

CHIP DESIGNERS' STRUGGLES TO PROVIDE TRIPLE-PLAY HD SERVICE TO TELEPHONE, CABLE, AND WIRELESS CUSTOMERS ARE CHANGING THE NATURE OF SOC ARCHITECTURE.



# Lessons from THE LAST MILE

BY RON WILSON • EXECUTIVE EDITOR

**T**he forces converging on the telecom and networking businesses have their roots in the changing desires of end users, and changing traffic patterns reflect those desires. For home-computer users, the mostly one-way HTML (hypertext markup language) traffic of Web browsing is gradually evolving into a rich mix of HTML, compressed HD (high-definition) video, interactive high-resolution graphics, and latency-intolerant HD audio. The heavily asymmetric traffic of Web browsing is becoming the more symmetric traffic of peer-to-peer networking.

Nowhere are these changes happening faster or with more public results than in the cellular-access networks, which are struggling to support new smartphones, such as the iPhone.

Mike Coward, chief technology officer at Continuous Computing, points out that mobile-broadband data traffic is doubling every nine months. "The

iPhone does 30 times the traffic of a conventional handset," Coward says. "But that's not the bad news. Netbook users appear to create 450 times the traffic of handsets. All the operators are running up against spectrum limitations."

The iPhone is not the end of the story, either. Handset designers are pressing ahead with plans for mobile devices that

can display and capture HD video. "Even with LTE [long-term evolution], there's not enough air bandwidth to give everyone HD video in their palm," Coward says. And a movie viewer in every palm is not the worst-case scenario. "Peer-to-peer traffic from netbooks and video sharing can be network breakers," he warns.

Mobile services must live within the physics of their air interfaces and thus face the most acute problem. Even cable- and telephone-service providers are under pressure, however. "While US-based broadband customers game, e-mail, and social-network over 384-kbps or 3-Mbps links, our counterparts in Korea, China, or Japan are real-time gaming and sharing video on 40- to 100-Mbps links," says Bruce Tolley, vice president of corporate marketing at Solarflare Communications. "A common deployment in Japan and China is IEEE 802.3ah

PON [passive-optical-network] fiber to the building, with 100-Mbps VDSL [very-high-speed-digital-subscriber-line] tails into the houses. This [bandwidth] is more than many of us have available in our corporate networks here in the United States.”

So US cable and telephone operators are scrambling to upgrade, driving optical fiber as close as possible to the customer premises and then bridging the so-called last mile with cable or twisted pairs. “Cable operators will be strong in triple-play [voice, data, and video] in the United States,” says Greg Fisher, vice president and general manager of Broadcom’s carrier-access business. “With new VDSL technology, the telephone-company operators should be able to provide 50 to

**AT A GLANCE**

- ▶ Triple-play use models are threatening today’s networks.
- ▶ Carriers are rushing to increase speed and to shape traffic.
- ▶ Traffic shaping and security require fast deep packet inspection.
- ▶ A new generation of silicon architectures is rising to the challenge.

100 Mbps on their copper for short distances.” That ability is a big deal for the carriers. Verizon believes it can charge more than \$100 per user per month for that kind of service. So nearly everyone is in the same boat. Sooner or later,

bandwidth limitations will keep customers from using the network as they wish.

Security is yet another issue lurking beneath the surface of the shift in network use. Network-application providers, ISPs (Internet-service providers), and carriers all must protect themselves from denial-of-service attacks and intrusion. And carriers must protect their subscribers. “People are not talking enough about mobile security,” Coward warns. “The first time there is a big intrusion into smartphones, users are going to blame their carriers.” The same argument could apply just as well to fixed-service providers.

**A NOSIER, SMARTER NETWORK**

According to networking experts, the solutions to both of these problems—

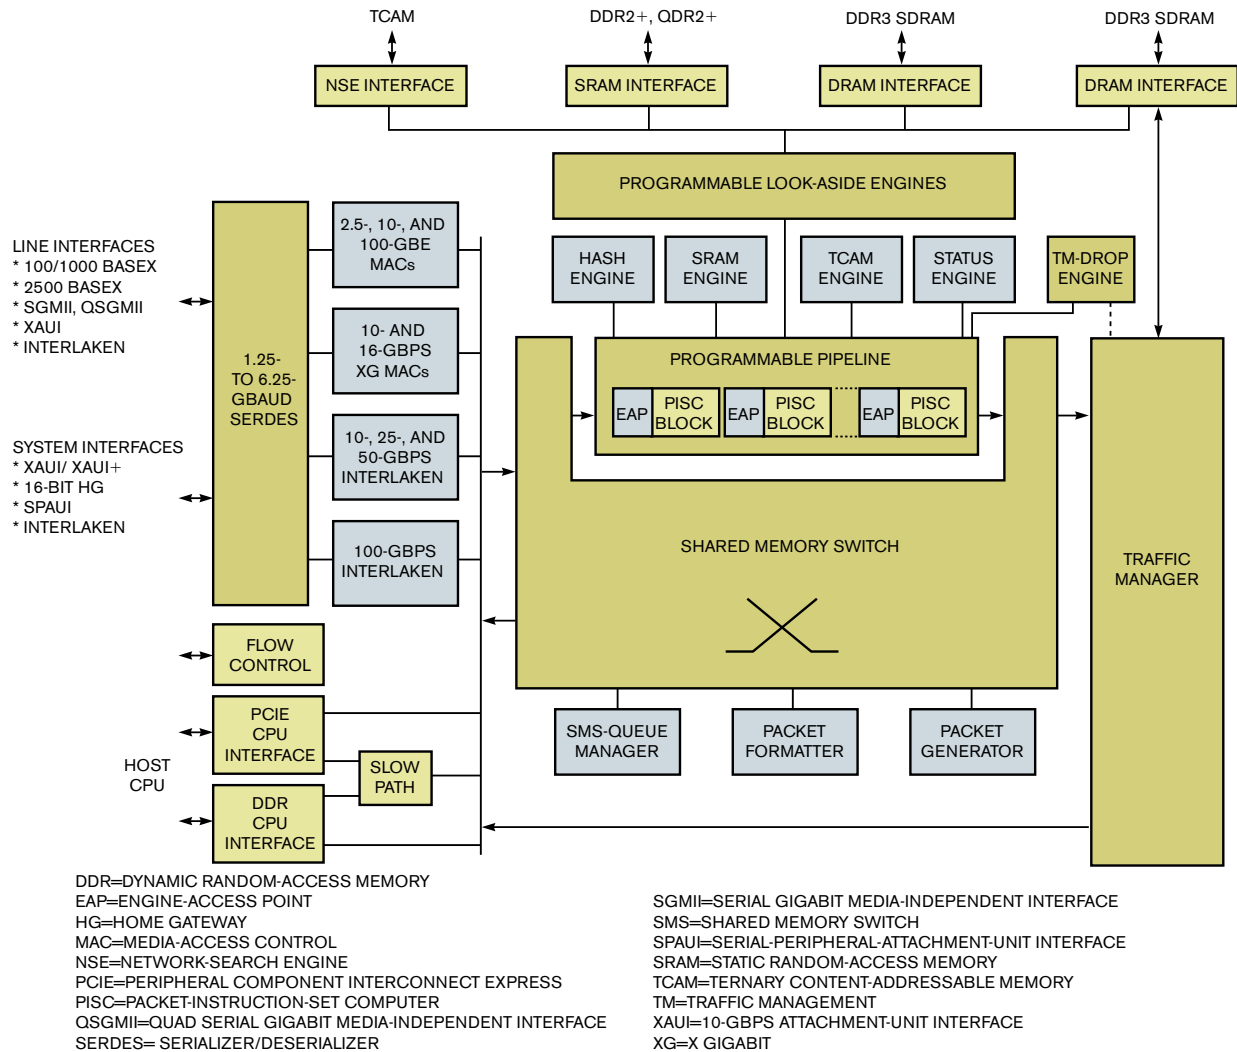


Figure 1 Xelerated’s HX330 is an evolution of the programmable-pipeline strain of architectural thinking.

running out of bandwidth and security—begin in the same place: with knowing what is in the packets traversing the network. To make the most of what bandwidth they have, carriers must shape the traffic that passes through their domains. And to protect themselves and their customers, carriers and service providers must identify and destroy pernicious packets. Both of these processes require inspection of the packets as they pass through switches, routers, and even, some argue, line cards. But where to perform this inspection, how deeply to look into the packet, and what to do with the resulting information are all debated issues, the resolutions to which are greatly influencing silicon design.

“Bandwidth gets very expensive in access networks,” says Kent Fisher, chief systems engineer at Freescale Semiconductor. “So there is a lot of incentive for carriers to parse the packet stream, identify the applications that are using the packets, and apply protocols and traffic shaping to get the most out of their bandwidth.”

DPI (deep packet inspection)—looking deep enough into a packet to identify its payload—has many attractions. DPI allows a switch or router to prioritize and schedule individual packets—for example, giving latency-intolerant audio packets an immediate departure, making



sure that the packet stream for an HD-video player gets its required minimum bandwidth, and scheduling HTML packets for a browser before data packets for a file swap. And detecting virus-bearing or denial-of-service traffic can also require looking at the payload. More controversial is the revenue aspect of the question. DPI allows a carrier to identify and charge extra for packets associated with premium services or to impede packets from rival services.

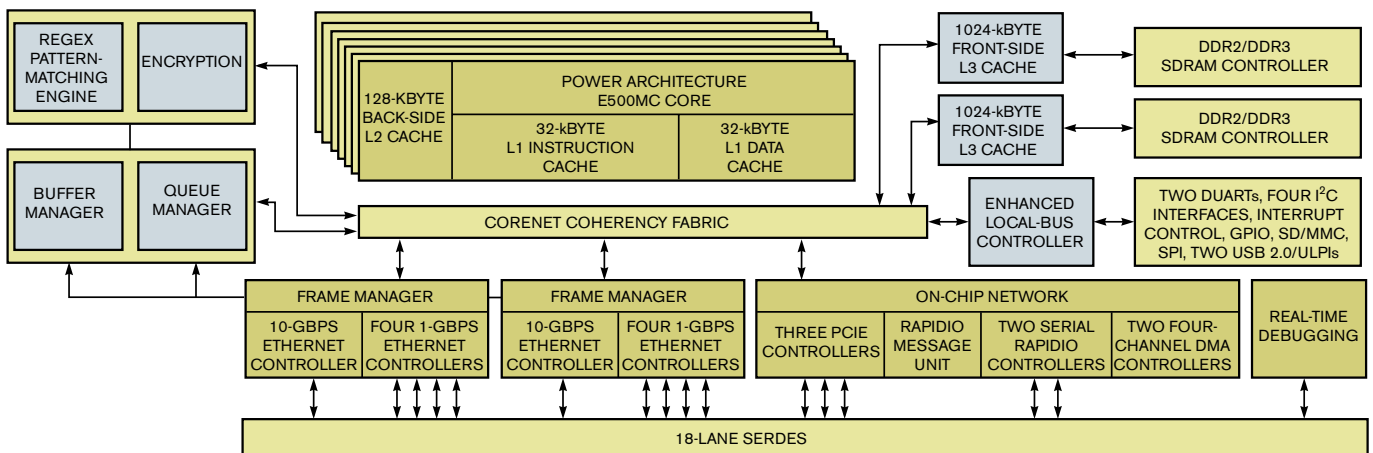
Other issues surround DPI, as well. If packets are encrypted and you can get the key, then inspection requires decrypting and re-encrypting each packet. “About half the time, there is no way to read encrypted traffic, so you have to rely on statistical techniques to guess what the packets are,” says Continuous Computing’s Coward.

And DPI is hard work. Instead of just breaking apart the header on each packet, you have to read the whole thing and, in the worst case, run it through a regular-expression processing algorithm to detect embedded patterns that can indicate data types or the presence of a virus. Particularly in software, that task takes a lot of cycles and a lot of energy. “With all of their requirements, mobile operators are asking us for 20 times more processing work per packet than in yesterday’s systems,” Coward says.

### WHO DOES THE WORK?

Who will do all this work is another difficult issue. “Classification and QOS [quality-of-service] processing have to happen from end to end of the network, even in the metro networks,” says Freescale’s Fisher.

“You don’t want to end up doing deep classification at really high bit rates,” however, says Syed Shah, a systems architect at the company. “It’s much more



DDR=DOUBLE DATA RATE  
 DMA=DIRECT-MEMORY ACCESS  
 DUART=DUAL UNIVERSAL ASYNCHRONOUS RECEIVER/TRANSMITTER  
 GPIO=GENERAL-PURPOSE INPUT/OUTPUT  
 L1=LEVEL 1  
 L2=LEVEL 2  
 L3=LEVEL 3  
 MMC=MULTIMEDIA CARD  
 PCIE=PERIPHERAL COMPONENT INTERCONNECT EXPRESS  
 SD=SECURE DIGITAL  
 SDRAM=SYNCHRONOUS DYNAMIC RANDOM-ACCESS MEMORY  
 SERDES=SERIALIZER/DESERIALIZER  
 SPI=SERIAL-PERIPHERAL INTERFACE  
 ULPI=USB 2.0 TRANSCIVER-MACROCELL INTERFACE/LOW-PIN INTERFACE  
 USB=UNIVERSAL SERIAL BUS

Figure 2 Freescale’s 4080 family processors bear a family resemblance to other heterogeneous multicore-processor architectures.

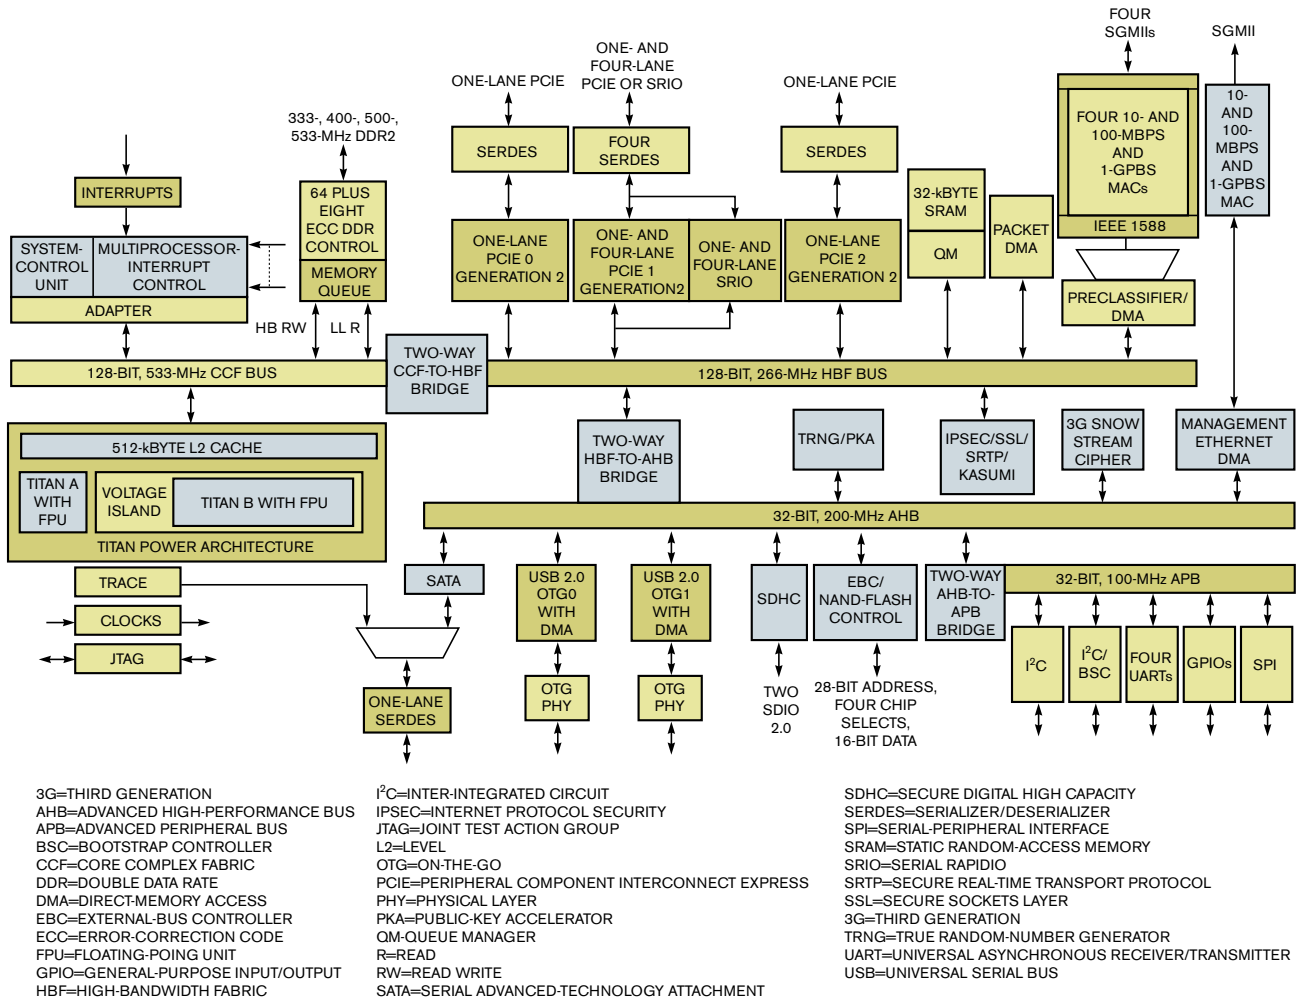


Figure 3 Applied Micro's chip architecture resembles nothing so much as the architecture of the networks in which it will find use.

feasible to inspect the packets in the access network.” Network architects recognized this situation years ago and came up with ideas such as MPLS (multiprotocol label switching) and the QOS bits in the IPv4 (Internet Protocol Version 4) header. In these schemes, a classification engine inspects each packet at or near its source and leaves a marker at Layer 2 or 3, indicating the priority the packet requires. Switches and routers deeper in the network then need not perform deep inspection.

“We see businesses trying to aggregate data and voice traffic from multiple ISPs and to route the packets using QOS bits or VLAN [virtual-local-area-network] tags,” says Michael Durant, vice president of engineering at Arcturus Networks. This approach can in principle keep most of the switching decisions at

Layer 2. “But browser vendors compete with each other on things like audio quality,” Durrant continues. “So they routinely set the QOS bits very high. That practice creates artificially stringent QOS demands.”

Even with all applications playing fairly, legitimate differences can exist in objectives between an application trying to impress a user, a base station trying to manage overloaded channels, a backhaul aggregator, and the metro network, for instance. So boxes deeper in the network may want to take a peek at suspect packets. “[Still,] I don’t think you really need DPI either in the line card or in the metro network,” says Thomas Eklund, vice president of marketing and business development at packet-processing chip vendor Xelerated. “Depending on regulations, inspection probably

makes the most sense integrated into the access-network fabric. There, you must classify each packet through Layer 4. Beyond that [layer], I would argue it isn’t really necessary.”

You must also consider the government regulations that Eklund calls Layer 8. The Federal Communications Commission’s sudden interest in network neutrality—the idea that the net, including carriers, should treat every packet the same—is of particular concern to equipment and silicon providers. Just what this doctrine means and how it might turn into regulation are areas of anxious debate. “Net neutrality appears to forbid DPI,” Eklund observes. “But security, user demands for QOS, and the carriers’ need to generate revenue may all require DPI.” Such conflicts typically lead to politically driven instability in

regulations and, hence, create a need for great flexibility in switches and routers throughout the network.

### ADDRESSING THE SILICON

From this statement of the problem, you can generalize about the kind of silicon that the next generation of access multiplexers, base stations, and carrier-Ethernet switches and routers will require. First, these chips will have to be fast. Wire speed for a VDSL2 twisted pair may be 100 Mbps. Deeper into the network, all transmission is optical, and a speed requirement of 10s of gigabits per second is not unusual. Switch and router boxes can't run below wire speed and depend on big buffers to make up the difference if carriers are succeeding in getting high channel usage because there would never be enough dead time in which to work through the buffer. And some new media types, notably audio conferencing and videoconferencing, are highly intolerant of the latencies big buffers would create.

The chip or chips must also be able to perform packet inspection. Just how



deep that inspection must go is a matter of great uncertainty. As a generalization, however, the closer to the edge of the network a chip will sit, the more likely it is to have to do DPI. After inspection, the hardware will have to classify the packet and place it in the right queue for export. Further, the system will have

to support a growing array of administrative, bookkeeping, supervisory, and error-recovery protocols.

What does all this mean for the silicon? In simpler times, the hardware was just a fast CPU with a lot of memory—sometimes, just an embedded PC. All the functions were in the software. As speeds and functions both grew, however, their product outran Moore's Law. At that point, the hardware architecture split into two planes. Sequential, control-oriented tasks stayed in a CPU in the control plane. The much faster but readily parallelizable packet processing moved to more specialized hardware in the data plane.

Under growing pressure, the data plane evolved further. As data rates grew too high for CPUs to keep up, some architects developed network processors—essentially, microcontrollers with tightly coupled hardware accelerators to handle the bottleneck tasks. Other design teams went in a different direction: a hardware pipeline. Fixed-function hardware engines could keep up with very high wire speeds; if the sequence of

tasks in packet processing remained the same, simple data flows between pipeline stages eliminated many of the loads and stores inherent in a CPU-centric architecture, saving time and power.

But as protocols grew more diverse and complex, the fixed functions and fixed topologies broke down. Pipeline stages began to look like programmable accelerators. "Life is too risky now for fixed-function pipelines," says Xelerated's Eklund. "Programmability is not necessary only in the access fabric. It has to go much deeper into the network." Pipelines also sprouted thickets of conditional bypass and feedback paths and, eventually, accelerators of their own until the pipeline became just the central engine in a network of processing elements (Figure 1).

## THE EVOLVING ENGINE

This growing complexity is erasing the distinction between the control and the data planes. At the same time, process migration is yielding less increase in circuit speed. As a result, some architects are returning to where they began: software on a CPU. This time, though, the CPU is a multicore cluster with both general-purpose processors and specific accelerators. Toby Foster, senior product marketing manager at Freescale, describes such a device (Figure 2). "The QorIQ chip family employs multiple e500 Power Architecture cores to cover applications from line cards to base stations and infrastructure," he says. "As the control and data planes merge, we see multicore chips with datapath accelerators—a queue manager, a crypto engine, a regular-expression matcher—encroaching on the traditional ASIC approach."

With all these cores, the traditional bus-based interconnect structure is failing, as well. To get the bandwidth the chip needs, architects provide each processing site, including the accelerators, with local caches, and they may tie everything together through a non-blocking switch fabric. If architects then provide hardware coherency across the caches and fabric, the programming model for the chip can approximate coding for a single CPU.

Even with good cache design, however, scheduling data movement under software control in such a chip involves a lot of work. "Traffic management in a multicore chip creates access issues,"

warns Satish Sathi, senior principal engineer at Applied Micro. "And these issues involve fairness, QOS, and conflicts for resources. You can resolve them in software, but that [approach] creates overhead."

Applied Micro's approach is hardware-based virtualization. In effect, Sathi explains, the control software sets up a route through the engines on the chip for each category of packets. A network of queues and a hardware-arbitration engine then steer the packets through the maze of engines, buses, and bridges (Figure 3). "The arbitration engine does dynamic arbitration based on actual end-to-end congestion on the chip," Sathi says. "Each packet gets inspected at the end of each task and routed to its next stop."

It's not a coincidence that this scenario sounds remarkably like a network—with nodes, routers, heterogeneous interconnect, and virtual channels. Increasingly, networking chip architectures are leaving behind the idea of a CPU core with accelerators on a bus and the concept of a CPU controlling a data-plane pipeline. Instead, the chips are becoming miniature models of the networks they will serve: heterogeneous collections of processing and routing sites, heterogeneous interconnect, virtual connections, and hardware-supported explicit routing of packet streams. The ideal we are approaching is the ability to define a virtual data-flow machine for each packet flow on an underlying fabric of programmable engines. Therein may lie the future not only of networking ICs but also of the SOC (system on chip) itself. **EDN**

## FOR MORE INFORMATION

**Applied Micro**  
[www.appliedmicro.com](http://www.appliedmicro.com)

**Arcturus Networks**  
[www.arcturusnetworks.com](http://www.arcturusnetworks.com)

**Broadcom**  
[www.broadcom.com](http://www.broadcom.com)

**Continuous Computing**  
[www.ccpu.com](http://www.ccpu.com)

**Freescale Semiconductor**  
[www.freescale.com](http://www.freescale.com)

**Solarflare Communications**  
[www.solarflare.com](http://www.solarflare.com)

**Verizon**  
[www.verizon.com](http://www.verizon.com)

**Xelerated**  
[www.xelerated.com](http://www.xelerated.com)

You can reach  
Executive Editor  
**Ron Wilson** at  
1-510-744-1263 and  
[ronald.wilson@reedbusiness.com](mailto:ronald.wilson@reedbusiness.com).

