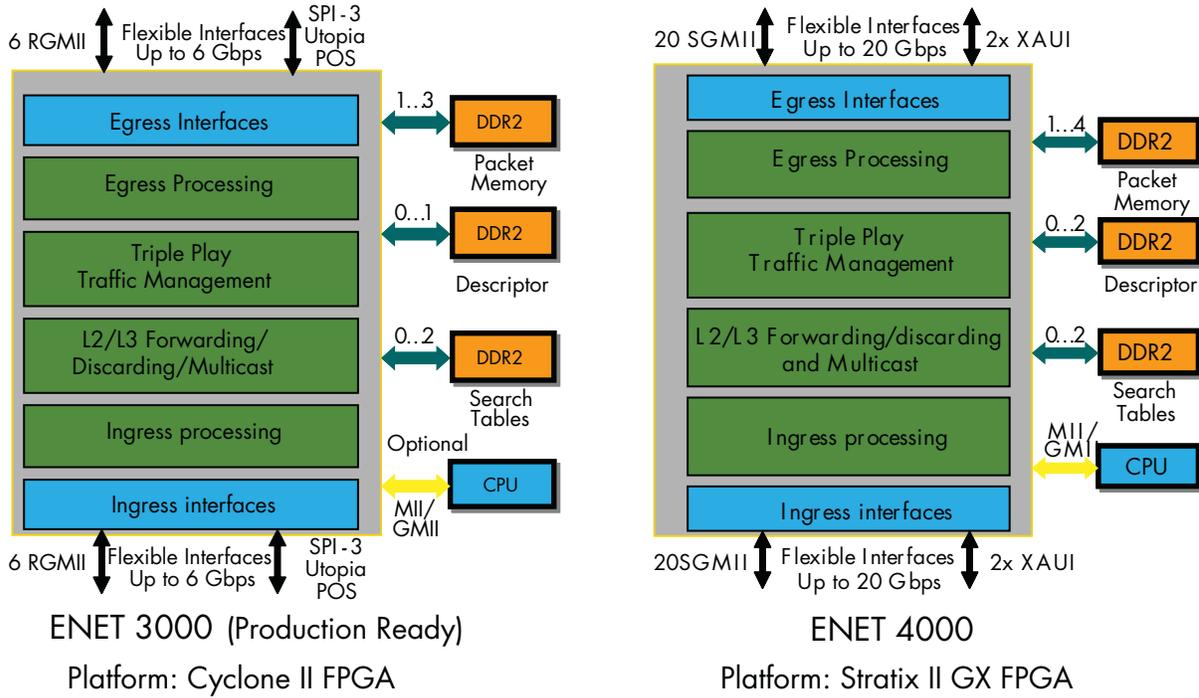
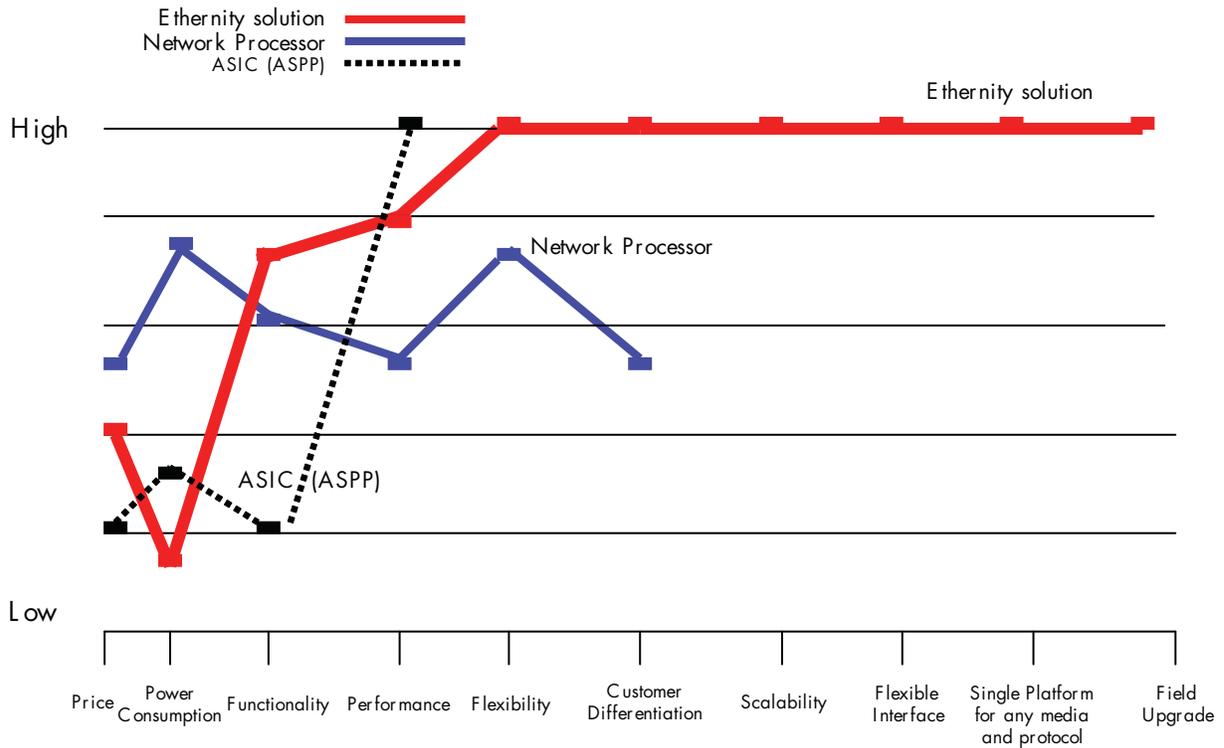


Figure 3. Ethernity Network Value Chart

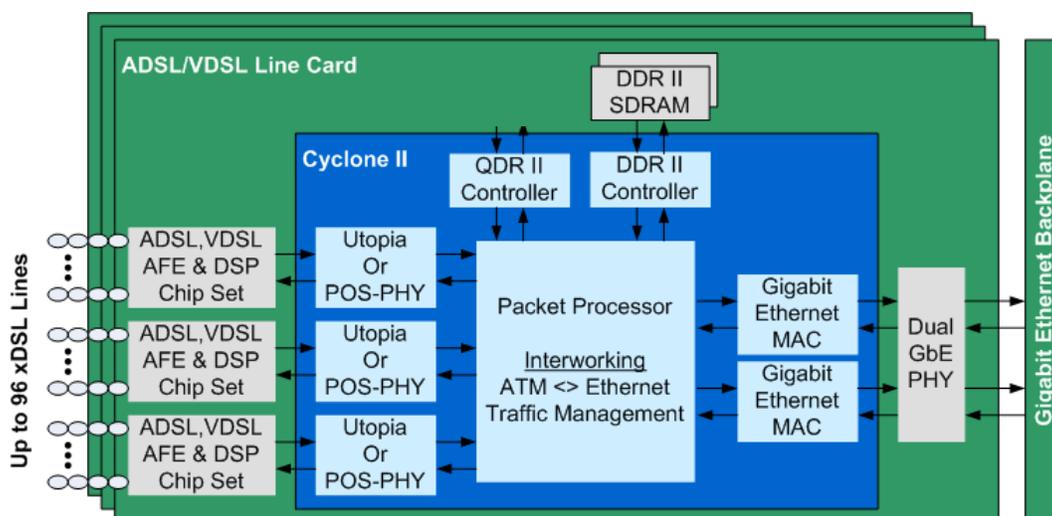


Performance

The architecture of the ENET3000 consists of a staged pipeline, with the tasks of parsing, classification, forwarding, switching and routing using up to 128K search entries, traffic management with up to 64K buffers, ATM SARing and flexible packet fragmentation, policing, and packet editing each handled by one or more engines in the pipeline. The data pipeline width is 64 bits and the engine runs at 100 MHz. Handling the worst-case packet size of 64 bytes, this results in a throughput of 7.5 MPPS or 5 Gbps, while the ENET4000 datapath width is 128 bits, with internal frequency of 167 MHz supporting up to 20 Gbps using two traffic manager blocks.

The ENET3000's (implemented in an Altera Cyclone II FPGA as shown in Figure 4) 5-Gbps network traffic performance is well beyond the capabilities of most network processors available for access applications such as DSLAM line cards. Network processors that are capable of this throughput are too expensive and consume too much power to be acceptable for typical DSLAM line card deployments.

Figure 4. Example ADSL Line Card Diagram



Scalability

In addition to being able to reach the high deterministic throughput of 5 Gbps and handle high-port-count line cards such as 96 dual latency ports, the ENET3000 can be scaled to lower port-count and lower throughput levels (Table 2). The same solution can be used in Metro Ethernet to support up to 24 SSSMII or 20 Gbps supporting 10G XAUI.

Table 2. Example ENET Family Configurations

Part Number	Throughput	Functionality Note	DDRs	Interfaces
ENET 3100 (in production)	<1.5 Gbps	256 flows - Cyclone II (1.3W)	2	1 Utopia/POS + 1 GbE
ENET 3200 (in production)	1.5-3 Gbps	512 flows - Cyclone II	3-4	Flexible interfaces up to 3 Gbps (2 Utopia/POS + 2 GbE)
ENET 3500 (in production)	3-5 Gbps	512 flows - Cyclone II 2K flows - Cyclone III	5-6	Flexible interfaces up to 6 Gbps (GbE, 3 Utopia/POS/SPI-3), or multiple FE (up to 24)
ENET4000x (Q1'08)	15-20 Gbps	16K flows - Stratix II GX	8	2 XAUI, 12 SGMII, RGMII

Low Power Consumption

An efficient architecture yields low power consumption. To process a given amount of network traffic, the optimized architecture of the ENET3000 uses much less logic and less switching relative to a RISC processor (even when the instruction set of the processor is tailored for networking applications). It requires well under 3W while handling 5 Gbps of network traffic, and can go as low as 1W in 1 Gbps of network traffic.

Flexibility

Implementing an AP in an FPGA has many advantages. One key advantage is the flexibility that is achieved on several levels. It is most important to understand that the ENET3000 is an inherently flexible AP—flexibility has been designed into the solution.

Network Traffic Processing

The device's many configuration registers and tables can be pre-configured or altered on the fly to handle a large range of network traffic types and traffic management/QoS requirements, including configurable support for packet fragmentation, and a wire-speed packet editor to enable the support of any Ethernet encapsulation in wire speed.

Interfaces

This FPGA-based solution offers the flexibility to customize interfaces. Unique customer-defined interfaces can be combined to suit the needs of the application, so the designer is not limited by the interfaces offered by an ASSP vendor.

Switching and Routing

The ENET3000 includes a flexible key generator that can extract up to 64 fields from a packet that then is used for classification, filtering, ACL, switching and routing. The use of DDR2 for searching the data base enables 128K entries to be used for classification, learning, filtering and forwarding, so that a different forwarding rule can be used per flow.

Integration

Such a solution provides the flexibility to add logic to the FPGA whereby miscellaneous board functions can be easily incorporated into it. It has been observed that many access cards designed with RISC NPUs also require FPGAs to implement features not supported by the NPU (i.e., framing or channel bonding) as well as implementing miscellaneous board logic. Unused logic is available to the user to implement various board logic such as GPIO, timers, or other customer specific functions.

Field Upgrades

FPGAs offer the flexibility to upgrade a product in the field. While products are often deployed with no definitive plans to upgrade or alter installed units, the option of altering the hardware in the FPGA offers important insurance against unforeseen events. This capability is important to carriers who want to future proof their systems through field upgradeability.

Cost

APs developed on 90-nm FPGAs have several time-to-market and cost-saving benefits, including design automation, flexibility, and reduced development risk. These benefits are used to accelerate AP development in the dynamic IP-DSLAM market and adapt to constantly evolving standards. Design flexibility is afforded by moving Altera's Stratix II FPGAs to the HardCopy structured ASIC migration path, which easily allows designers to create customized designs according to individual customer specifications. Line cards for multi-service platforms, including 96-port dual latency non-blocking ADSL2+, non-blocking 32-port 100/100 VDSL2, multi-port Fast Ethernet, Gigabit Ethernet, and Fiber-to-the-Home (FTTH), can be supported.

Conclusion

There are numerous requirements for line termination access packet processors that must be met to fulfill the design challenges for DSLAM, OLT, and MSAN broadband access line cards. An FPGA-based access flow processor, like the ENET3000, combines the reconfigurability, performance, and scalability of FPGA and structured ASIC technologies with the best aspects of a customized, access packet processor architecture on a single chip. This combination provides the scalability, flexibility, and low cost that designers need to deliver a wide range of programmable access packet processor solutions offering optimized product differentiation, customized packet processing, and rapid time to market.

Further Information

- Point-to-Point Protocol over ATM (PPPoA):
http://en.wikipedia.org/wiki/Point-to-Point_Protocol_over_ATM
- Internet Protocol over ATM (IPoA) and Internet Protocol over Ethernet over ATM (IPoEoA):
www.h3c.com/portal/res/200706/01/20070601_108465_ATM%20Introduction_195627_57_0.pdf
- Point-to-Point Protocol over Ethernet over ATM (PPPoEoA):
www.techweb.com/encyclopedia/defineterm.jhtml?term=PPPoA
- Q-in-Q:
<http://en.wikipedia.org/wiki/Q-in-Q>
- Subnetwork Access Protocol (SNAP):
http://en.wikipedia.org/wiki/Subnetwork_Access_Protocol
- DSL Forum's Working Text (WT-101):
www.ieee802.org/1/files/public/docs2004/WT-101v1.pdf

Acknowledgments

- Sri Purisai, Strategic Marketing Manager, Communications Business Unit, Altera Corporation
- David Levy, CEO, Ethernity Networks



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.