For more information on MAN, visit:
www.xilinx.com/esp/optical/index.htm

Xilinx Sees Bright Future for Metropolitan Area Networks

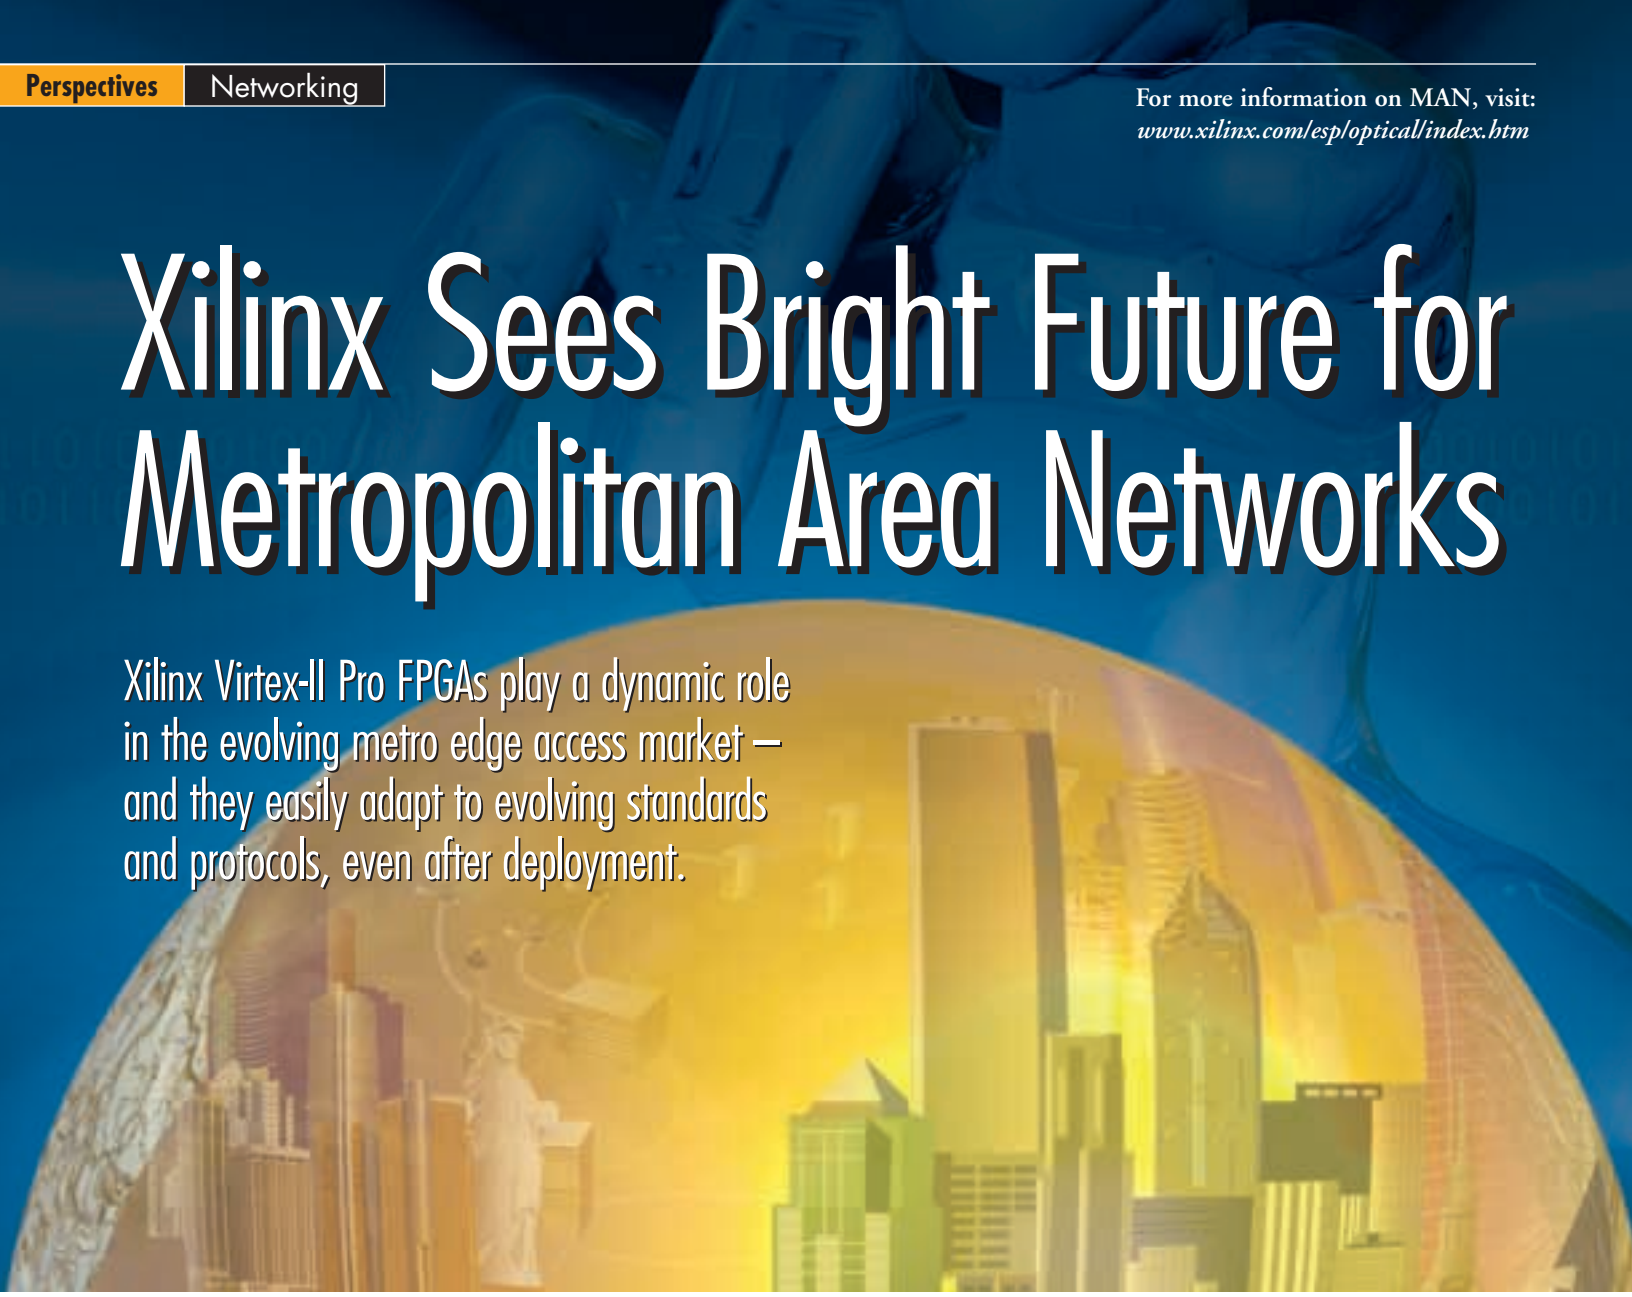

Xilinx Virtex-II Pro FPGAs play a dynamic role in the evolving metro edge access market – and they easily adapt to evolving standards and protocols, even after deployment.

by Robert Bielby, Senior Director
Strategic Solutions Marketing
Xilinx, Inc.
robert.bielby@xilinx.com

The dynamics of the metropolitan area network (MAN) are undergoing a fundamental transformation. The explosion of bandwidth in local area networks (LANs), the deployment of Gigabit Ethernet, and the growth of dense wave division multiplexing (DWDM) in long-haul, wide area networks (WANs) have all served to fuel the demand for networks capable of servicing significantly more data traffic.

All Roads Lead to the MAN

Today, most LAN and WAN traffic converges at the MAN – a transport technology comprising a series of fiber-optic rings that typically encircle major metropolitan areas. Within the MAN, Synchronous Optical Network/Synchronous Digital Hierarchy (SONET/SDH) is the principal networking protocol. Initially deployed for voice traffic, where a typical line, such as T1 (1.54 Mbps), was sufficient to transport multiple voice channels, SONET is very inefficient when it comes to handling IP-based traffic. In addition, SONET is not highly scalable. So, while corporate LANs are moving to 10 Gigabit Ethernet, and WANs are moving to line speeds up to 40 Gbps, the interface between the two networks can easily be 1,000 times slower than the slowest technology in the network.

According to industry estimates, 80% of today's telecommunications traffic consists of data. Even though this percentage is expected to increase further, service providers are still focusing on legacy voice services, because they provide the revenue base that will allow carriers to build out their new service models. To increase the efficiency and effectiveness of their investments, service providers are turning to the MAN.

In addition to their efforts to manage existing data traffic more effectively in the MAN, carriers are adding new services – voice and video over IP (Internet Protocol), virtual private networks (VPNs), 3G (3rd generation) wireless access, wholesale Ethernet delivery, and transparent LAN services – to create new revenue sources.

Another benefit of this increased revenue-generating traffic is that it allows carriers to take advantage of the currently over-built Internet backbone, helping to offset the heavy investments carriers already have in WANs.

Although there remains much discussion and indecision regarding the eventual MAN technology winner, Resilient Packet Ring (RPR) and Metro Ethernet Forum (MEF) are the prime candidates for providing higher performance and more efficient data transport within the MAN. Besides being able to handle voice and data traffic more efficiently, technologies such as RPR can also significantly reduce a service provider's costs by realizing the benefits of a single, converged network. These benefits are far-reaching in their ability to reduce significantly the cost of operations, accounting, management, and provisioning (OAMP) – which typically constitute up to 49% of a service provider's network costs. In short, many factors make it clear that the high-growth area in the telecommunications market is centered squarely in the MAN.

Requirements of the New MAN

For the MAN market to take off, the equipment must do two things – provision and billing for services. This is tougher than it sounds because the metro edge will be a primary point where multiple traffic types – with varying traffic requirements – will converge. To provide basic provisioning and billing, successful interface equipment at the metro edge must meet all the following requirements:

- Deliver provisioning and bandwidth consistent with Service Level Agreement (SLA) policies, regardless of subscriber location on the network.
- Feature a highly scalable architecture capable of servicing thousands of endpoints while supporting a broad range of applications.
- Provide reliability on the order of 99.999% uptime with support for redundant hardware and ring topologies, fiber protection, and restoration capabilities.
- Support services requiring deterministic and predictable performance, such as real-time voice and video applications. These services should deliver minimal latency and jitter.
- Converge voice and data services seamlessly.
- Be optimized for a ring topology and incorporate service protection.
- Be agile and flexible enough to support a wide range of services.
- Be cost-effective to operate.

These requirements place extreme demands on areas such as packet processing, network traffic management, and backplane technologies. Developing a product that delivers this broad range of features and capabilities – while remaining flexible enough to accommodate a variety of traffic types, specification changes, and/or enhancements – requires high-performance, leading-edge technologies.

ALTHOUGH THERE REMAINS MUCH DISCUSSION AND INDECISION REGARDING THE EVENTUAL MAN TECHNOLOGY WINNER, RESILIENT PACKET RING (RPR) AND METRO ETHERNET FORUM (MEF) ARE THE PRIME CANDIDATES FOR PROVIDING HIGHER PERFORMANCE AND MORE EFFICIENT DATA TRANSPORT WITH IN THE MAN.

MANaging Traffic

Networking solutions for the MAN must be able to cost-effectively support a high density of customers using multiple traffic types. Moreover, the density of aggregated traffic leading into a single network area requires that traffic management provide highly effective throughput, while supporting such services as multicasting.

Source address filtering, which helps reduce traffic congestion by identifying traffic that can be "touched" (as opposed to traffic that should not be), requires operation at layers 2 and 3 of the Open System Interconnection (OSI) model. At line rates of 10 gigabits per second (OC-192), filtering poses significant processing challenges for most silicon technologies. Furthermore, while traffic management places extreme packet processing requirements on the data plane, traffic management also requires the ability to prioritize traffic on the control plane.

Because differentiated services are critical to revenue generation, advanced traffic management must be applied to both on-ramp and off-ramp access points. This is the only way to ensure that customers receive the services and bandwidth that they are paying for. Conversely, traffic management must also make sure that customers aren't receiving more bandwidth than they are paying for.

This traffic contract is typically enforced at the on-ramp, or ingress side, of the network. Enforcement is usually based on "leaky-bucket" policing algorithms that drop arriving packets that are outside the scope of the provisioned service contract.

Typically, carriers use weighted fair queuing scheduling in conjunction with shaping to ensure that bandwidth guarantees are supported. To deliver the low-jitter services required to support voice and video traffic, weighted fair queuing scheduling is usually applied on a per-flow basis.

Placing these system requirements in the context of a router that supports OC-768 (40 Gbps) line speeds requires enqueue/dequeue packet processing at rates of greater than 100M packets per second (PPS), with peak scheduling decisions of up to 100M PPS. To be competitive, the router must support in excess of 100K unique flows, with each flow spanning a wide range of granularity, from 64 Kbps to 40 Gbps.

Suffice to say, supporting this level of processing performance imposes significant demands on the basic characteristics of semiconductor technologies. To make matters worse – as mentioned earlier – two principal Layer 2 technologies are vying for acceptance as the standard for MAN applications: Resilient Packet Ring and optical Metro Ethernet Forum – and both are still being defined.

With RPR, for example, there is still considerable consternation regarding traffic fairness – which details how ring traffic is added and dropped. This lack of agreement has caused a fundamental split across the industry, spawning three derivatives of the RPR specification.

RPR, officially referred to as 802.17, is expected to become a formal standard in 2003, but early versions of the RPR specification have already been shipped to carriers. Long-term compatibility between equipment shipped today and the final 802.17 specification is far from certain. Indeed, compatibility is most likely impossible – unless that equipment has been implemented in a flexible, high-performance fabric.

The Virtex-II Pro "Killer App"

On March 4, 2002, Xilinx redefined the programmable logic landscape – again. The new Virtex-II Pro™ FPGAs herald an astonishing breakthrough in system-level solutions. With as many as four IBM PowerPC™ 405 processors immersed in the industry's leading FPGA fabric, Xilinx/Conexant's high-speed serial I/O technology, and Wind River System's cutting-edge embedded design tools, Xilinx delivers a complete development platform of infinite possibilities. The inherent system-level performance and feature sets are a perfect match for the performance demands and diverse requirements of equipment for the emerging MAN.

Each PowerPC core runs at 300+ MHz, delivering 420 Dhrystone MIPS, and is supported by IBM CoreConnect™ bus technology. The unique Xilinx IP-Immersion architecture makes it easy to harness the power of high-performance processors, and to integrate soft IP (intellectual property) easily into the industry's highest-performance programmable logic.

The Xilinx XtremeDSP solution is the world's fastest programmable DSP solution. With up to 556 embedded 18 x 18 multipliers, 10 Mb of embedded block RAM, an extensive library of DSP algorithms, and tools that include System Generator for DSP, Xilinx ISE, and Cadence SPW, XtremeDSP is the industry's premier programmable solution for enabling tera-MACs per second applications. This level of high-performance DSP is critical to supporting the computation of packet transmit schedules.

The first programmable devices to combine embedded processors along with 3.125 Gbps serial transceivers, the Virtex-II Pro series addresses all existing connectivity requirements as well as those associated with emerging high-speed interface standards. Xilinx Rocket I/O™ transceivers offer a complete serial interface solution, supporting 10 Gigabit Ethernet with XAUI, 3GIO, SerialATA, and a host of other protocol technologies. Xilinx SelectI/O™-Ultra supports 840 Mbps, LVDS, and high-speed single-ended standards such as XSBI and SFI-4.

In a single off-the-shelf programmable device, systems architects can take advantage of microprocessors, the highest density of on-chip memory, multi-gigabit serial transceivers, digital clock managers, on-chip termination, and more. This will result in a dramatic simplification of board layout, a reduced bill of materials, and unbeatable time to market.

Additionally, systems designers can partition and repartition their systems between hardware and software at any time during the development cycle – and even after the product ships. The overall system can thus be optimized – guaranteeing that performance targets are achieved in the most cost-efficient manner – and hardware and software can be debugged and observed simultaneously at speed. This capability is critical, as traffic types typically are not known until the product is out in the market and because, even when possible, evaluating "corner cases" is either impossible or too time-intensive.

Optimized for the PowerPC, Wind River's industry-proven embedded tools are the premier support for real-time microprocessor and logic designs. The Virtex-II Pro FPGA is driven by the lightning-fast Xilinx ISE 5.1i software, the most comprehensive, easy-to-use development system available.

Conclusion

Clearly, the metropolitan area will be next arena of growth in the telecommunications market. While many factors contribute to this potential for growth, factors such as the current lack of standardization are impeding broad-based deployment. Bridging the bandwidth gap between the LAN and the WAN – while simultaneously supporting a host of new applications and corresponding traffic patterns – will place unprecedented demands on semiconductor technologies.

The introduction of the Virtex-II Pro FPGA family heralds a new era of programmable solutions that will provide the performance, features, capabilities, and flexibility to address the extreme demands of the metro market. In short, these FPGAs promise to be a key technology in catalyzing the growth of this new market.