OpenSystems Publishing

NETWORK DATA CENTERS…IP SWITCHES…TETRA…INFINIBAND

CompactPCI® and AdvancedTCA® SYSTEMS

OCTOBER 2005
VOLUME 9 NUMBER 8
www.CompactPCI-Systems.com
www.AdvancedTCA-Systems.com

The Magazine for Developers of Open Communication, Industrial, and Rugged Systems

Highly parallel silicon: Tackling complexity

Network Processors Product Guide

PICTURED
THE FLEXCOMPUTE ATCA-PMC40 NETWORK PROCESSOR BLADE FROM CONTINUOUS COMPUTING CORPORATION.

VOLUME 9 • NUMBER 8
OCTOBER 2005

www.compactpci-systems.com
www.advancedtca-systems.com

CompactPCI and AdvancedTCA SYSTEMS

The Magazine for Developers of Open Communication, Industrial, and Rugged Systems

COLUMNS

E-LETTER

October: www.compactpci-systems.com/eletter
Leveraging the AdvancedTCA platform for VoIP media gateways and media servers
By Alan Percy, AudioCodes, Inc.

November: www.compactpci-systems.com/eletter
A look at network processor development environments and programming
By Curt Schwaderer

COVER
As more communications and networking system vendors seek to exploit the new highly parallel silicon such as network processors, a software development complexity crisis is brewing. See page 38.

Cover photo courtesy IP Fabrics.

FlexCompute ATCA-PMC40 network processor blade photo courtesy Continuous Computing.

Published by: OpenSystems Publishing™

FEATURES

EVENTS

SDR '05
November 14-18, 2005
Anaheim, CA
www.sdrforum.org/sdr05/main.html

Productronica
November 15-18, 2005
Munich, Germany
www.productronica.com

WEB RESOURCES

Subscribe to the magazine or E-letter at:
www.opensystems-publishing.com/subscriptions

Industry news:
Read: www.compactpci-systems.com/news
Submit: www.opensystems-publishing.com/news/submit

Submit new products at:
www.opensystems-publishing.com/vendors/submissions/np

"Great in a lab,
so how about
in a UAV?"

CompactPCI and AdvancedTCA SYSTEMS

AN OPENSYSTEMS PUBLICATION

Communications Group

- CompactPCI and AdvancedTCA Systems
- CompactPCI and AdvancedTCA Resource Guide
- CompactPCI and AdvancedTCA E-letter

Editorial Director Joe Pavlat
jpavlat@opensystems-publishing.com

Assistant Managing Editor Anne Fisher
afisher@opensystems-publishing.com

Senior Editor (columns) Terri Thorson
tthorson@opensystems-publishing.com

Technology Editor Curt Schwaderer
cschwaderer@opensystems-publishing.com

Associate Editor Jennifer Hesse
jhesse@opensystems-publishing.com

European Representative Hermann Strass
hstrass@opensystems-publishing.com

Art Director Steph Sweet

Senior Web Developer Konrad Witte

Graphic Specialist David Diomede

Circulation/Office Manager Phyllis Thompson
subscriptions@opensystems-publishing.com

OpenSystems Publishing

Editorial/Production office:
Fountain Oasis, Ste 203, 16872 E. Ave of the Fountains
Fountain Hills, AZ 85268
Tel: 480-967-5581 ■ Fax: 480-837-6466
Website: www.opensystems-publishing.com

Publishers John Black, Michael Hopper, Wayne Kristoff

Vice President Editorial Rosemary Kristoff

Embedded and Test & Analysis Group

Editorial Director Jerry Gipper
Editorial Director Don Dingee
Technical Editor Chad Lumsden
Associate Editor Jennifer Hesse
Special Projects Editor Bob Stasonis
European Representative Hermann Strass

Military & Aerospace Group

Group Editorial Director Chris Ciufo
Managing Editor Bonnie Crutcher
Assistant Editor Sharon Schnakenburg
Senior Editor (columns) Terri Thorson
European Representative Hermann Strass
European Bureau Chief Stefan Baginski

ISSN #1098-7622 ONLINE ISSN #1550-0381
Publication Agreement Number: 40048627
Canada return address: WDS, Station A, PO Box 54, Windsor, ON N9A 615

CompactPCI and AdvancedTCA Systems is published nine times a year by OpenSystems Publishing LLC., 30233 Jefferson Ave., St. Clair Shores, MI 48082. Subscriptions are free, upon request in writing, to persons dealing with or considering *CompactPCI and AdvancedTCA Systems*. For others inside the US and Canada, subscriptions are $45/year. For 1st class delivery outside the US and Canada, subscriptions are $90/year (advance payment in US funds required).

POSTMASTER: Send address changes to *CompactPCI and AdvancedTCA Systems* Fountain Oasis, Ste 203, 16872 E. Ave of the Fountains, Fountain Hills, AZ 85268

The Power of Choice

Subscribe or read online now!

The Most 3U CompactPCI CPU Options

Features include processing speeds up to 1.8 GHz, temperature range to +70 °C, extensive rear I/O options, power consumption down to 6 watts, with or without active cooling. COTS or custom, SMA offers a wide variety of digital and analog I/O, special function and communication boards, plus enclosures, backplanes and power supplies. Choose from a comprehensive line of fieldbus products including Profibus, DeviceNet, CANopen, Interbus, ASi, and MVB, and all major OS's.

For all your 3U CompactPCI needs, choose one partner. Choose SMA.

www.SMAcomputers.com

SMA Computers
9550 Warner Ave. #250
Fountain Valley, CA 92708
Phone +1 714.593.2338

www.SMAcomputers.com

SMA Technologie AG
Hannoversche Strasse 1–5
34266 Niestetal, Germany
Phone +49 561 95 22-0

www.SMA.de

Emerging technologies and opportunities

By Joe Pavlat

CompactPCI & AdvancedTCA

My father used to say, "The number of things I know nothing about is increasing at an alarming rate." That's truer than ever, and this month's authors give us a glimpse into the complex engineering behind modern communications systems.

Much has been written about the transition of the global telecom network from circuit switched technology to packet-based architectures like those found in data centers and enterprise networks. But the differences between data center and telecom technologies are broader than that. Traditionally, telecom equipment exhibits a very high level of service, with redundant resources and self-healing characteristics. Enterprise data centers that support business and Internet communications have generally not been designed to strict telecom requirements, but that is now changing. In *Introduction to Network Data Centers* Kevin Bross takes us through some of the new thinking going on in the enterprise space and the movement towards telecom-grade reliability for critical functions. Kevin also details some of the enhancements being proposed for the AdvancedTCA standard that position AdvancedTCA as a great platform for these Network Data Centers.

In this month's *Technology in Europe* column, editor Hermann Strass updates shows and activities in Europe. He also details a sophisticated communications standard that provides digital and IP-based communications over a robust radio network combining cellular, trunked, and packet radio technologies. Highly available and reliable, the central control of the TETRA network uses CompactPCI processors specially built to work over very wide temperature ranges. In addition to serving critical business applications, TETRA is ideal for emergency communications networks that must be set up quickly and is being deployed in Venezuela, Brazil, Korea, Siberia, Kazakhstan, and other locales.

John Gudmundson from PLX Technology explains new developments in PCI Express technology for this issue's Technology feature. He explains how new chips offload many functions formerly done by a host processor to intelligent PCI Express chips. These chips do more than initiating connections and aggregating data. They also support multiprocessor configurations through what is commonly known as Non-Transparent Bridging. As intelligence and processing power move out into PCI Express I/O subsystems, ever-higher throughput can be achieved.

We all know you can never be too thin or too rich. In the communications-computing environment it's probably safe to add that you can never have too much memory or too much bandwidth. For our Guest feature, Joe McDevitt explains some of the issues and myths about fabric technologies and their real costs and throughput. While many in the industry believe Ethernet will be the predominant fabric used in AdvancedTCA systems and Advanced Switching (AS) will be used in places where Ethernet just isn't good enough, Joe makes a convincing case that InfiniBand combines the best attributes currently available with solid 10 Gbps performance, low latency, and lowest overall system cost. And, unlike 10G Ethernet and AS, InfiniBand is available now.

Matthias Lim and Sorin Vasile address the issue of AdvancedTCA bandwidth from another perspective in their article *Maximizing backplane switching performance of AdvancedTCA systems.* They explore some of the problems that result in inefficient use of bandwidth – partially filled cells, suboptimal routing, and data errors. They also discuss some of the electrical and backplane issues associated with going yet faster, introducing the concept of adaptive link technology and optimized backplane routing.

In *Next generation multimedia applications standards*, Thomas Howe takes us through some of the complex issues associated with multimedia applications, including voice, video, and data. The "disruptive technology" of Voice over IP (VoIP) is rattling traditional carriers who can no longer maintain their control over voice communications. Low cost providers like Vonage – and free providers like Skype – are making it increasingly difficult for providers to offer differentiated services, causing commoditization that threatens profits. Thomas explains how open industry standards are a help, not a hindrance, because they allow vendors to develop value added yet interoperable products.

Drilling down even deeper to some of the technologies associated with VoIP, Kevin Graves explains in his introduction to this month's *Network processors product guide* the importance of Session Border Controllers to VoIP communications. He examines their role in a network environment that requires ever better security and mobility. As the complexity of these network elements increases, development times and costs rise, and the rapidly changing world of standards can threaten investment. Kevin explains how programmable network processors can provide flexible yet powerful solutions aimed at developing a wide range of wire speed packet processing products efficiently.

Joe Pavlat

Editorial Director

Secure, reliable communication

By Hermann Strass

CompactPCI boards are being used as the core of telecommunications systems designed to serve emergency response personnel and other organizations that must keep lines of communication open even where radio cell access is not possible.

European applications

TErrestrial Trunked RAdio (TETRA) is a standard for digital, cellular, and trunked radio networks defined by the European Telecommunications Standards Institute (ETSI). This digital radio system is superior to the Internet or fixed phone nets in privacy, audio quality, and speed in professional applications. Teltronic (Zaragoza, Spain) has been developing and manufacturing radio communications equipment systems for professional use since 1974. Their mobile and digital trunk radio systems are deployed in many European and other countries ranging from South America to Asia (Korea, Kazakhstan) to serve police forces, fire brigades, ambulance services, electricity companies, and other organizations. Teltronic uses CompactPCI boards from Kontron (Germany) for their Central Network Controller (CNC) and for the Network Management System (NMS), which are part of their TETRA systems. The NEBULA system from Teltronic merges individual TETRA systems in the field with an Ethernet/IP backbone. Teltronic is one of the first providers to support Ethernet/IP services such as network management with software upgrades over File Transfer Protocol (FTP), system configuration over Hypertext Transfer Protocol (HTTP), or system monitoring with Simple Mail Transfer Protocol (SMTP) on TETRA systems. Figure 1 shows equipment racks for a TETRA installation by Teltronic.

Teltronic's TETRA product range is continuously expanding. The NEBULA architecture in support of TETRA provides high redundancy and system availability. Its features include distributed intelligence and switching and it does not require Global Positioning System (GPS) synchronization. TETRA is approximately three times faster (28.8 kbps) than GSM (9 kbps). What's more, it only takes 0.3 seconds to complete a connection (GSM = 5 seconds). TETRA offers phone and packet data services, which makes it easy to use for Internet connections. Data and subscriber identities are encoded and use dynamic keys to inhibit movement patterns or subscriber identities being analyzed. TETRA can operate in areas where no radio cells are installed or if they are destroyed. An ad hoc network without base stations can be set up in an emergency over limited distances, as with walkie-talkies.

The CNC in Teltronic's TETRA infrastructure uses Kontron's CP604 system/peripheral CPU (6U CompactPCI) from Kontron as the heart of each trunk radio network installation. Each CP604 manages up to 256 base stations and transmits user conversations within two miles. It operates within a temperature range of –22 °C to +55 °C. Depending on network size and traffic load, the TETRA infrastructure may include up to four CP604 boards, which are based on low power Pentium III microprocessors at 700 MHz or 400 MHz clock rate. There are two Ethernet ports available for network connections. Teltronic uses the Intelligent Platform Management Interface (IPMI) for optimal system management and remote maintenance. Users in Brazil, Venezuela, and Siberia are among those deploying TETRA systems.

The TETRA NMS in a NEBULA system is based on a Kontron 6U CP605 CompactPCI board. It uses the Pentium 4 microprocessor with 2.0 MHz or 2.4 MHz clock rate. It also features two Gigabit Ethernet and one Fast Ethernet port for network connections.

Both CompactPCI boards from Kontron may be replaced in future systems by newer, higher performance CP6011 and CP6000 CPU boards with an extended temperature range of -40 °C to +85 °C. In systems such as NEBULA-TETRA with worldwide deployment, selecting products takes more than just datasheet figures and prices. When Teltronic started to select suppliers for CompactPCI components the list of potential suppliers quickly shrank to four. The supplier must be European based because of geographical, time zone, and language requirements. In a previous installation in South America it took about three months to replace a failed part due to language and other problems with an Asian manufacturer. Kontron has local support in Asia and the Americas as well as Spanish-speaking personnel in Spain. They offer a wide range of 3U and 6U CompactPCI boards as well as technical support on all levels. Kontron also makes a customer specific front panel to match Teltronic's corporate design. Teltronic and Kontron are truly international companies, both based in Europe, which makes cooperation easy. Another important criterion was long-term product support. Kontron must guarantee the delivery of boards for a period of 10 years.

Figure 1

European events

LASER2005, held June 13 to 16, 2005 in Munich, Germany, attracted approximately 23,000 visitors and 985 exhibiting companies. One interesting new product was a fast holography system able to digitize and create a 3-D picture of a human face with extreme accuracy (0.4 mm). Even skin pores and individual hairs can be visualized. This supports facial surgery and forensic science applications. A green laser pulse of 35 ns duration is used because it is eye safe and does not require special treatment of the surface of the skin. With this extremely short exposure time, holograms of live subjects can be made without any motion artifacts. The holographic system was developed by the Center for Advanced European Studies and Research (CAESAR) in Bonn, Germany.

Comlase (Sweden) demonstrated their Nitrel surface treatment technique. The treatment protects delicate semiconductor and glass surfaces from oxidation, contamination, and moisture to improve the performance and lifetime of laser diodes and avalanche photodiodes.

Two conferences, FiberComm2005 and World of Photonics (WoP), were co-located with Laser2005. WoP was extremely international with more than 80 percent of its' visitors coming from outside Germany, mainly from the US and Asia.

The 2nd MicroCar2005 conference was held June 21 to 22, 2005 in Leipzig, Germany. Major European car companies and supporting industries presented 100 papers and 13 workshops, with topics including:

- Microtechnology and nanotechnology materials and analytics
- Mechatronics
- Microsecurity for car applications

Professor Michel of the Fraunhofer Micro Materials Center at the Fraunhofer IZM in Berlin initiated the European Center for Microreliability and Nanoreliability (EUCEMAN).

The MicroCar2005 conference took place in the same location as the Zuliefermesse (subcontractor fair for automotive products) and Mechatronic fairs. The subcontractor fair had about 31 percent non-German exhibitors (out of 462 exhibitors). More than 4,500 visitors came to both the Zuliefermesse and Mechatronic fairs. This is an increase of about 23 percent over the last event. Fair organizers actively arranged for several hundred business-to-business meetings in preparation for the event. This year the Czech Republic was the partner country of these fairs.

European products

Goepel Electronic (Germany), founded 1991 in Jena, develops and produces customer-specific test and automation systems. They are a pioneer of PCI eXtension for Instrumentation (PXI) technology, starting in 1998 with their first PXI product and expanding their current PXI product range of 23 different boards continuously. PXI was originally developed by National Instruments, Austin, Texas. It is now a PICMG standard (PICMG 2.8). All of Goepel's boards are supported with drivers and Dynamic Link Library (DLLs) for Windows NT and Windows 2000/XP and with Virtual Instruments (VI) libraries for the LabVIEW programming environment from National Instruments. Goepel uses PXI boards and the LabVIEW software environment in their turnkey customer-specific test and automation systems.

LV Power (Israel) has introduced an AdvancedTCA board level power supply for 200 W nominal output in a variety of output voltage configurations. The power supply operates on 36 V to 75 VDC input (AdvancedTCA nominal is 48 V) at an input current of less than 6 A at 36 VDC. This power supply card is equipped with overvoltage and overcurrent protection. An IPMI controller handles thermal protection and other control functions.

For further information, contact Hermann at: hstrass@opensystems-publishing.com.

Next generation now

By Susan M. Miller

The Alliance for Telecommunications Industry Solutions (ATIS) and other global communications standards bodies are driving toward consensus on a converged interconnected network architecture – one that delivers exciting new applications and services, supported by global industry standards.

Responding to the industry's need for a common vision for Next Generation Networks (NGNs), ATIS released in November 2004 its *NGN Framework, Part I: NGN Definitions, Requirements and Architecture*. The ATIS Framework is the first comprehensive set of North American NGN requirements delivered internationally, the result of direct input from the senior executives and officers from the industry's leading service providers, manufacturers, and software developers.

The ATIS Framework defines a high level architecture for NGN that meets the business and wide scale deployment needs of telecommunications companies. Its creation signals the arrival of industry wide consensus on the business requirements, architecture, and technology path for NGN. Ultimately, the ATIS Framework itself will serve as a blueprint for the development of needed industry standards. Figure 1 shows the ATIS NGN network architecture.

ATIS's role as the source for an NGN Framework is not surprising given its diverse membership and unique standing within the industry. ATIS is a technical planning and standards development organization of more than 300 leading global communications companies that provides a forum for service providers, equipment manufacturers, and hardware and software developers to collaborate on the technology that delivers tomorrow's communication infrastructure. Together, ATIS member companies pursue the priority technical work and develop supporting standards that drive the business of communications and information technology.

The business requirements for NGN

The transition to a fully converged NGN network means that service providers need a standards-based, service-enabled network architecture that is ready and able to deliver value-added services. Many of today's legacy services are vertically integrated, which inhibits integration with other applications. The ATIS Framework calls for an NGN target architecture that separates services from transport and enables efficient interworking between applications to support innovative converged services. This interworking promises to add value to NGN services by integrating them into the larger convergence of media and access modes. It will allow service providers to roll out the kind of customized and convenient advanced services that segments of their markets are already asking for today. By putting the disparate parts of the communications puzzle together (such as wireline and wireless services, switched and IP networks, voice and other media, and access modes of all types), service providers will have the flexibility to create the right combinations of services for their markets, and deploy them to the benefit of end customers.

In order to fulfill these needs, the ATIS NGN target architecture offers support for:

- Standards-based interfaces allowing plug-and-play integration of any number of applications.
- A standardized session control function through which application servers can signal. This will allow for full convergence of services over a number of access modes. For example, blending instant messaging with unified communications and Voice over Internet Protocol (VoIP).
- A logical subscriber database that holds all customer profile data, in a manner that it can be consistently accessed from anywhere in the NGN. This concept, similar to the Home Location Registers and Visitor Location Registers of the wireless world, would make both service rollout and ongoing administration manageable and scalable.
- A set of access-independent application and service creation capabilities, supporting adoption of converged applications and services to any device the end user chooses and delivered with consistency.

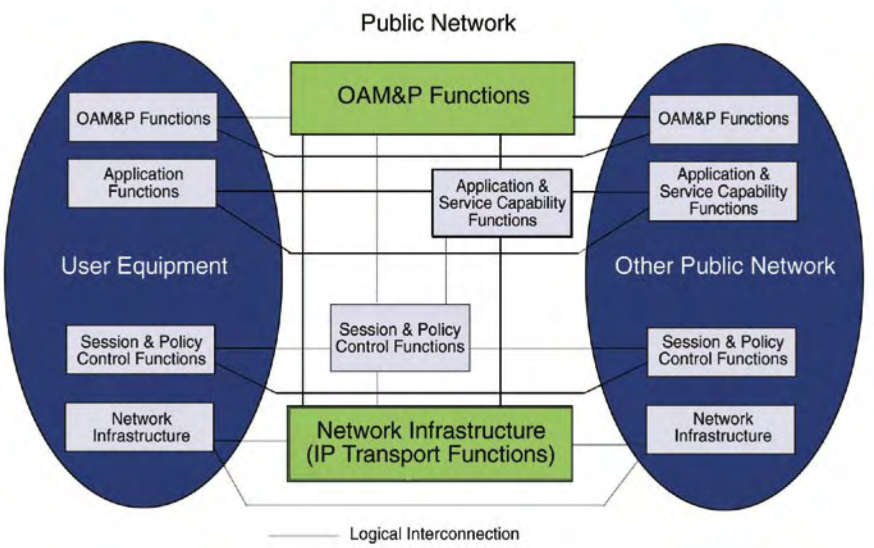

Figure 1

4U 8-slot CompactPCI Subsystem

The cPCIS-6418U is a high-performance CompactPCI platform with a 64-bit, 8-slot 6U backplane in a 4U enclosure. The 8-slot backplane dedicates one system CPU slot and seven peripheral slots. Dual AC inlet and three redundant power supply units provide 500W+250W power redundancy. Featured with high availability and ample room for expansion, the cPCIS-6418U is compliant with the PICMG 2.5 computer telephony integration standards and is ideal for VoIP and network equipment applications.

For more info, go to:
www.adlinktech.com/products

An Entire Family of High-Speed, Low-Power Universal Intel® Pentium® M Processor Boards

The cPCI-6840 family combines Intel's latest high performance low power processor the Pentium® M with the 855GME MCH and two different I/O controller hubs. The Pentium® M has a highly efficient instruction execution unit and coupled with a large L2 cache, it can outperform other processors at the same clock speed. The 855GME has an integrated 250MHz 32-bit 3D/2D graphics engine capable of supporting simulations CRT and flat panel displays. The ICH packs a plethora of peripherals include Serial ATA 150, PATA, USB 2.0, watchdog timer and serial port.

For more info, go to:
www.adlinktech.com/products

2-CH 65MS/s 14-Bit Digitizer with SDRAM

ADLINK's PCI-9820 and PXI-9820 are 2-CH 65 MS/s, high-resolution PXI digitizer with 512MB SODIMM SDRAM memory. They features flexible input configurations including programmable input ranges and user selectable input impedance. With lots of on-board acquisition memory, these digitizers are not limited by the PCI 132MB/s bandwidth and can record the waveform for a long period of time. The PCI-9820 and PXI-9820 are ideal for high-speed waveform capturing, software radio applications, and signal digitizing applications which require ample memory.

For more info, go to:
www.adlinktech.com/products

8-CH High-Power Relay Outputs & 8-CH Isolated Digital Inputs Card

The PCI-7260 is a high-power relay output card for automation and machine control. It comes with eight channels switches 5A currents at 250VAC or 30VDC. The front panel connector and PCB are implemented for carrying high currents. The removable terminal block is convenient for easy wiring. The PCI-7260 also offers eight isolated digital input channels with debouncing circuits to eliminate signal interference.

For more info, go to:
www.adlinktech.com/products

A Variety of Intel® Pentium® M based Single Board Computers

3U PXI CPU Module
with VGA, GbE & HotSwappable CF

3U CompactPCI CPU Module
with Dual GbE & USB 2.0

Half-Size Low Power SBC
with Dual GbE, LVDS & CF

ETX Module with High-End Graphics
& Optional HDTV I/F

Call us toll-free at (866) 4-ADLINK
or email info@adlinktech.com

PXI Systems Alliance

PICMG

Intel® Communications Alliance Associate Member

ADLINK TECHNOLOGY INC.

www.adlinktech.com

The ATIS NGN architecture is also designed to support a variety of business models, including wholesale, retail, Virtual Network Operator, and Virtual Services Operator. In developing their NGN strategies, market participants may consider a variety of business model aspects. These include:

- Market segment addressing
- Insourcing versus outsourcing NGN infrastructure
- Electronic supply chain
- Merchandising model
- Technology considerations
- Development and deployment constraints
- Operating and maintenance models
- Regulatory issues

Additionally, accommodation of existing key services and infrastructures through integration and/or interoperability is an important business aspect. The NGN must leverage the deployed base while supporting new services and technologies efficiently. Paramount to network integration is a common protocol that supports applications across different networks and around the globe.

> The ATIS Framework calls for an NGN target architecture that separates services from transport...

The role of IMS

ATIS advocates the IP Multimedia System (IMS) session-based architecture and views it as the appropriate technology choice to consistently support new value-added services and the NGN business requirements of service providers. The Third Generation Partnership Project (3GPP) developed IMS, which is being adopted globally. IMS establishes packet data connections between Session Initiation Protocol enabled devices, making possible the exchange of all types of packets, including voice, video, content, and others. An international consensus on the use of IMS creates the opportunity to build an open IP-based service infrastructure that will enable an easy deployment of new rich multimedia communication services mixing a wide range of telecom and data services. Ultimately, the NGN will allow the secure flow of high quality data packets to any device, any time, over both fixed (PSTN, ISDN, WLAN) and wireless (GSM, CDM, WCDMA) networks.

The global NGN efforts

ATIS is sharing its NGN Framework with other global standards organizations, including the 3GPP, of which ATIS is an organizational partner, the European Standards Telecommunications Institute (ETSI), and the International Telecommunication Union (ITU). ATIS is collaborating with these standards bodies in an effort for the industry to arrive, to the furthest extent possible, at a consistent global view of the NGN. These actions will ensure that the deliverables produced from global standardization efforts support the requirements of the North American marketplace. To this end, on April 1, 2005, ATIS and ETSI TiSPAN conducted a symposium where both organizations shared details of their respective NGN standardization activities.[1] Both organizations will pursue further collaboration in several areas, to include standards for:

- Interconnection requirements
- Emergency communications
- User identification
- Number portability
- ENUM and directory services
- Security requirements
- Access agnostic solutions

ATIS' vision is to drive industry wide support for the creation of an end-to-end solution for the Next Generation Network. This can be accomplished through a coordinated approach to standards development that results in suites of standards that support next gen network interconnection, interoperability, and reliability and that deliver new and exciting services across multiple networks.

As a follow up to its NGN Framework document, in August 2005, ATIS produced a release strategy supporting the creation of standards for a defined set of network architecture capabilities that enable the introduction of new NGN services. Among the network architecture capabilities, or enablers, that are the primary focus of ATIS' NGN standards activity are network security, Quality of Service, media gateway functions, multicast, presence, digital rights management, and the decoupling of services from access technology.

The ATIS NGN Framework is the roadmap that provides an industry-wide definition and architecture for NGN, and addresses the numerous technical and standardization requirements as determined by industry. With the framework squarely in place, it's now time to move forward, create the needed industry standards, and build tomorrow's converged network. For a copy of the ATIS NGN Framework visit www.atis.org/topscouncil.shtml.

***Susan M. Miller** is president and chief executive officer of the Alliance for Telecommunications Industry Solutions. Previously, Susan served as ATIS' vice president and general counsel. Before joining ATIS, Susan practiced telecommunications law with Weil, Gotshal & Manges. She received her Juris Doctorate from Catholic University's Columbus School of Law, and her Bachelor of Arts from Dickinson College.*

[1] Telecoms & Internet converged Services & Protocols for Advanced Networks

For further information, contact:

Alliance for Telecommunications Industry Solutions
1200 G St. NW
Washington, DC 20005
Tel: 202-434-8850
Fax: 202-393-5453
E-mail: atispr@atis.org
Website: www.atis.org

The case for InfiniBand in AdvancedTCA

By Joe McDevitt

AdvancedTCA, at its core, is a fabric agnostic architecture. One of the fabrics defined early on was InfiniBand via the AdvancedTCA 3.2 subspecification. As AdvancedTCA systems move towards production it is likely that the majority of these systems are AdvancedTCA 3.1 or Ethernet-based systems, but there are limitations in this kind of deployment. Others propose Advanced Switching (AS) as the "King of All Fabrics," but truly there may be no king of fabrics in AdvancedTCA. Time to market, bandwidth, latency, and even system level costs are some areas in which InfiniBand may prove a surprisingly good fit for AdvancedTCA.

Some target telecommunication engineers are outright hostile to using AdvancedTCA in their companies' applications. The chief reason for that hostility remains bandwidth. Engineers report that 1 Gbps Ethernet was slower than their current proprietary solutions and that 10G Ethernet was not the timely solution they needed. On the InfiniBand side, one x4 InfiniBand port can provide 10 Gbps of bandwidth utilization, which would fully utilize one AdvancedTCA channel, providing the timely bandwidth demanded. Other fabrics have plans for these bandwidth levels; AdvancedTCA 3.1 Ethernet via Option 9 and AdvancedTCA 3.4 AS via Option 3 can match this raw performance. The difference is availability. With InfiniBand, the full bandwidth available on a simple FR4 AdvancedTCA backplane is realized today, without limitations on node boards or switches. Both Ethernet and AS solutions suffer from the lack of availability of proper AdvancedTCA centric node and switch solutions.

Another key factor affecting performance is packet latency, the time it takes a single packet to leave the host and reach the destination. InfiniBand offers high bandwidth with low latency. Packet construction time is an area where InfiniBand excels, thereby reducing latency. The InfiniBand Host Channel Adapter (HCA) does this directly without host support. HCA does not carry the overhead of Ethernet, which relies on a complex software protocol stack to ensure reliable delivery. Achieving low latency also requires a well-designed packet switch. InfiniBand answers this with a switch that can cut through a packet from input port to output port by looking at the first eight bytes of the header, before the packet has even fully arrived. Other switch architectures will buffer the entire packet before routing, increasing latency. This cut-through routing does mean forwarding a bad packet, but that doesn't cause a fabric or application failure because the packet will be checked at its destination.

> InfiniBand allows processors to do what they were purchased to do, rather than munging packets.

Shell games

One highly discussed area revolves around fabric cost. Interesting shell games occur when this subject is discussed. Ethernet does have a clear message about per port hardware cost being lower than other fabrics. However, per port charges are only partial solutions and therefore only a partial story of the cost. 1G Ethernet needs multiple ports simply to approach half the bandwidth of InfiniBand. Alternatively, some fabrics rely on future technology that is not available in the scale and port counts needed for AdvancedTCA. What this future cost will be is not known – it can't be extrapolated nor inferred from "sister" technologies. A real world example based on existing AdvancedTCA nodes and switches, an *apples to apples* comparison, of system cost is possible. If a fully loaded InfiniBand AdvancedTCA 3.2 Option 1 (two switches and 12 CPU nodes) system is compared to the exact same configuration of a loaded Ethernet AdvancedTCA 3.1 Option 4 system, the total system cost of the InfiniBand system is around $1000 greater than the Ethernet system or approximately two percent of the total system cost. This minor increase comes with a 10X increase in fabric bandwidth and similar performance increase in fabric latency. When one looks to CPU utilization or packet construction offloading, a conservative estimate of 30 percent of available CPU cycles are used in packet construction in this real world, Ethernet system. A more liberal estimate of 10 percent CPU cycles for packet construction on the InfiniBand side translates into about $1500 savings from wasted CPU cycles. That savings occurs again in the form of more power for applications running on these CPUs. InfiniBand allows processors to do what they were purchased to do, rather than munging packets. Certainly link aggregation and offload engines increase the likelihood of matching performance and limiting CPU usage, but these performance enhancements belay the apparent cost per port benefit of Ethernet. Additionally, the silicon cost of InfiniBand switches is admittedly more expensive, mainly because of market size differences. However, a major difference is that InfiniBand management software is open source. The ease and general familiarity customers required on the Ethernet side requires a complicated laundry list of ITEF RFC support. Frequently, the management costs of the Ethernet are not included when raw per port hardware costing is used, and this is one reason the system level cost of Ethernet is not realized from its per port lead. The result is "The Great Fabric Shell Game." Instead of finding the ball under the shell, we are asked to forget cost additions of offload engines, and of the multiple port aggrega-

tions needed to argue performance limitations. What's more, we are asked to buy into wild speculation on future enhancements and ignore development Ethernet software management cost. All this activity is an effort to show Ethernet is the low cost fabric of choice. Per port cost saving is a partial solution. Therefore one must look toward total system cost and the performance gained. Today, InfiniBand is the clear winner when timely system level price for performance is taken into account.

Advanced Switching proponents point to the economies of scale due to the reuse of the physical and link layer technologies to predict lower cost. While there is value in reusing existing technology, it is clear that these economies will not be achieved. AS will have much higher prices than volume PC-based PCI Express products such as motherboard chipsets and add-in cards. AS devices will sell in much lower volume and will have much higher performance, reliability, and support designed into them and will therefore demand higher cost. AS and PCI Express are different technologies that will be deployed on different classes of devices. PCI Express is not at all suited for system level AdvancedTCA construction via AdvancedTCA redundant links. AS requires management software, thereby changing the market, node architecture, and chip design as well as adding cost. PCI Express does not require management software, is a great technology for desktop systems, and even enhances InfiniBand by allowing reduced BOM costs via memory free HCA solutions. The AS "economies of scale" argument is simply not valid, it is not PCI Express. Additionally, current port counts on AS switches mean that some blocking will occur on early deployment of AS on AdvancedTCA, thus hurting the performance argument.

Conclusion

Ultimately, no AdvancedTCA fabrics king may emerge. Ethernet, Advanced Switching, InfiniBand, and RapidIO, among others, can and likely will all coexist in AdvancedTCA. Pitting one fabric against another is fun and the heated discussions that arise are interesting. AdvancedTCA is designed to its core as a fabric agnostic architecture. As engineers and technologists, it is our job to ensure that once all the tradeoffs have been studied, the best solution is chosen, and for some AdvancedTCA applications that is assured to be InfiniBand. Figure 1 shows the ATC5232 InfiniBand-based AdvancedTCA Node Board and Figure 2 shows the ATS2148 InfiniBand-based AdvancedTCA Switch Board from Diversified Technology, Inc.

Figure 1

Figure 2

Joe McDevitt *is vice president of technical development at Diversified Technology, Inc.*

To learn more, contact Joe at:

Diversified Technology, Inc.
476 Highland Colony Pkwy.
Ridgeland, MS 39158
Tel: 601-856-4121
Fax: 601-898-4185
E-mail: JoeM@dtims.com
Website: www.dtims.com

Handles

Industry's largest selection of handle styles

Classic, ergonomic, telco and long-telco styles

Strongest latching button

Patented alignment pin

Panels

Blue LEDs meet IEEE 1101.10 hot-swap specifications

Surface mount LEDs designed for IEEE front panels

Switches and test jacks designed for IEEE front panels

Take the Elma Latch Challenge

One touch and you'll never go back to non-latching cPCI and VME64x handles

At Elma we're sure you'll prefer our line of ergonomically designed latching handles. Why? Because our patented 2-step extraction process eliminates problems. Unlike conventional handles, ours won't pinch your fingers. Or accidentally trip, causing costly downtime. Elma's hotswap and non-hotswap latching handles are perfect drop-in replacements for all popular front panel products. Our full line of cPCI and VME64X handles and front panels are available off the shelf or can be adapted to your specifications. But don't just take our word that these are safer, better feeling handles. Test one yourself.

Enclosures & Components | Backplanes | System Platforms | Switches, Knobs & LEDs | Cabinets

ELMA
Your Solution Partner

USA Elma Electronic Inc.

Phone: 510.656.0606 Fax: 510.656.8008 E-mail: sales@elma.com Web: www.elma.com

RSC# 19 @www.compactpci-systems.com/rsc

Next generation multimedia applications standards

By Thomas Howe

Many next generation applications require rich media collaboration, including not only voice, but video and data as well. In order for these applications to be developed quickly and inexpensively, all vendors and system integrators must adopt industry standards. Thomas identifies the most important standards for collaboration, especially as they relate to the convergence of digital media servers and next generation packet communications.

Today's communication applications are true multimedia experiences. Only a few years ago, simply adding video or data to a common voice application would qualify as multimedia. Today, we have the next generation of multimedia integration, where voice, data, and video combine with presence and chat to create compelling solutions. Tomorrow, we will see *bots* added to this application mix as well, as computers will communicate to us directly through voice or chat. As an example, imagine a conference that has, as a participant, another application that listens and records the conversation. In the middle of the call, you could ask it, "Can you play back what I said about customer focus?." The bot would then search back through the conference recording, find the appropriate keywords, and play that audio back into the conference. Eventually, multimedia applications will become *any media* applications. They will truly focus on the problem to be solved for a customer.

A changed landscape

This movement towards rich, multimedia communication applications is based in economic reality. The advent of technologies such as Voice over IP (VoIP) has inalterably changed the landscape for every carrier in the world. For the first time, customers have the option of choosing any service provider for their service. It is as easy for a broadband subscriber in the US to get service from ATT CallVantage as from Skype. Thanks to the global nature of the Internet, service now crosses borders as well, as evidenced by a significant portion of Vonage subscribers living overseas. This internationalization of carriers comes with a price: commoditization. It is increasingly difficult for service providers and carriers to provide services that are different from each other. If carriers cannot provide differentiating services, they will be caught in a commodity market where they will be forced to compete on price and join the dreaded *race to the bottom*.

The key driver that enables these new applications is standards. Standards allow independent vendors to contribute components of the solution in competitive yet interoperable ways. The key standards bodies that are driving next generation applications are the Internet Engineering Task Force (IETF) and the Third Generation Partnership Project (3GPP). The IETF is the moving force behind Session Initiation Protocol (SIP), the current standard for session initiation on the Internet. The 3GPP is a consortium of vendors and carriers that defines how multimedia applications will work for the next generation of mobile phone and Internet endpoints. Others contribute through computer industry collaboration, one example being the Speech Application Language Tags (SALT) Forum, founded by Microsoft, Intel, Cisco, and Philips. Another example in the hardware world would be the efforts made by PICMG, a consortium of 450 companies who collaboratively develop open specifications for high performance telecommunications and industrial computing applications.

Of all these standards efforts, SIP is arguably the most important. SIP is the de facto standard in the telecommunications world for enabling VoIP. Long before VoIP, the Multiparty Multimedia Session Control working group created SIP to manage multimedia sessions over the Internet. As such, SIP is built from the ground up to be multimedia aware. SIP handles voice, video, data, and presence with equal adroitness. In its role as an umbrella standard SIP references many other standards, such as RTP, the transport protocol for real time applications, and the SIP Instant Messaging and Presence Protocol (SIMPLE).

From an applications standpoint, the addition of presence information is the most groundbreaking technology in the last five years. Presence is the vehicle by which user preferences may be communicated to applications in real time. Anyone with experience with IM programs such as Microsoft Instant Messenger or Yahoo Instant Messenger understands presence. Your IM program will notice that you've moved your mouse or typed on your keyboard, and mark your status (or presence) as being online. If you walk away from your keyboard, it will set your presence to away after a few minutes. You can set it up so that it reads your calendar, and when you have a meeting, it will set your presence to busy. Perhaps when you pick up your phone, it will set your presence to *on the phone*. Using technologies like SIMPLE, all of this can be forwarded to an application and then used to make smart phone decisions. If your status is set to *busy*, maybe it will bring the call straight to voicemail. For another, it might be better instead if the call beeped in on the other line, especially if the caller was a family member. By including presence, the application has the opportunity to make smarter decisions in multimedia applications.

Not all is rosy

The standards front is not all rosy, however. One glaring omission from the SIP suite of protocols is one targeted towards media server control. Media servers are network elements that provide audio processing services such as conferencing, voice recording, and prompt playback. Today, this control interface is not standardized, and many proprietary implementations exist. One approach, called Media Sessions Markup Language (MSML) is a recent Internet draft from industry leader Convedia. MSML is used to control and invoke many different types of services on IP media servers, including audio and video streams.

As designers of next generation applications and networks, our success is predicated on open standards and protocols. Fortunately for us, there's never been more reason for vendors and carriers of all sizes to drive these standards to widespread adoption. Success will be yours if you, like many others, choose to compete on your special talents, and leverage the intelligence of the entire industry to complete your solution.

***Thomas Howe** brings over 15 years of telecommunications engineering experience to Versatel Networks. Prior to joining the company as CTO, he was CEO and founder of Tangerine Inc., a leading maker of call agents for small IP telephony service providers. Thomas also founded Howe Digital, a VoIP engineering design consulting organization. At PictureTel/ Polycom, he contributed to the audio and video software of the first PC-based video conferencing system, the first H.323 implementation, and the program that became Microsoft NetMeeting.*

To learn more, contact Thomas at:

Versatel Networks
80 Jean-Proulx Street
Gatineau, Quebec Canada J8Z 1W1
Tel: 819-771-0011 • Fax: 819-771-0031
E-mail: thowe@versatelnetworks.com
Website: www.versatelnetworks.com

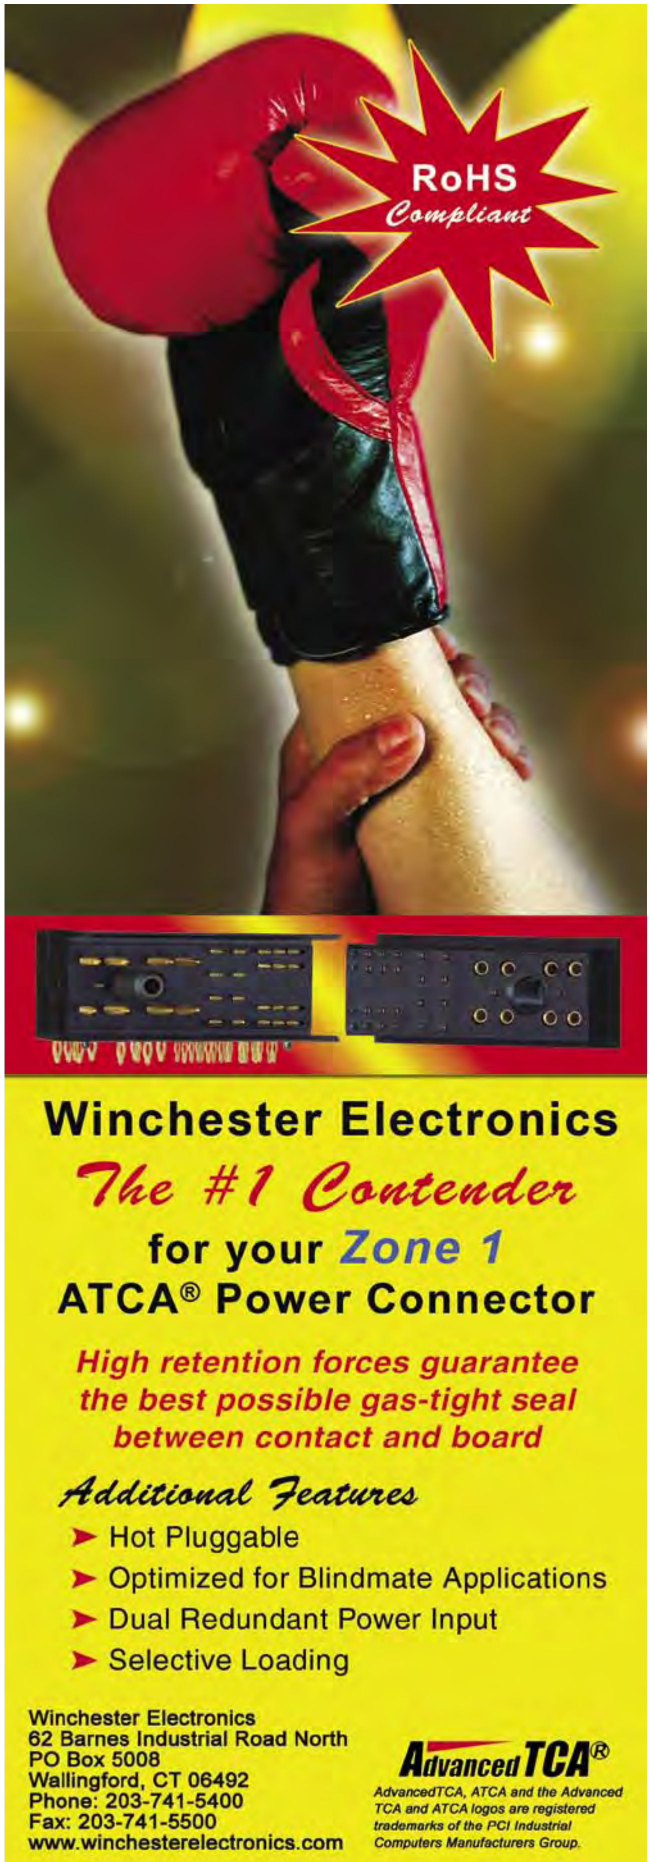

Introduction to Network Data Centers

By Kevin W. Bross

Network Data Centers are an increasingly attractive potential home for certain classes of communications equipment. The accelerating migration from circuit- to packet-switched networks facilitates the development of value-added services with differing levels of location dependence.

Today's high speed networks no longer require distributing core capabilities throughout the network in order to maintain acceptable latencies. Instead, a Home Subscriber Service (HSS) handles such core elements as subscriber identification and access information, subscriber location functions (for example, when roaming), and subscriber authentication capabilities. An HSS network element is essentially an IP Multimedia Subsystem (IMS) database. Designers are centralizing functions such as HSS in a few redundant locations to serve an entire network. Some service providers are choosing to optimize performance capabilities for certain network elements if they can find a reliable environment that frees equipment from some of the constraints typically imposed by standard Central Office (CO) environments and codified in Telcordia's Network Equipment Building Systems (NEBS) family of specifications.

During the Internet buildup in the late 1990s, many new data centers were built to handle the various emerging web hosting business models. Unfortunately, not all data centers were created equal. Some data centers were little better than improved warehouse space with generators and improved air handling.

However, some companies designing data centers built much more robust facilities, which are sometimes referred to as Network Data Centers. Network Data Centers typically have the following characteristics:

- Electrical feeds from two or more electrical grids
- N+1 generator capability, in addition to extensive battery banks
- Connections to six or more backbone carriers
- Hardened perimeter and extensive security
- Highly redundant Heating, Ventilation, and Air Conditioning (HVAC) systems with particulate control
- On-site network operations center for 24x7x365 support

In short, Network Data Centers are carrier grade environments designed to house mission-critical systems. Table 1 shows the attributes of a Network Data Center maintained by Sterling Network Services. The company's Phoenix, Arizona complex totals more than 75,000 square feet in the Downtown Phoenix Technology Exchange. An example layout of a Network Data Center from Sterling Network Services is shown in Figure 1.

Building	Blast-resistant concrete structure with 18-inch thick steel reinforced walls
Electrical feeds	Dual power entry vaults from Arizona Public Service's downtown Phoenix grid which is fed by multiple substations
Utility reliability	Zero critical power interruptions since business inception
Total input power	16,000 A @ 480 V
Backup power	Multiple diesel generators (1 MW and 1.5 MW)
Power distribution	AC (480/220/110 V) or DC (–48 V)
HVAC system	Liebert HVAC units (16 30-ton units in one data center alone)
Particulate control	Yes
Ceiling height	18 feet
Carriers	More than 40 service providers, including AT&T, MCI, Sprint, Qwest, SBC, Cox, Level (3) Communications, XO Communications, Global Crossing, AboveNet, WilTel Communications, InterNAP, Time Warner Telecom, and others

Table 1

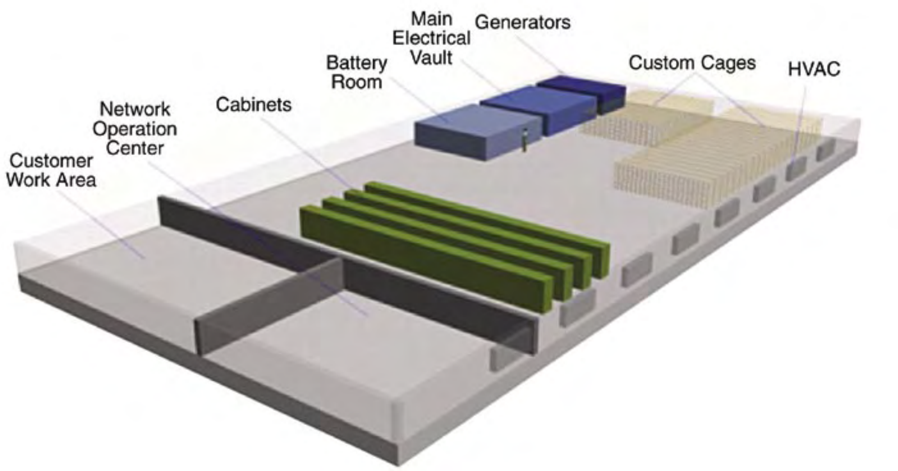

Figure 1

If we compare the typical environment of a CO versus a Network Data Center, we find some critical differences as depicted in Table 2. Note that data in the center column is derived largely from NEBS GR-63-CORE specification. Although CO requirements are clearly defined in Telcordia's NEBS specifications (and similar European Telecommunications Standards Institute [ETSI] specifications), the requirements of Network Data Centers are not as codified. To address this issue, Telcordia has recently initiated a new subcommittee to propose standardized requirements for Network Data Centers. The data listed in Table 2 is based on discussions with various data center operators over the last five years (mostly in the US).

The lower ambient temperature, broader temperature rise allowance, more tolerant acoustic limits, and increased airflow unfettered by local filters all conspire to provide increased cooling capability in each shelf. Combine these factors with a deeper cabinet geometry and higher power dissipation per square foot, and equipment in Network Data Centers has the potential to noticeably outperform products constrained by CO requirements.

Attribute	Central Office	Network Data Center
Normal ambient temperature	5 °C to 40 °C	10 °C to 35 °C
Short-term ambient temperature (for example, during HVAC failure)	–5 °C to +55 °C	10 °C to 40 °C
Maximum thermal rise (exhaust – inlet)	12 °C max	Unspecified
Filtering required at cabinet/ shelf level?	Yes	No
Power feeds	–48 VDC (wide ranging)	220/110 VAC, optional –48 VDC
Acoustic limits at normal temperature	60-65 dBA	Unspecified other than OSHA limits
Maximum cabinet depth	600 mm (about 24 in.)	600 mm, 750 mm, or 900 mm (about 24 in., 30 in., or 36 in.)
Maximum dissipation per square foot	50-120 W/sq. ft.	100-250 W/sq. ft.

Table 2

Some network elements such as HSS systems fit well with a centralized model. However, some network elements are naturally deployed in both centralized and distributed fashions. For example, the IMS Call Session Control Function (CSCF) capabilities typically divide into three distinct subelements:

- Serving CSCF (S-CSCF)
- Proxy CSCF (P-CSCF)
- Interrogating CSCF (I-CSCF)

Describing each element's role is beyond the scope of this discussion, but it is instructive to note that the S-CSCF is commonly centralized[1], while the P-CSCF and I-CSCF may be distributed or colocated with other network elements such as a gateway. In order to leverage common code bases, common middleware, validation efforts, spares, and application portability between these various elements, it is advantageous if a single architecture can be efficiently utilized in both the Central Office and Network Data Center environments.

Moving into the Network Data Center environment

Recognizing this issue, the subcommittee revising the PICMG 3.0 specification has adopted several modifications intended specifically to allow AdvancedTCA products to move into the Network Data Center environment.

One significant step in this direction is an allowance for AdvancedTCA front boards to dissipate up to 400 W per slot, rather than the previous limit of 200 W per slot. From the beginning, the PICMG 3.0 specification allowed for double-wide front boards to draw up to 400 W through a single connector. Allowing up to 400 W to be dissipated in a single slot is an acknowledgement that AdvancedTCA products are moving beyond the CO into environments where increased power dissipation may be acceptable.

However, it is important to note that changing a specification does not alter the laws of physics. Just because the PICMG 3.0 specification allows boards to draw more than 200 W for a given slot, it doesn't mean that shelves will be able to cool more than 200 W in that slot, depending on the power and thermal constraints imposed by the environment. To this end, the draft specification encourages board developers to stay within the 200 W limit if possible, stating explicitly that cooling more than 200 W per slot may not be practical in some thermal environments.

In some cases, a single blade design could yield products above and below the 200 W line of demarcation. Products drawing more than 200 W per slot could be simple derivatives of their lower thermal cousins, in the manner of blades designed for both full- and low-voltage processors. Other products may offer different I/O or peripheral capabilities, such as supporting higher powered mezzanines or more memory modules in the higher power configuration. Still other products may use advanced silicon functions such as voltage and frequency scaling to adjust to the maximum power/thermal limitations defined at the shelf level.

Another Network Data Center related change in the current PICMG 3.0 draft specification requires that board developers define their airflow requirements at both 40 °C and 55 °C ambient temperature. The 40 °C requirement provides worst-case thermal guidance for systems integrators targeting Network Data Centers. Some high power blades may not be usable in a given chassis if a 55 °C ambient temperature must be survived, but that same chassis may be able to cool those blades if the maximum ambient temperature remains below 40 °C.

A third change in the current draft of the PICMG 3.0 specification deals with shelf level airflow characterization. Based on this draft, shelf vendors will now have the option to characterize the airflow through each slot in their shelf both with and without air filters installed. If a shelf is going to be used in a Network Data Center, air filtering at the shelf or cabinet level is typically not needed. Removing the dense air filter designed to meet 80 percent dust arrestance could have a measurable impact on the airflow through the shelf. The current draft specification leaves it up to the shelf vendor to decide whether to provide this data. In some cases, a lower density air filter or some other form of airflow straightener may be needed in place of the dense air filter in order to achieve optimal airflow. These draft changes in the PICMG 3.0 specification are expected to be ratified by PICMG around the end of the year.

Many new communications services have elements that can be efficiently centralized, and this centralization often occurs in highly reliable Network Data Centers rather than more traditional COs. Figure 2 shows cabinets in a Network Data Center from Sterling Network Services. Network Data Centers generally have more forgiving environmental conditions than CO settings, opening the door for higher power equipment to be used in the Network Data Center. However, Telecommunications Equipment Manufacturers (TEMs) and service providers often need to leverage the same application software, middleware, and hardware in both the CO and Network Data Center environments in

Figure 2

order to leverage their substantial investments in development and validation of equipment.

AdvancedTCA equipment (built on the PICMG 3.0 standard) is already being used in CO equipment today. Recent draft changes adopted by the PICMG subcommittee overseeing changes to the PICMG 3.0 specification are designed specifically to allow the AdvancedTCA architecture to expand into the Network Data Center. These changes include:

- An increase in the maximum power draw per front board
- Specification of board cooling requirements at data center thermal limits (in addition to Central Office thermal limits)
- Use of AdvancedTCA shelves without air filters installed

These changes will allow higher performance AdvancedTCA bladed solutions to be built and deployed in Network Data Centers, fostering the development of new applications and services for the communications industry in an industry-standard architecture.

__Kevin W. Bross__ is a modular systems architect in the Modular Communications Platform Division in the Intel Communications Infrastructure Group. He has held a variety of engineering and marketing roles during 16 years at Intel. Kevin was a key contributor to several portions of the PICMG 3.0 specification and has been instrumental in the development of Intel's AdvancedTCA SBC, chassis, and management products. Kevin earned a Bachelor of Science degree in computer science and in mathematics and Bachelor of Arts degree in business administration from Principia College.

References

[1] Centralized functions should not be confused with centralized computing. In other words, functions such as S-CSCF can be efficiently distributed across multiple dual-processor blades, rather than needing a single large server such as an 8- or 16-way machine. It's the network function that is centralized in a common location, but that function can be accomplished by multiple systems working together in parallel.

For further information, contact Kevin at:

Intel Corporation
15400 NW Greenbriar Pkwy., CO5-156
Beaverton, OR 97006
Tel: 503-696-1411
E-mail: kevin.bross@intel.com
Website: www.intel.com/go/atca

Applying PCI Express switches to network adapter cards

By John Gudmundson

The introduction of many types of products based on PCI Express (PCIe) technology is moving forward at an accelerating pace. Solutions from root complex chipsets to switches, bridges, and endpoints for PCI, PCI-X, InfiniBand, Fibre Channel, Ethernet, and other standards are now available. Processors, both network and fully programmable, are being developed with native PCIe interfaces. Within the realm of switches, high port and lane count solutions are available today. These devices are well suited for applications within storage, server, and communications for backplane switching fabric designs. New generation switches coming to market will expand this portfolio and include lower lane and port count designs targeted at many mass market high-volume applications. These markets include laptop docking stations, printer and other peripheral device controllers, XMC and AMC mezzanine cards, and network adapter cards. Such PCIe switching solutions will need a wide variety of features to fill such diverse applications. One market in particular that can utilize PCIe technology is that of intelligent network adapters. Given the nature of these card architectures, lower port and lane count switches are highly beneficial for I/O aggregation and processor isolation.

Migration to intelligent adapters

Intelligent network adapter devices are rapidly moving to expand offload engine capability. With a local processor, these adapters can:

- Initiate and terminate TCP/IP connections
- Encode and decode SSL encrypted sessions
- Provide iSCSI protocol processing
- Manage RAID controller functionality and many other functions

It is imperative as the line rates reach and exceed Gbps speeds that these functions be offloaded from the host processor. Innovations within server and PC operating systems are providing compatibility with adapter card software drivers to allow these adapters to be practical. Intelligent adapters will need to be low cost, yet provide high throughput. In order to accomplish a low cost per bit of data transferred and provide high bandwidth, adapter cards are moving to incorporate several ports per card. Multiple ports can enhance bandwidth or add redundancy and to do so will require a method to aggregate multiple flows through a single upstream port to the host.

Adapter I/O data aggregation

One method to fan in data from multiple ports is through the use of an aggregation switch. Using switches based on PCIe technology provides low cost, high bandwidth, Quality of Service (QoS), low latency, and other crucial benefits (see Figure 1). The PCIe switch combines two I/O endpoints and allows aggregation to the PCIe four-lane upstream port. Switches allow peer-to-peer support for data transfer to the local CPU for offload-engine processing before being transferred to the x4 port. This prevents the bottleneck of having data sent upstream to the host and back to the local CPU. I/Os could be based on Gigabit Ethernet, Fibre Channel, InfiniBand, or other outside the box type communications standards. The PCIe devices shown in Figure 1 are bridges from one protocol to PCIe technology. In this design, the one lane of PCIe to the I/O is more than sufficient as this allows a net bidirectional flow of 2 Gbps per direction. The PCIe links are aggregated through the switch to the upstream PCIe host. The local CPU or a network processor handles the offload engine computations. The x2 link to the CPU provides the bandwidth to allow wire-speed operation from two ports. In addition to the bandwidth, these switches offer low latency. Switches will be available with latency from ingress to egress ports under 150 ns, even for packets with larger data payloads. This low latency is not possible with many competing standards but is absolutely required in communications systems.

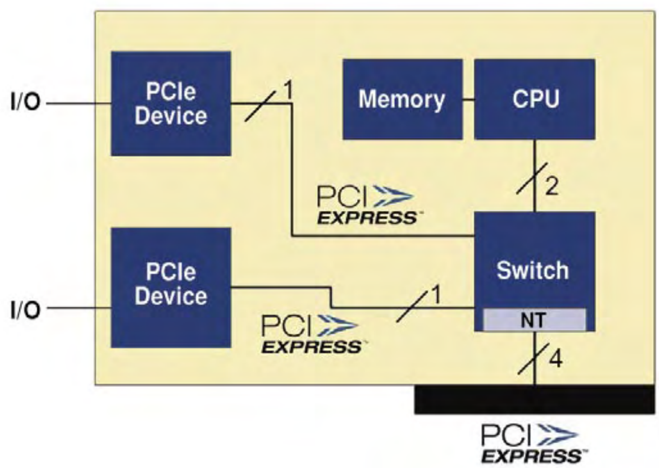

Figure 1

Processor isolation

Other critical switch features include QoS, end-to-end data integrity, and processor domain isolation. PCIe switches coming to market will provide for multiple Virtual Channels (VCs) along with the mandatory eight types of traffic class. Each class of data is mapped to the VCs available on a link by link basis. The switch offers a number of selectable VC arbitration techniques to provide the VC prioritization. PCIe technology offers numerous data integrity techniques. One important optional method is based on end-to-end cyclic-redundancy check packet protection. Switches for such adapter applications will provide this capability to ensure reliable transport through an extended series of PCIe links from the adapter through the host side to the root complex. Perhaps the most significant need in these switches is the ability to provide processor isolation. With local intelligence and a system host, system discovery and configuration can result in memory mapping conflicts with these multiprocessor systems. Switches will be enabled with selectable Non-Transparent Bridging (NTB) ports to provide processor isolation. The switch's non-transparent port will prevent the system processor from enumerating and configuring the downstream endpoint devices by making the switch appear to be an endpoint using type "0" configuration status registers. Also provided for NTB is address translation for both address and requestor ID routed packets between these processor domains. Switching through NTB ports does not result in any significant added latency. Next generation systems, soon to be on the market, will take advantage of these advanced PCIe switching solutions.

John Gudmundson *is a senior marketing manager for PLX Technology, Inc. Prior to PLX, he held marketing positions at Sunrise Telecom, Nortel Networks, Integrated Telecom, and Ascend Communications. John is an active member of the Advanced Switching Interconnect Special Interest Group (ASI SIG) and is a member of the IEEE and American Marketing Association. He holds a BSEE from the University of California, Los Angeles, a BSBA from the University of California, Berkeley, and an MBA from University of Southern California.*

For further information, contact John at:

PLX Technology, Inc.
870 Maude Ave.
Sunnyvale, CA 94085
Tel: 408-328-3500
Fax: 408-774-2169
E-mail: jgudmundson@plxtech.com
Website: www.plxtech.com

What can you say about Convergence in a Box?

How about Wow.

Features

- Scalable from 4 to 256 Low Speed Links
- Supports SS7 HSL and ATM T1/E1
- Hardware configurations from 1U simplex to 4U High Availability
- High density and redundancy from a single point code
- Supports SIGTRAN M2PA, M2UA, M3UA and SUA and SS7 interworking
- Common hardware platform with easy management software
- Gateways upgradable without having to change the application

The Adax Signaling Gateway

Adax, the company you've come to trust to deliver high-performance signaling solutions across broadband, narrowband, and IP networks, now offers even greater choice. The Adax Signaling Gateway supports SS7/IP switching, routing, tunneling, and backhaul. Fully redundant options with no single point of failure are available today. These products enable a simple and straightforward migration of existing SS7 nodes to IP transport, saving the costs associated with leasing or provisioning dedicated long haul SS7 circuits. The same solution meets the new demands for IP signaling and SS7/IP interworking in converging IP, circuit, and wireless networks.

Building Blocks

Integrated Blades

Complete Gateways

adax inc (510) 548 7047 sales@adax.com www.adax.com
adax europe ltd +44 (0) 118 952 2800 sales@adax.co.uk

adax

Maximizing backplane switching performance of AdvancedTCA systems

By Matthias Lim and Sorin Vasile

AdvancedTCA is a wide reaching standard that covers a broad array of applications for next generation communications and computing systems. However, in some high capacity switching applications, system designers are encountering bandwidth demands that strain the resources provided by the AdvancedTCA standard. Limited connectivity between the node and hub boards often constrains backplane throughput for many high capacity packet-based applications. Various standards bodies and component vendors are aggressively pursuing 6 Gbps to 10 Gbps backplane serial links. Although AdvancedTCA implementations can employ the resulting incremental bandwidth when available, this is still only part of a scalable solution. System designers recognize that truly efficient system implementations maximize the use of backplane bandwidth, directly translating improved link speed to scalability advantages. This type of efficiency is rarely demonstrated in today's switch fabric technologies. Building more robust, efficiently scalable, and reliable 10/20/40 Gbps per Node Board AdvancedTCA systems for higher capacity applications demands a more advanced switching architecture. Without this, much of tomorrow's serial link improvements will not increase effective system capacity.

AdvancedTCA benefits

Broad adoption of the AdvancedTCA standard has ushered in a new era of collaboration among system and component players in the communications and computing markets. Even at this early stage in its evolution, it has had an undeniable impact on numerous applications across a wide range of modular platform applications. Considerable effort was required to ensure that a common form factor, power supply, airflow, and backplane connectivity produced a suitable and cost-effective standard across a common set of interchangeable hardware components. The clear advantage of this approach is to ensure economies of scale for device and subsystem vendors while system designers enjoy a competitive environment populated by numerous vendors supplying similar or interchangeable parts.

AdvancedTCA backplane connectivity limitations

As is often the case with any broad standard, challenges and limitations emerge when attempts are made to extend towards implementations outside the bounds of the initial applications. Today's communications system designers are struggling to adapt the existing AdvancedTCA platform to their critical system requirements. These needs include Node Board scalability extending from 10G to beyond 40G of user port capacity. User bandwidth requirements are often supplemented by additional internal system overhead consisting of tagging and forwarding information. Communications applications require robust QoS mechanisms to support emerging applications with latency and bandwidth guarantees for video, voice, and data – often delivered simultaneously to the same end user. These service capabilities must also be provided in a fashion that ensures carrier-class reliability through fully redundant hardware mechanisms.

Today's base AdvancedTCA specification, PICMG 3.0, does not define the backplane switching mechanisms required to provide these capabilities. Instead, a single channel (four bidirectional ports as shown in Figure 1) is defined between each Node Board and Hub Board in the Zone 2 Fabric interface, the primary interconnect intended for data traffic. Even mesh implementations, which offer additional interconnect bandwidth, do not add to the single channel between any two boards in the widely adopted 14- and 16-slot systems.

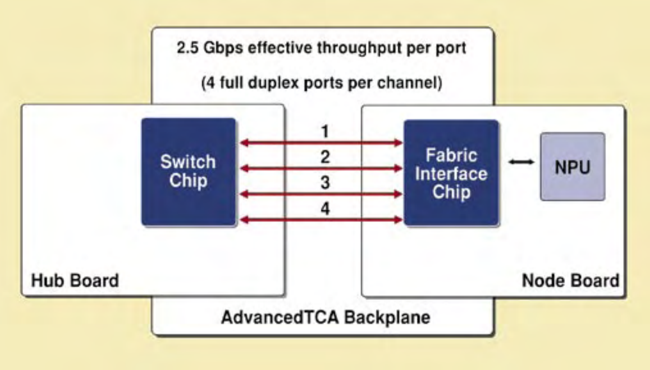

Figure 1

Taking 3.125 Gbps (2.5 Gbps actual data throughput without 8b/10b coding overhead) as the maximum raw throughput rate per Zone 2 Fabric interface port, a single channel presents a serious limitation to system designers who need adequate throughput for 10G or more of carrier-class user bandwidth per Node Board. Though the base specification does not restrict increases in backplane rate beyond 3.125 Gbps, most currently available backplane switching technologies support only 2.5 Gbps to 3.125 Gbps of raw throughput per port. In addition numerous tagging and checksum bytes encapsulating their backplane payload often burden high performance communications switching platforms, increasing the demands placed upon the system backplane (Figure 2).

Led by various semiconductor vendors and standards initiatives, significant and promising efforts are underway to increase the raw capacity per AdvancedTCA port through 6 Gbps to 10 Gbps and beyond. Depending upon the switching implementation, however, this increase in raw throughput may not fully solve the scalability demands of these high capacity applications.

Backplane interconnect options

Today's system designers are presented with a dizzying array of alternatives for AdvancedTCA backplane interconnect. Two categories of solutions have emerged to meet their needs. One group has eyed the opportunities offered by open standards. These solutions have been reviewed, debated, and ratified by the PICMG committee. Another group that finds its way into the AdvancedTCA market comprises companies that build self-contained, off-the-shelf switch fabric technology. They couple proprietary switching architectures with integrated serial link technology. Both groups have their merits and compete effectively in the market.

A number of semiconductor vendors support the PICMG 3.x standards and provide the hope, if not the promise, of future Node Board interoperability. These vendors form the first group.

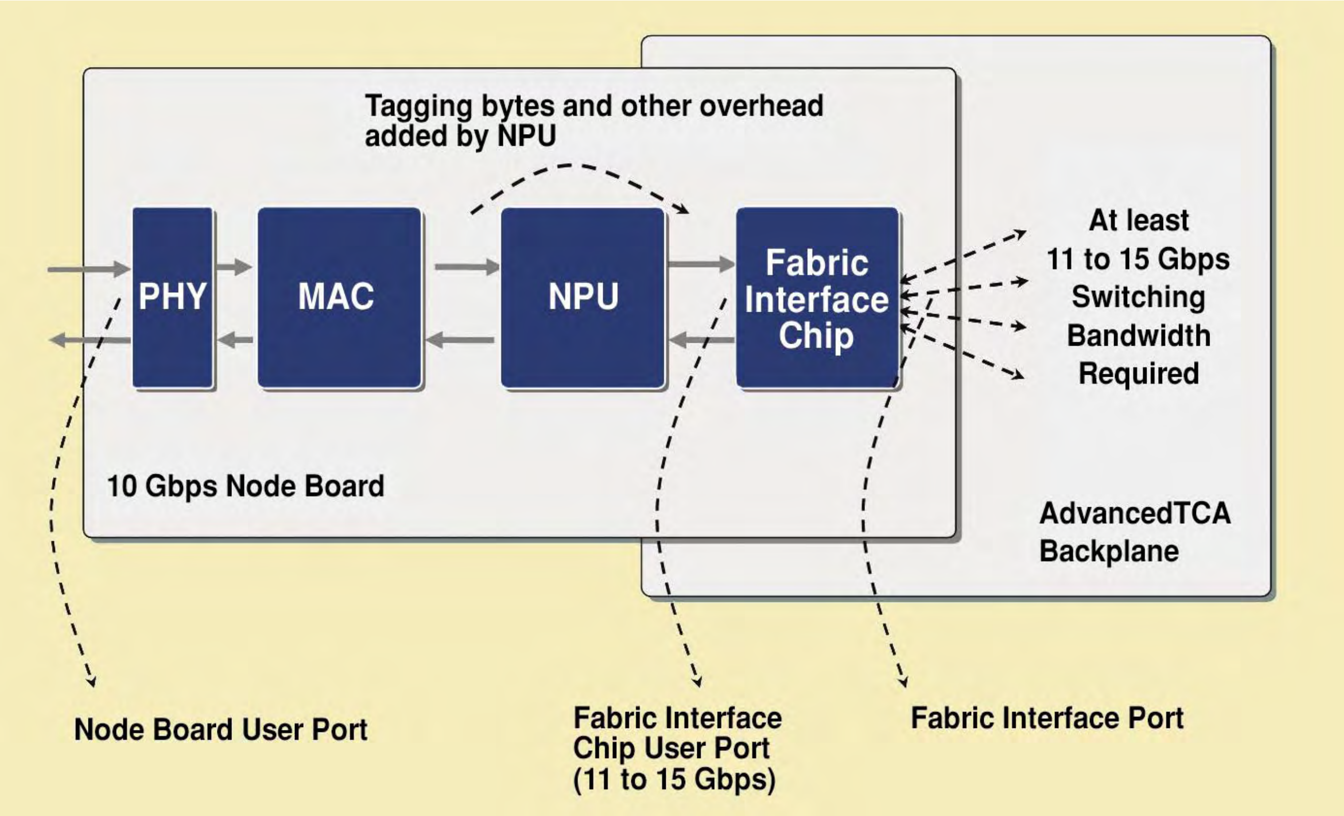

Figure 2

The standards include Ethernet, Fibre Channel, InfiniBand, StarFabric, PCI Express, Advanced Switching, and RapidIO, which are defined in PICMG 3.1, 3.2, 3.3, 3.4, and 3.5 respectively. Each has its own advocates, ecosystems, and best-fit applications. Despite varying degrees of success for each standard, growing demands for carrier-class 10 Gbps user throughput per Node Board have created opportunities for alternative switching technologies.

Proprietary switch fabric technology, freed from the constraints of industry consensus, holds the greatest promise for pushing the envelope of AdvancedTCA Node port user capacity to 20G/40G. Here again, today's implementations suffer from lagging serial link rates, often coupled with inefficient cell-based internal switching mechanisms.

Neither of these solutions offers the market true 10G Node user throughput with full redundancy, nor the promise of 20G/40G given the constraints of today's AdvancedTCA channel limitations.

Solutions for the AdvancedTCA Node Board bottleneck

Fundamentally, the issue to be addressed is the Node Board backplane bandwidth constraints imposed by the limited number of ports between each board as defined in the PICMG 3.0 specification. Potential solutions that would increase this bandwidth are likely to encompass one or both of two characteristics – increased efficiency of the port and channel bandwidth as well as maximization of the throughput or link rate of each port. Maximizing the Node user throughput of next generation AdvancedTCA Node Boards requires solutions that combine both approaches.

Switch fabric architectural efficiency issues

Backplane efficiency can be generally defined as the amount of user traffic carried in the backplane divided by the raw backplane throughput. In an ideal system, every bit of raw throughput can be used to carry a bit of user traffic. Real world systems, however, must allocate a portion of raw bandwidth for purposes such as framing, DC-balancing CODECs, and error detection. Reliable operation requires these critical functions, and little improvement is expected beyond today's high rate CODECs. Beyond this, many scalable implementations devote considerable raw backplane bandwidth to provide a speedup factor to account for partially filled cells. Minimizing this speedup can maximize backplane efficiency.

Switch speedup solves an otherwise intractable problem. When user traffic and any associated system overhead bytes are divided into fixed-size cells certain inefficiencies occur. Most notably, the division of variable sized packets or frames into a number of cells almost invariably results in frequent partially filled cells. In the extreme, pathological traffic distributions can leave as many as half of the backplane cells largely unfilled – an enormous waste of bandwidth. The traditional approaches to resolve this issue range from brute force (massive over-provisioning of backplane bandwidth) to clever (minimizing the cell size and combining partially filled cells wherever possible). These solutions fail to address the underlying problem of cell-based backplane transport and invariably result in less-than-optimal backplane efficiency.

In order to reduce the inefficiency caused by speedup, backplane interconnects must depart from cell-based switching and instead employ packet-based switching. In this switching methodology, the packet payload exactly matches the packet size. Packet-based switching directly transports the entire user traffic packet, with any associated system tagging bytes, without being disassembled into two or more transport cells. This reduces the backplane speedup to the barest minimum, often close to none, and maximizes the backplane efficiency.

The Enigma switch fabric chipset, for example, implements this approach and thus brings the speedup factor very close to 1x.

Accelerating the backplane

Overcoming the speedup issue is only part of the puzzle in maximizing AdvancedTCA backplane use. With growing bandwidth demands, one inevitable push is to extend serial link speeds beyond 3.125 Gbps. Today's efforts push backplane link speeds to 6.25 Gbps or even 10 Gbps to fully capitalize on these scarce interconnect resources. To be accepted in critical applications, a link must be fast, robust, power-efficient, observable, manufacturable, and of course, readily available in silicon. To meet these requirements, an advanced link technology should include the following:

- Comprehensive multitap transmit and receive equalization
- Continuous decision feedback based receive equalization
- Continuous transmit equalizer adaptation
- Waveform observability of the post-equalized receive signal
- Voltage- and time-based Bit Error Rate (BER) testing mechanisms

Additionally, world-class solutions permit the system designer flexibility in adjusting both serial link rates and CODECs to optimize the serial link subsystem to match the system environment. These characteristics make it possible to deploy the same advanced serial link technology throughout a fully populated system, accounting for inter-link crosstalk equally well for both short traces and long.

Such high speed serial links, however, are very sensitive to signal attenuation and impairments associated with all the components


AdvancedTCA®
CompactPCI®
CONEC®
Technology in Connectors
USA
Tel: (919)460-8800
Fax: (919)460-0141
Email: info@american.conec.com
CANADA
Tel: (905)790-2200
Fax: (905) 790-2201
Email: info@conec.com
EUROPE
Tel: 049-2941-765-0
Fax: 049-2941-765-65
Email: info@conec.de
WWW.CONEC.COM

used to construct the backplane channel. Therefore, designers must apply a more stringent set of rules for PCB layout of the backplane, the node and hub boards, and package to minimize the impact of these factors. In addition, environmental factors such as humidity and temperature changes affect link performance, as do geometry manufacturing variations and component and trace material composition. A smart serial link technology is needed to adapt to these factors.

This type of state-of-the-art adaptive serial links technology is integrated into Enigma's switch fabric chipset. For instance, if temperature or humidity fluctuates, the continuously adaptive link adjusts to reduce the impact from these environmental elements and maintain the lowest possible bit error rate. The other benefit of a continuous adaptive technology is the minimization of serial link maintenance through self-calibration. Calibration for a large number of serial links in a system can be an enormous manufacturing challenge and very inefficient in practice. This serial link technology makes the system highly scalable by achieving higher speed, but not at the expense of system reliability and availability.

Deploying advanced link technology in a real backplane requires a careful balance of design parameters. The Advanced Differential Fabric (ADF) backplane connectors now available are capable of operation beyond 6 Gbps under some circumstances. The backplane design itself, however, is fraught with risk. Potential impairments include:

- Signal crosstalk
- Reflections from via stubs
- Dielectric losses
- Skin effect losses

Much has been written about methods to reduce and balance these effects, but the application of advanced backplane design techniques to the AdvancedTCA specification introduces unique complications. The routing geometry proposed in the AdvancedTCA specification is a very good solution for backplanes with a dual star topology. For more complicated topologies – dual-dual star and full mesh – it is still possible to use this geometry, but this leads to a board thickness of up to 8 mm in a full mesh topology. Further, the placement of hub slots in some configurations leads to very long trace lengths leading to increased signal attenuation.

In response, Elma Trenew's R&D department has optimized the routing geometry and material such that a board thickness of only 3.2 mm was possible. In addition to reduced stub effects, this also permits additional flexibility in selecting from a wider range of PCB vendors that are able to produce PCBs that are 3.2 mm thick. Elma also changed the placement of the hub slots to the middle of the backplane, reducing the length of the traces and therefore dielectric losses. This has resulted in commercially available AdvancedTCA dual-dual star and full mesh backplanes capable of considerable improvements in serial link performance.

Conclusion

Today's widely accepted AdvancedTCA platform imposes a significant bottleneck on system designers pushing to maximize Node Board user bandwidth. Inefficient architectures coupled with 3.125 Gbps link technology fail to deliver true 10G Node ports, particularly when system overhead such as tagging bytes are included and full redundancy is required. However, more efficient and minimum speedup switching architectures coupled with advanced serial link technology and backplane techniques are easily capable of delivering carrier-class 10G performance in a current AdvancedTCA platform. These technologies are also well suited for pushing this industry standard to 20G, 40G, and beyond.

Matthias (Wai-Leng) Lim *is a member of Enigma Semiconductor's Applications Team. His main focus is in system switching applications for modular systems. He previously worked at QuickLogic as an applications engineer for six years specializing in DSP applications, PCI 33/66 MHz, and high speed interconnects such as HyperTransport in the FPGA environment. Matthias holds a MS in Engineering Management and Leadership from Santa Clara University and a BS in Electrical Engineering from the University of Arizona.*

Sorin Vasile *is currently a senior project manager at Elma Trenew Electronic GmbH. He has eight years of experience working as a backplane design engineer and as a technical project manager. He focuses now on the management of the backplane projects and the development of new products implementing switched fabric architectures. Sorin holds a BS degree in Electronic Engineering from the Politehnica University of Bucharest (Romania).*

For further information, contact the authors at:

Matthias Lim
Enigma Semiconductor Inc.
2975 Scott Blvd., Suite 110
Santa Clara, CA 95054
Tel: 408-567-0536 • Fax: 408-567-0527
E-mail: matthias_lim@enigmasemi.com
Website: www.enigmasemi.com

Sorin Vasile
Elma Trenew Electronic GmbH
Stuttgarterstr. 11, D-75179 Pforzheim
Germany
Tel: +49 (0)72319734-610 • Fax: +49 (0)72319734-81
E-mail: sorin.vasile@elma.de
Website: www.elma.de

Programming network processors: Solving the complexity crisis

By Kevin Graves

While the activity in the networking and telecom space isn't near what it was back in the good old days of the late 90's, those of us in the business of designing networking equipment have seen the complexity of applications and the number of challenges posed by the development process in this area continue to rise. Very predictably, this additional complexity has resulted in compromises being made in product functionality, lengthened development cycles, or both, while rapidly evolving standards have fueled increased risk.

Among the macro factors driving the increased complexity of networked systems are:

- Heightened security concerns, resulting in an increased number of network elements with security functionality along with the combination of what have traditionally been discrete security devices into multifunction devices, such as Unified Threat Management systems
- Increased mobility, resulting in significant increase in network endpoints (for example, handsets) as well as contributing to IP conversion and an increase in traffic
- The general increase in traffic requiring higher line speeds, higher throughput, and more network elements
- Convergence and standardization on IP and eventual migration to IPv6, resulting in an increased number of gateways (including media gateways and Packet Cable)

Despite the convergence and standardization of IP, there's still an incredible amount of change with new protocols, additional standards, new uses, and new threats emerging at an alarming rate. While each one of these elements drives increased function, capacity, and/or performance in networking equipment, let's focus on just one driving force, security, for the purpose of illustration. Maintaining information security is one of the biggest challenges affecting the Internet's use for communications. If we narrow the scope a bit more, we are witnessing Voice over Internet Protocol (VoIP) rapidly gaining popularity as a low cost alternative to traditional service, but implementing it securely exposes major challenges for communications systems designers.

Let's drill down a bit further and look at a key element of VoIP systems, Session Border Controllers (SBCs). As Figure 1 illustrates, these devices enable VoIP media to traverse enterprise firewalls and Network Address Translation (NAT), as well as supporting other advanced features such as Secure Real-time Transfer Protocol (SRTP). Examples of complex SBC features include *network hardening* to prevent denial-of-service attacks, *deep packet processing* such as signature detection for intrusion detection, and bandwidth-theft protection.

SBC system providers, like other communication system vendors, are rising to the challenge by providing equipment with the performance to analyze packets of data on the fly at wire speed. To do this, designers and developers (SBC vendors are a perfect example) have typically implemented programmable processor-based systems (as opposed to purely ASIC- or FPGA-based designs), because such systems maximize the flexibility necessary to keep pace with the rate of change, typically incorporating the following:

- High degrees of parallelism, either through multicore general purpose processors or network processors
- Fixed-function accelerators, such as encryption/decryption, hash units, and TCAMs
- Various memory types, each with potentially different access models and timings
- Various buses/interconnects to other silicon, again, each with different access models and timings
- Interaction with other system elements or architectural planes such as a control plane or management plane

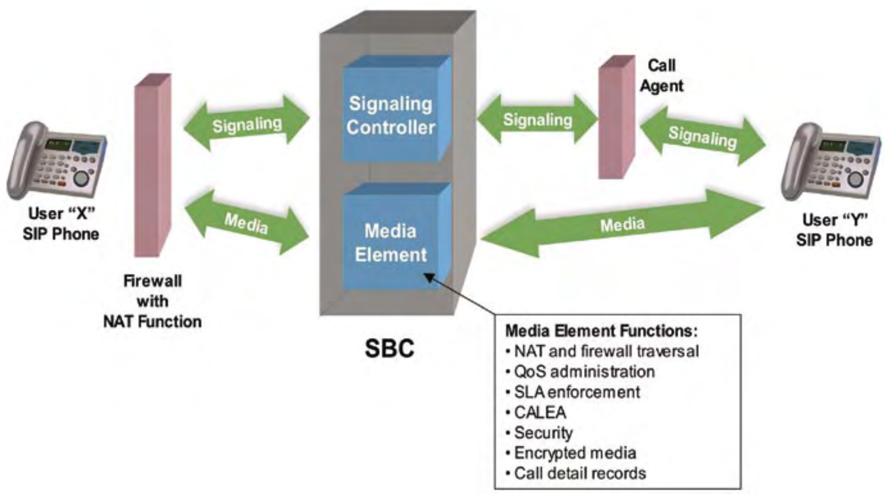

Figure 1

The most obvious impact of this complex architectural environment is increased development and lifecycle software costs...

Despite pressure from carriers to minimize the number of unique boxes in their networks, SBCs have strategically managed to establish a foothold as standalone network elements. Why? It's simple – SBCs have stayed ahead of the innovation curve and have been flexible enough to incorporate new features much faster than other network elements. It should be noted that SBCs have a key architectural advantage over many other pieces of networking gear in that they have visibility to both the signaling and media content. This ideally positions them to implement a multitude of new features.

But all of this comes at the cost of drastically increased software development complexity. The reason is simply that meeting the performance goals has typically required a new and complex architecture for use within the products. A good example is the proliferation of a relatively new class of processor called the *Network Processor (NP).* NPs are optimized for packet processing because they generally incorporate a high degree of parallelism through pipelining or superscalar processor arrangements:

- Multithreading
- Multiple memory types
- Dedicated-function accelerators (such as hash units and CRC generators)

The most obvious impact of this complex architectural environment is increased development and lifecycle software costs compared to using general-purpose processors due to the steep learning curve and lengthened development, debug, and test phases. What's more, the lengthened development cycle typically results in a functional prototype not being available until late in the project cycle, thus delaying integration with other system components and overall system performance modeling. How many of us have worked on projects that were cancelled because the functional prototype system wasn't ready for a key customer demonstration? How many engineer-hours have been spent optimizing or redesigning what was only *thought* to be the performance bottleneck, only to have the true bottleneck discovered elsewhere in the system? Another byproduct of the high complexity is that many times designers are hesitant to modify or enhance working designs due to the high risk of change and lengthy debug cycle, therefore negating one of the strongest benefits of NP-based designs, flexibility.

In our SBC example, we can clearly see the importance of flexibility and extensibility. Reviewing the history of SBCs reveals the original reason for their existence: The VoIP signaling protocols, namely Session Initiation Protocol (SIP), H.323, and Media Gateway Control Protocol (MGCP), are "poorly behaved," meaning they transfer key information in layers above OSI Layer 4. This behavior prevented the VoIP media traffic from traversing the various Network Address Translation (NAT) devices and firewalls between endpoints. To solve this, SBCs monitored both the signaling traffic (to "sniff out" the parameters of new VoIP calls) and the media traffic (to "open up pinholes" and enable the media to pass through the various security devices).

But SBC vendors didn't stop innovating. As VoIP traffic increased, other deficiencies in the existing networking and telecom equipment manifested themselves, and the SBC became the natural "clearing house" for dealing with such issues. Once it was clear that having visibility to both the signaling and media traffic provided a significant architectural advantage over other boxes in the network, many additional features were added, such as Quality of Service (QoS), route optimization, sophisticated call detail records, and encryption/decryption. None of this innovation would have been possible without a very flexible underlying implementation that allowed for extensibility and scaling.

Growing complexity

The task of implementing such systems is becoming more complex with each product generation, however. As more commu-

nications and networking system vendors seek to exploit the new highly parallel silicon such as NPs, a software development complexity crisis is brewing, demanding innovative thinking to provide a solution. After decades of communications system design where system optimization relied on low-level code or microcode optimization to squeeze out maximum performance, the task of optimizing the newest NP applications has become too complex for this technique to be accomplished in a timely and cost-effective manner.

As an example of SBC application complexity, consider the relatively simple task of programming a NP to implement pinhole firewall, NAT, and encryption. In somewhat simplified terms, the software engineer must program the processor to first determine if the incoming packet is part of the VoIP media (as opposed to some other control or management packet), and if so, perform a flow/session lookup to determine if the packet is part of an established and accepted session. This lookup is often based on the 5- or 6-tuple of IP and Layer 4 header fields with an optional VLAN identifier. In addition, the lookup must be performed on a table capable of tracking the maximum number of supported concurrent connections (often tens or hundreds of thousands). Due to the size of the table, engineers usually employ hashing or other search algorithms rather than linearly searching the entire table. If the packet isn't part of an established session, it is then processed as an exception, and some combination of logging, dropping, and alerting is performed on the packet. If the packet is part of an established session, NAT is performed by updating the packet header fields using data structures indexed by the unique flow index returned from the lookup. Other state variables are inspected to determine if this is a secure flow, and if so, encryption keys and cipher state (for example, which cipher and/or hash is to be used) are accessed through similar data structures, and the packet is encrypted or decrypted. Lastly, packet values are updated to account for any padding and updates to UDP checksum and the IP header checksum. The packet is then forwarded to the appropriate output port.

Expressing the aforementioned logic using assembly language code would typically require thousands of lines of code and would not only involve coding the packet logic described previously, but also dealing with the intricacies of the underlying hardware, such as creating a functional pipeline, then partitioning and synchronizing the function across pipelined processors.

Implementing this in C language would not be much easier, because of the lack of any type of parallel operating system or library support and the need to interact with the underlying NP at the "hardware" level. The time to develop and debug the code from scratch could easily take months. Even with reference code available, integrating and testing the resultant application would certainly take weeks. If changes or upgrades are required, the changed packet logic must be debugged and tested, as well as repartitioned and remapped to the underlying hardware.

Then, if new features need to be added, it's back to the drawing board – not only to implement the new features, but also to repartition the entire application across the parallel engines.

Calling for a new approach

Employing powerful software abstraction techniques offers a far better approach. These techniques enable application programmers to focus on adding packet-processing and application features rather than on dealing with the underlying parallel hardware bits and bytes. One way to do this is via a very high level, application-specific programming model that supports programming in a *functional* language that could naturally be implemented as a Virtual Machine (VM) atop the NP (see Figure 2). Implementing a VM in software atop the NP provides the programmer an architecture-independent environment with the potential to be completely portable and scalable. Other significant benefits include offering superior robustness by building in logic to perform bounds checking, null pointer/handle checking, and other exception handling. What's more, a VM approach allows for advanced concepts such as dynamic compilation.

Referring back to the SBC example discussed earlier, use of a programming model that provides a robust set of built-in algorithms such as a high performance n-tuple table lookup (typically required for NAT) or a standard way to encrypt/decrypt or authenticate packets (required for implementing secure media) would decrease the effort of programming and debugging the application by more than an order of magnitude. Tailoring the programming model to a specific class of applications (for example, packet processing) allows the model to include a very focused set of data types, control mechanisms, and built-in values, all of which would be useful in dealing with VoIP RTP packets/streams. Such tailoring further abstracts the programming task. IP Fabrics is promoting such an approach and has implemented it with its Packet Processing Language (PPL) software.

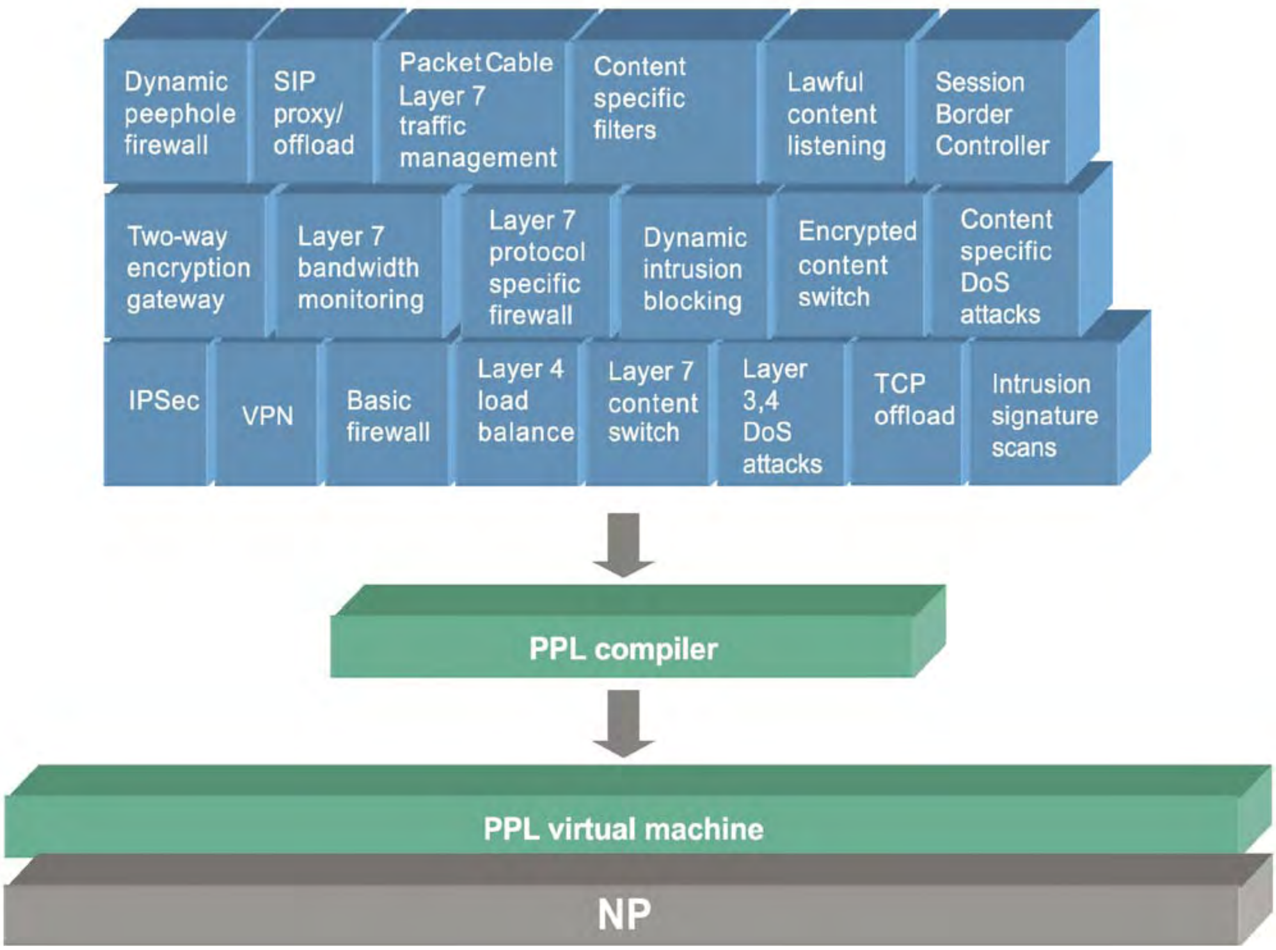

Figure 2

PPL provides these functions and more for the Intel IXP28XX and IXP23XX NP families.

IP Fabrics' PPL is *packet-centric*, meaning the primary features and benefits center around primitives and complex algorithms that process IP packets (it should be noted that PPL is suitable for processing entire Layer 2 frames – from the Layer 2 headers up through Layer 7 content). In the PPL programming model, the fundamental object is the packet to which the PPL virtual machine provides a rich set of packet-handling functions such as header insertion/stripping, connection/session tracking, content inspection, and rate monitoring. Additionally, the PPL software has a number of lower level built-in functions such as automatically calculating key packet state values. These values include the offset and length of the various headers and whether or not the packet is part of a fragment. The software also allows common packet fields to be accessed symbolically. But probably the most important aspect of the IP Fabrics approach is that it makes it possible for users to express their logic in this very high level manner, but then automatically maps this logic onto a parallel, complex network processor, yielding a very high performance implementation.

The concept of masking complexity by abstraction is not a new idea in itself, and skeptics might say that NP virtual machine abstraction must come at a price – most often a real or perceived performance penalty when compared to ideal performance. By focusing on a particular application-specific domain (for example, packet-processing) however, the VM can be implemented with highly optimized, best-of-breed algorithms and state machines used in that domain. These hand-tuned application building blocks will often outperform equivalent functionality implemented by NP application designers. The overall VM architecture can be optimized to the application's domain-specific processing and data flow characteristics. IP Fabrics PPL, for instance, uses through pipelining and parallelism as the underlying architecture. And lastly, one could argue that the final system performance will be greater when all components of the system are available as early in the development cycle as possible, thus allowing more time to analyze and optimize in real-world conditions. As the complexity of designs increase, it will become more difficult to perform any low-level tuning and optimization.

By contrast, NP experts have the background and can take the time required to heavily optimize the VM implementation. VM users can then leverage these optimizations via the high-level programming language.

One final point to make is that abstraction isn't the solution to all of today's networking problems. There are very clear cases for highly tuned and optimized designs for some network elements, such as core routers and switches. However, the complexity crisis is very real and needs to be addressed. And, one proven approach to address this is with an application-specific abstraction layer enabling developers to focus their attention on application-level issues and optimize functionality rather than spending precious design time dealing with the underlying hardware architecture.

Kevin Graves *is a member of IP Fabrics' founding team, and has been the company's CTO since 2002. In total, Kevin has 21 years of experience, with most of his pre-IP Fabrics experience at IBM where he had roles in engineering and marketing management for communications adapters. Recently he was senior director of engineering of the Telecommunications Division of RadiSys, where he had responsibility for the Oregon, Houston, and Boston design centers. He has a strong technology and customer focus in packet networks, signaling, and media gateways. He has a BS in computer science from Pennsylvania State University.*

For further information, contact Kevin at:

IP Fabrics
14964 NW Greenbrier Parkway
Beaverton, OR 97006
Tel: 503-444-2400
Fax: 503-444-2401
E-mail: kevin.graves@ipfabrics.com
Website: www.ipfabrics.com

PRODUCT GUIDE

NETWORK PROCESSORS

COMPANY NAME/ MODEL NUMBER	AdvancedTCA	CompactPCI	PICMG 2.16	1G Ethernet	10G Ethernet	T1/E1/J1	T3/E3/J3	OC-3	OC-12	PMC sites	PrPMC sites	PTMC sites
ACT/Technico www.acttechnico.com												
RaidStor			•	•								
ADLINK Technology www.adlinktech.com												
cPCI-6240-2			•	•						1	1	1
cPCI-6841		•		•								
cPCI-6842		•		•								
Advantech www.advantech.com												
SF-420		•										
AMCC www.amcc.com												
nP3710 ATCA	•			•				•	•	1		
Artesyn Communication www.artesyncp.com												
SpiderWareNP			•					•				1
Katana750v			•	•								2
CES www.ces.ch												
TNS 4880			•			•	•	•	•			1
Continuous Computing www.ccpu.com												
FlexCompute ATCA-PMC40	•			•						2		4
FlexPacket ATCA-BC10	•			•								1
FlexPacket cPCI-IXP2406			•							2		
PACKETblade BC10	•			•								2
GarrettCom www.garrettcom.com												
Magnum Blade DS12		•										
Interphase www.interphase.com												
iNAV 4000			•	•		•		•	•	2		
IP Fabrics www.ipfabrics.com												
Double Espresso				•								
PPL for Intel IXP28XX/23XX	•											
Motorola www.motorola.com/computing												
5385		•	•	•								
CPIP5365			•							2		
cPCI-920/922			•								1	

continued on page 46

PRODUCT GUIDE

NETWORK PROCESSORS

COMPANY NAME/ MODEL NUMBER	AdvancedTCA	CompactPCI	PICMG 2.16	1G Ethernet	10G Ethernet	T1/E1/J1	T3/E3/J3	OC-3	OC-12	PMC sites	PrPMC sites	PTMC sites
Performance Technologies www.pt.com												
FlexNAT NAT/LSNAT			•	•								
PMC101 Processor Module		•		•								
RadiSys Corp www.radisys.com												
Promentum ATCA-7010	•			•								
Spectrum Signal Processing www.spectrumsignal.com												
aXs.1110		•				•	•	•				
TeraChip www.tera-chip.com												
TeraChip TLC-NIC101	•				•							
TriEMS www.triems.com												
TRL3613		•						•				
Voiceboard www.voiceboard.com												
TDM to Packet Processor			•			•						
Xalyo Systems www.xalyo.com												
XS-AMC1	•			•				•	•			

A look at network processor development environments and programming

By Curt Schwaderer

CompactPCI & AdvancedTCA

In today's technology world, it's a foregone conclusion that a circuit board with any kind of intelligence will use a microprocessor. Designers put CPUs on boards without any obstacles or development worries. But some of you may remember when designers commonly developed circuit boards with fixed-function logic chips and the introduction of something called a microprocessor loomed on the horizon. Early on, microprocessors lacked widespread use because of the programming complexities involved. Engineers had to be skilled in assembly or machine language. Little or no software development tools existed. Of course, we all know how the microprocessor story played out – the wide variety of programming languages, run-time software, and development environments make it easy and natural to develop systems that perform as required while having flexibility to update and enhance the capabilities of the system.

Today, network processors have reached a new era. Network processor technology has moved from the hot new technology on the block to an industry-wide understanding of what a network processor is and the product benefits and development challenges involved. As with all technologies, though, companies are working hard at eliminating challenges, hoping to make network processor usage as ubiquitous as the microprocessor.

In this month's Software Corner, we'll take a look at the advancements in the tools, languages, and run-time software available for network processors. A number of companies involved with network processor technology are referenced here to bring you up-to-date with what's going on in the world of network processor programming. Four network processor vendors were interviewed for this article as well and my thanks goes out to their representatives for taking the time and effort to provide much of the content here. If you haven't looked at network processors lately, I think you're going to be very interested and excited about the new software development capabilities of today's leading network processors.

EZCHIP

EZchip

Eyal Choresh, Vice President of Software, Systems, and Support for EZchip, mentions programming network processor-based boards can appear complex to designers at first. Customers often have complex and processing-intensive control and management plane applications that must be integrated with – and make optimal use of – available network processor cycles. The integration of the control plane software with the data plane processing of the NPU is just as critical to the success of a project.

EZchip addresses software development complexity starting with the NPU architecture. The generic *microengine* term comes in the form of a *Task Optimized Processor* or TOP. There are four types of TOPs in the EZchip NPU:

- TOPparse – Parsing packets for key fields to pass onto the TOPsearch
- TOPsearch – Takes the key fields found by TOPparse and performs table lookups
- TOPresolve – Performs traffic metering and policing along with packet forwarding tasks
- TOPmodify – Performs packet content modification as the packet is being transmitted over the network

The language to program these TOPs is somewhere between assembly and a high level language. Basic instructions perform low level tasks, with higher level macros to accomplish

Continued on page 50

AMCC

AMCC

Key aspects of AMCC's product architecture and strategy include the core architecture, coprocessor, and traffic management competencies coupled with software expertise in tools and run-time software. AMCC has also developed systems knowledge they use to steer the NPU architecture.

Steve Klinger, Senior Manager of Software & Solutions Marketing in the Communications Business Unit of AMCC described AMCC NPUs as a set of multithreaded nPcores, each core supporting 24 contexts with an instruction set optimized specifically for packet processing. See the figure on page 52. Steve notes, "Each context is a *run to completion* programming model – a packet arriving on an I/O port is handled start to finish by a single context. Around the core is all the logic that accomplishes multithreading. We've found that this provides an optimal way to reduce programming complexity while achieving nearly 100 percent utilization of the nPcore cycles in the NPU."

All contexts of all nPcores execute from the same program store in the chip. So, the data plane code structure tends to be a front-end lookup based on the header contents of the incoming packet with a jump through a jump table of packet processing options based on the lookup results. As a result, when a context is handed a packet that context will jump to the relevant code block to execute the rest of the processing to be performed for that packet.

Continued on page 52

AGERE

Agere

Key network processing functions where typical network processors excel are pattern matching, policing, data manipulation, queue/buffer management, and statistics gathering. The Agere NPU family supports these important NPU functions with a combination of silicon acceleration and software components. In addition, Agere's NPU family supports fine-grained traffic management to enable revenue generation from QoS and customized services on a per subscriber basis.

The figure shows the Agere NPU block diagram. The Agere NPU includes an ARM CPU that can serve as a control plane processor or interface with an external CPU and manage the data plane processing part of the silicon. The data plane processing components include functional blocks that comprise the packet processing path through the NPU.

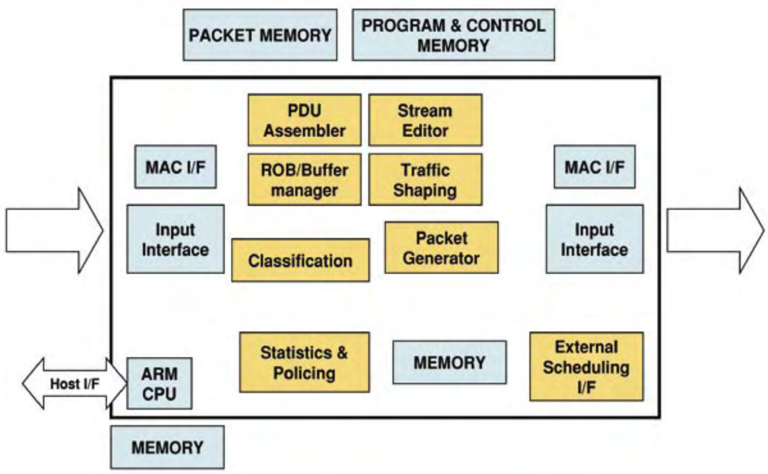

Agere's NPU software strategy has been to reduce the number of lines of software code that their NPU customers must write to make their network equipment function. Writing fewer lines of code enables these customers to reduce product development costs, gain more time to focus on developing differentiating features of their products, and boost overall equipment reliability. Fewer lines of code also typically translates to fewer mistakes.

To reduce the number of lines, Agere implemented a *Functional Programming Language (FPL)* strategy with programming model highlights that include:

- Single-threaded programming model hides the underlying multi-threaded operation
- Hardware hides all aspects of management of parallelism such as synchronization and scheduling
- Hardware isolates the programmer from housekeeping exercises such as acquiring and freeing memory, loading and storing PDUs, manipulating pointers, building linked lists, and creating and maintaining lookup tables

Continued on page 56

INTEL

Intel

Intel's approach to network processors includes a microengine building block with an instruction set tuned for packet processing that can scale across multiple parts in the product line. Also included in the part is an Intel XScale processor that can be used to manage various aspects of the microengine environment, or in some cases, serve as a control and/or management plane processor. Their programmable network processors start with the Intel IXP2300, which contains an XScale processor with four microengines, up to the IXP2850. The latter contains an XScale, 16 microengines, and integrated hardware acceleration for encryption. The figure illustrates a simplified architecture diagram of the Intel IXP NPU family.

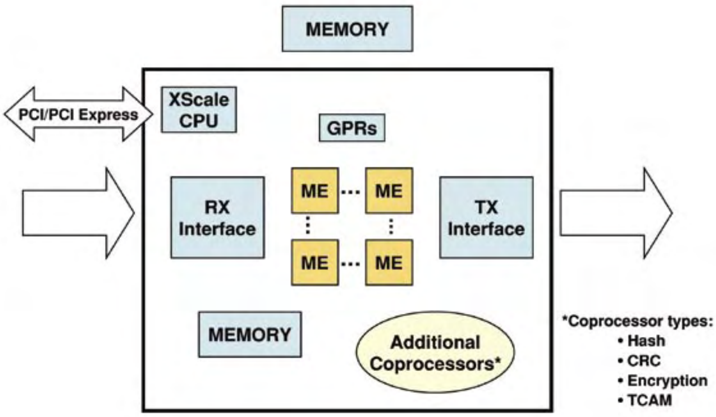

Each microengine is designed to be *generically programmable* meaning that the same instruction set applies to all microengines in the NPU. However high speed interfaces between adjacent microengines provide for small-cycle-transfer of data or control information between adjacent microengines. Each microengine can be programmed to provide four or eight contexts, and each has its own control store where the data plane program resides. The contexts within a microengine execute from this control store. So in general, programming the Intel product line of NPUs consists of designing the algorithms that make up the application, determining how many microengines and/or contexts within microengines to apply to each algorithm, and organizing the processing such that functions that communicate often are situated in adjacent microengines. This approach capitalizes on the high speed communication mechanisms provided in the silicon. This fully programmable architecture provides equipment manufacturers with broad flexibility to implement differentiated services while supporting lines rates up to 10 Gbps.

Any number of hardware assist components may be available as well, depending on the part. Onboard hash units and TCAMs exist for special processing. The IXP2850 has an on-chip crypto that can perform AES, MD-5, and SHA-1 cryptography under microengine program control. There are also up to three types of memory. Scratchpad memory is on-chip and is typically the fastest access time. Programmers employ SRAM

Continued on page 57

Continued from page 48

more complex tasks such as checksum computation. Packets flow through the NPU, being pipelined from TOPparse through TOPsearch and TOPresolve, until finally being transmitted by TOPmodify. This pipeline is controlled in hardware, removing the complexity of a previous stage needing to find an available TOP in the next stage to process its packet. The figure shows a general block diagram of the part.

In general, the user writes four separate data plane programs, one for each type of TOP. The NPU itself deals with the parallelism and pipelining of packets through each pool of TOPs. Programmers use the EZchip development environment to simulate their applications, then load and run the application on the NPU. Each class of TOP executes from the same program store (labeled *memory* in figure). Each TOP processes packets according to the program, then passes the result to the next TOP pool in the pipeline stage. Each TOP pool forms one stage of the processing pipeline. Passing packets and control information between TOPs is also hardware controlled. This eliminates the program from having to worry about details involving the number of processing engines running the program or how to send data and control information to one of the TOPs in the next pipeline stage (associated blocking/buffering and related issues).

Development environment

EZchip's IDE, EZdesign, is a set of tools bundled in a GUI-based development environment on Windows. The environment contains a code editor, preprocessor, and compiler, as well as a clock-accurate simulator for modeling the application. In addition, EZdesign can generate packets and lookup tables used by the simulation environment. A generic program template can be used as the starting point for writing the data plane processing for a given TOP class.

Once the program for each TOP in the pipeline is written, the tools allow you to point-and-click your way to a configuration output file that's in an EZchip format called Network Processor Script Language (NPSL). A piece of run-time library code that is linked with the user's control plane application interprets the NPSL configuration file. So, from the viewpoint of the control plane application, it's simply compiling in some C files that perform the specialized initialization, configuration, and program downloading to the NPU.

Once the program is loaded into the NPU, EZdesign can set breakpoints, and probe registers and memory, among other tasks. It's also possible to look into each internal stage of the NPU pipeline to see what's going on with each TOP. Programs can be run in interactive or batch mode, which generates a log file. The log file provides a comprehensive dump of what's going on inside the NPU so you can run things real time and get the interaction and debugging information afterwards. These capabilities are present for hardware debugging or working with the simulation environment. A remote debug library for the host CPU does the linkage between the debugging environments into the target system. These libraries use the PCI to talk to the network processor itself, like an in-circuit emulator.

Writing data plane code for the Ezchip NPU

The EZchip TOP instruction set is a mixture of assembly and high level macros Since TOPs perform specific functions, each type of TOP has a different set of high level macros based on what functions they perform.

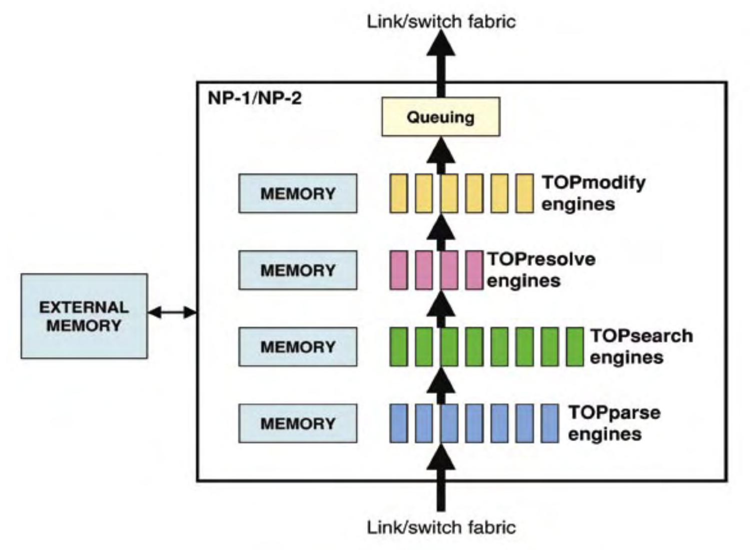

EZchip NPU instructions have the format:

<opcode> <destination> <source(s)>
<Operation type> <flags>

Packets enter the TOPparse stage and go into Frame Memory. The TOPparse has the ability to look anywhere in the Frame Memory for its processing. The final stage (TOPmodify) then has access to Frame Memory that can incorporate areas of the packet that the previous stages have requested to be modified while the packet is transmitting. Internal to the part and hidden from the program is the queuing between stages of the pipeline. Some stages may be very processing-intensive, while others are relatively simple, so you can imagine that input queues for a given stage may get filled. The previous stage would have nowhere to queue their control information. In these cases on-chip backpressure occurs within each stage until it reaches the TOPparse stage where flow control input signaling can throttle the input for the part.

Control plane software

The EZdriver kit contains a set of libraries with high level C APIs that give access to the network processor data plane. APIs exist to download microcode, create table structures, and provide debug services. These APIs abstract the specific table format used by TOPsearch by requiring that table information be added through the APIs. Then libraries within EZdriver generate the tables in the right format for the TOP. The EZdriver library uses the PCI interface for communication with the NPU. So, the user supplies callback read and write routines that the EZdesign libraries use to communicate with the NPU. There are also APIs available for sending and receiving of packets between the NPU and the control plane processor (for example, misses in the data plane lookup and route updates). The control plane APIs are not Network Processing Forum compliant but do represent similar functionality for controlling the NPU.

The EZdriver product supports Linux and VxWorks. Customers have ported to QNX, OSE, and other operating systems. EZdriver contains an OS abstraction layer that most people have been able to port in less than a week. All code is in ANSI C, and users get full source.

Available solutions and services

Available data plane solutions for the EZchip NPUs are Layer 2 switching with VLAN, VPLS, Layer 3 routing with IPv4/IPv6/

MPLS, load balancing through URLs, firewall and network address translation, and Martini draft.

The TOP opcodes are centered around packet processing, but don't have any specific support for WAN frames. So, things like ATM segmentation and reassembly or ATM Adaptation Layer (AAL) processing must be done external to the chip.

EZchip is able to write data plane solutions for customers from startups to large companies with big projects. Some companies that do control plane software that runs on the EZchip NPUs include Futuresoft and IP Infusion. However, the customers did the porting themselves.

There are no current formal partnerships with other control plane software vendors, but cooperation between these vendors occurs as customers require partnerships in a more informal way.

What's next?

The main area EZchip continues to focus on is making the part easy to use. For example, maybe a GUI development environment where a customer looks at graphical representations of packets and with a few clicks and GUI assistance the engineer can specify what the processing is and 50 percent to 70 percent of the data plane code will be generated automatically. They are also looking at porting the tool chain to other non-Windows environments.

EZchip product line overview

EZchip is a spin-off from LanOptics. LanOptics was formed in 1990, developing token ring hubs that were widely used in the European market. LanOptics went public in 1992 and branched out into Ethernet hubs and switches, controller cards, and eventually network subsystems. Leveraging its competencies in networking and ASIC design, LanOptics formed EZchip with the goal of doing everything the LanOptics product did in their network system product on a chip.

- The EZchip NPU silicon
- EZdesign – tools for data plane software development
- EZdriver – control plane interface software and libraries
- Evaluation platforms
- AdvanncedTCA based blade with the EZchip NPU which provides 10 1 Gigabit Ethernet interfaces
- EZsystem with a pizza box with the EZchip NPU, which provides 10 1 Gigabit Ethernet interfaces along with an Ethernet control interface and serial connection
- Four-day training on the EZchip software environment
- Data plane design and implementation services

Continued from page 48

The NPU also may contain a pool of coprocessor units such as an internal Ternary Content Addressable Memory (TCAM), hash unit, and traffic manager. Silicon FIFOs and arbitration logic make these coprocessor units available to any of the nPcores, freeing users from worry about programming in arbitration and contention issues between contexts except in extreme cases.

The NPU also has LA-1 and quad-data-rate SRAM bus external interfaces for connection to external TCAMs and external memory. Two buses to DRAM – one for context memory and one for packet memory – allow the memory references to process a packet to be performed in parallel, speeding up the overall processing of a given packet.

Lookup is performed using an internal TCAM-based *policy engine* to determine which data plane code block in the control store to jump to, then further lookups and classification-based information deeper in the packet finally perform transformations or other processing like transmit/drop decisions.

Traffic policing, queuing, and scheduling are well developed in the AMCC NPUs – there are various levels of granularity to the scheme in the NPU. At the highest level, there is the ability to enforce a bit rate on a given interface for WAN edge and access applications. Designers can further implement policing and statistics based on a particular type of software defined flow. Algorithms are available at a high level such as weighted round robin. For incoming flows, a traffic manager block handles admissions control (Weighted Random Early Discard or WRED and dynamic thresholding) with support for up to 128K per-flow queues. As discussed previously, this traffic management block also provides bit rate scheduling per flow, the ability to bundle flows into a *pipe* with each pipe representing a specific Service Level Agreement (SLA). The traffic manager block allows the software to define up to five levels of hierarchical groupings. Traffic processing is symmetrical, so software processing can occur right after packets arrive and/or just before they are transmitted.

On-chip packet transform engine

As the data plane program makes changes to the packet, these changes aren't directly performed in packet memory. The changes are fed to an on-chip *packet transform engine*, which adds, removes, and modifies contents within a packet as the packet is being transmitted out the interface.

The architecture is similar and consistent across the NPU product line – the family consists of faster, more/less nPcores, or different numbers and types of coprocessors.

Steve sees ease of programming NPUs as a major barrier to widespread NPU usage, but says this has and continues to be aggressively addressed and the market has grown to the point of being an interesting business. "Before the technology bubble burst in 2000, there were a large number of start-ups aggressively using NPU technology to get a better product faster to market more cheaply than using ASICs. The business has now evolved to the point where Tier I OEMs are widely adopting NPUs. The key to widespread adoption is to achieve performance goals while making things easy to program. Efficient data plane coding still

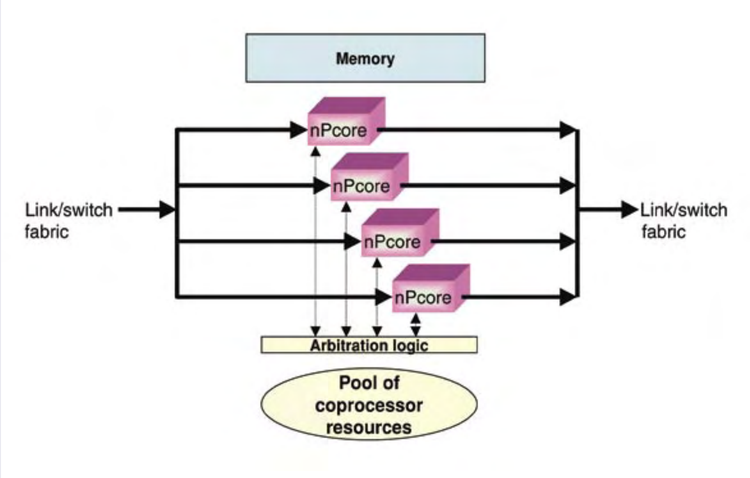

requires knowledge of the part and how the silicon components interact. Customers are looking for tools and architectures that allow them to run high performance data plane applications without having to deal with every register and bit inside the NPU," he adds.

Development environment

The AMCC development environment is called the nPsoft Toolkit. The Toolkit consists of an integrated development environment, code generation tools, simulation, and debugging environments.

Code generation tools include a macro assembler and linker as well as C language data plane code. Once the application is developed, it's linked with the appropriate application libraries and the nPkernel. The nPkernel foundation services enable queuing and interaction between units of the parts, and macros for such tasks as adding headers or issue classification requests to a policy engine. These services abstract functions to a high level, but implementation is highly optimized for the part. The nPkernel also ensures that code is easily ported from one device to another. The simulation environment allows for:

- Command line interface or graphical interface execution
- Validation and bug fixing
- Some performance analysis

The simulation environment allows the programmer to step through the processing of each packet and includes built-in facilities for packet generation and storing/logging. Using the simulator programmers set breakpoints and track what code gets exercised, determine how many lines for processing per given packet, and gather performance statistics.

The command line driven debugger enables developers to peek and poke registers on the chip, similar to an operating system ROM debugger for NPUs, and includes instruction traces, set breakpoints on context, and on-chip read/modify registers. While it's not a graphical debugger, the debugger's simplicity has the benefit of being able to be easily ported to custom designs. This allows the same debugging environment on a development platform to be present on customer hardware, a significant aid to hardware bring-up.

Continued on page 54

Continued from page 52

Writing data plane code for IPv4/MPLS

Data plane code is written in C, as are data structures and syntax. Many of the NPU-specific function calls come in the form of nPkernel macros. The structures that the data plane coder deals with, C language and the macros along with the assembler and linker, convert the syntax into an executable image to be run on the NPU. The source code in Example 1 illustrates the AMCC NPU syntax.

```
MPLS/Ipv4 Packet Classification Example
// Classify the Packet
search_db (PKTCLFY_DB, classifyResult,
classifyKey0, classifyKey1, classifyKey2);
// Process the packet based on the classification result
switch ( classify_pkt(PKTCLFY_DB, classifyResult) ) {
        case 0:
                goto ingressDropFrameSPU;
//port administratively disabled
        case 1:
                goto enetDix8023xPause;
//pause frame
        case 2:
                goto enetDixInvalid;
//null MAC-SA
        case 3:
                goto enetDixInvalid;
//multicast/bcast MAC-SA
        case 4:
                goto enetDixArp;
//bcast arp
        case 5:
                goto enetDixArp;
//unicast arp (ie arp reply)
        case 6:
                goto enetDixIpv4ExpiredTTL;
//TTL = 0 or 1
        case 7:
                goto enetDixLerIpv4ChecksumWrap;
//special checksum case
        case 8:
                goto enetDixLerIpv4ChecksumCarry;
//special checksum case
        case 9:
                goto enetDixLerIpv4;
//IPv4 with no exception cases
(etc.)
```

Example 1

Writing in C provides a level of familiarity for the programmer. Using macros supplied with the nPkernel, customers write a smaller amount of code, from low hundreds of lines for any particular processing path up to a high end of a few thousand for a complete integrated application with multiple flows being processed different ways. Steve states, "AMCC uses a specialized microcode library together with a specific NPU silicon architecture. This removes many NPU overhead complications for data plane program development."

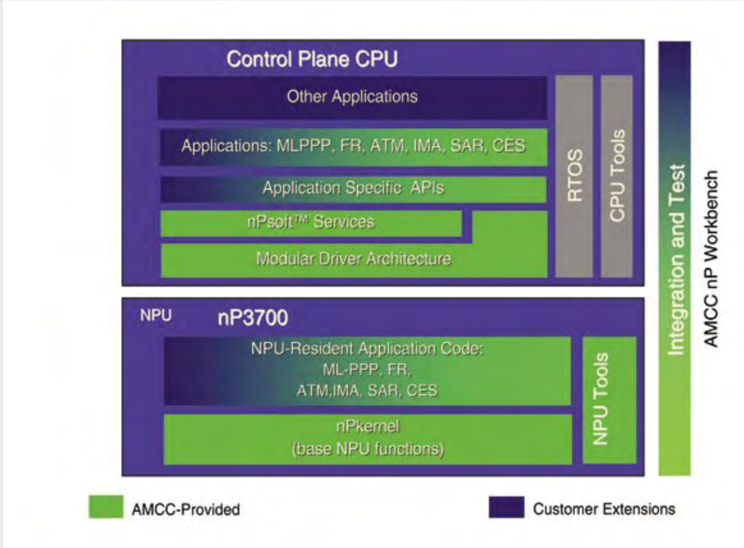

Control plane software

The control plane software (see figure above) is a separate domain with an independent set of tools, but AMCC does provide some control plane interface software. AMCC also offers PowerPC processors for use as control plane processors, but the AMCC NPUs are not tied to a specific control plane CPU.

AMCC's control plane driver handles microcode download to the NPU. Data plane program store initialization between the control plane CPU and the NPU comes in the form of AMCC *applets*. An applet is a "special" type of packet that the NPU knows to DMA into the control store memory. This allows the run-time environment to update data plane programs during operation of the part. The control plane interface also has the ability to pass packets between the control plane and the data plane. Additional services include statistics gathering and table updates.

These control plane APIs are similar to, but not Network Processing Forum compliant. Steve mentions that the NPU architectures are quite varied at this point, making it difficult to provide enough standardized, specific set of functional APIs. While it's disappointing that standardization of these APIs has lagged, Steve is optimistic that standards will continue to evolve and become an important part of the industry as NPU usage matures.

AMCC uses VxWorks for bring-up of the development boards, but OS abstraction within the control plane software environment makes it easy to support a variety of operating systems.

Available solutions and services

In the early days, most of what was provided was example code. Customers looked at the code, but ended up rewriting the entire application. AMCC is now providing rigorously designed and tested application software, which they call nPsoft Application Service Libraries. These applications are also provided in full source.

Among the many application libraries available are ATM Layer 2, AAL5 processing and Ethernet/IP, and MPLS switching and routing. Other applications that may not be as obvious include multiservice switching, multilink PPP, Inverse Muxing over ATM (IMA), and Frame Relay interworking solutions.

The standard tools purchase provides the NP kernel, host driver, and a basic framework used as the starting point for data plane development along with a sample data plane solution for IPv4. The commercial application software is a separate product.

Formal partnerships with control and management plane vendors haven't been financially rewarding in the past, leading AMCC to handle integration and cooperation on a customer-by-customer basis. Steve does think AMCC is likely to add a services partner in the future to provide additional services beyond what they do themselves. VioSoft has been a key technology contributor to development environment. So, AMCC has utilized vendors for the NPU tool chain. Other possible partnerships include more value-added tools and development of additional application libraries.

What's next

Historically, AMCC NPUs perform Layers 2-4 processing within the communications infrastructure and telecommunications markets. AMCC has heavily penetrated the Metro Area Network (MAN) application space, gaining a dominant share of this arena. Over the next few years, AMCC plans to move into deeper packet processing and target new markets. Some of the areas mentioned are mobile infrastructure, broadband access (DSL, PON, cable), and wireless LAN. These new initiatives will begin to move past the telecom space into the enterprise as well.

Steve believes that the important aspects of quality software development environments for NPUs are to make the chips fundamentally easier to program and to couple that effort with effective development tools and flexibly written data plane applications.

As far as a possible emergence of a unified software platform analogous to an operating system for CPUs, Steve thinks the widely varying nature of NPU architectures makes one unified software environment difficult to achieve right now. He does believe the future holds room for innovation and as the market matures, a common software platform will become viable.

AMCC product line overview

The AMCC network processor product line grew from an acquisition of MMC Networks, a company widely credited with inventing the concept of a network processor. AMCC is now shipping their fifth generation of NPU product. Key aspects of the product architecture and strategy include the core architecture, coprocessor, and traffic management competencies coupled with software expertise in tools and run-time software. AMCC has also developed systems knowledge they use to steer the NPU architecture.

- NPU silicon
- nPsoft services – core data plane software for NPU programming
- nPsoft toolkit – code generation, simulation, and debug
- nPsoft applications
  - Ethernet L2 switching
  - IP routing/forwarding
  - MPLS LSR and LER
  - POS, ATM, PPP, FR, Ethernet
  - VLAN services
  - ATM L2 switching with OAM
  - ATM AAL5 SAR
  - Policing
  - Multiservice interworking
  - Multilink PPP, IMA, ML-FR
  - IEEE 802.1 P/Q
  - Pseudo-wire Circuit Emulation Services (SAToP)
- Evaluation platforms
  - AdvancedTCA
  - Modular line card
- Design support
  - Training, architecture, and implementation assistance
- Control Plane CPU – PowerPC

Continued from page 49

- Programming languages are highly optimized for network processing solutions, with all programming in high-level languages and no microcode programming
- Multistage operation-specific functional pipeline achieves high performance, predictable throughput, and efficient resource utilization, enabling deterministic behavior
- Programmer retains low level control of the data stream and data constructs, allowing code and application optimization

Recognizing that classification is similar to database searching, Agere employs and supports FPL as a high level *pattern matching* language. FPL is a functional language that specifies mappings or transformation of packets and describes the task to be performed on a packet, not how the task is to be performed. Each line of FPL code is a rule that specifies a pattern associated with a resulting action. So, if the pattern is matched against the packet being processed, the specified action for that rule is executed.

Writing data plane code for the Agere NPU

A rules-based language such as FPL hides the complex parallelism of the network processor implementation, enabling programmers to specify actions on classes of packets without worrying about parallelism or pipelining issues.

An example of a rule programmed in FPL is:

```
ClassB: 10b net host action($1, $2)
```

The *ClassB* labels the rule. The next three parameters describe the pattern to match on the packet. If the packet matches, then the actions listed in the action() parameter are executed.

Programming of all other data plane functions, such as buffer management, policing, traffic shaping, and packet modification, are accomplished by writing small functions (or scripts) in an enhanced subset of the C language, called C-NP.

With the Agere NPU architecture and programming model, a given functionality can be programmed in 30-40 times fewer lines of code than implementing the same functionality on network processors using a multiRISC engine approach. For example, Agere sites that less than 10,000 lines of data plane software implements a complete DSLAM application including support for multiple encapsulations (PPPoEoA, IPoEoA, IPoA, PPPoA), Layer 2 bridging, Layer 3 routing, VLAN tagging, Virtual MAC addressing, ATM AAL5 processing with per flow traffic management, IGMP, SNMP, and DHCP. 100 percent of the data plane code is in high level languages (FPL and C-NP).

Development environment

Along with the programming languages, Agere also provides a comprehensive tool suite integrated in a single framework, called the Software Development Environment (SDE). The SDE includes support for system simulation (multiple processors and I/O interface devices). It includes FPL and C-NP compilers, debugger, code profiler for test coverage analysis, traffic generator and analyzer, a simulator with comprehensive set of performance metrics, and a traffic plotter to view simulation results. Developers use this environment for all the data plane program development and optimization. The tool suite generates a single image, including configuration information as well as compiled FPL and C-NP programs, which can be directly loaded onto the hardware using a single API call by the application. No additional tools are required for any aspect of data plane development.

For target access applications, Agere's complete application software packages include the complete data plane software and an application level, very high level API. The company provides FPI software packages for DSLAM and Wireless Access applications. The control plane portion of FPI is built over the object level API, and each package includes an additional 100,000+ lines of C code. With Agere's programming model, the entire data plane for a complete application is under 10,000 lines of code. Additionally, source code availability and built in extensibility features let customers create differentiated products by modifying the software to include proprietary features.

Among the data plane software sample applications, or application notes, Agere makes available for reference are:

- IPv4 over Ethernet processing
- ATM processing
- MPLS, IPv4/v6, PPP, and ATM AAL interworking applications
- Cell-based policing scripts and traffic management
- IPv4 with Diffserv
- Access Control List filtering
- IP/MPLS router
- Radio Network Controller

Control plane software

All Agere network processors come with the Run-Time Environment (RTE), a software package with an API that hides device level details from the application developer. About 175,000 lines of portable ANSI C code, the RTE package supports API to map to functional processing objects in the network processor such as classification tables, queue tables, and port tables. The RTE package includes API to:

- Create tables, and add, modify, and delete one or more table entries
- Enable the host application to add and delete routes dynamically as connections are established and terminated
- Support statistics gathering and OAM

The RTE also includes several debugging and diagnostics tools useful during board bring-up.

Hardware development platforms

Agere's complete hardware development platforms include network processor cards, various I/O interface cards, and host processor. Designers can perform complete system integration on these early platforms before the target hardware is available. Such systems are available for all target applications: multiservice edge and core, DSLAM access, and wireless access.

A flexible, extensible, chassis-based system called Festino supports the high performance (5G) network processors for multiservice edge and core applications. I/O cards supporting ATM, packet over SONET, and Ethernet interfaces are available. In addition to a control processor on each network processor card, a separate system controller card is available for control of the entire system.

Compact systems in the Hy-Span family support the 2G network processor family, with variations for wireline DSLAM and wireless access applications. DSL modem and GE I/O cards are available for the DSLAM platform. For wireless, I/O cards that support NxT1/E1, channelized and unchannelized OC-n, and Ethernet interfaces are available.

All hardware platforms come with Linux operating system and associated Board Support Packages.

- Cell-based policing scripts and traffic management
- Diffserv
- Access Control List filtering
- Other software packages
- IP/MPLS router
- Radio Network Controller

Agere product line overview

Agere historial competencies in telecom networks and systems make it particularly well positioned in the network processor market. The company used to be part of Lucent Technologies, and before that part of AT&T. As a result, Agere has a rich portfolio of telecom system and network intellectual property.

- NPUs
- Festino development environment
- Hardware development environments for Ethernet and ATM
- FPL & C for data plane programming
- Code editor, simulator, debugger
- Application environments
- MPLS, IPv4/v6, PPP, and ATM AAL interworking applications

Continued from page 49

and DRAM buses to connect to external memory typically used for packet and control plane information. I/O buses are also available. The IXP2350 uses a SPI-4 bus and also has a PCI Express interface for connection with a control plane processor.

Development environment

Software development begins with the Intel Software Development Kit. This is a GUI integrated development environment containing a code editor, compiler, debugger, and simulation environment for the IXP NPU family. The development environment runs on Windows or Linux systems and the debugger can be used with the simulator (called the Intel Transactor) or with target hardware.

The development kit also comes with packet generation and capture capabilities with templates for packet formats such as IPv4, IPv6, and VLAN-tagged packets. The simulator is set up to generate a set of output log files, one per port, containing each packet sent out that interface.

The Intel Transactor is a cycle-accurate simulator that enables the programmer to write data plane microcode, specify which code blocks go into which ME control store, and initialize the control plane environment through scripts executed by the Transactor. The Transactor software uses a *foreign model* interface that allows the user's control plane software to interface with the data plane software being run in the simulator. In this way, as complete a system simulation environment as is possible can be achieved.

The debugger within the Intel SDK environment allows for setting breakpoints, data watch points, probing memory, and stepping/tracing/modification of variables within the debug environment. The entire environment is controlled within the graphical context of the development environment.

Writing data plane code for the Intel NPU

The developer can write the data plane software in Intel microengine assembly or in C. For more demanding applications that require every available cycle, microengine assembly is typically the language of choice. For applications wanting to reuse code blocks and/or provide more complex processing, the C language is used. Intel claims that the performance of microcode written in C is within 90 percent of the performance of the same application written in microengine assembly.

Intel's latest advancement in programming data plane software comes in the form of IXP-C – a new tool that allows the developer to write data plane processing in C as one large program. Next, the tool partitions the code between the microengines for a given NPU type. So in this case, the developer inputs what NPU is being used, the line rate desired, and the C program doing the processing. The IXP-C tool then determines how many microengine contexts are required to execute the data plane code on the NPU being used at the line rate specified.

In addition to the languages, two basic programming models should be considered. One extreme is to write a single data plane program and run it across all contexts of all microengines for a highly parallelized programming option. Another choice is to write a series of data plane processing stages, then pipeline the processing between contexts and MEs, a highly pipelined approach to programming the NPU. In most cases, the data plane program mixes these two approaches to pipeline the processing where packet latency is least worrisome and to provide parallel contexts processing packets where more complex processing of each packet is performed.

The Intel SDK also defines the concept of a *microblock*, a self-contained algorithm implemented in microcode assembly or C and

Continued on page 58

Continued from page 57

controlled by a piece of C code that runs on the XScale within the NPU for exception processing and/or statistics. These microblocks can be chained together to implement a given data plane application. Intel provides microblocks that implement some of the more common processing like receive and transmit functions, IPv4/v6 routing and forwarding, and traffic management algorithms.

Control plane software

The Intel SDK provides data plane control and management and control plane interface software in the form of a *Resource Manager (RM)* and a set of core components. The Resource Manager is the main software component that manages the data plane environment. It manages memory allocation, packet buffer pools, queues between control and data plane, and message passing. Core components are C code blocks that run on the XScale processor and manage a data plane (microcode) portion of a microblock. This infrastructure allows the user to write programs to download the microengines, packet buffer, and control information memory. In addition, the infrastructure makes it possible to provide a message and packet interface between the microcode environment and the control plane.

Intel has been an active participant in the Network Processing Forum, and the Intel SDK includes a set of NPF-compliant APIs in the form of the Control Plane Platform Development Kit (CP-PDK). With these APIs any NPF-compliant control plane application can be interfaced to the Intel NPU data plane environment.

Available solutions and services

Since Intel's entry into network processor offerings, they have been developing internal solutions as well as an ecosystem of software, silicon, and hardware vendors that offer products and services relating to the Intel eXchange Architecture (IXA) product family. The Intel program is called the IXA Developer's Forum. For example, IP Infusion is a control plane protocol software company and a member of the IXA Developer's Forum. They have demonstrated their NPF-compliant routing stacks to be interoperable with the Intel SDK. Solutions range from protocol stack and operating systems to silicon and development platforms.

Intel also provides example applications in the SDK for IP and ATM applications, DiffServ, MPLS, routing, and forwarding applications. The user can use these applications to include the standards-based functionality within their application while using the tools, third-party offerings, and/or internal development to complete the required functionality for their product.

Intel services focus mainly on Tier I communications companies. So within the IXA Developer's Forum community, there are product and consulting companies that offer similar services for non-Tier I companies.

What's next

From a hardware perspective, Intel's process technology and manufacturing capabilities enable the company to scale its NPUs to smaller geometries. This allows Intel to implement a range of design options such as integrating more powerful control plane cores with the data plane, increasing the number of microengines, or incorporating additional specialized acceleration features.

Intel product line overview

Key competencies Intel brings to network processor development are their processes and a leadership position in processor technology. Intel has also made very significant investment in development of software tools and an ecosystem that provides a variety of hardware and software capabilities that can be used with their IntelInternet eXchange (IXP) product line. Over the past 10 years, Intel has also made acquisitions such as Dialogic in the communications products space and Trillium, a protocol stack software company, to gain additional competencies in communications software and product design.

- IXP network processor family
  - The Intel IXP4XX (core CPU with fixed-function coprocessors for edge applications < 1 Gbps)
  - IXP23xx (XScale CPU with 2-4 microengines for edge applications for edge applications < 4 Gbps)
  - IXP28xx (XScale CPU with 16 microengines and optional security coprocessing for up to 10 Gbps)
- Intel SDK
  - Code editing, compiler, debugger
  - Cycle accurate simulator with packet generation and output logging
  - Example data plane applications with control plane I/F software
  - Control Plane Platform Development Kit – NPF compliant APIs.
- Evaluation platforms
  - Multiple development platforms for AdvancedTCA and SBC for all NPUs in the family
  - Design and consulting services for Tier I customers
  - IXA Developer's Network (third-party products and services)
  - Additional platforms
  - ADI CompactPCI platforms commercial product
  - RadiSys AdvancedTCA commercial board products
  - IP Fabrics PCI Express IXP2350 development platform
  - Silicon classification and coprocessing
  - IDT classification coprocessor
- Software products, tools, services
  - IP Infusion routing and control plane stacks
  - Teja Technologies Teja NP software development tools
  - IP Fabrics Packet Processing Language software platform
  - Silicon & Software Systems consulting services

Other notable developments

If you think only the NPU vendors are advancing software capabilities for network processing, think again. Software companies have emerged with software products that support network processors. For example, IP Fabrics, Inc. (see page 38 in this issue) and Teja Technologies are two companies with software products available today targeted at making data plane programming for network processors even easier.

Teja Technologies (www.teja.com) has a product called Teja NP that provides a software platform for multiprocessor System-on-Chip architectures. It's a graphical environment where the programmer can specify the sequence of packet processing for a given application, set some parameters that tell the builder metrics that help it decide where to partition the functionality, then build the data and control plane software for a given chip. The tool outputs microcode for the data plane processing specified and C code for the control and management plane parts of the application. The programmer then loads the outputs onto a target system and performs test and validation. Applications like IPv4 forwarding, IPv6/IPv4 interworking, and TCP/IP acceleration have been developed by Teja and provide accelerated time to market for customers needing these functions in their product. Teja has products for the Intel NPU family and the Broadcom SiByte processor that shows signs that a common software platform may indeed be viable across multiple NPU and multiprocessor environments.

It will be interesting to see if the evolution of network processors will also prompt the emergence of standardized languages and run-time foundations like the CPU did for operating systems and eventually, Java.

Powerful silicon, available software

I hope the latest updates from these silicon and software vendors has opened your eyes as to how far programming network processors has come over the past few years. Dramatic advancements in tools, programming environments, services, and available software products mean no longer having to be afraid when adopting NPUs for your products. Among the silicon and software now available for NPUs, there is a solution that fits just about every network application possible.

For further information, contact Curt by e-mail at: cschwaderer@opensystems-publishing.com

AVIONICS

SBS Technologies, Inc.

Website: www.sbs.com
Model: AVC-cPCI-3008 **RSC No:** 20563
Rugged, yet highly flexible COTS computing platform • Suited for mission applications requiring real-time high-capacity aircraft data acquisition and processing • Functions as a flight-ready 1553 bus data monitoring, collection, and analysis system • Can be fitted into new aircraft, or retrofitted inexpensively into older aircraft • Can store data in crash survivable memory, and in removable memory for routine maintenance • At the heart of the AVC-cPCI-3008 lies the SBS PASS 3200 1553 data bus analyzer application • An onboard robust data analysis tool, PASS 3200 provides time-stamped real-time data in engineering form that can be analyzed onboard, or downloaded to a PC-based ground station for immediate analysis

BACKPLANE

ELMA Bustronic Corporation

Website: www.elmabustronic.com
Model: ATCA Backplane **RSC No:** 22176
A 14-slot AdvancedTCA backplane • Compliant to PICMG R3.0 draft 0.65 specification • Topology: Dual star base and fabric interface • Two hub slots and 12 node slots • 12-layer PCB • 3.2 mm • FR4 material • Power studs for redundant -48 VDC

RSC 22176

BACKPLANE ACCESSORIES

Kontron

Website: www.kontron.com
Model: CP-ASM7-P47-RIO **RSC No:** 22708
Supplied with up to four Positronic47 power supplies, 8 or 16 slot CompactPCI backplane, and a 1U hot swap fan tray concept, this 7U enclosure is designed to meet the requirements associated with rugged environments and varying ambiance • EMI protected (IEEE1011.10, IEC 1587 V1) • Fully tested • High-quality machined components • Anodized/chromed aluminum • Built-in 8 slot backplane with IPMI and hot-swap support • Optional device bay for 3.5-inch and 5.25-inch devices • With all slots Rear I/O • Three cooling fans with sensored outputs

BACKPLANE: FULL MESH

ELMA Bustronic Corporation

Website: www.elmabustronic.com
Model: 2-slot ATCA Backplane **RSC No:** 21240
A two-slot AdvancedTCA backplane • Point-to-point mesh topology • 10-layer stripline design with optimized routing through signal integrity studies • Includes connectors for direct plugging of the IPM Sentry shelf manager from Elma, or wired connections to other shelf managers • Switchless 15x replicated mesh topology with point-to-point links between slots • Standard full mesh can also be implemented on the backplane • Compliant to the PICMG 3.0 Specification

BRIDGE: PROCESSOR-TO-PCI

Tundra Semiconductor Corporation

Website: www.tundra.com
Model: Tsi108 **RSC No:** 20221
Host bridge for Freescale and PowerPC processors • Suitable for printer controller board, processor PMC, base station, and media gateways applications • DDR2 memory support for up to 50% memory power savings compared to DDR • Supports PCI-X, Gigabit Ethernet, and Flash memory • 2.5 W typical power consumption • Integrated Clock Generator with optional Spread Spectrum capability • Low latency nonblocking internal switch fabric

RSC 20221

CARRIER BOARD: AMC

CorEdge Networks, Inc.

Website: www.coredgenetworks.com
Model: CEN-RC3 **RSC No:** 21698
An AdvancedTCA/AdvancedMC reprogrammable carrier • Three general purpose single-width/extended full-height AdvancedMC bays • Full hot swap support • Radial IPMI system management • Radial JTAG distribution and test engine • Radial synchronization clocks 1, 2, and 3 • 200 W power maximum • 0 °C to +55 °C operating temperature range • Carrier IPMC Renesas Hitachi H8 • Base channel programmable multi-protocol fabric switch • Optional Fat Pipe programmable multi-protocol fabric switch

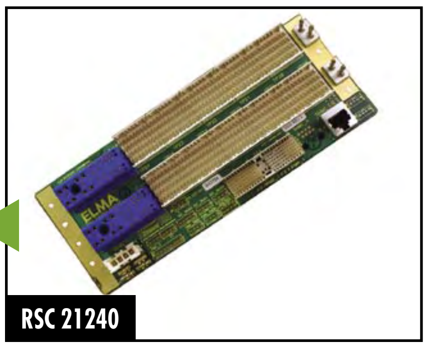
RSC 21240

For further information, enter the product's RSC# at www.compactpci-systems.com/rsc

CARRIER BOARD: PMC

Critical I/O

Website: www.criticalio.com
Model: XGE4022-PMC **RSC No:** 22882
Dual channel, hardware based TCP/IP Offload Engine (TOE) • 442 MBps sustained data rate • Standard sockets interface • Minimal host processor overhead • TCP, IP, UDP, iSCSI, and RDMA protocol offload • Zero copy/direct data placement • Software library and driver support • Hardware Built in Test (BiT) • Commercial or rugged (conduction cooled) • Copper or fiber optic

TEK Microsystems, Incorporated

Website: www.tekmicro.com
Model: JazzFiber PMC **RSC No:** 20299
Protocol-agnostic fiber optic PMC I/O module • Optimized for both streaming I/O and signal processing applications • 64-bit 133 MHz PCI-X host interface • Supports ANSI/VITA 17.1 Serial FPDP • Four fiber optic transceivers operating at up to 3.125 Gbps each, which can be configured as four independent interfaces or combined into a single 4x link • Supports both PCI and PCI-X protocols at up to 133 MHz • Provides over 1 GBps of throughput to support four 247 MBps Serial FPDP streams in a single module • Full GB of onboard DDR SDRAM allows deep buffering of streaming

data at the full 1 GBps data rate • Fully compatible with a wide range of VME and CompactPCI carrier cards including Tekmicro's JazzStream and PowerRACE I/O controllers • Fully integrated into Tekmicro's JazzStore data recording and playback systems architecture


RSC 20299

CARRIER BOARD: PMC (INTELLIGENT)

N.A.T. GmbH

Website: www.nateurope.com
Model: NVPT1001 **RSC No:** 20300
Intelligent carrier for PCI Mezzanine Cards (PMC) • Motorola MPC8280 PowerQUICC II at 333 MHz or 450 MHz • PCI interface and compliance: Intel 21555, 64 bit/ 66 MHz • Hot swappable, PCI Rev. 2.2 • PMC: PLX PCI 6150, 32 bit/33 MHz, PCI Rev. 2.2, P1386.1/Draft 2.4a • H.110 Bus (PICMG 2.5) Agere T8110 H.110 controller, on CompactPCI J4 connector • IP-Backplane (PICMG 2.16) • AMD AM79C973 10/100Base-T Ethernet controller, on CompactPCI J3 connector • 32-256 MB SDRAM (PC-100, 64 bit) installed in a SODIMM slot • 16/32 MB Flash PROM (16-bit), separate boot Flash (512 KB) • Two 32-bit/33 MHz PCI Rev. 2.2, P1386.1/Draft 2.4a compliant PMC slots, signals of PMC I/O connector (P14) routed to FPGA for open interconnect • Serial I/O: two RS-232 compatible serial ports on the front panel • Intel 82544 Gigabit Ethernet controller on standard RJ-45 or optical connector on the front panel • Routing Pool: Altera 1K100 FPGA, central switching resource • Operating System support and firmware: OK-1, VxWorks, Linux, TCP/IP, SNMP, SS7, ISDN, and others • Power consumption: 2 A at 3.3 V and 1 A at 5 V (prel.) • Temperature (operating): 0 °C to +60 °C with forced air cooling • Temperature (storage): -40 °C to +85 °C • Relative humidity: 10% to 90% (noncondensing)

CONNECTOR: OTHER

Hypertronics

Website: www.hypertronics.com
Model: D02 Circ. Plastic Con. **RSC No:** 21160
Suitable for critical medical applications including electronic catheters, patient monitors, and wearable therapeutic devices • Available with a combination of 1 coaxial contact (rated at frequencies to 18 GHz) and up to 9 signal contacts (rated at 2.5 A) • Alternatively, the coax could be replaced with a 25 A power contact, in the same housing • Also available with from 7 to 25 signal contacts, three 8 A power contacts, and in combinations of signal and power with two 8-amp power contacts and up to seven 2.5 A signals in the same connector • Incorporate the Hypertac hyperboloid contact system, which provides high cycle life, low insertion force, low contact resistance, and immunity to shock and vibration


RSC 21160

DATACOM: ATM

SBS Technologies, Inc.

Website: www.sbs.com
Model: Telum 1001-012M/S **RSC No:** 21483
AMC.1 and PCIe Rev 1.0 compliant • 1 port full duplex OC-12 interface with optional Automatic Protection Switching (APS) • 8 MB local memory; supports up to 16,000 VCCs • Hot Swap compliant • Intelligent Platform Management Interface (IPMI) • Carrier Grade Linux, VxWorks

DATACOM: WAN

SBS Technologies, Inc.

Website: www.sbs.com
Model: Telum 624/628-TEJ **RSC No:** 21482
4 ports T1/E1/J1, AMC.0 compliant extended height • 128 DSO channels • AMC Hot Swap compliant • i-TDM-to-IP converter • Onboard CSU/DSU • Carrier Grade Linux, VxWorks

DEVELOPMENT TOOLS

Agilent Technologies

Website: www.agilent.com
Model: E2943A **RSC No:** 21495
Active probe board • Plugs into the AdvancedTCA backplane to emulate real ASI designs or topologies • Offers analyzer and exerciser support for the AMC standard • Suitable for bring-up and debug, validation, and compliance testing of AdvancedTCA form factor-based ASI boards and backplanes

ENCLOSURE + CARD RACK + POWER SUPPLY

Asis-Pro

Website: www.asis-pro.com
Model: 5U ATCA AC shelf **RSC No:** 22854
Compact design allowing 5 slots + 5 standard 8U RTM blades • Integral AC power input: two redundant power supplies, cooled independently, 1200 W each • Cable free chassis design with integral cable tray for optimum cable management • Easily front accessed, field removable, and replaceable air filter tray, fan trays, and shelf management board • Fully AdvancedTCA PICMG 3.0 compliant • Meets NEBS requirements for cooling, noise, and shielding • Dedicated slots for two Asis Management carrier modules • Front connectors include serial ports, LAN ports, and Telco alarm ports • Redundant "no tools" field replaceable AC power supplies • Comprehensive power filtering and overcurrent protection eliminate need for rack-level Power Distributed Units (PDUs) • Rear access for serial and LAN connection • Full mesh backplane • 8 fans (N+1 redundancy) for cooling the high density/high performance computing environment • Provides cooling for 200 W + 15 W (RTM) per slot • Provides a platform for building carrier grade telecom applications that support 5-nines (99.999) availability by decreasing Mean Time Between Failure (MTBF) and requiring much less than five minutes Mean Time to Replacement (MTTR)

ELMA Electronic

Website: www.elma.com
Model: Conduction Cooled ATR **RSC No:** 22660
Commonly used in applications where the enclosure is sealed or in high altitudes • Conduction-cooled ATRs have a machined card cage

NEW PRODUCTS

accommodating wedge locks that transfer heat • Compliant with ARINC 404A specifications • Optimized for cooling via thermal simulation studies • -40 °C to +70 °C operating temperature range • Optimized for low weight, suitable for avionics applications where weight is a critical issue • Designed for 3U cards, the chassis has backplane options in CompactPCI, VME, or customized versions • Versions for 6U cards available • All Elma conduction-cooled ATRs meet the IEEE 1101.2 specification • MIL 38999-type power connector with an integrated line filter • Power supply and line filter combination are optimized to meet MIL-STD 461E • Rugged aluminum dip-brazed construction is designed to meet MIL-STD 810F for shock and vibration • Elma also offers a line of standard convection-cooled ATRs that come in 1 and 1-ATR tall long formats per ARINC 404A specifications

Hybricon Corp.

Website: www.hybricon.com
Model: RME1021M **RSC No:** 20836
10 U rack-mount high power enclosures • Ruggedized construction and a compact stackable design for vertically mounted cards • Designed to cool extremely dense CPU and DSP boards • Dual DC impellers powered by a dedicated power supply output provide 16.1 CFM per slot, sufficient cooling for 125 watts per slot • Meet MIL-STD-461 EMI radiated and conducted emissions and susceptibility standards as well as MIL-STD-810, MIL-S-901, and MIL-STD-167 environmental requirements • Front and rear cover panels with honeycomb inlet and exhaust and provisions for shielded connectors for I/O • Available with 21-slot CompactPCI, VME64x, VME, or VXS backplanes • Up to 2400 W embedded power • Closed loop fan speed control has programmable settings for minimum and maximum speeds and temperature break points • Fully compatible with IEEE 1101.10/11 packaging standards • Patented CoolSlot air deflecting card guides improve airflow • Mounting for two internal hard drives • Thermal simulation of enclosure (thermal report available) • Low smoke wire • Honeycomb EMI air filters on front and rear panels • Front panel LCD display with monitoring of voltages, fans, and temperature with RS-232, Ethernet, and SNPM support

Inova Computers, Inc.

Website: www.inova-computers.com
Model: Rail-Man **RSC No:** 22849
Just 18TE wide, Inova's Rail-Man is a suitable choice for DIN-Rail mounted CompactPCI • Complete with a neat and purposefully designed, scratch and corrosion resistant enclosure, 70 W AC/DC or 60 W DC/DC power supply, single slot CompactPCI backplane, and high-performance Intel Pentium III or Celeron M CPUs • The Rail-Man combines the simple DIN-Rail assembly with its efficient space and cost saving approach with open CompactPCI standards and associated modularity and vendor independence

Kontron

Website: www.kontron.com
Model: XL8500 **RSC No:** 22711
Suitable for software or hardware developers wanting to use AdvancedTCA CPU blades together with a hub blade • Can also be used to exploit AdvancedTCA blades with the full mesh backplane in a small form factor chassis of only 5U • With the redundant shelf manager and the full rear I/O possibilities, no AdvancedTCA feature will suffer any disadvantage due to the smaller form factor • The system comprises a 5U/84HP rack with either a 750 W wide range AC power supply or a 48 VDC input and enough airflow to cool all blades

FABRICS: SWITCHED FABRIC

ZNYX Networks

Website: www.znyx.com
Model: ZX6000/7000 **RSC No:** 20403
24-port (ZX6000) or 48-port (ZX7000) Gigabit Ethernet nonblocking switch fabric (upgradable) • 16-slot chassis support for base interface • PICMG 3.1 Option 1 & 2 support • Four (ZX6000) or two (ZX7000) PTMC Option 5 slots • Configurable front and rear panel egress • OpenArchitect 3.3 switch management with Linux kernel • Wire-speed Layer 2-7 packet classification • Three ports out-of-band management • CompactFlash • Real-time clock • USB • Enhanced protocol package includes OSPF, RIP, EGP, BGP, IP Multicast, VRRP, and more

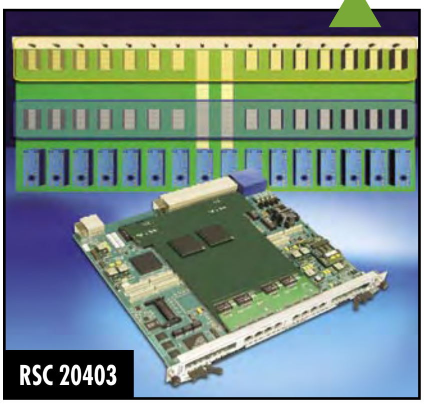
RSC 20403

GATEWAYS

AudioCodes Ltd

Website: www.audiocodes.com
Model: Mediant 1000 **RSC No:** 20318
Converged wireline and wireless VoIP media gateway • AudioCodes' VoIPerfect technology • Scalable • E1/T1/J1 and analog (FXS/FXO) interfaces • Cost efficient for low density gateways • Life-line fallback to PSTN in case of power failure • PSTN fallback for assured connectivity

I/O: FPGA

Acromag, Inc.

Website: www.acromag.com
Model: PMC-DX502/2002 **RSC No:** 22925
32 bi directional RS-422 differential I/O lines • Front or rear I/O connection • Customizable FPGA with 500,000 or 2,000,000 gates (Xilinx Virtex-II XC2V500 or XC2V2000 • FPGA code loads from PCI bus or Flash memory • 256K x 36-bit SRAM memory • Supports dual DMA channel data transfer to CPU • Supports both 5 V and 3.3 V signalling • Extended temperature option (-40 °C to +85 °C)

NEW PRODUCTS

MIL-STD-1553

Data Device Corp. (DDC)

Website: www.ddc-web.com
Model: BU-65570T/72T **RSC No:** 22838
CompactPCI dual tester/simulator card • One or two 1553 buses • PXI certified • PICMG certified • Variable amplitude transceivers • 128 K-words of shared RAM per bus • Simultaneous emulation of BC, 31 RTs, and MT for each bus • IRIG-B Interface • DMA data transfers via PCI master operation • Software configuration of bus coupling and termination • PCI plug-and-play compatible • 32-bit time tag with 1 µsec resolution • Replay of previously recorded bus traffic via menu and run-time library

POWER SUPPLY

Performance Technologies

Website: www.pt.com
Model: ZT 6301 **RSC No:** 21711
Highly reliable modular package designed for AC power input systems • Highly dense, redundant supply is suitable for telecommunications, industrial automation, and a variety of embedded computer applications utilizing the CompactPCI 3U form factor • Hot-swap N+1 load sharing • DIN input/output connector • Protection features: overvoltage, short circuit; all outputs,overtemperature protection • Status LEDs: Fault, Input OK • Status output signals: (DEG#), (FAL#) • Main output remote sense (+3.3 V, +5 V) • Built-in EMI filter • Eurorack-compatible module • Front drawer-style handle

Picor Corporation

Website: www.picorpower.com
Model: QPI-5 Active EMI Filter **RSC No:** 20211
14 A rating • 40 VDC (maximum input) • 100 VDC surge 100 ms • >60 dB CM attenuation at 250 kHz • >80 dB DM attenuation at 250 kHz • 707 VDC hipot hold off to shield plane • -40 °C to +100 °C PCB temperature • Efficiency >99 percent at full load • 1.0" x 1.0" x 0.2" System-in-Package (SiP) • SMT Land Grid Array (LGA)

PROCESSOR: MULTIPLE MPU

CWCEC: Processing

Website: www.cwcembedded.com
Model: Compact Champ AV IV **RSC No:** 20219
PICMG 2.16 multiprocessor card, CompactPCI form factor (no CompactPCI interface) • Four PowerPC7447A/7448 (AltiVec enhanced) CPUs up to 1.5 GHz • 64 KB L1 and 1 MB L2 internal caches operating at core processor speed • 48 GFLOPs peak computational power • Up to 512 MB DDR-250 SDRAM per processor • (2 GB total) • PowerPC 8540 at 500-800 MHz for control functions • 128 or 256 Mbytes DDR SDRAM • QuadFlow architecture with 3.2 GBps peak onboard throughput • Five Gigabit Ethernet (GbE) ports, one per processor • Support for two 64-bit, 100 MHz PCI-X mezzanine modules (PMC-X) Two EIA-232 serial ports • Support for StarFabric PMC modules with differential routing to backplane • VxWorks BSP • IXLibs-AV optimized AltiVec DSP function library • Verari VSI/Pro Image VSIPL DSP library • Air-cooled Level 0 ruggedization

PROCESSOR: PENTIUM M

RadiSys Corporation

Website: www.radisys.com
Model: Procelerant CE **RSC No:** 20302
PICMG COM Express 1.0 compliant • Basic form factor (95 mm x 125 mm) • Intel Pentium M or Celeron M • Intel 915GM chipset • Up to 1 GB DDR2 SDRAM • Intel Integrated Graphics • Dual independent displays • Analog VGA • LVDS • Dual SDVO • 10/100/1000Base-T Ethernet • High-definition audio • 8 USB 2.0 ports • 4 SATA ports • 1 ATA100 port • 3x1 PCI Express links • PCI 32/33 compliant 32-bit/33 MHz PCI Bus • ACPI compliant power management • Phoenix BIOS • 8 GPIO lines

RADAR/SONAR

CWCEC: Real-Time Video and Graphics

Website: www.cwcembedded.com
Model: ProWare PMC-440 **RSC No:** 20558
A new rugged FPGA PMC card for the capture, processing, and output of data derived from high-speed sensors such as electro-optical/infrared (EO/IR) and radar systems • Onboard FPGA delivers up to 20 billion operations/sec performance for FFT and digital filter DSP functions • Can be configured with either of two versions of the Xilinx Virtex-II Pro FPGA: the XC2VP20 (9,280 logic slices/88 18x18 multipliers) or the XC2VP40 (19,392 logic slices/192 18x18 multipliers) • 64-bit, 66 MHz PCI interface with support for PCI-X • Legacy bus interface for C40 comm link, SHARC link, and other proprietary DSP interfaces • Available in both air-cooled and conduction-cooled versions • I/O options include four front-panel RocketIO transceivers with aggregate peak throughput up to 1.6 GBps and four RocketIO transceivers accessible via the module's Pn4 connector that support throughput rates up to 1 GBps • 30 LVDS/LVTTL ports via the front panel and 48 LVDS/LVTTL ports via the Pn4 connector • Front panel I/O options available on air-cooled versions • 256 MB of 32-bit DDR266 SDRAM • Temperature sensor • Current sensor • Four indicator LEDs • User-programmable with CWCEC's ProWare Design Kit

ROUTERS/SWITCHES

Kontron

Website: www.kontron.com
Model: CP932 **RSC No:** 22697
Gigabit Ethernet switch with integrated NIC • Supports 10/100/1000Base-T Ethernet on five ports with RJ-45 connectors, each with two status LEDs • Optional sixth port connected through a Gigabit Ethernet NIC to the Master CPU of the CompactPCI backplane segment • Broadcom BCM5388 Gigabit Ethernet switch

RadiSys

Website: www.radisys.com
Model: ATCA-2100 **RSC No:** 19739
Single slot AdvancedTCA PICMG 3.0/3.1 compliant module • Highly integrated fabric module: GbE base interface, Fibre Channel fabric, shelf manager, two PMC sites • 28-port GbE PICMG 3.0 Base fabric switch: 14 GbE links to user slots, 1 GbE link to redundant switch blade, 8 GbE links for external connections (uplinks), 4 GbE links to onboard 2 PMC sites, 1 GbE link to off-board shelf manager • Wire-speed L2 and L3 switching: support for 4K IEEE 802.1Q VLANs, packet classification/filtering, link aggregation, QoS, CLI, and SNMP management • Fibre Channel switch: 14 FC links to user slots,17-port 2 Gbps FC PICMG 3.1 fabric switch, 1 FC link local on board with PCI interface, 2 FC links for external connections (uplinks),supports FC-AL-2 protocol, hub or switch • Integrated Shelf Management: PICMG 3.0 compliant ShMC/ShMgr, HPI 1.1 compliant programmatic interface, CLI, and SNMP management, 2 PMC sites can be used for System Mgmt PrPMC and network timing PMC or intelligent WAN uplinks

SHELF AND MECHANICAL COMPONENTS

SBS Technologies, Inc.

Website: www.sbs.com
Model: AMC-7S **RSC No:** 21490
A 2U 19-inch AdvancedMC standalone chassis that can be controlled by a processor AdvancedMC plugged directly into its backplane • Passive backplane system will allow a computer to be built exclusively from AdvancedMC modules

SOUTHCO

Website: www.southco.com
Model: AdvancedTCA Faceplate **RSC No:** 21498
Ejector handles • Safety ground pins • Captive retention screws • Meet all requirements of the AdvancedTCA PICMG 3.0 R2.0 standards for faceplate mounting and fastening • Ejector handle for the standard AdvancedTCA application features required 500N (112 lb/f.) load capacity and a secondary catch • Catch firmly locks into the faceplate, where it operates a microswitch that permits hot swapping • Ejector handle capable of withstanding more than 250 inject/eject cycles as specified in the AdvancedTCA standard • Hardened captive screws in the required M3 thread size and #2 Philips headstyle are available in both press-in and flare-in mounting styles • Bullet-tip design and an extremely smooth finish • All metallic parts of Southco AdvancedTCA compliance mounting hardware are free of hexavalent chrome and fully RoHS compliant

SOFTWARE-DEFINED RADIO

Pentek, Inc.

Website: www.pentek.com
Model: 7131 Multiband DR **RSC No:** 20561
A 16-channel multiband digital receiver with dual 14-bit, 105 MHz A/Ds • Four quad multiband digital receiver chips driven by the samples from

both A/Ds • Equipped with either XC2V1000 or XC2V3000 FPGA devices from the Xilinx Virtex-II family, with logic densities of one or three million gates, respectively • Operating temperature range of -20 °C to +65 °C • Qualified for 20 *g* shock and 2 *g* sine vibration • Optional conformal coating extends operating range to 100 percent relative humidity and protects the board from environmental contaminants

RSC 20561

TEST SYSTEMS

Advint LLC

Website: www.advint.com
Model: Puma Test Sys. **RSC No: 22840**
Ruggedized, portable, and field deployable • Small size and weight facilitates compact packaging in rugged enclosures to support mobile operations with a reduced logistics footprint • 18-slot PXI and 4-slot SCXI chassis containing wide variety of digital, analog, electro-optics, RF, and switch instrumentation for functional testing of modern electronics and avionics • System provides power distribution, Unit Under Test (UUT) power supplies, system cooling, and UUT isolation • LabWindows/CVI, LabVIEW, and Visual Basic test programming environments under the control of TestStand test management software • IVI instrument drivers • Graphical User Interfaces (GUIs) are provided for direct control of system resources; allows the operator to manually control the test system to facilitate debug, troubleshooting, and repair activities • XML and SQL test results available • MacPanel's Series L2000 Interface Test Adapters (ITAs) are used for a reliable, high density interface to UUTs • UPS power pack allows for operation from generators or any international power standard • Fully automated system self-test accomplished via signal wrap-around of stimulus instruments to measurement assets to ensure comprehensive tester operational assurance • Field/field calibration via portable NIST traceable standards • Field calibration capability allows the system's test instruments to be calibrated in place at the user's site within the normal system configuration • Open architecture allows for easy and cost effective technology insertion and customization to meet other test requirement needs

THERMAL MANAGEMENT

SOUTHCO

Website: www.southco.com
Model: 5T Heat Sink Screws **RSC No:** 22700
A consistent set preload up to a maximum of 5 pounds per foot (22 N) per screw to maximize surface contact and thermal conductivity throughout the range of chip expansion • Can allow for a smaller heat sink, conserving space and costs • Designed to mount at the outer edges of the heat sink • Cause minimal or no interference with heat sink fins and do not compromise their function • Hardened-steel captive screws with stainless steel springs • Standard #2 Phillips/slot headstyle • Free of hexavalent chrome • In salt spray corrosion testing have exceeded 144 hours of treatment • Multiple screw thread sizes and lengths are available

VIDEO PROCESSOR

Mango DSP, Ltd.

Website: www.mangodsp.com
Model: Mango Harrier **RSC No:** 22851
15 DSP TMS320C6415 @ 600 MHz • Up to 72 GIGA instructions per cycle • 256 MB SDRAM @ 133 MHz per DSP • 32/66 MHz Primary Side PCI Bus • Up to 64/66 MHz Secondary Side PCI Bus • 133 MHz system clock • 5 Altera Stratix EP1S25F1020 FPGA • 2 Camera Link IN 68 bit @ 85 MHz • 2 Camera Link OUT 68 bit @ 85 MHz • DSP to FPGA: EMIFA bus 64 bit @ 133 MHz, EMIFB bus 16 bit @ 133 MHz • B/W FPGA: 2 buses 32 bit @ 85 MHz DDR • Intel 64-bit 66 MHz 21154 PCI bridge • Extended 6U 233 x 260 mm • +3.3 VDC, 5 Amp typical • +5 VDC, 4 Amp typical

VOICE: VOIP

AudioCodes Ltd

Website: www.audiocodes.com
Model: MediaPack 112/114/118 **RSC No:** 20315
Analog VoIP media gateways • Fax, voice, and modem support • Toll quality voice compression • MWI, long haul, metering, CID, and outdoor protection • Integration with leading PBXs, IP-PBXs, softswitches, H.323 gatekeepers, and SiP servers • Spans a range of 2 to 8 FXS analog ports • Selectable, multiple LBR coders per channel • T.38 compliant • Echo canceller, jitter buffer, VAD, and CNG • Complies with MGCP, H.323 (V4), and SiP control protocols • Comprehensive support for supplementary services • Web management for easy configuration and installation • EMS for comprehensive management operations (FCAPS)

NMS Communications Corporation

Website: www.nmscommunications.com
Model: CG6500C1 **RSC No:** 20320
High density PSTN, VoIP, and voice processing platform • Channelized T1/E1 interfaces for up to 480 DS0 streams • 8 or 16 digital PSTN trunks, software selectable between T1 and E1 (120 ohm) • A-law or µ-law PCM encoding • Dual 10/100Base-T interfaces either on rear panel or connected to PICMG 2.16 backplane • Onboard RTP/RTCP processing by fast PowerPC • Voice processing DSP resource pool of up to 9,600 MIPS • Supports Fusion VoIP gateway and media server development environments • Low latency media streaming • Single-slot 6U CompactPCI solution • Supports all popular signaling protocols • Full speed H.110 bus with 4,096 timeslots to support interoperability with other boards in open architecture, high capacity systems • Uses MVIP-95 switching model • Supports full PICMG hot swap standard • Feature rich Natural Access • Software development kits support Microsoft

For further information, enter the product's RSC# at www.compactpci-systems.com/rsc

CompactPCI and AdvancedTCA Systems

ADVERTISER INFORMATION

Page/RSC#	Advertiser/Product description
59	ACT/Technico – NAS RAIDStor Solution
30	Adax – Adax Signaling Gateway
13	ADLINK Technology – Pentium M based SBCs
2901	AdvancedTCA Summit – AdvancedTCA Summit December 6-8, 2005
2101	Advanet/Radic – cPCI & VME boards
6	AEI – Server Class Fast Ethernet & Gigabit Ethernet Cards
51	AP Labs – FS-5973 3U cPCI Conduction Cooled ATR
5	Arrow – COTS by Motorola
26	AudioCodes – TP-12610 ATCA VoIP Media Gateway Board
35	Conec Corporation – CompactPCI/AdvancedTCA Connectors
62	Digital Power – Hot-Swap/Redundant Power Supplies
15	Diversified Technology – 12U AdvancedTCA/CompactPCI System
40	ELMA Bustronic – ATCA Backplanes
19	ELMA Electronic – Elma Handles and Panels
14	Embedded Planet – 64-bit IBM 970FX
11	Excalibur Systems – Dragon/Mace/Lance/Camelot
2	GE Fanuc Automation – CPCI-7808/CPCI-7506
47	General Standards – Data Acquisition Boards
3301	Interface Concept – PowerPC-based SBCs
9	Kontron – AdvancedTCA/AMC modular solutions
25	MEN Micro – VME and CompactPCI FPGA Solutions
36	Micrel – High-Voltage Hot-Swap Power Controllers
17	Motorola – Motorola VMEbus Technologies
45	National Instruments – High-Density Switch Systems
23	One Stop Systems – CompactPCI/PCI Express/Max Express
39	Performance Technologies – Advanced Managed Platforms
41	Performance Technologies – Integrated Slot Components
43	Performance Technologies – Software and Middleware
4601	Pinnacle Data Systems – AdvancedTCA Solutions
53	Positronic Industries – Zone 1 Connectors
4602	Positronic Industries – Connector Solutions
18	Red Rock Technologies – Mass Storage Modules
44	SBE – SBE Interface Cards
67	SBS Technologies – PSL09 PrPMC/CR4
24	Schroff – AdvancedMC Mechanical Hardware Kits
34	Simon Industries – Conduction-Cooled Heat Frames
7	SMA – 3U CompactPCI CPU Options
61	Sundance – SMT3000/SMT300/SMT7008
63	Sundance – SMT791/SMT787/SMT795
65	Sundance – SMT6050/Diamond RTOS
31	Technobox – PMCs & PIMs
2902	Technobox – Async I/O PMCs
3302	Technobox – Adapters and Tools for PMC
3	VadaTech – Customized CompactPCI Solutions
68	VMETRO – PCI Express Protocol and Link Analyzer
2102	Winchester Electronics – Zone 1 ATCA Power Connector
27	XTech – Customized CompactPCI/AdvancedTCA Front Panels

www.compactpci-systems.com/rsc

OpenSystems Publishing

Advertising/Business Office

30233 Jefferson Avenue
St. Clair Shores, MI 48082
Tel: 586-415-6500 ■ Fax: 586-415-4882

Vice President Marketing & Sales
Patrick Hopper
phopper@opensystems-publishing.com

Business Manager
Karen Layman

Communications Group

Patrick Hopper
Vice President Marketing & Sales
phopper@opensystems-publishing.com

Christine Long
Print and Online Marketing Specialist
clong@opensystems-publishing.com

Embedded and Test & Analysis Group

Dennis Doyle
Senior Account Manager
ddoyle@opensystems-publishing.com

Doug Cordier
Account Manager
dcordier@opensystems-publishing.com

Military & Aerospace Group

Tom Varcie
Account Manager
tvarcie@opensystems-publishing.com

Andrea Stabile
Advertising/Marketing Coordinator
astabile@opensystems-publishing.com

International Sales

Stefan Baginski
European Bureau Chief
sbaginski@opensystems-publishing.com

Reprints and PDFs

Call the sales office: 586-415-6500