Embedded COMPUTING DESIGN

Connecting Silicon, Software, and Strategies for Intelligent Systems

Available on the App Store

FOLLOW US:

DECEMBER 2012
VOLUME 10 • NUMBER 8

www.embedded-computing.com

INTERCONNECTS HIT HIGH SPEEDS

What's Hot in 2013 in Embedded?

Annapolis Micro Systems

The FPGA Systems Performance Leader

High Performance Signal and Data Processing in Scalable FPGA Computing Fabric

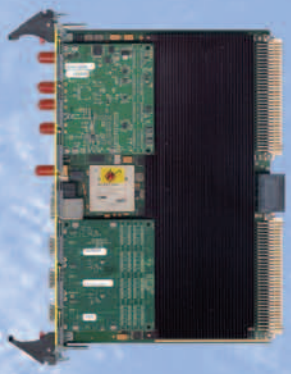

GEOINT, Ground Stations, SDR, Radar, Sigint, COMINT, ELINT, DSP, Network Analysis, Encryption, Image Processing, Pattern Matching, Oil & Gas Exploration, Financial Algorithms, Genomic Algorithms

Direct Seamless Connections with no Data Reduction
Between External Sensors and FPGAs
Between FPGAs and Processors over IB or 10GE
Between FPGAs and Standard Output Modules
Between FPGAs and Storage Arrays

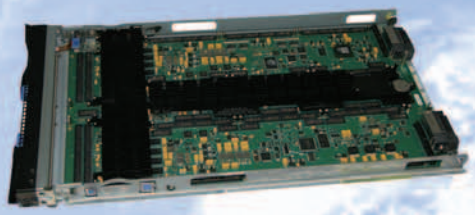

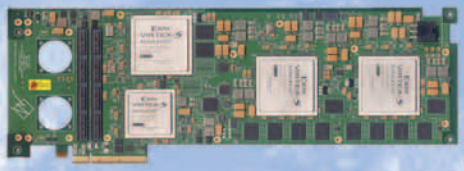

Annapolis Micro Systems

Ultimate Modularity
From 1 to 8 Virtex 4, 5 or 6 FPGA/Memory Modules
Input/Output Modules Include:
Quad 130 MSPS thru Quad 550 MSPS A/D
1.5 GSps thru 5.0 GSps A/D, Quad 600 MSps D/A,
Dual 1.5 GSps thru 4.0 GSps D/A
Infiniband, 10G, 40G or 100G Ethernet or SFPDP

VME/VXS/VPX, IBM Blade, PCI-X/PCI Express, PMC/XMC, MicroTCA

No Other FPGA Board Vendor Streams This Volume of Data Real Time Straight Into the Heart of the Processing Elements and Then Straight Back Out Again

190 Admiral Cochrane Drive, Suite 130, Annapolis, Maryland USA 21401
wfinfo@annapmicro.com USA (410) 841-2514 www.annapmicro.com

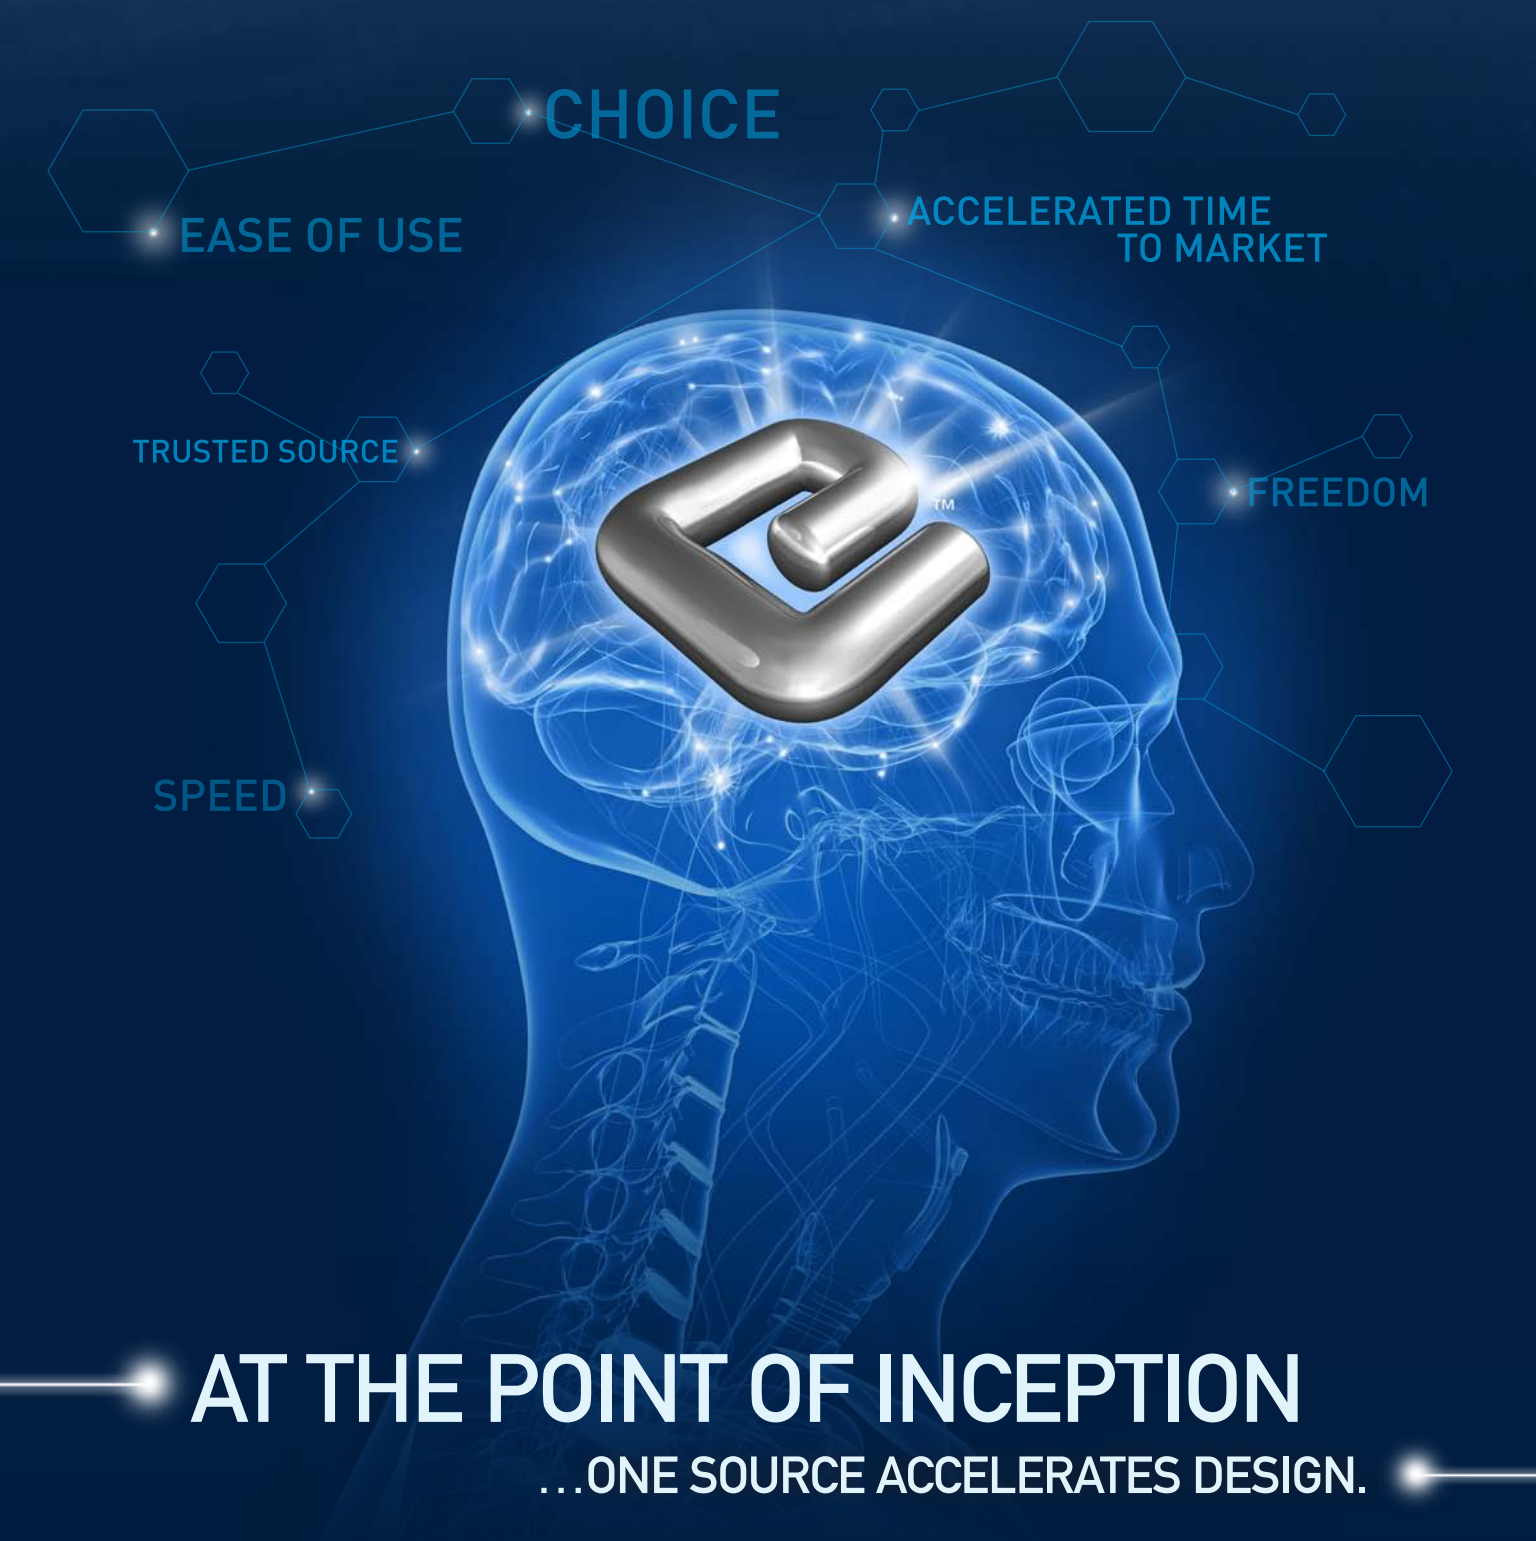

AT THE POINT OF INCEPTION

...ONE SOURCE ACCELERATES DESIGN.

From inception forward, the Embedded Software Store increases the ease, speed, and simplicity of getting your designs to market. Developed jointly by Avnet Electronics Marketing and ARM, ESS provides an embedded marketplace, bringing together developers and software vendors to collaborate and provide solutions for customers.

embedded software store™

Your Gateway to Accelerated Design™

www.EmbeddedSoftwareStore.com

Brought to you by

AVNET electronics marketing ARM

Embedded COMPUTING DESIGN
www.embedded-computing.com
December 2012 | Volume 10 • Number 8

ON THE COVER

High-speed interconnect technology has become a necessary component to optimize the performance of high-bandwidth embedded systems. Read updates on some of the leading interconnect standards in our Silicon section, and don't miss our special advertising section exploring the question of "What's Hot in 2013 in Embedded?"

ISSN: Print 1542-6408, Online: 1542-6459

E-CASTS

Building Intelligent Therapeutic Devices
December 11, 12 p.m. EST
Presented by: Wind River

Improving Healthcare through Standards and Interoperability
December 12, 11 a.m. EST
Presented by: IBM, Freescale Semiconductor, Wind River

PERFORMANCE ~~OR~~ *and* SWaP

You *Can* Do More with Less.

XPand6000 Series Small Form Factor (SFF) Systems from X-ES support the highest-performance Freescale™ QorIQ™ and 3rd generation Intel® Core™ i7 processors in the smallest SFF Systems available, at less than 72 in^3 and 3.5 lbs. With COM Express processor modules, PMC/XMC I/O, and SSD storage, application-ready SFF Systems enable you to focus on your area of expertise. **Call us today to learn more.**

X-ES

Extreme Engineering Solutions

608.833.1155 www.xes-inc.com

Tight on space and high in performance. That's Extreme.

Advertiser Information

Page	Advertiser/Ad title
2	**Annapolis Micro Systems, Inc.** – High performance signal and data processing
12	**EMAC, Inc.** – Atom D525 touch PC
44	**AMD** – Innovate now!
28	**ATP Electronics** – ATP industrial grade DRAM and flash products
3	**Avnet** – At the point of inception
19	**Elma Electronic Inc.** – Conveniently designed with a small footprint
38	**Elma Electronic Inc.** – Smart, reliable platforms for embedded computing
39	**Embedded Software Store** – A marketplace for software developers
41	**embedded world** – It's a smarter world
34	**Flexera** – The solution company business model
27	**Interface Concept** – Switches and IP routers
35	**Lilee Systems** – Seamless M2M connectivity
33	**Micron Technology** – Semiconductor memories for the automotive electronics industry
36	**Microsemi Corporation** – FPGA technology to deliver new capabilities for systems requiring security, low power, reliability, and integration
28	**PC/104 Embedded Consortium** – Maintaining and distributing PC/104 specifications
37	**SafeNet** – Monetizing embedded software
25	**Supermicro** – Embedded building blocks
17	**Technologic Systems** – Industrial controllers
31	**VIA Technologies, Inc.** – VIA system on module solution
40	**VIA Technologies, Inc.** – VIA Technologies embraces both x86 and ARM
43	**Wind River Systems, Inc.** – Innovation is 100 percent reliability 65 million miles from home. Again.
22-23	**WinSystems, Inc.** – A 1.66 GHz Intel Atom industrial SBC of EPIC proportions
5	**X-ES** – Performance and SWaP

Reader Resources

Subscribe to magazines and E-letters:
www.opensystemsmedia.com/subscriptions

E-casts: ecast.opensystemsmedia.com

TechChannels: tech.opensystemsmedia.com

White Papers:
whitepapers.opensystemsmedia.com

Videos: video.opensystemsmedia.com

Get news via RSS:
www.embedded-computing.com/news/rss.xml

Embedded COMPUTING DESIGN

OpenSystems media®

ECD Editorial/Production Staff

Warren Webb, Editorial Director
wwebb@opensystemsmedia.com

Jennifer Hesse, Managing Editor
jhesse@opensystemsmedia.com

Brandon Lewis, Associate Editor
blewis@opensystemsmedia.com

Steph Sweet, Creative Director
ssweet@opensystemsmedia.com

Sales Group

Tom Varcie
Senior Account Manager
tvarcie@opensystemsmedia.com

Rebecca Barker
Strategic Account Manager
rbarker@opensystemsmedia.com

Eric Henry, Strategic Account Manager
ehenry@opensystemsmedia.com

Ann Jesse. Strategic Account Manager
ajesse@opensystemsmedia.com

Christine Long
Vice President, Online Business
clong@opensystemsmedia.com

International Sales

Elvi Lee, Account Manager – Asia
elvi@aceforum.com.tw
Gerry Rhoades-Brown
Account Manager – Europe
gerry.rhoadesbrown@husonmedia.com

Regional Sales Managers

Barbara Quinlan, Southwest
bquinlan@opensystemsmedia.com

Denis Seger, Southern California
dseger@opensystemsmedia.com

Sydele Starr, Northern California
sstarr@opensystemsmedia.com

Ron Taylor
East Coast/Mid Atlantic
rtaylor@opensystemsmedia.com

Reprints and PDFs

republish@opensystemsmedia.com

OpenSystems Media Editorial/Production Staff

Warren Webb, Editorial Director
Embedded Computing Design
wwebb@opensystemsmedia.com

Jennifer Hesse, Managing Editor
Embedded Computing Design
jhesse@opensystemsmedia.com

Jim Turley, Embedded Moderator
Embedded Computing Design
jturley@opensystemsmedia.com

Mike Demler, Editorial Director
DSP-FPGA.com and EDA Digest
mdemler@opensystemsmedia.com

John McHale, Editorial Director
Military Embedded Systems
jmchale@opensystemsmedia.com

Joe Pavlat, Editorial Director
CompactPCI, AdvancedTCA, & MicroTCA Systems
jpavlat@opensystemsmedia.com

Jerry Gipper, Editorial Director
VME and Critical Systems
jgipper@opensystemsmedia.com

Sharon Hess, Managing Editor
VITA Technologies
sharon_hess@opensystemsmedia.com

Monique DeVoe
Assistant Managing Editor
PC/104 and Small Form Factors
DSP-FPGA.com
mdevoe@opensystemsmedia.com

Brandon Lewis, Associate Editor
CompactPCI, AdvancedTCA, & MicroTCA Systems
blewis@opensystemsmedia.com

Curt Schwaderer
Technology Editor

Steph Sweet
Creative Director

David Diomede
Art Director

Joann Toth
Senior Designer

Konrad Witte, Senior Web Developer

Matt Jones, Web Developer

Editorial/Business Office

Patrick Hopper, Publisher
Tel: 586-415-6500
phopper@opensystemsmedia.com

Subscriptions Updates

Karen Layman, Business Manager
www.opensystemsmedia.com/subscriptions
Tel: 586-415-6500 ■ Fax: 586-415-4882
30233 Jefferson, St. Clair Shores, MI 48082

Rosemary Kristoff, President
rkristoff@opensystemsmedia.com

Wayne Kristoff, CTO

16626 E. Avenue of the Fountains, Ste. 201
Fountain Hills, AZ 85268
Tel: 480-967-5581 ■ Fax: 480-837-6466

Tracking Trends in Embedded Technology

By Warren Webb

@warrenwebb wwebb@opensystemsmedia.com

Bandwidth demands redefine interconnect technology

As high-performance embedded devices stretch the limits of technology, subsystem interconnection strategies must be constantly updated to match the growing data transfer requirements. Medical instrumentation, military systems, communications installations, and industrial automation are just a few areas where soaring bandwidth, increased processing requirements, and escalating application complexity continue to expand the interconnection technology envelope.

Embedded designers are using the latest high-speed interconnects to not only boost data rates, but also dynamically optimize performance, bypass failed subsystems, and coexist with legacy components. Although several different architectures still vie for universal acceptance, there is little doubt that high-speed interconnect technology has become a necessary element of most high-bandwidth embedded systems.

In the Silicon section of this issue of *Embedded Computing Design*, industry experts bring you up-to-date on several high-speed interconnect standards that could be a part of your next embedded design. Sam Fuller, founder and executive director of the RapidIO Trade Association, explains how the RapidIO architecture has been optimized for interconnect applications where high reliability, low latency, and deterministic operation are required. In addition, Rick Wietfeldt, chair of the Mobile Industry Processor Interface (MIPI) Alliance Technical Steering Group, and Al Yanes, president and chairman of the Peripheral Component Interconnect Special Interest Group (PCI-SIG), answer questions on new techniques and methods to scale PCI Express performance for mobile applications. Continuing the high-speed connectivity topic, Craig Wiley, chairman of the Video Electronics Standards Association (VESA), discusses the latest updates to the Embedded DisplayPort (eDP) standard.

This month's Software section highlights how one of the hottest Operating System (OS) platforms – Android – can be deployed in one of the hottest market segments in embedded design – In-Vehicle Infotainment (IVI). As automotive operational and entertainment systems transition from a few traditional in-dash components to individual multimedia and information clusters for the driver and each passenger, the corresponding embedded software burden skyrockets. Designers must integrate GPS navigation, cellular communications, and automotive controls with multiple audio and video systems operating simultaneously. In the near future, vehicles will be permanently connected to the Internet, giving developers access to real-time cloud data services, telematics, and video/audio streaming. Andrew Patterson, business development director for Mentor Graphics Embedded Software Division, presents the advantages of developing these evolving vehicle-based applications using the Android OS.

> "Subsystem interconnection strategies must be constantly updated to match the growing data transfer requirements."

Along with interconnect and software selections, embedded designers must also choose a processor architecture that maximizes performance while reducing power and cost. In the Strategies section, technology experts examine the benefits of the growing trend toward multiple processing elements. In that vein, AMD's Frank Altschuler, senior product manager, and Jonathan Gallmeier, senior member of the technical staff, explore multicore image processing using a mix of CPU and Graphics Processing Unit (GPU) elements. In a Q&A, Mark Brewer, president and CEO of Typesafe, presents the advantages of general-purpose programming languages like Java and Scala for embedded development, along with the implications for multicore platforms. Concluding the Strategies section, David Sonnier, technical fellow in system architecture for the Networking Solutions Group of LSI Corporation, evaluates multicore System-on-Chip (SoC) architectures for the next-generation communications infrastructure.

The hardware and software topics delivered in this issue of *Embedded Computing Design* are on the cutting edge of today's embedded technology, as are the products and services showcased in our special advertising section seeking to answer the question of "What's Hot in 2013 in Embedded?" With a new year upon us, we can expect continued development of performance-enhancing tools, plus a few new innovations that could completely change your design direction. Stay tuned.

If you have ideas for future articles and coverage that would help in your design efforts, please let us know. We are always interested in contributed technical articles or videos that would be of interest to other embedded designers. Contributed articles are a great way to expose your technology or expertise to the embedded community, so if you have an idea, send along an e-mail with a short abstract.

Q & A

PCI Express over M-PHY: Proven technologies scale I/O performance to meet mobile needs

Q&A with Rick Wietfeldt, Chair, MIPI Alliance Technical Steering Group, and Al Yanes, President and Chairman, PCI-SIG

The MIPI Alliance and PCI-SIG recently announced an agreement to deliver an adaptation of the PCI Express (PCIe) architecture over the MIPI M-PHY physical layer technology to provide scalable I/O functionality for mobile devices. Rick and Al describe how this collaboration combines the interoperability of the PCIe standard with the power efficiency of the M-PHY physical layer to create a low-power, high-performance technology for the mobile market.

ECD: ***What is the MIPI Alliance, and how does it help embedded device developers?***

WIETFELDT: MIPI Alliance is an organization with global membership focused on developing interface standards for mobile devices. These interface standards include high-speed gigabit-per-second interfaces for performance applications such as wireless modems, cameras, and displays, as well as low-speed and multipoint interfaces for other applications such as audio and control. Key MIPI objectives include achieving low-power operation for battery-constrained mobile devices and minimizing the diversity of interfaces that speed OEM time to market. Embedded device developers benefit from a well-known and optimized set of industry-standard interfaces.

ECD: ***What is the PCI-SIG, and how does it help embedded device developers?***

YANES: The PCI-SIG is the industry organization chartered to develop and manage the PCI Express (PCIe) standard. With a global membership community, the PCI-SIG's principal charter is to evolve the PCIe architecture to meet the current and emerging needs of its members and the industry. A board of directors comprising nine people, each elected by the membership, leads the PCI-SIG. The PCI-SIG fulfills its charter by continuing to promote innovation and evolving through interoperability testing, technical support, seminars, and industry events.

As PCs become lighter and thinner and tablets and smartphones become more functional, consumers want seamless, always on/always connected functionality from their computing devices. To respond to these market expectations, device manufacturers need efficient, intelligent I/O technologies. The PCIe architecture satisfies all of these requirements, and with the adaptation to operate over the M-PHY specification, it can deliver consistent high performance in power-constrained platforms such as Ultrabooks, tablets, and smartphones. By delivering this technology, the PCI-SIG is meeting the emerging needs of its members and the industry.

ECD: ***What is M-PHY physical layer technology, where is it currently deployed, and what can we expect in the near future?***

WIETFELDT: The M-PHY physical layer was developed to satisfy the anticipated needs of mobile devices including high performance, low power, and high Electromagnetic Interference (EMI) immunity. High performance is achieved through three "gears" of operation including Gear 1 at 1.45 Gbps/lane, Gear 2 at 2.9 Gbps/lane, and Gear 3 at 5.8 Gbps/lane. Each gear includes two closely spaced frequencies of operation providing the capability to select operating frequency based on optimal mitigation of EMI issues in the mobile device, such as ensuring adequate radio receive sensitivity.

M-PHY provides key bandwidth scalability via support for varying lane widths – typically one, two, four, or eight lanes – and bandwidth asymmetry via differing numbers of lanes in the transmitter/receiver directions. M-PHY was also designed to accommodate a variety of overlay protocols for different applications, such as USB 3.0 or PCIe for chip-to-chip connections, Camera Serial Interface to camera sensors, and Display Serial Interface to (internal) displays. The flexible support for differing protocols provides a key foundational element for broad use of MIPI M-PHY technology in a number of different applications within a mobile device.

M-PHY has been broadly accepted within key industry organizations developing interfaces for specialized applications. JEDEC has adopted M-PHY for its Universal Flash Storage (UFS) interface as an evolution to the embedded Multi-Media Card (eMMC) interface for embedded storage. USB-IF has selected M-PHY for its SuperSpeed Inter-Chip (SSIC) interface for transporting the USB 3.0 protocol. And most recently, the PCI-SIG has chosen M-PHY for use as a low-power PHY to carry the PCIe protocol.

ECD: ***What was the impetus behind this collaboration to develop an adaptation of PCIe over M-PHY, and how will this technology be implemented in mobile applications?***

YANES: PCI-SIG is interested in adapting its ubiquitous PCIe architecture to operate on the MIPI M-PHY to provide a robust I/O architecture for use by its members interested in ultra-low-power solutions. MIPI has succeeded in defining a best-in-class low-power, scalable PHY (M-PHY). The layered architecture of PCIe allows for an easy way to replace the physical layer with a new PHY. This design uses the M-PHY layer combined with the PCIe data link and transaction layer.

WIETFELDT: The collaboration to develop the PCIe on M-PHY solution stemmed from the need in increasingly powerful mobile computing devices to merge the ubiquitous PCIe protocol on the emerging M-PHY physical layer. Similar to the joint MIPI/USB-IF development of the SSIC specification, known informally as "USB 3.0 over M-PHY", the joint MIPI/PCI-SIG development of the "PCIe over M-PHY" technology (see Figure 1) delivers high performance and broad ecosystem support of applications in the PCIe market, as does SSIC in the USB market and application space.

ECD: ***How can PCIe technology reduce product development schedules and cost?***

YANES: As a broadly adopted technology standard, PCIe benefits from several decades of innovations with universal support in all major operating systems, a robust device discovery and configuration mechanism, and comprehensive power management capabilities that few, if any, other I/O technologies can match. PCIe technology has a flexible, layered protocol that enables innovations to occur at each layer of the architecture independent of the other layers.

PCIe architecture meets the software support and testability needed for today's power-constrained platforms; no new development or investment in the architecture is needed. In this way, power-efficient PHY technologies such as MIPI M-PHY can be integrated with the familiar and highly functional PCIe protocol stack to deliver best-in-class and highly scalable I/O performance in mobile devices such as Ultrabook, tablet, and smartphone devices.

ECD: ***Can this new technology be used in traditional embedded applications other than mobile?***

YANES: PCIe is currently deployed in a number of embedded platforms that run the gamut from enterprise to consumer applications and form factors. With the adaptation to support the M-PHY, this ubiquitous I/O technology has been rendered suitable for adoption in Ultrabook, tablet, and smartphone platforms almost overnight. In addition, it is expected that this technology will be adopted in future storage applications in various topologies due to the anticipated migration of storage attach points from SATA to PCIe technology. As a power-efficient, general-purpose, load-store I/O architecture, component and device designers can implement this technology in other I/O expansion usage models of their choice.

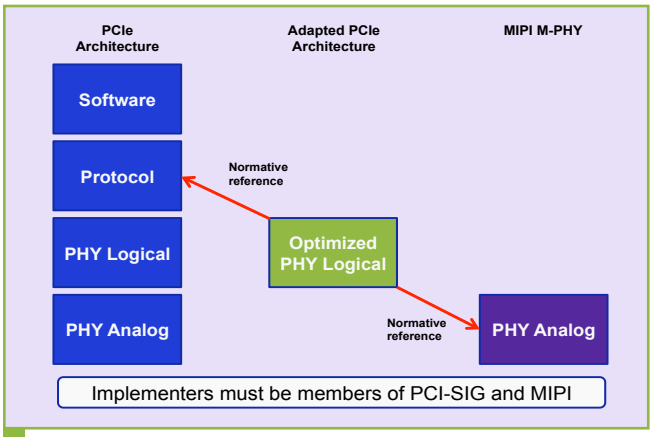

Figure 1 | The proposed PCIe over MIPI M-PHY technology leverages the layered architecture of PCIe and the extensible protocol stack of M-PHY to deliver scalable I/O functionality for mobile devices.

WIETFELDT: While the MIPI Alliance develops interface standards primarily for mobile devices such as smartphones, tablets, and other consumer electronics, MIPI interfaces can be used in products not employing wireless communications. The primary difference is nontechnical and relates to MIPI's licensing model, which involves Reasonable and Non-Discriminatory – Zero Cost (RANDZ) terms for mobile devices and Reasonable and Non-Discriminatory (RAND) terms for other devices. This licensing model is very common in the high technology industry.

ECD: ***Does the MIPI Alliance or PCI-SIG offer any software tools, libraries, and educational materials to help embedded designers get started with these standards?***

YANES: PCI-SIG offers a range of technical support for members, including test specifications and procedures, compliance workshops around the world, and annual developer conferences to help members design and implement PCIe standards.

WIETFELDT: MIPI Alliance offers an online learning center, which includes general interface tutorials, specification tutorials, white papers, webinars, and other presentations/videos. MIPI Alliance also provides a suite of testing resources, which includes test suite documents, test tools, and multivendor interoperability events. ECD

Rick Wietfeldt, *PhD, is chair of the Technical Steering Group at MIPI Alliance and senior director of technology at Qualcomm.*

Al Yanes *is president and chairman of the PCI-SIG and a Distinguished Engineer for IBM in the Systems & Technology Division.*

RapidIO: Optimized for low-latency processor connectivity

By Sam Fuller

Interconnect architectures reflect the problems they are designed to solve. Focusing too heavily on raw bandwidth metrics often misses more important attributes such as reliability, determinism, and latency, which continue to be critical factors in embedded computing design. Comparing interconnects based on these criteria demonstrates the advantages of using RapidIO for connecting today's high-performance processors and the tightly coupled computing systems of the future.

As Moore's Law has continued to drive the performance and integration of processors, the need for higher-speed interconnects has continued to grow as well. Today's interconnects commonly sport speeds ranging from 10 Gbps to 80 Gbps and have roadmaps leading to hundreds of gigabits per second.

In the race to faster and faster interconnect speeds, some topics are not often discussed, including the types of transactions supported, communications latency and overhead, and what sorts of topologies can be easily supported. Designers tend to think of all interconnects being created equal and having a figure of merit based solely on peak bandwidth.

Reality is quite different. Much as there are different forms of processors optimized for general-purpose, signal-processing, graphics, and communications applications, interconnects are also designed and optimized for different connectivity problems. An interconnect typically solves the problems it was designed for and can be pressed into service to address other applications, but it will be less efficient in these applications.

RapidIO design objectives

It is instructive to look at RapidIO in this context. RapidIO was designed to serve as a low-latency processor interconnect for use in embedded systems requiring high reliability, low latency, and deterministic operation. It was designed to connect different types of processors from different manufacturers in a single system. Because of this, RapidIO has found widespread use in wireless infrastructure equipment, where there is a need to combine general-purpose, digital signal, FPGA, and communication processors together in a tightly coupled system with low latency and high reliability.

The usage model of RapidIO required providing support for memory-to-memory transactions, including atomic read-modify-write operations. To meet these requirements, RapidIO provides Remote Direct Memory Access (RDMA), messaging, and signaling constructs that can be implemented without software intervention. For example, in a RapidIO system, a processor can issue a load or store transaction, or an integrated DMA engine can transfer data between two memory locations. These operations are conducted across a RapidIO fabric, where their sources or destination addresses are located, and typically occur without any software intervention. As viewed by the processor, they are no different than common memory transactions.

RapidIO was also designed to support peer-to-peer transactions. It was assumed multiple host or master processors would be in the system and that those processors needed to communicate with each other through shared memory, interrupts, and messages. Multiple processors (up to 16K) can be configured in a RapidIO network, each with their own complete address space.

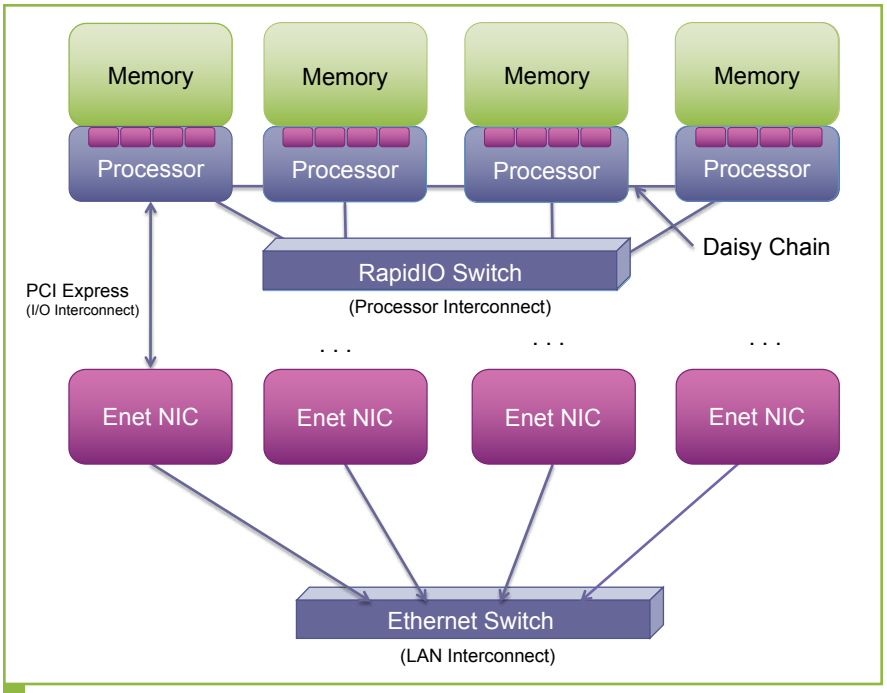

Figure 1 | RapidIO, PCI Express, and Ethernet provide different options for connecting processors, I/O, and systems.

RapidIO also provides a clean dividing line between the functionality of switches and endpoints. RapidIO switches only make switching decisions based on explicit source/destination address pairs and explicit priorities. This allows RapidIO endpoints to add new transaction types without requiring changes or enhancements to the switch devices.

Comparing interconnects

As more and more of the system is incorporated onto a single piece of silicon, PCI Express (PCIe) and Ethernet are being integrated within Systems-on-Chips (SoCs). This integration, however, has not changed the nature of the transactions provided by these interconnects (see Figure 1).

PCIe does not natively support peer-to-peer processor connectivity. Using PCIe for this sort of connectivity can be exceedingly complex because it was designed as a peripheral component interconnect (hence PCI). It was intended to connect peripheral devices, typically slave devices like I/O and graphics chips, to the main host processor. It was not designed as a processor interconnect, but rather as a serialized version of the PCI bus. Building a multiprocessor interconnect out of PCI requires a step beyond the base PCI specification to create new mechanisms that map address spaces and device identifiers among multiple host or root processors. To date, none of the proposed mechanisms to do this – Advanced Switching (AS), Non-Transparent Bridging (NTB), or Multi-Root – I/O Virtualization (MR-IOV) – have been commercially successful.

For systems where there is a clear single host device and other processors and accelerators operate as slave devices, PCIe is a good choice for connectivity. However, for connecting many processors together in more complex systems, PCIe has significant limitations in topology and support for peer-to-peer connectivity.

Many developers are looking to leverage Ethernet as a solution for connecting processors in systems. Ethernet has evolved significantly in the past 35 years. Similar to the increase in computer processing speeds, its peak bandwidth has grown steadily. Currently available Ethernet Network Interface Controller (NIC) cards can support 40 Gbps operating over four pairs of SERDES with 10 Gbps signaling. Such NIC cards contain significant processing on their own to be able to transmit and receive packets at these speeds.

The sending and receiving of Ethernet packets by a NIC is a long way from being a solution to tightly coupled interprocessor communications. The overhead associated with both PCIe and Ethernet transaction processing (both stacks must be traversed in a NIC), plus the associated SERDES functions and Ethernet media access protocol and switching add latency, complexity, and higher power dissipation as well as cost to systems where much more direct connection approaches could be utilized (see Table 1).

Using Ethernet as an integrated embedded processor interconnect requires significant transaction acceleration and enhancements to the Ethernet Media Access Controller (MAC) as well as the Ethernet switch devices themselves. Even with these enhancements, RDMA operations should be limited to large block transactions to amortize the overhead of using Ethernet.

Standards that have been deployed to solve this problem include the iWARP RDMA protocol from the Internet Engineering Task Force and RDMA over Converged Ethernet (RoCE). Both iWARP and RoCE are typically implemented through acceleration coprocessors. Despite this acceleration, RDMA

Bandwidth and per-port economics		
System requirement	**10 GbE Ethernet**	**RapidIO**
Switch-per-port performance raw data rate	10 Gbps	20 Gbps (4x)
End-to-end packet termination	>10 µs	~1-2 µs
Messaging performance	Software	Hardware – directly coupled to processor
Reliability	Drops packets on error or congestion	Integrated error detection and retry

Table 1 | A comparison of Ethernet and RapidIO shows the benefits of a more direct connection approach.

transactions must still be carefully managed to reduce communications overhead. The reason is that although Ethernet offers high bandwidth, especially in 10 GbE and 40 GbE implementations, it also has high transaction latencies that are typically measured in microseconds.

Current RapidIO applications

The RapidIO value proposition has been widely acknowledged in the embedded market for many years. This same value proposition can now be extended to more mainstream data-processing markets, which are evolving to demand many of the same system attributes that communications networks have long required.

One well-known application wherein RapidIO is utilized is the wireless base station. This application combines multiple forms of processing (DSP, communications, and control) that must be completed within a very short time frame. Communications between processing devices should be as quick and deterministic as possible to ensure that real-time constraints are achieved.

For example, in 4G Long-Term Evolution (LTE) wireless networks, frames are sent every 10 milliseconds. These frames, which contain data for multiple concurrent mobile sessions, are distributed across multiple subcarriers, which are supported by multiple DSP devices. Communications between the DSP and general-purpose processing devices must be deterministic and low latency to ensure that a new frame is ready for transmission every 10 milliseconds. At the same time, the receive path must support the data arriving from the mobile devices connected to the network. On top of this complexity, the system must track the location of the mobile device in real time and manage the device's signal power.

Another example of a RapidIO application is semiconductor wafer processing. Similar to the wireless infrastructure application, semiconductor wafer processing has real-time constraints, including a control loop of sensors, processing, and actuators. Leading-edge systems often have hundreds of sensors collecting information, with sensor data processed by tens to hundreds of processing nodes. The processing nodes generate commands that go to actuators and AC and DC motors to reposition the wafers and wafer imaging subsystems. This is all performed in a recurring control loop with a frequency of up to 100 kHz or 10 microseconds. Systems like this benefit from the lowest-latency communications possible between devices.

The future of high-performance computing

The introduction of virtualization, ARM-based servers, and highly integrated SoC devices is paving the way for the next stage in the evolution of high-performance computing. This evolution is moving toward more tightly coupled clusters of processors that represent processing farms built to host hundreds or thousands of virtual machines. These processor clusters will be composed of up to thousands of multicore SoC devices connected through high-performance, low-latency processor interconnects. The more efficient this interconnect is, the better the system's performance and economics will be.

Technologies like PCIe and 10 GbE are not going away anytime soon, but they will not be the foundation for these future tightly coupled computing systems. PCIe is not a fabric and can only support the connectivity of small numbers of processors and/or peripherals. It can simply serve as a bridge to a fabric gateway device. While 10 GbE can be used as a fabric, it has significant hardware and software protocol processing requirements. Its widely variable frame sizes (46 B to 9,000 B for jumbo frames) drive the need for fast processing logic to support several small packets and large memory buffers to support large packets in endpoints and switches, thus hiking up silicon cost. The use of PCIe or 10 GbE will either restrict the topologies and connectivity available or add cost and overhead to the system. These drawbacks create opportunities both for proprietary fabrics and for open solutions such as RapidIO in this interesting new market.

Implementing integrated server, storage, and networking systems presents an opportunity for OEMs to innovate. A key component of that innovation will be the internal system connectivity. RapidIO is a mature, well-proven technology with the attributes required for success in this market. As was the case for wireless infrastructure, where RapidIO went from early innovation to become the *de facto* base station interconnect standard, RapidIO's biggest challenge in server, storage, and high-performance computing will be to cross the chasm from today's innovators and early adopter markets to mass-market proliferation. ECD

***Sam Fuller** is a founder and executive director of the RapidIO Trade Association.*

RapidIO Trade Association
sam@rapidio.org
www.rapidio.org
Follow:

Embedded DisplayPort: Increased flexibility and power savings render greater display efficiency

By Craig Wiley

While mobile device display performance continues to increase, system chip processes geometries continue to shrink, resulting in a greater proportion of system power consumed by the display and its high-speed interface. The new Embedded DisplayPort (eDP) v1.4 standard offers several new features that maximize system power efficiency, further consolidate the display interface, and address a wide range of system profiles to satisfy the growing demand for power optimization in the embedded display system.

The Video Electronics Standards Association (VESA) first released the Embedded DisplayPort (eDP) standard in 2009 as an extension of the DisplayPort standard for use with embedded displays. VESA developed eDP to replace the aging Low-Voltage Differential Signaling (LVDS) standard, and today eDP is used widely in notebook computers and all-in-one systems. The primary benefits of eDP over LVDS include the reduction of signal wires due to its higher data rate, compatibility with submicron chip processes, decreased interference with wireless services, and its ability to accommodate new features.

Since its initial release the eDP standard has gone through a series of revisions adding new features not shared with DisplayPort, as eDP has evolved to target battery-operated embedded display systems. For example, in 2010 eDP version 1.2 was published, adding control of display and backlight features over the auxiliary channel. In 2011 eDP version 1.3 introduced Panel Self Refresh (PSR).

While eDP was originally designed for notebooks and all-in-one systems, it is becoming increasingly optimized for smaller form factor systems including tablet and PC smart phone applications. Released in December 2012, eDP version 1.4 adds new optional features developed to address this broader range of form factors and further reduce system power.

The importance of lowering display-related system power

Today, mobile devices are a major driving force in the electronics industry. Every year new mobile devices are introduced with increased processing capability, better displays, a smaller and lighter form factor, and extended battery life. Taking into account typical CPU idle time, an average display consumes about 75 percent of system power. While system chip power reduction is accomplished through shrinking semiconductor process geometries, display power reduction comes through improvements in backlight and LCD technologies, as well as new pixel structures.

However, the recent trend toward brighter, higher-resolution displays is driving up display power. The second-generation iPad had a 1024x768 display and a 25 watt-hour battery, while the latest iPod has a 2048x1536 display (a 400 percent pixel increase) and a 42.5 watt-hour

battery (a 70 percent power increase), both delivering a 10-hour battery life. The higher-resolution display requires additional pixel-driving circuitry and a higher data rate display interface, as well as faster Graphics Processing Unit (GPU) rendering and display image processing circuitry.

This display power challenge has led to many new architectural developments at the platform level. Reducing display power means longer battery life and less battery capacity requirement and therefore smaller, lighter, and less expensive systems. Rather than being treated as a simple rendering device, display deployment has become more integrated into the overall system design. The new eDP v1.4 brings many of these concepts together, as explained in the following discussion.

Panel Self Refresh (PSR)

Introduced in eDP v1.3, the PSR function provides a means to lower display-related system power by allowing portions of the GPU and display interface to enter a low-power state during a static display image condition (see Figure 1). When the system enters PSR mode, the remote frame buffer built into the LCD timing controller (Tcon) performs the routine display refresh task, offloading the GPU and display interface. The local frame buffer can then be updated later by the GPU with a new static image (such as when a flashing cursor is turned on or off), or the system can exit from PSR mode to display constantly changing images (such as when playing a video). Because most displays are refreshed 60 times per second, PSR allows the GPU circuits and display interface to remain in a low-power state for most of the system operating time, resulting in significant power savings.

PSR with selective frame update

The latest version of eDP enhances the PSR mechanisms by enabling updates to selected regions of the video frame in the Tcon frame buffer. With eDP v1.3, the entire frame must be updated every time any portion of the image changes. With eDP v1.4, only the new portion of the frame requires updating (see Figure 2). This lets the GPU display interface remain active for a shorter duration, further reducing system power.

Additional eDP v1.4 features that support PSR

Advanced link power management

The ability to manage the display interface is improved in eDP v1.4 by greatly reducing the wake-up time from the low-power state, which is important to maximize the efficiency of the PSR selective update process. Previous versions of eDP required more than 100 microseconds to wake up and retrain the data link. With eDP v1.4, if the new PSR mode is supported, wake-up from standby is a maximum 0.5 microseconds, and wake time from the lower-power sleep state is 20 microseconds.

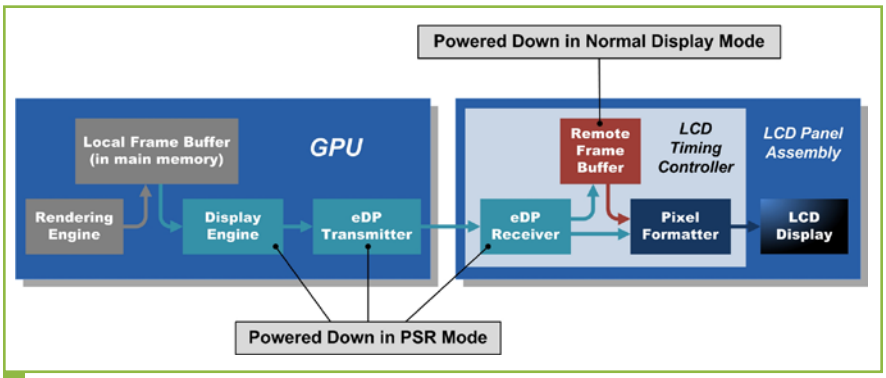

Figure 1 | The PSR function helps reduce power by allowing portions of the GPU and display interface to remain in a low-power state during static display image conditions.

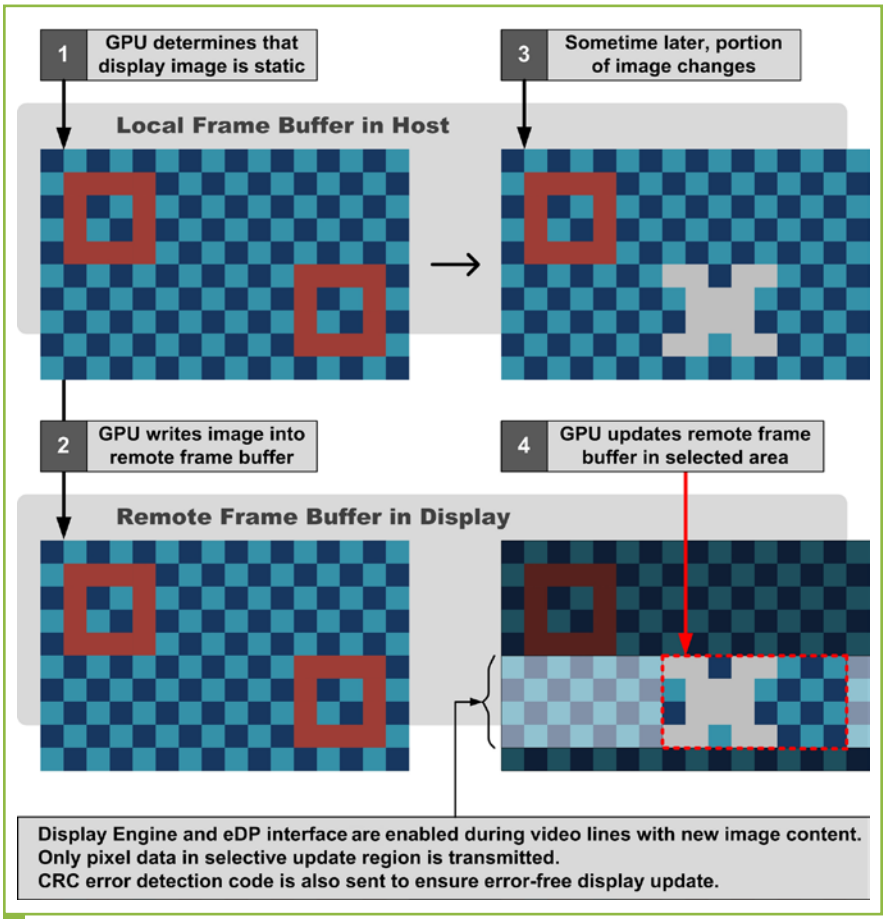

Figure 2 | The latest version of eDP enables a PSR selective update process that further decreases system power.

Video timing synchronization

In DisplayPort and earlier versions of eDP, display-to-GPU synchronization is performed over the DisplayPort main link, the high-speed data channel that carries video data using special control codes. In eDP v1.4, frame synchronization is also provided over the lower-speed auxiliary channel. This enables quick bursts of eDP main link operation to perform timely selective updates of the remote frame buffer.

Display stream compression

Display interfaces including DisplayPort and earlier versions of eDP send uncompressed pixel data across the display interface. Image compression is traditionally limited to media delivery to the system or for storage. In the interest of further power savings, eDP v1.4 introduces a display stream compression algorithm that reduces the data rate across the display interface. In contrast to typical video or image compression algorithms, the display stream compression algorithm is optimized for high data throughput, low latency, and low gate count and targets low-compression ratios in the 2x to 5x range, depending on the image type. By using a minimum 2x compression ratio configuration, for example, the display interconnect bandwidth can be cut in half, with typically no loss in image content. Display stream compression can also be used when updating a display's PSR frame buffer, providing further power savings. It can also be used to support display resolutions beyond the uncompressed main link capability.

Regional backlight control

When eDP v1.2 was released in 2010, it introduced the capability to control LCD backlight modulation frequency and brightness through the auxiliary channel, thus eliminating the need for an extra backlight control interface. In eDP v1.4, regional backlight control provides the option to independently set different portions of the display backlight region. This allows the GPU to actively darken selective portions of the display based on display contents, increasing power savings. Up to 15 backlight regions can be controlled with a single auxiliary transaction.

> "eDP v1.4 will help propel eDP into new applications and establish it as a universal embedded display interface."

Multitouch over auxiliary channel

Touch-sensitive displays are common in many smaller embedded display systems and will become more prevalent in PCs. The latest version of eDP adds the ability to transport multitouch data from the display to the host through the auxiliary channel. The multitouch data transport uses a framework that is compatible with the USB Human Interface Device specifications. By eliminating the dedicated USB interface commonly used for this purpose, both electrical connections and power are saved.

Reduced differential voltage swing

Prior to eDP v1.4, the standard utilized the same interface signal voltage amplitudes as DisplayPort. Today, eDP v1.4 reduces the main link's minimum differential voltage amplitude levels from 400 mV to 200 mV, which is suitable for the short transmission distances in small embedded form factors. This reduces interface drive power by as much as 75 percent, as power is proportional to the square of amplitude. Additional flexibility is also added to the link training amplitude step sizes, helping increase suitability for different transmission media including chip-on-glass.

Increased link rate flexibility

Another addition in eDP v1.4 is greater flexibility in the main link data rate. Previous versions of eDP limited link rate selection to 1.62 Gbps, 2.7 Gbps, and 5.4 Gbps per lane (main link data channel), which is the same as DisplayPort. The current version of eDP now includes seven standard rates, enabling further power efficiency.

With DisplayPort, increasing a 1080p 60 Hz display format from 24-bit color to 30-bit color would require changing a two-lane 2.7 Gbps configuration to either a four-lane 2.7 Gbps configuration or a two-lane 5.4 Gbps configuration. Ether case represents a 100 percent increase in interface bandwidth and power, with only a 25 percent increase in data content (additional dummy bits would be added to make up the difference). With eDP v1.4 in the same application, the rate could be bumped up from 2.7 Gbps to 3.24 Gbps, resulting in only a 20 percent link rate increase.

Customizable link rates are also supported, with a new set of registers defined for the display to declare this capability, as well as the custom link rate frequency. In addition to further power optimization, this increases application diversity as link rates can be adjusted based on particular system timing requirements. It also enables link rate adjustment to mitigate radio frequency interference with system wireless services.

System deployment with eDP v1.4

Offering increased flexibility and reduced system power, the latest release of eDP v1.4 will help propel eDP into new applications and establish it as a universal embedded display interface. Given the current cycle of eDP standard publication and adoption, it is expected that systems using some of the new features in eDP v1.4 will appear on the market in two to three years. Higher-resolution notebook PCs and tablets will likely be the first candidates to leverage this power-optimized interface. ECD

***Craig Wiley** is chairman of the Video Electronics Standards Association (VESA) and senior director of marketing for Parade Technologies.*

VESA
moderator@vesa.org
www.vesa.org
Follow:

Developing Android-driven applications for automotive infotainment

By Andrew Patterson

Google's Android mobile Operating System (OS) has quickly become the dominant platform for smartphone devices. Can Android be equally effective in the In-Vehicle Infotainment (IVI) sector? Despite Android's many functional advantages, software developers must take a careful look at Android's strengths and weaknesses before adopting the OS in modern IVI systems.

Developers today can make a strong case that Android is now the most successful portable OS of all time. In terms of recent smartphone sales, according to research firm IDC, Android devices represent a 68 percent share of the global market (quarter ended September 2012). This compares to a 17 percent market share held by Apple. Within 12 months it is expected that more than 1 billion Android devices will be in use – an achievable goal considering that nearly 700,000 new smartphones are activated every day. Easy access to software and development tools for Android means just about anyone from an individual engineer to the largest R&D department at a large corporation can get involved.

For corporations selling online services, it's almost a limitation not to have an appropriate Android smartphone app. During the past five years, user expectations have migrated from seeing a good website to seeing a good mobile website to having an Android or iPhone app available. According to AOL Tech, the download rate for Android apps in 2012 is 1.5 billion installs per month, with a total of nearly 20 billion installs to date.

Unfair to compare

It was inevitable that individuals who own both a smartphone and a vehicle equipped with an IVI system would compare and contrast the two. The functionality of a typical infotainment system has evolved in the past 10 years, restricted by the lengthy development cycles of automakers and their traditionally conservative approach to product development. Quality and reliability are paramount, as well as the overwhelming need to keep costs low to ensure the final product stays competitive.

At the recent Paris Motor Show, several automakers announced their latest models embodying the concept of the always connected automobile. One such system was Renault's Android-based R-Link infotainment system featuring built-in Android apps such as navigation, multimedia, and phone support via an online store for Renault-approved apps. Despite these and other IVI enhancements, any driver today can look at a contemporary smartphone and find much more functionality and personalization on that device compared to an IVI system. Carmakers are becoming increasingly desperate to incorporate this level of functionality and flexibility into a vehicle without compromising its safety and security. Using Android, there are several ways to accomplish this endeavor, each with its own set of advantages and disadvantages.

Bring Your Own Device (BYOD) to your vehicle

If the Android smartphone can be considered the ultimate infotainment device, then why not have it connected inside the car? This is the approach taken by the Car Connectivity Consortium, an industry alliance established to allow smartphone screens to be displayed on

the infotainment head unit. Several infotainment platform providers, including Mentor Graphics, offer this approach, whereby the head unit acts as a thin-client display, with apps running directly on the smartphone. Connectivity is provided through USB cable today, but Wi-Fi connections are emerging. Bluetooth 3.0 may also offer sufficient bandwidth for video streaming between smartphones and IVI systems.

The advantage of this approach is that the phone connectivity technology will not become obsolete as the car ages, which is an important factor given that the typical smartphone enjoys a higher refresh rate over its lifetime. The concept of random IVI software updates is seen as too risky for the more permanent car-based system; OEMs want to keep the process under their strict control. Looking 10 years ahead, this means that the infotainment system can still be current and relevant because its functionality is based on the smartphone at the time.

This approach also presents a cost advantage, as the permanently fixed infotainment system costs less for an OEM or Tier 1 developer to design and maintain. Another benefit is related to shared or rented vehicles – the smartphone immediately personalizes the vehicle to which it is tethered, without needing to learn a new user interface every time. One example showing the advantages of integrating a smartphone into the infotainment system is Android car mode, which turns an Android phone into a better driving companion by providing quick access to key applications such as GPS navigation, voice-activated commands, and a phone's contact list.

The main disadvantages of allowing smartphone screens to be displayed on the infotainment head unit are loss of control and marketability of the infotainment system as a car feature. High-end automobile manufacturers are now differentiating themselves through the sophistication of an infotainment system; they are not willing to pass that advantage over to phone makers. There is also the unknown security risk lurking in terms of the potential for someone to hack into the vehicle system via a smartphone.

Considerations for building in an Android OS

Many designed-in infotainment systems such as the Renault's R-Link build Android directly into the vehicle and pre-load a number of approved and tested apps. This offers a pre-built, tested, state-of-the-art infotainment system to a prospective car buyer. The idea here is that it's now possible for the vehicle owner to download additional Android apps from an online store managed by the manufacturer. The Android OS is kept isolated from other vehicle functions and apps are only offered from manufacturer-approved repositories to help protect the system from malware. However, as Generation Y Android users start to dominate the car buyer population, they will want the freedom to download their favorite apps and won't be happy with a pre-defined mix decided for them.

Looking at this from an OEM perspective, adopting Android as a base OS poses a few significant business risks. Some OEMs are nervous about the omnipresence of Google as the owner and licensor of the Android OS platform. Because Google manages the release schedule and content of Android, many automotive strategists are wary about Android changes affecting their product release cycle. What would happen if the license or terms of use suddenly changed?

The original Android OS was designed exclusively for mobile smartphones, and it must be modified to handle the wide variety of audio streams in the vehicle with signals coming from reversing sensors, radio, DVD player, navigation, phone, and external sources. The middleware in Android that covers audio stream routing has proved difficult to modify and re-test; the intended infotainment system has to link in at several points including the audio flinger (mixer providing a single output at a specified sample rate), underlying audio

hardware, and audio manager. Some developers are questioning why they should commit to the technology when it's possible to dock a smartphone into the vehicle.

Embedded Android architectures

Developers can choose from several possible approaches to implement Android into a vehicle. Some vehicle manufacturers use Android as the core OS for the infotainment system, deeming that it's secure and mature enough to fulfill this role. For designers who are not as bold and want to stick with Linux, Android can still be included as a guest OS in a "container" (see Figure 1). Using the Linux Container (LXC), resources can be assigned by the Linux host to the Android guest, which includes memory available for apps, access privileges, services available, and interaction with other domains. The container is intended to be a secure environment, so users can potentially download entrusted apps into this area.

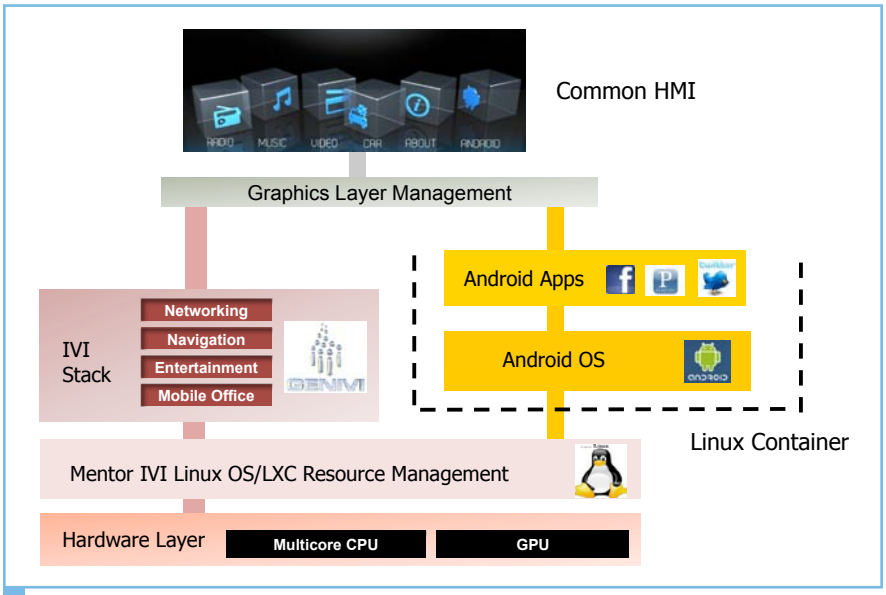

Figure 1 | When running Android in a Linux Container, privileges and permissions can be tightly controlled.

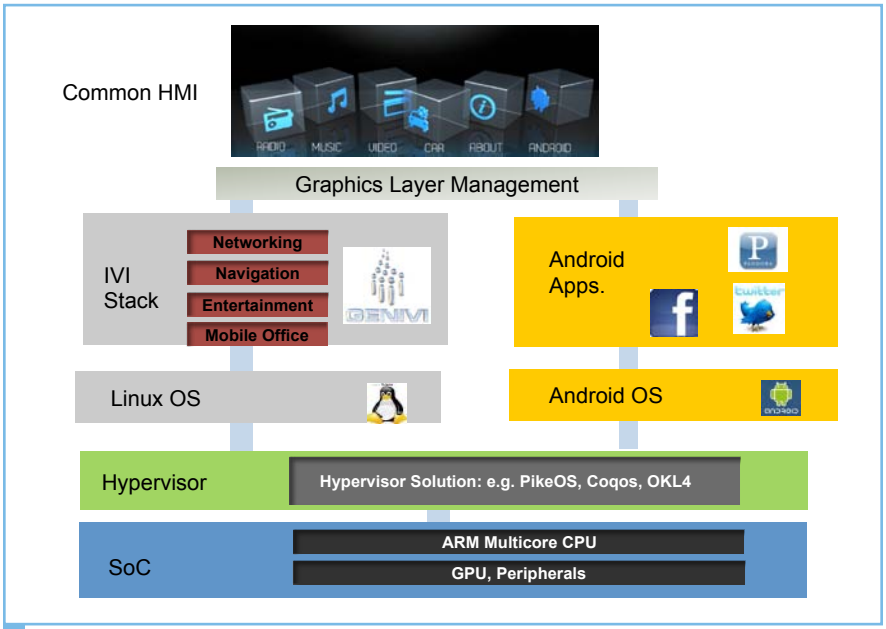

Figure 2 | Android and Linux can run simultaneously on a virtualization layer or hypervisor.

Another technique for including Android in an IVI system is to use a hardware or software virtualization layer (see Figure 2). In this scenario, each OS or domain runs on a dedicated virtual machine, and the hardware resources available from the underlying host platform are shared. Communication is allowed in a controlled way between the different domains, and boot-ups may be independent, allowing safety-critical features running on a dedicated domain to be available more quickly than the infotainment or Android systems.

Several hardware platform providers provide isolated domains in hardware. Software virtualization is available using proprietary software from providers such as SYSGO, OpenSynergy, and Open Kernel Labs. These virtualization layers consume a small amount of overall resource (typically 1 percent to 4 percent) and allow a high degree of domain isolation and safety.

In a few years, all drivers will expect their vehicles to be permanently connected to the Internet. This will allow access to cloud data services, telematics, video and audio streaming, and apps download. It is no longer a question of if this happens, but rather when all of this becomes available to the general public. The explosion of Android in smartphones has ensured that Android apps will need to be accessible in vehicles, and users will decide if these are built in or accessed via a BYOD solution. **ECD**

Andrew Patterson *is business development director for Mentor Graphics Embedded Software Division, specializing in the automotive market.*

Mentor Graphics
Embedded Software Division
Andrew_patterson@mentor.com
www.mentor.com
Follow:

Conveniently Designed with a Small Footprint

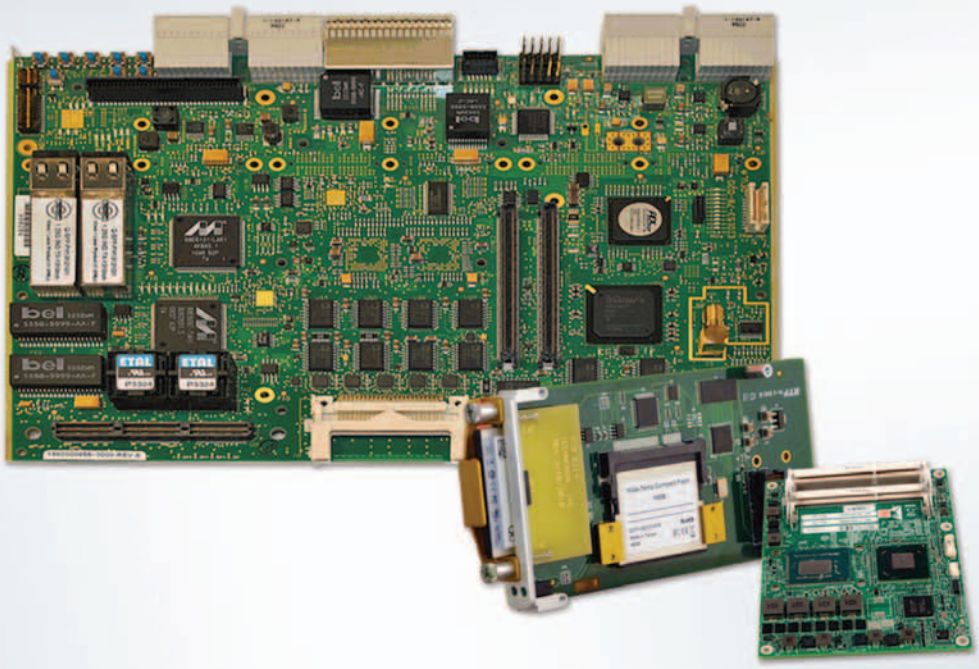

Radisys Type 6 COMe module

- Engineering time significantly reduces open source standards based designs
- High capacity for application-specific I/O on the carrier
- Carrier can be designed to fit any form factor
- Tailored to your requirements with COM Express modules & mezzanines
  - Large COTS ecosystem of COM Express & XMC/PMC modules for processor and additional I/O
- Meets most environmental requirements

Talk to Elma Bustronic about your next carrier card designs for COM Express. We will work with you as an extension of your team to provide the right carrier for your project's footprint.

ELMA
Your Solution Partner

Learn why the smallest details matter to Elma Bustronic by visiting **www.elmabustronic.com.**

Or by contacting us at **510.490.7388** or **sales@elmabustronic.com**.

bus tronic
Elma Bustronic

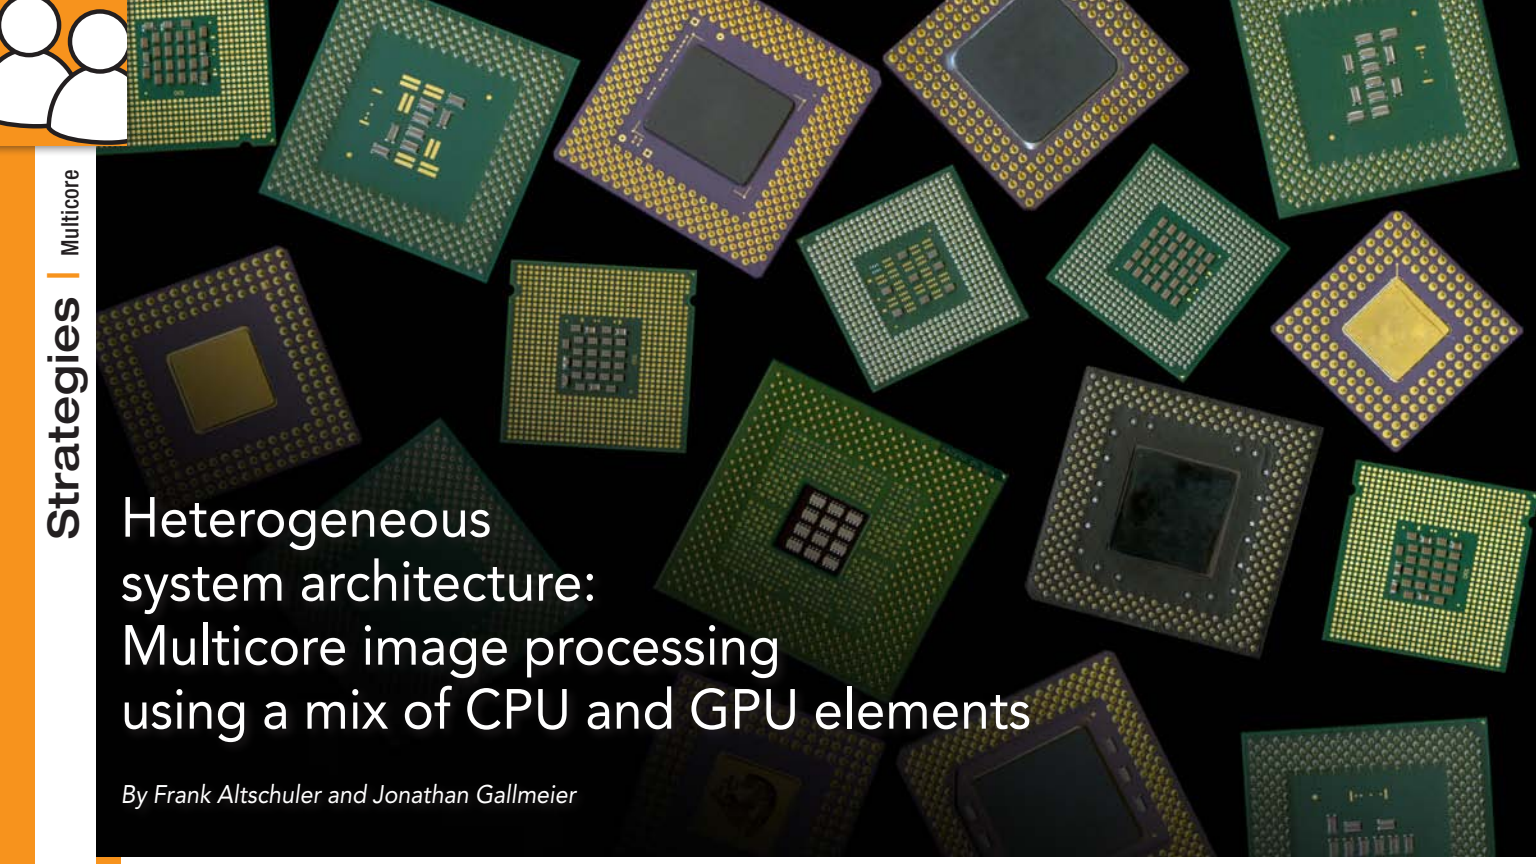

Heterogeneous system architecture: Multicore image processing using a mix of CPU and GPU elements

By Frank Altschuler and Jonathan Gallmeier

Image processing is computationally intensive, requiring immense resources in CPU and memory throughput. Parallelism through multiple CPU cores helps, but even with today's dual-, quad-, and higher-count core devices, multimedia tasks demand either a great deal of power or the addition of dedicated image processing hardware. Adding a pool of parallel processing elements as a programmable accelerator to the CPU creates a proper balance of general-purpose processing, high performance, and low power, as demonstrated in an example showing 2x-3x performance and performance-per-watt improvements.

Although multimedia processing has benefited tremendously from the increased availability of multicore processors, until recently, the tendency has been to consider multicore from the standpoint of largely homogeneous CPU architectures. A typical approach offers dual, quad, and higher numbers of cores arranged in Symmetric Multi-Processor (SMP) or cache-coherent Non-Uniform Memory Architecture (ccNUMA) groups. These devices sometimes come with additional hardware acceleration to offload particularly difficult tasks, such as the Context-Adaptive Variable-Length Coding (CAVLC) or processor-intensive Context-Adaptive Binary Arithmetic Coding (CABAC) entropy encoding functions used in H.264 video processing.

Many companies now offer integrated graphics in their devices, largely as a way to save board space and bill of material costs. An additional benefit of having an integrated Graphics Processing Unit (GPU) is that it can provide another path to attacking challenging computational tasks. As an example, AMD's R-Series processors integrate the Bulldozer CPU architecture with a discrete-class AMD Radeon GPU to produce an Accelerated Processing Unit (APU).

Figure 1 shows a high-level view of the architecture. The GPU is an array of Single Instruction Multiple Data (SIMD) processing elements. Note the common memory controller between the array and x86 cores. This architecture is an early step in the movement to AMD's Heterogeneous System Architecture (HSA), which breaks significantly from the traditional multicore approach by treating integrated graphics processing elements as a pool of computational resources that can be applied to nongraphics data. It is similar to the way that an attached hardware accelerator can be used, but with the advantage of being a large cluster of programmable devices – up to 384 cores at the high end – with many hundreds of GFLOPS worth of computational throughput.

Table 1 illustrates the basic operating parameters of an AMD desktop part with the same underlying architecture as the AMD R-Series APU. The Bulldozer-based CPU's floating-point performance at 121.6 GFLOPS is dwarfed by the GPU, which offers 614 GFLOPS to applications that can utilize the resource effectively. The key to leveraging the added GPU processors in imaging is to consider the time needed to transfer a block of data to the GPU, process it, then return the results to the main thread.

Performance analysis in a video transcoding application

Although current-generation AMD R-Series processors require a buffer transfer between the GPU and CPU, data-parallel applications can achieve an improvement in throughput and power efficiency that significantly outweighs the cost of the data transfer.

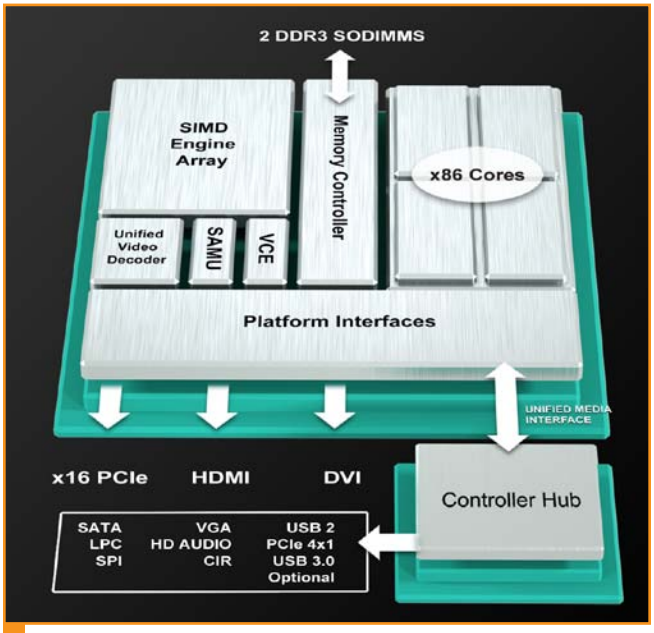

Figure 1 | A high-level view of the AMD R-Series APU architecture shows a common memory controller between the array of SIMD processing elements and x86 cores.

Handbrake A10 5800K (100 W TDP)	CPU	GPU
Frequency (GHz)	3.8	0.8
Cores	4	384
FLOPS/core	8	2
GFLOPS	121.6	614.4
A10 5800K total GFLOPS	736	
A10 5800K power (W TDP)	100	

Table 1 | A comparison of operating parameters demonstrates CPU versus GPU performance in an AMD desktop part.

One example of a data-parallel application is Handbrake, a widely used video transcoding utility that runs on a variety of operating systems and hardware platforms. The utility is notable in that it has well-integrated multicore support and therefore serves as an appropriate platform to study the effectiveness of multicore strategies. Handbrake uses x264, a full-featured H.264 encoder used in several open-source and commercial products that provide adequate multicore support.

OpenCL has been used in both of these projects to maximize the GPU compute resources available within the APU. In x264, the quality optimizations associated with the lookahead function were ported completely to OpenCL on the GPU. This can be used in transcoding applications to improve the output video quality and normally consumes up to about 20 percent of CPU time. Handbrake uses x264 directly as the video encoder. In addition, Handbrake uses OpenCL to perform video scaling and color space conversion from the RGB to YUV color space. AMD's hardware video decoder (UVD) was also used in Handbrake to perform some of the video decoding tasks.

Data was produced on a Trinity 100 W A10 desktop machine; however, the underlying architecture is identical to that used in the embedded product.

To simplify analysis of the APU's complex power management architecture, the top Thermal Design Power (TDP) was evaluated instead of breaking out the power consumed by each subsystem. The data displayed in Table 2 shows that when processing is performed purely in the CPU, a 14.8 fps frame rate is achieved, and processing is accomplished in 62 seconds. Involving the GPU and UVD blocks in the processing increased the frame rate to 18.3 fps, reducing processing time and hence power consumption by 14 seconds or 22 percent.

The floating-point operations total in Table 2 is the product of the floating-point throughput from Table 1 and the processing time, from which can be drawn a figure of merit: operations per watt. Looking at the total floating-point operations and normalized operations per watt, it becomes clear that processing resources are being left on the table. CPU-only throughput is 75 operations per watt, while it is boosted to 353 operations per watt when the GPU is added for an improvement in theoretical floating-point throughput of nearly 5x (4.7).

Twenty-two percent is nontrivial savings, but while this generation of APU makes significant progress in the direction of HSA, all memory accesses are not uniform from the standpoint of latency and throughput. When the kernel accesses memory, it does so directly. Such operations occur at a theoretical rate of 22 GBps or 16 to 18 GBps in typical applications. When a buffer is created in host memory to access GPU data, however, that data goes through a different path where the effective

Handbrake A10 5800K (100 W TDP)	*High preset, 1080p input transcoded to 720p 6 Mbps output*		
	No GPU	GPU used + UVD	Difference
	14.8 fps	18.3 fps	3.5 fps
	62 sec	48 sec	-14 sec
Floating-point operations total (x10^6)	7,539.2	3,5328	27,788.8
Operations/watt	75.392	353.28	277.888

Table 2 | An analysis of a Handbrake application on an A10 desktop machine shows the performance difference when using a CPU alone versus using a GPU and UVD blocks for processing.

A 1.66GHz Intel® Atom™ Industrial SBC of EPIC Proportions

Fanless Extended Temperature All-in-One Single Board Computer

The EPX-C380 is expandable with PC/104 I/O modules.

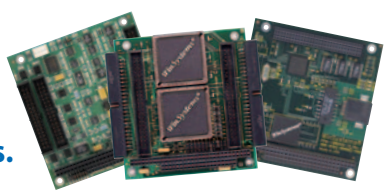

The EPX-C380 is a rugged single board computer that provides a powerful platform for industrial applications.

Powered with either a single or dual core processor, it has a wealth of onboard I/O plus expansion options. Also it supports Windows®, Windows® Embedded, Linux, and other x86 real-time operating systems.

Features include:

- 1.66GHz Intel® Atom™ N450 single core or D510 dual core processor
- Embedded Gen 3.5+ GFX video core supports CRT and LVDS flat panels simultaneously
- Custom splash screen on start up
- Optional 1MB of battery backed SRAM
- Two, Gigabit Ethernet ports
- Two SATA channels
- Eight USB 2.0 ports each with over-current protection
- Four serial RS-232/422/485 channels
- 48 bi-directional digital I/O lines
- CompactFlash (CF) card supported
- MiniPCIe and PC/104-*Plus* expansion connectors
- High Definition audio supported (7.1)
- WDT, RTC, LPT, status LEDs, and beeper
- +5 volt only operation
- EPIC sized: 4.5" x 6.5" and RoHS compliant
- EBX sized SBC also available
- -40°C to +85°C operation
- Responsive and knowledgeable technical support
- Long-term product availability

Contact us for additional information, custom configurations, and pricing. Our factory applications engineers look forward to working with you.

Call 817-274-7553 or Visit www.winsystems.com/EPX-C380E2

Ask about our product evaluation program.

WinSystems®
An Employee-Owned Company

715 Stadium Drive • Arlington, Texas 76011
Phone 817-274-7553 • FAX 817-548-1358
E-mail: info@winsystems.com

Linux

Extended Temperature

RoHS compliant

epic™

SBCS
ENCLOSURES
I/O MODULES
PANEL PCS
SUMIT
COMPACT FLASH
POWER SUPPLIES
EMBEDDED PCS

bandwidth is closer to 8 MBps. This means that the gap in memory bandwidth between what the kernel can expect locally and what it can expect if the programmer sets up a local buffer to GPU memory is nearly 3:1. This is symmetrical. If the GPU needed data residing in CPU memory, the same relationship would hold true.

Future devices adhering to the full HSA architecture will have a uniform memory model, eliminating the need to copy data between the CPU and GPU memory regions. Algorithms such as the lookahead functions of x264, when moved into the GPU's processing domain, will have near parity from the standpoint of memory bandwidth.

Performance analysis in a benchmarking application

A second example of a data-parallel application is LuxMark, a graphics-oriented benchmarking package that uses OpenCL and illustrates what can be achieved when differences in memory access times are removed from the equation. The default integrated benchmark was run on the AMD Embedded R-464L APU, whose characteristics are shown in Table 3.

R-464L (35 W TDP)	CPU	GPU
Frequency (GHz)	2.3	0.685
Cores	4	384
FLOPS/core	8	2
GFLOPS	73.6	526.08
R-464L total GFLOPS	599.68	
R-464L power (W TDP)	35	

Table 3 | A comparison of operating parameters demonstrates CPU versus GPU performance in an AMD Embedded R-464L APU.

LuxMark v. 2.0	AMD R-464L (35 W)
Default integrated benchmark	
OpenCL GPU	198
OpenCL CPU + GPU	230
OpenCL CPU	146

Table 4 | Results from running the LuxMark 2.0 benchmark on the AMD Embedded R-464L APU indicate how the GPU outperforms the CPU in this processing task.

The results in Table 4 clearly demonstrate that the GPU outperforms the CPU by more than 25 percent for this class of processing task. With a full HSA implementation, one would expect to see the addition of the CPU plus GPU to be close to the sum of GPU plus CPU: 344. In practice, the resulting number is 230, which is a nontrivial 63 percent boost but not reflective of fully utilized resources. The gap is largely explained by the overhead of the CPU parsing data and transferring the data back and forth between the CPU and GPU at the lower speeds described earlier. Such features will enable the utilization of both the CPU and the GPU with minimal overhead, allowing the combined performance of both CPU and GPU to be significantly better than either one alone.

Maximizing computational resources

Newly released APU products provide significant computational resources to designers of embedded computing products. Taking advantage of those resources requires a solid understanding of the underlying hardware architecture, insight into the data flow issues within the target application, and familiarity with the latest tools such as OpenCL. ECD

Frank Altschuler is senior product manager at AMD.

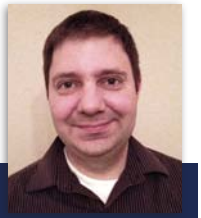

Jonathan Gallmeier is a SMTS/CAS engineer at AMD.

AMD
www.amd.com

Follow:

SUPERMICRO®

Embedded Building Blocks

Modular Computing Solutions for Embedded Applications

Small Form Factor
X7SPE-H(F)-D525, X7SPA-H(F)-D525
X9SCV Series

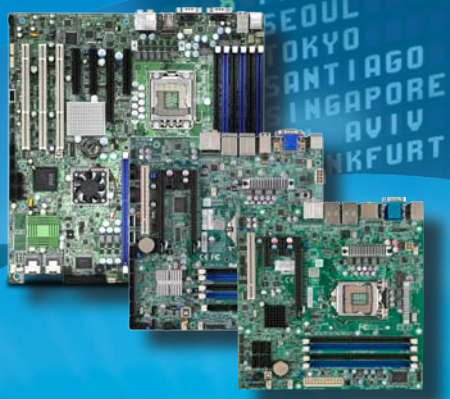

High-Performance Application-Optimized
X8DA6, C7SIM-Q, C7Q67

Compact Form Factor Short-Depth
SuperServer® 5015A-EHF-D525
SuperServer® 5017C-LF

- Supports Intel® Xeon®, Intel® Core™ i7/i5/i3 and Intel® Atom™ Processors
- Energy Efficient, Low Power and Fanless Solutions
- High Performance, Scalable Systems
- High-Density, Compact Form Factors
- Open Standard and Multi-OS Support
- Ruggedized and Industrial Grade Chassis
- 7- Year Life Cycle Support

Versatile, Whisper Quiet and Highly Configurable
SuperServer® 5037C-i/T

Industrial PC Short-Depth
SuperServer® 6046T-TUF

intel Intelligent Systems Alliance

Digital! We Power the Cloud

21 dB Whisper-Quiet

www.supermicro.com/Embedded

Q & A

A modern software platform for the era of multicore and cloud computing

Q&A with Mark Brewer, President and CEO, Typesafe

Created from the ground up to address multicore and parallel computing, the Scala programming language smoothly integrates features of object-oriented and functional languages, enabling developers to be more productive while retaining full interoperability with Java. Mark explains how Scala-based middleware technology can maximize modern multicore hardware and cloud computing software by raising the abstraction level for building multithreaded applications.

ECD: ***What are the advantages of using general-purpose programming languages like Java and Scala for embedded development?***

BREWER: Scala is a general-purpose programming language designed to express common programming patterns in a concise, elegant, and type-safe way. Scala smoothly integrates features of object-oriented and functional languages, enabling developers to be more productive while retaining full interoperability with Java and taking advantage of modern multicore hardware.

Scala is also a functional language. Inspired by the long tradition of functional programming, Scala makes it easy to avoid shared state so that computation can be readily distributed across cores on a multicore server and across servers in a data center. This makes Scala an especially good match for modern multicore CPUs and distributed cloud computing workloads that require concurrency and parallelism.

Scala is equipped with an expressive type system that detects and avoids many kinds of application errors at compile time. At the same time, a sophisticated type inference capability frees developers from the redundant type information "boilerplate code" that is typical of Java.

Because the code sizes are typically reduced by a factor of two to three when compared to an equivalent Java application, Scala is well suited for an embedded environment due to its lightweight and concise nature, as demonstrated in the following code samples:

Java class definition

```
public class Person {
  public final String name;
  public final int age;
  Person(String name, int age) {
      this.name = name;
      this.age = age;
  }
}
```

Equivalent Scala class definition

```
class Person(val name: String, val age: Int) {}
```

Java code

```
import java.util.ArrayList;
...
Person[] people;
Person[] minors;
Person[] adults;
{  ArrayList<Person> minorsList = new ArrayList<Person>();
   ArrayList<Person> adultsList = new ArrayList<Person>();
   for (int i = 0; i < people.length; i++)
       (people[i].age < 18 ? minorsList : adultsList)
       .add(people[i]);
   minors = minorsList.toArray(people);
   adults = adultsList.toArray(people);
}
```

Equivalent Scala code

```
val people: Array[Person]
val (minors, adults) = people partition (_.age < 18)
```

Scala protects investments in existing Java libraries, tools, and developer programming skills. Scala programs are compiled directly to Java bytecode that runs on the mature Java Virtual Machine (JVM), leveraging its robust just-in-time compilation, garbage collection, and well-understood deployment techniques. The operations team doesn't see a difference. Developers keep working with their familiar tools, but they're writing code that's shorter, faster, more scalable, more correct, and maybe even more fun.

ECD: ***How does the latest wave of multicore processors affect software development with object-oriented programs?***

BREWER: The Java programming language was created in 1995, so it was suited to handle the first generation of Internet applications, with object-oriented programming models, multiplatform runtime, and network orientation. However, with the advent of cloud computing and interactive applications that demand near real-time capabilities, the Java language and traditional Java middleware have begun to show their age when faced with the equally significant hurdles of large-scale distributed applications and multicore platforms.

In most languages, the key to utilizing the full power of multicore CPUs is by writing concurrent multithreaded applications. Considering the shared state, state visibility, threads, locks, concurrent collections, and thread notifications involved, writing this type of application is difficult, even for experienced developers. These concepts are by nature error-prone and often result in deadlocks or application crashes.

In an attempt to keep up with the times, Java has evolved, but has become bulky and cumbersome in the process. Java and Java Enterprise Edition (JEE) application servers are inherently difficult to scale predictably; scaling up is extremely hard to accomplish, and the commercial licensing terms offered by most vendors prohibit the economical scale-out that customers so desperately need as an alternative to this model.

Another essential component of building scalable applications for embedded devices is middleware. Akka, built upon the scale afforded by the Scala programming language, is a message-oriented programming model for building multithreaded applications. Akka raises the level of abstraction so that developers only need to worry about messages and business logic, instead of dealing with the low-level plumbing required in Java.

Play! builds upon Akka to deliver a Model-View-Controller (MVC)-style Web framework with a development experience much like that enjoyed by Rails developers. Play-mini is a subset of the Play! framework that consists of a REST layer on top of the Netty non-blocking I/O socket server. It offers the ability to deploy Akka applications for service-layer jobs that don't need the rest of the Play tools (for example, the MVC/interface layer).

ECD: ***How does advanced middleware technology enable developers to build better software for the cloud?***

BREWER: Akka is an event-driven middleware framework implemented in Scala for building reliable, high-performance distributed applications. It raises the abstraction level for the developer, removing the need to worry about the low-level plumbing required to create highly concurrent applications in languages such as Java.

Raising the abstraction level is key for building better software in the cloud. Developers can focus on implementing business logic and adding value, not spending time worrying about implementing low-level features such as high-availability services and state/memory management. With Akka, the developer gets an ideal fabric for the cloud that is:

- Elastic and dynamic, with the ability to expand and contract based on the actual load the system is experiencing
- Fault-tolerant and self-healing, where the system detects failures automatically and can restart individual components or entire servers based on the business requirements

- Customizable and adaptive load-balancing, where software components can route messages based on system load or any other user-definable criteria
- Configured so that clusters can rebalance the load in failover situations via actor migration, where actors are moved between systems dynamically

The bottom line is that it's very easy to build loosely coupled and dynamic systems that can almost organically change and adapt at runtime.

ECD: ***Briefly explain Typesafe's technology and the current applications for it in an embedded computing environment.***

BREWER: Typesafe's vision is to enable development of concurrent, fault-tolerant applications with a single unified programming model, managed runtime, and binary compatible distribution.

Typesafe was founded in 2011 by the creators of the Scala programming language and Akka middleware, who joined forces to create a modern software platform for the era of multicore hardware and cloud computing workloads. The company provides an easy-to-use packaging of Scala, Akka, Play!, and developer tools through both an open-source stack and a commercial stack that provides commercial support, maintenance, and operations tools via the Typesafe Subscription. In conjunction with its partners, Typesafe also provides training and consulting services to accelerate the commercial adoption of Scala, Akka, and Play!

Companies that use the Typesafe Stack in an embedded way typically rely on its low latency, high throughput, and resiliency design points. It is used as for handling millions of messages per second across networks and network devices.

ECD: ***What challenges are your customers dealing with right now?***

BREWER: Performance at scale is a main driver for customers to consider the Typesafe Stack. We've found that customers who have tried the traditional development paradigms using JEE application servers, PHP, and Ruby hit a performance or efficiency wall. In addition to being highly efficient, the components of the Typesafe Stack are compact from both a disk and memory standpoint, especially when compared to the traditional JEE application server, and lend themselves well to the embedded space. ECD

Mark Brewer *is president and CEO of Typesafe.*

Typesafe
mark.brewer@typesafe.com
typesafe.com

Follow:

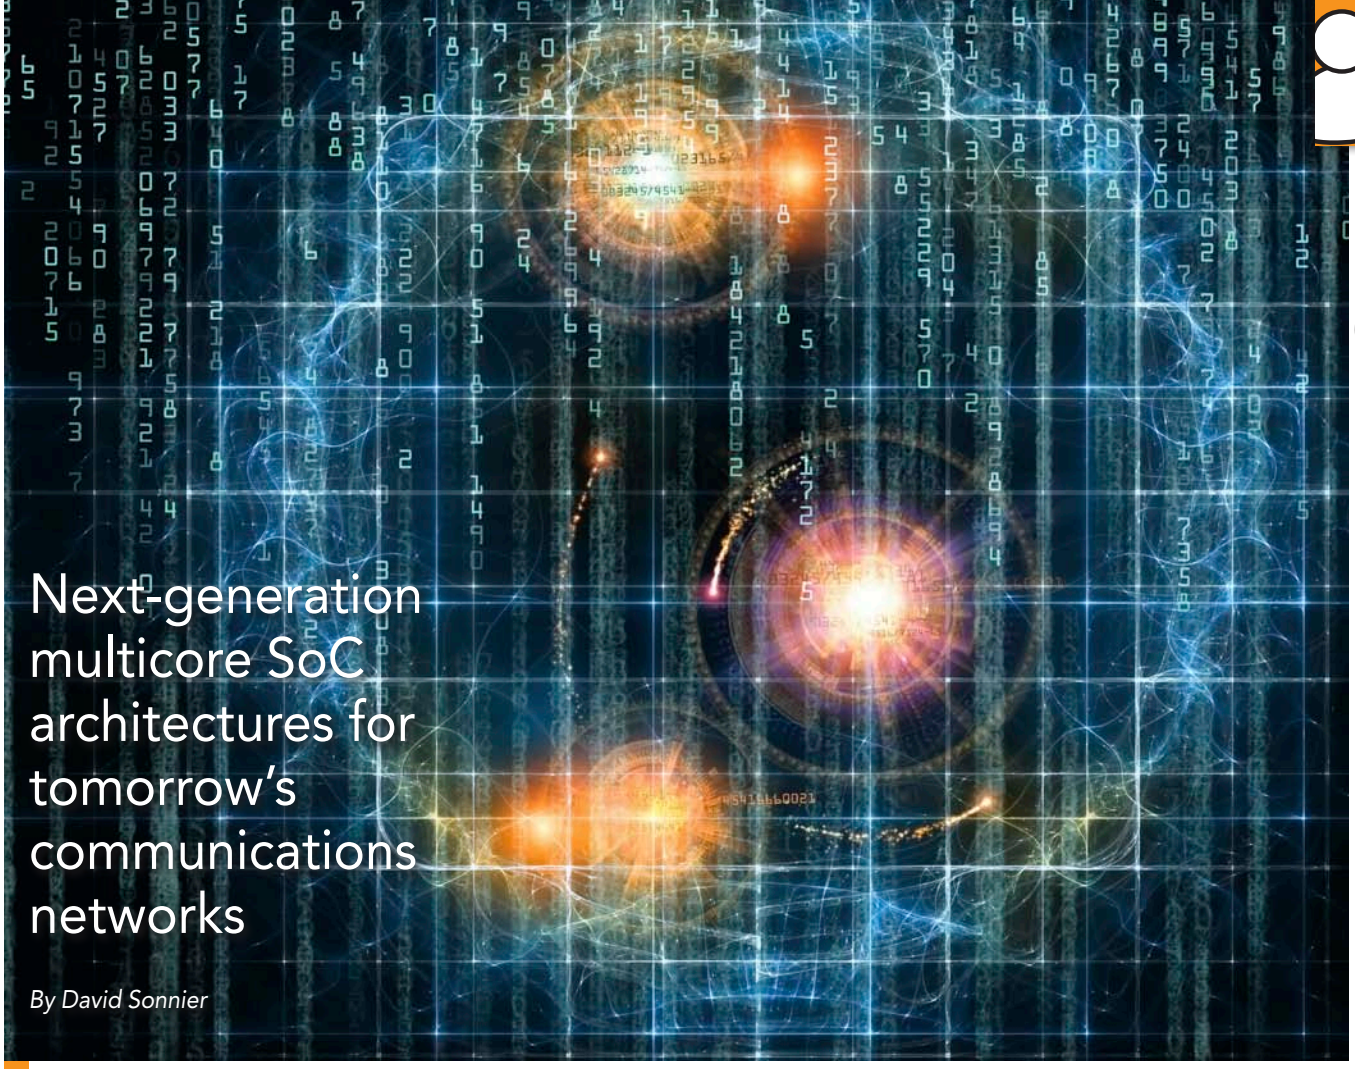

Next-generation multicore SoC architectures for tomorrow's communications networks

By David Sonnier

IT managers are under increasing pressure to boost network capacity and performance to cope with the data deluge. Networking systems are under a similar form of stress with their performance degrading as new capabilities are added in software. The solution to both needs is next-generation System-on-Chip (SoC) communications processors that combine multiple cores with multiple hardware acceleration engines.

The data deluge, with its massive growth in both mobile and enterprise network traffic, is driving substantial changes in the architectures of base stations, routers, gateways, and other networking systems. To maintain high performance as traffic volume and velocity continue to grow, next-generation communications processors combine multicore processors with specialized hardware acceleration engines in SoC ICs.

The following discussion examines the role of the SoC in today's network infrastructures, as well as how the SoC will evolve in coming years. Before doing so, it is instructive to consider some of the trends driving this need.

Networks under increasing stress

In mobile networks, per-user access bandwidth is increasing by more than an order of magnitude from 200-300 Mbps in 3G networks to 3-5 Gbps in 4G Long-Term Evolution (LTE) networks. Advanced LTE technology will double bandwidth again to 5-10 Gbps. Higher-speed access networks will need more and smaller cells to deliver these data rates reliably to a growing number of mobile devices.

In response to these and other trends, mobile base station features are changing significantly. Multiple radios are being used in cloud-like distributed antenna systems. Network topologies are flattening. Operators are offering advanced Quality of Service (QoS) and location-based services and moving to application-aware billing. The increased volume of traffic will begin to place considerable stress on both the access and backhaul portions of the network.

Traffic is similarly exploding within data center networks. Organizations are pursuing limitless-scale computing workloads on virtual machines, which is breaking many of the traditional networking protocols and procedures. The network itself is also becoming virtual and shifting to a Network-as-a-Service (NaaS) paradigm, which is driving organizations to a more flexible Software-Defined Networking (SDN) architecture.

These trends will transform the data center into a private cloud with a service-oriented network. This private cloud will need to interact more seamlessly and securely with public cloud offerings in hybrid arrangements. The result will be the need for greater intelligence, scalability, and flexibility throughout the network.

Moore's Law not keeping pace

Once upon a time, Moore's Law – the doubling of processor performance every 18 months or so – was sufficient to keep pace with computing and networking requirements. Hardware and software advanced in lockstep in both computers and networking equipment. As software added more features with greater sophistication, advances in processors maintained satisfactory levels of performance. But then along came the data deluge.

In mobile networks, for example, traffic volume is growing by some 2,600 percent per year, owing mostly to the increase in video traffic. This is already causing considerable congestion, and the problem will only get worse when an estimated 50 billion mobile devices are in use by 2016 and the total volume of traffic grows by a factor of 50 in the coming decade.

In data centers, data volume and velocity are also growing exponentially. According to IDC, digital data creation is rising 60 percent per year. The research firm's Digital Universe Study predicts that annual data creation will grow 44-fold between 2009 and 2020 to 35 zettabytes (35 trillion gigabytes). All of this data must be moved, stored, and analyzed, making Big Data a big problem for most organizations today.

With the data deluge demanding more from network infrastructures, vendors have applied a Band-Aid to the problem by adding new software-based features and functions in networking equipment. Software has now grown so complex that hardware has fallen behind. One way for hardware to catch up is to use processors with multiple cores. If one general-purpose processor is not enough, try two, four, 16, or more.

Another way to improve hardware performance is to combine something new – multiple cores – with something old – Reduced Instruction Set Computing (RISC) technology. With RISC, less is more based on the uniform register file load/store architecture and simple addressing modes. ARM, for example, has made some enhancements to the basic RISC architecture to achieve a better balance of high performance, small code size, low power consumption, and small silicon area, with the last two factors being important to increasing the core count.

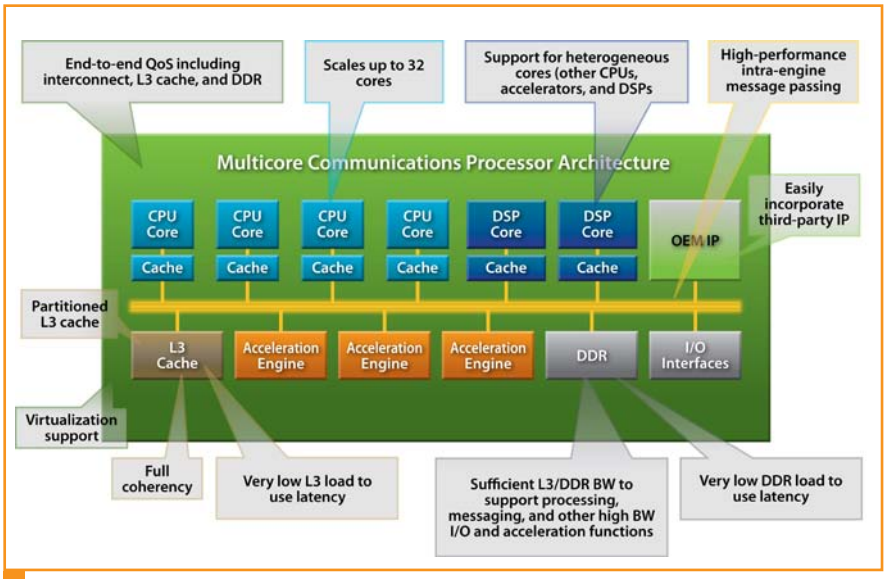

Figure 1 | SoC communications processors combine multiple general-purpose processor cores with multiple task-specific acceleration engines to deliver higher performance with a lower component count and lower power consumption.

Hardware acceleration necessary, but …

General-purpose processors, regardless of the number of cores, are simply too slow for functions that must operate deep inside every packet, such as packet classification, cryptographic security, and traffic management, which is needed for intelligent QoS. Because these functions must often be performed in serial fashion, there is limited opportunity to process them simultaneously in multiple cores. For these reasons, such functions have long been performed in hardware, and it is increasingly common to have these hardware accelerators integrated with multicore processors in specialized SoC communications processors.

The number of function-specific acceleration engines available also continues to grow, and more engines (along with more cores) can now be placed on a single SoC. Examples of acceleration engines include packet classification, deep packet inspection, encryption/decryption, digital signal processing, transcoding, and traffic management. It is even possible now to integrate a system vendor's unique intellectual property into a custom acceleration engine within an SoC. Taken together, these advances make it possible to replace multiple SoCs with a single SoC in many networking systems (see Figure 1).

In addition to delivering higher throughput, SoCs reduce the cost of equipment, resulting in a significant price/performance improvement. Furthermore, the ability to tightly couple multiple acceleration engines makes it easier to satisfy end-to-end QoS and service-level agreement requirements. The SoC also offers a distinct advantage when it comes to power consumption, which is an increasingly important consideration in network infrastructures, by providing the ability to replace multiple discrete components in a single energy-efficient IC.

The powerful capabilities of today's SoCs make it possible to offload packet

processing entirely to system line cards such as a router or switch. In distributed architectures like the IP Multimedia System and SDN, the offload can similarly be distributed among multiple systems, including servers.

Although hardware acceleration is necessary, the way it is implemented in some SoCs today may no longer be sufficient in applications requiring deterministic performance. The problem is caused by the workflow within the SoC itself when packets must pass through several hardware accelerators, which is increasingly the case for systems tasked with inspecting, transforming, securing, and otherwise manipulating traffic.

If traffic must be handled by a general-purpose processor each time it passes through a different acceleration engine, latency can increase dramatically, and deterministic performance cannot be guaranteed under all circumstances. This problem will get worse as data rates increase in Ethernet networks from 1 Gbps to 10 Gbps, and in mobile networks from 300 Mbps in 3G networks to 5 Gbps in 4G networks.

Next-generation multicore SoCs

LSI addresses the data path problem in its Axxia SoCs with Virtual Pipeline technology. The Virtual Pipeline creates a message-passing control path that enables system designers to dynamically specify different packet-processing flows that require different combinations of multiple acceleration engines. Each traffic flow is then processed directly through any engine in any desired sequence without intervention from a general-purpose processor (see Figure 2). This design natively supports connecting different heterogeneous cores together, enabling more flexibility and better power optimization.

In addition to faster, more efficient packet processing, next-generation SoCs also include more general-purpose processor cores (to 32, 64, and beyond), highly scalable and lower-latency interconnects, nonblocking switching, and a

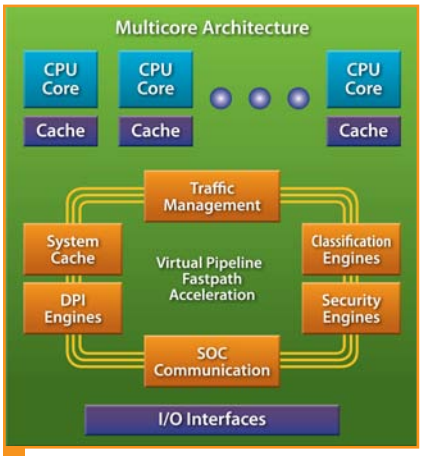

Figure 2 | To maximize performance, next-generation SoC communications processors process packets directly and sequentially in multiple acceleration engines without intermediate intervention from the CPU cores.

wider choice of standard interfaces (Serial RapidIO, PCI Express, USB, I2C, and SATA) and higher-speed Ethernet interfaces (1G, 2.5G, 10G, and 40G+). To easily integrate these increasingly sophisticated capabilities into a system's design, software development kits are enhanced with tools that simplify development, testing, debugging, and optimization tasks.

Next-generation SoC ICs accelerate time to market for new products while lowering both manufacturing costs and power consumption. With deterministic performance for data rates in excess of 40 Gbps, embedded hardware is once again poised to accommodate any additional capabilities required by the data deluge for another three to four years. ECD

***David Sonnier** is a technical fellow in system architecture for the Networking Solutions Group of LSI Corporation.*

LSI Corporation
david.sonnier@lsi.com
www.lsi.com

Follow:

What's Hot in 2013 in Embedded?

Embedded COMPUTING DESIGN
What's Hot in 2013 in Embedded?

The coming year promises strong potential for growth in many market segments that are increasingly relying on embedded technologies to manage power, optimize performance, and run any number of essential processes for fast and reliable system operations. In particular, the trend toward deploying intelligent systems connected by smart devices is driving productivity and ingenuity in the embedded space.

In this special advertising section, we approached embedded technology companies seeking their perspectives and opinions on what the New Year holds for our industry, summed up in the question "What's Hot in 2013 in Embedded?" Leading vendors spanning all of our coverage areas – Silicon, Software, and Strategies – present their viewpoints on what they believe will fuel the fires of innovative development in 2013. Hot topics run the gamut, from automotive infotainment to Machine-to-Machine (M2M) connectivity, FPGA and System-on-Chip (SoC) integration to software monetization, and standardized embedded platforms with flexible processor options to e-Commerce-based support for tackling the challenges of software design complexity.

These trends represent opportunities for embedded designers to achieve faster performance, improved power efficiency, heightened security, and a range of features that address the current and future demand for intelligent devices and systems.

FLEXERA SOFTWARE

Microsemi

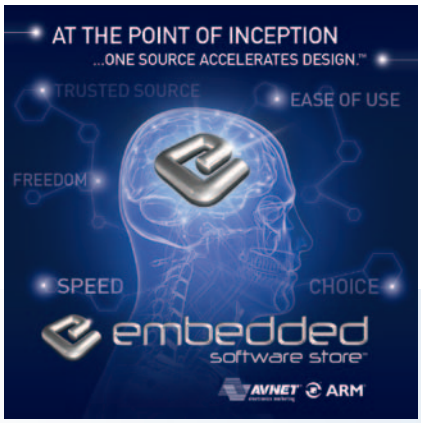

Semiconductor Memories for the Automotive Electronics Industry

Micron Technology

Advancement and innovation in the automotive industry are what's hot – from solutions that help keep drivers' eyes on the road with heads-up displays, to wireless solutions that allow cars to talk to each other, to systems that improve night vision through infrared sensors and safety features that provide visibility to protect pedestrians.

The key megatrends that are fueling the pervasiveness of microelectronics in cars and driving memory consumption include:

- **Infotainment:** According to industry analysts, many buyers today care more about the infotainment technologies embedded in the dashboard than what's under the hood. High-end electronics provide drivers and passengers with in-car navigation and entertainment systems, as well as information delivered over a wireless network. As a result, the global automotive embedded telematics market is expected to grow at a Compound Annual Growth Rate (CAGR) of more than 20 percent from 2013 to 2015.
- **Safety:** Automotive safety is moving from passive-only systems for accident response to active systems for accident prevention, which can take control of the car and support the driver in emergency situations. Some Advanced Driver Assistance System (ADAS) features are already well-known and provide welcome convenience and safety. These include adaptive cruise control, blind spot monitoring, lane departure warning, and night vision. Microelectronics is the fundamental enabling technology for all advanced safety features in today's cars.
- **Environment:** Restrictive government regulations aimed at reducing emissions drive continuous improvement in engine and transmission control applications. Microelectronics will play a major role in the quest for the ultimate environmentally friendly car engine. Fuel-saving engines and hybrid drives have one thing in common: they make use of intelligent electronics to reduce fuel consumption.

The megatrends fueling the growth of automotive electronics have a significant impact on the market demand for semiconductor memories. On average there are 1.5 memory devices in each car, and in some scenarios, up to seven per car. With more than 121 million memory devices shipped into the automotive market in the past year, Micron is at the forefront of developing memory solutions that are driving the power of progress in the automotive industry to a new level.

Car telematics and infotainment applications that require additional storage space for rich multimedia data and advanced software and applications are driving explosive growth for both volatile and nonvolatile memories. Fully digital dashboard instrumentation increases the needs for advanced solid-state storage solutions. And new ADAS features rely on embedded intelligent control units with significantly large memory systems for program and data.

Infotainment applications alone are – and will continue to be – the highest memory consumers, representing more than 60 percent of total market value in 2012, followed by ADAS and body control at 34 percent, and then power train solutions.

Memory Matters: High Quality, Innovation, Broad Memory Offering

Micron has one of the broadest portfolios of leading memory technologies for the automotive market. From best-in-class DRAM, NOR, NAND, and e•MMC devices for infotainment; to NOR and DRAM for driving assistance, advancement safety systems, and user-friendly dashboard instrumentation; to highly reliable NOR for power train and transmission – Micron's product offerings enable customers to rely on a single source for all their memory needs.

Micron Technology
www.micron.com

The Solution Company Business Model

Flexera Software

FlexNet® Licensing® and FlexNet Operations®

In 2013, hardware commoditization will continue to negatively challenge the profit margins of intelligent device manufacturers. With globalization continuing to transform the marketplace, device manufacturers can no longer rely on predictable sales of devices to drive growth.

In addition, the intelligent device marketplace will continue to see growth as leading manufacturers and service providers formulate and execute their Machine-to-Machine (M2M) strategies. As more companies move into the M2M space, device manufacturers will continue to experience increased competition. Because customer requirements and expectations will constantly shift and grow, to remain relevant and competitive, device manufacturers must adapt.

To drive growth in the face of a demanding marketplace, device manufacturers must differentiate themselves by rethinking their business models and product strategies. One trend that will continue to gain more traction in 2013 is device manufacturers' transformation of business models to become more like solution companies, where software drives differentiation and value. The solution company business model empowers manufacturers to transform their product lines by harnessing the power of embedded software (already provisioned on their devices) to differentiate themselves from their competitors.

Solution-centric business models allow device manufacturers to deliver more value to customers in the form of competitive differentiation, easy capacity and capability provisioning and upgrades, protection of Intellectual Property (IP), added revenue streams, and targeting of multiple geographies without increasing manufacturing Stock Keeping Units (SKUs). The result is increased customer satisfaction, increased profit margins, and significantly improved business performance.

Licensing and entitlement management is critical to the success of the transformation. Flexera Software has helped many device manufacturers evolve from using an exclusively hardware-based business model to one that optimizes software revenues. Flexera Software's FlexNet Licensing and FlexNet Operations solutions enable device manufacturers to monetize intelligent devices, equipment, and software.

FlexNet Licensing and FlexNet Operations benefit device manufacturers in a number of ways:

- Protect and monetize IP, enabling subscription licensing, maintenance/upgrade fees, usage-based pay-for-use or pay-for-burst, metering, and many other models. These solutions support the full software licensing spectrum, from strict enforcement to usage-based "trust, but verify" models.
- Allow device manufacturers to efficiently create lower-end products at lower price points, mid-market products at moderate price points, and premium products at higher price points – all with the same basic code and hardware – by using electronic licensing to turn device capabilities and capacities on and off.
- Enable device manufacturers to offer customers a way to immediately activate newly purchased product capabilities by simply entering a license key. Device manufacturers can also offer these capabilities on a "try before you buy" basis.
- Allow device manufacturers to enter new markets by protecting IP that runs in virtual and cloud (private and public) environments. This gives the customer the benefit of a lower cost by eliminating expensive hardware while preserving profit margin. Software-only virtual appliance solutions can also be delivered under a Software as a Service (SaaS) and/or external cloud model, representing a rapidly growing worldwide market.

Flexera Software
www.flexerasoftware.com/ecm

Seamless M2M Connectivity

Lilee Systems

Lilee Systems TransAir Intelligent Connectivity Solution

Making connectivity a reality for an increasing number of application areas: that's what's hot. From our personal lives to the business world, having key data available at our fingertips is essential, and knowing that the information is always up-to-date and easily accessible is becoming the norm. The expectation that the massive amounts of data collected can make the world a better place also is sought.

This is the age of the Internet of Things (IoT) and increased intelligent Machine-to-Machine (M2M) implementation. The means by which intelligent systems enable IoT is set to be the challenge that every IT team around the world is looking to achieve both with company-owned tools and systems, as well as the growing Bring Your Own Device (BYOD) network addition. And, of course, the desire is that this can be achieved with remote management and without compromising security or adding painful complexity to the overall network topology.

This is a challenge that Lilee Systems is knee-deep in tackling. For Lilee Systems, it is all about the science of keeping M2M assets connected and communicating seamlessly no matter where they are, whether they are stationary or mobile, how fast or slow they are moving, how many endpoints they need to talk to, whether they utilize wireless technologies cost-effectively, and no matter what the environmental condition. The result beyond seamless M2M connectivity will be advancements in situational awareness (GPS, video, voice), proactive information distribution, improved safety, optimized resources, and much more.

Connected transportation is one of the three hottest areas that Lilee Systems predicts will receive a great deal of attention in 2013. The other two are smart energy and environmental monitoring. Serving the M2M connectivity requirements of these areas requires deep understanding of networking and wireless technologies. The applications within these three market areas touch lives and require complete solutions that are achieved by way of a recipe that combines hardware and software – a recipe that is found in the Lilee Systems TransAir intelligent connectivity M2M solution.

The architecture of the Lilee Systems TransAir intelligent connectivity solution has two parts: the messaging layer and the core network layer. Requirements for link abstraction and message routing along with user-specific applications are provided for in the messaging layer. The core network layer in the TransAir end-to-end solution consists of both wired and wireless connectivity. It enables the components of the messaging layer to communicate with one another. For the end user, this means that setting up new elements within growing M2M networks is simplified with respect to the time required for implementation and the amount of staff training needed. With remote configuration and maintenance enabled, IT teams get the added assurance of maximized system uptime. M2M solutions need built-in flexibility such that costs for power usage and data collection/transmissions can be kept in budget. Look to see further advancements here in 2013.

Lilee Systems sees that the solutions for these three hot markets can be used to enable intelligent connectivity requirements across additional markets, including industrial automation, government and aerospace, eHealth, retail, and finance.

Lilee Systems
www.lileesystems.com

FPGA Technology to Deliver New Capabilities for Systems Requiring Security, Low Power, Reliability, and Integration

Microsemi Corporation

By Paul Ekas, VP of Marketing, SoC Products Group

FPGA suppliers have made impressive gains in functional integration. Virtually all major vendors offer devices combining an FPGA fabric with System-on-Chip (SoC) functionality, including built-in math blocks, high-speed serial interfaces, and microprocessor cores. As we move into 2013, we will see the addition of new capabilities for industrial, military, aviation, communications, medical, and other applications that require even more integration, along with advanced security, low-power operation, and improved reliability.

Design security is important for protecting FPGA configurations from cloning, reverse-engineering, or tampering. The greatest advancements here will come from making base-level security easier to use and adopt. Security should be inherent, which requires the use of SoC FPGAs built on nonvolatile flash technology since, unlike SRAM-based FPGAs, they don't expose a programming bitstream at start-up.

SoC FPGAs must go beyond simply providing design security. During 2013, we will see an increasing emphasis on protecting application data that the FPGA is processing. This includes hardware protection from differential power analysis attacks, the inclusion of non-deterministic random bit generators, and hardware firewalls to protect the integrated microcontroller core.

Low power consumption is also a critical focus area, especially in military designs that must support SWaP requirements. While microcontrollers have included power-saving modes for years, SoC FPGA low-power modes have only just become available. For systems that operate reactively (in standby mode, awaiting a trigger event to begin processing) or periodically (processing on a recurring basis), placing SoC FPGAs in ultra-low-power mode between activity bursts can reduce power by 95 percent or more. Today's solutions can enter into and exit from low-power mode within 100 milliseconds, without affecting register and SRAM states or the FPGA I/O. These new SoC FPGAs reduce standby power by 100x over previous-generation devices.

Reliability is an equally important SoC FPGA focus area for 2013. The biggest issue is Single Event Upsets (SEUs), the most common cause of system failures. SEUs occur when alpha particles in packaging or neutrons in the atmosphere deliver enough charge to change one or more configuration bits in an SRAM-based FPGA. SEUs are the most common cause of system failure, which can lead to injury in safety-critical industrial and medical systems. SRAM-based FPGA suppliers have attempted SEU mitigation techniques, but there are still significant periods of time after an SEU event when the system is operating in a corrupted state. In contrast, flash-based SoC FPGAs are immune to SEU effects.

We expect to see growing adoption of SoC FPGAs that deliver higher levels of integration, combining the inherently reliable flash-based FPGA fabric with numerous hard IP blocks on a single die. This architecture maximizes functional resources in smaller silicon space while improving performance and enabling design customization. Embedded processor cores remove the complexity and logic consumption of soft processor cores implemented in the FPGA fabric, and the tight coupling of peripherals and subsystems further improves size and performance.

The latest SoC FPGAs integrate a full-featured microcontroller subsystem with a powerful embedded processor, plus advanced security processing accelerators, DSP blocks, SRAM, embedded nonvolatile memory, and multigigabit serial communication interfaces. Because these devices use a flash process, they can include embedded nonvolatile memory for data and code storage, significantly improving security. The result is reduced board component count, system cost, and power consumption. As we move into 2013, designers of mission- and life-critical systems can use SoC FPGAs with inherently better reliability and security provided via low-power operation in a highly integrated solution.

Microsemi Corporation
www.microsemi.com

Monetizing Embedded Software

SafeNet

Software Monetization Solutions

SENTINEL SOFTWARE MONETIZATION SOLUTIONS

Embedded developers are evolving as embedded systems increasingly become off-the-shelf hardware running solution-specific software. Because of this transition, software is becoming a more important part of an embedded system's intellectual property value. Hardware vendors are finding that they are now software vendors, and this evolution is forcing an operational change that some vendors are finding very challenging and others are seeing as an opportunity.

Beyond the transition to software-focused embedded solutions, developers are facing problems with deliberate and unintentional misuse of their technology, tampering and theft of IP. Usage enforcement, copy protection, and theft prevention can guard against these threats, but they are only the first step toward unlocking the potential value of IP, which requires a focus on both revenue and profitability.

To monetize IP, intelligent device manufacturers need to leverage the four aspects of a successful software monetization strategy – control, packaging, management, and tracking. When software monetization strategies are implemented successfully, intelligent device manufacturers are able to simultaneously offer a better customer experience and a more profitable solution.

Access and Usage Control

Controlling IP is the foundation of software monetization, as intelligent device manufacturers face problems with deliberate and unintentional misuse of their software, product and feature overuse, competitive IP theft, reverse-engineering, and code tampering. The first step is controlling who has access to the device's software, when they're granted access, and to what extent. By effectively controlling access to source code, device vendors can protect revenue and safeguard the integrity of their brands and products.

Vendors must control the use of their software at the product and feature level to prevent deliberate or unintentional overuse of their offerings and ensure that they are being fairly compensated. Feature-based licensing helps device vendors minimize manufacturing costs while achieving greater packaging flexibility and enabling sophisticated pricing strategies. Feature-based licensing and entitlement management open the door to a world of marketing and sales tools and allow device manufacturers to develop and maintain a single, feature-rich application installed on all devices. The functionality of the device is then controlled through licensing. This enables software vendors to ship the same product with different functionality to different customers at varying price points and upgrade products remotely with lower support and fulfillment costs, thus delivering a better customer experience.

Tracking Usage

Tracking product activation and usage down to the feature level gives intelligent device vendors information to drive product packaging decision-making, roadmap investment, and sales and marketing strategies. With the right knowledge, product management and engineering teams can discontinue unpopular feature combinations and create software packages with the most valued features. In addition, marketing and sales teams can utilize reports to better determine what, when, and how products are being used, and leverage this invaluable data to plan, launch, and execute more effective sales and marketing activities.

Embracing the Change

Shifting from a hardware focus to a software focus typically doesn't happen overnight. Intelligent device manufacturers who embrace the transition and use software monetization solutions to overcome challenges will minimize threats to their IP, reduce manufacturing and inventory costs, expand their product lines, and bring innovative devices to market with greater efficiency, making them more nimble and competitive for the future.

SafeNet
www.safenet-inc.com/software-monetization/embedded-applications

Smart, Reliable Platforms for Embedded Computing

Elma Electronic Inc.

ELMA
Your Solution Partner

Data Processing and Storage

As video streaming and processing have become ubiquitous in many applications, the demand for platforms that can handle this data is increasing. OpenVPX is proving to be a strong candidate in video processing applications; while it has made strong progress in the defense space due to its rugged design, it is also of interest for critical industrial applications. Elma has been working tirelessly with multiple companies in the VPX ecosystem to show that interoperability is feasible and contributes to a wider range of solutions and vendors, thereby reducing risk of adoption.

Storage has progressed in leaps and bounds over the past few years, making rugged, solid-state options more affordable and reliable. Elma has created an entire family of storage products to address the growing need for storing the overwhelming amounts of data being processed, such as in radar or unmanned surveillance systems or industrial inspection equipment. Requirements such as high read/write rates, removability, RAID arrays, and secure data are all supported by our storage products.

The Need for Small Form Factor (SFF) Platforms

Elma is releasing more application-specific solutions to address the growing demand for harsh environment, small form factor computing in support of industrial control, energy exploration, and infrastructure. More and more functionality must be placed in increasingly less space, making environmental issues such as heat dissipation ever more challenging. Elma is building on our 40-plus years of packaging experience to developing innovative, leading embedded solutions that address Size, Weight, and Power (SWaP) constraints in deployed environments. We are also leveraging our success in legacy standards to build intelligent tailored solutions for our customers.

High-Speed Optical a Closer Reality

Optical backplanes have been long anticipated in the embedded computer marketplace and have been seen by some as the ultimate solution for higher-bandwidth interconnections. With the release of two ANSI-VITA standards (66.x) addressing backplane optical interfaces, the practicality of backplane-based optical solutions has moved closer to reality. These standards are laying the groundwork for migration to optical technologies within the embedded computing industry.

Communications/Telephony

The increased demand for faster and greater access to mobile data is driving the wider communications markets into new infrastructure investments and small cell site networks. Elma believes that Voice over IP (VoIP) implementations will increase as companies move away from traditional analog phone systems, necessitating more equipment build-outs.

Worldwide adoption of IPv6 will also drive the need for embedded computing platforms. The last IPv4 IP addresses were assigned this past year, and as more mobile devices come online there will be an increased need for the new IP addresses, spurring the need for new equipment.

Elma sees these next-gen advancements as a positive force for embedded computing suppliers. To date, we have increased our efforts directed at this market space and have invested in new product offerings that span several PICMG standards, including an entire family of MicroTCA and new CompactPCI Serial platforms, innovative AdvancedTCA backplanes, and rugged mobile systems.

Elma Electronic Inc.
www.elma.com

A Marketplace for Software Developers

Embedded Software Store

Avnet Electronics Marketing and ARM celebrate one year of the popular Embedded Software Store.com – an online information and e-Commerce-based website dedicated solely to the embedded design community. Joining forces to help solve the industry challenge of escalating software complexity via one, simplified source, the collective efforts of these two visionary leaders provides embedded developers across the globe access to proven solutions with competitive pricing, licensing terms, support, compatibility, and interoperability.

By consolidating a large number of software options that support the ARM architecture within a single domain, the Embedded Software Store helps developers tackle the challenges presented by increased software and SoC design complexity.

Over the past year, the Embedded Software Store (ESS) has grown to include products from 24 vendors, while increasing their online partner support and product selection to nearly 1,000 products.

ESS was launched at ARM's TechCon 2011 to provide an e-Commerce website and online information targeting embedded designers. It took the industry by storm and established a comprehensive one-stop online store that enables embedded designers to evaluate and purchase leading software products to accelerate their ARM processor-based design through features such as:

- Pass-through Licensing Agreements
- Software/Hardware Associative Sells
- Order Processing via credit card or Avnet Account (Americas-only currently)
- Interactive Online Support Community
- Vendor Support

The Embedded Software Store offers developers a broad range of foundation software for their embedded design. Examples include:

- Software tools – the Embedded Software Store sells a broad range of the ARM Keil® MDK software development toolkit.
- RTOSs and the ESS offer leading RTOSs software through vendors such as Micrium, CMX, Code Time Technologies, and RoweBots. Customers have a broad choice based on functionality and memory requirements.
- Customers who are exploring open source can find a less time-consuming solution and simplified approach offered by Timesys, which provides a tool, LinuxLink, that helps a developer leverage open source code for specific silicon target boards.

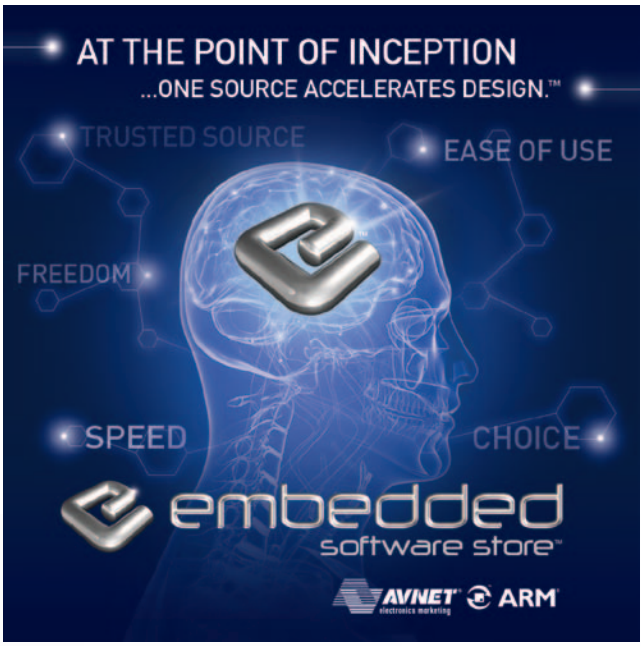

Your next great design is within reach. At the point of inception, visit EmbeddedSoftwareStore.com to collaborate, create, and connect.

Embedded Software Store
www.EmbeddedSoftwareStore.com

VIA Technologies Embraces Both x86 and ARM

VIA Embedded

By Cliff Moon, VIA Technologies

Board-Level Integration Allows ISA-Independent Interfaces

The processor wars will heat up even more during 2013, spawning fierce battles between the proponents for ARM-based systems and their instruction-set adversaries in the x86 camp. Instead of fighting the instruction-set architecture battle, VIA Technologies designs embedded microprocessors using both instruction sets to deliver a broad portfolio of chips and subsystems for embedded markets. The company's Nano, Eden, and QuadCore product families integrate VIA's 64-bit, x86 CPU core to target performance-driven applications that need full compatibility with the PC ecosystem. For applications that need the absolute lowest power consumption, VIA integrates ARM's 32-bit CPU cores into embedded SoCs. System designers now have more freedom to optimize their products to take advantage of the best aspects of either CPU architecture while working with a semiconductor company that has almost two decades of experience building boards and modules for embedded applications.

With standard interfaces such as USB, SATA, HDMI, and PCI Express, embedded system designers can implement a board-level product that integrates all of the processor-specific components into a module or single board computer. This integration allows the system to be differentiated at the software-application level, since the CPU at the heart of the system connects through the same standard hardware interfaces with full software driver support for most embedded operating systems. The system designer can select embedded boards that use either x86 or ARM, based on the requirements – or just an arbitrary preference – of the software developers. By selecting different boards or subsystems, the design can scale to meet system-level constraints for performance and power consumption. For example, as the figure below illustrates, designers using the Pico-ITX form factor (the size of a credit card) can target high performance using as many as four x86 CPU cores or target ultra low power using a single-core ARM processor.

Since the embedded processors are integrated at the board level, VIA Embedded can offer a broader range of features and performance by designing with semiconductor components from multiple vendors – not just chips from VIA Technologies. For example, the VIA VAB-800 expanded the Pico-ITX family by using a low-power Cortex-A8 (i.MX) processor from Freescale. VIA also uses its own ARM-based processors in products such as the ARM DS system for digital signage. By designing ARM and x86 processors, VIA has a much deeper understanding of microprocessor-based systems than vendors who only build boards. As the battles between ARM and x86 continue to intensify, VIA Embedded is well-positioned in both ISA camps and will continue to expand its product offerings at the silicon, board, systems, and platform levels.

VIA VAB-800

VIA EPIA-P910

VIA Embedded offers board-level products with either x86 or ARM processors, using common form factors such as Pico-ITX. For example, EPIA-M910 supports the high-performance VIA QuadCore, while the VIA VAB-800 uses a low-power ARM processor.

VIA Embedded
www.viaembedded.com

Nürnberg, Germany

we are the future

February 26 – 28, 2013

embedded world 2013

Exhibition&Conference

... it´s a smarter world

Register now for your free entrance ticket:

embedded-world.de

Programmed for success.

With some 900 exhibitors, embedded world is the top pioneering exhibition for embedded technologies worldwide. Make a note of the date now!

Exhibition organizer
NürnbergMesse GmbH
Tel +49 (0) 9 11. 86 06-49 12
visitorservice@nuernbergmesse.de

Conference organizer
WEKA FACHMEDIEN GmbH
Tel +49 (0) 89. 255 56-13 49
info@embedded-world.eu

Media partners

MEDIZIN+elektronik
Fachmedium für Elektronik in der Medizintechnik
elektroniknet.de
computer-automation.de
energie-und-technik.de
MEDIZIN-UND-elektronik.DE
DESIGN&ELEKTRONIK
KNOW-HOW FÜR ENTWICKLER
Markt&Technik
Die unabhängige Wochenzeitung für Elektronik
Elektronik
Fachmedium für industrielle Anwender und Entwickler
Elektronik automotive
Fachmedium für professionelle Automobilelektronik
ENERGIE&TECHNIK
Fachmedium für Energieeffizienz
Computer&AUTOMATION
Fachmedium der Automatisierungstechnik

NÜRNBERG MESSE

E-community Beat

Keeping the embedded conversation going

By Jennifer Hesse

www.embedded-computing.com

BLOGS | MARKET STATS | SURVEYS | VIDEOS | SOCIAL NETWORKING

ARM TechCon 2012 demo: Springboard for embedded development

To unleash the power of high-performance mobile processors like the Qualcomm Snapdragon S4 Plus APQ8060A, embedded designers need a flexible development platform that can provide a quick path to ship end products to market. Awarded the Best in Show for Hardware Design at ARM TechCon 2012, the DragonBoard Development Kit from Intrinsyc offers a complete package for starting embedded designs, including a System-On-Module (SOM) with the Snapdragon dual-core processor plus Wi-Fi and Bluetooth capabilities and a carrier board with a full complement of I/O features that can be tailored to specific application requirements. Developers can find support for designing intelligent devices and systems at the community forum http://mydragonboard.org.

Watch the DragonBoard demo at http://opsy.st/TuJV0T.

See more videos in our library: http://video.opensystemsmedia.com.

Man embraces machines in M2M-connected devices

Forget raging against the machine. In today's society, the machine, or more specifically, Machine-to-Machine (M2M) communications is becoming all the rage, with forecasts predicting the market will hit 400 million units by the end of 2017, up from more than 110 million at present, according to a new report from Juniper Research. While smart metering was once the dominant field for M2M deployment, the consumer electronics and telematics sectors are quickly becoming the anchor industries for the technology, thanks to the increased demand for embedded devices spurred by eReaders and telematics systems promising improved driver efficiency and cost management.

The price of M2M modules is expected to decline, especially for 3G modules, as automotive and consumer electronics systems require better bandwidth and latency, and as 2G infrastructure is retired in some markets. While 4G chipsets are currently shipping in low volumes, they are projected to see growth in the automotive industry and specific applications such as live video monitoring.

Read more: http://opsy.st/TJgilW

EDA Digest TechChannel

The premier issue of EDA Digest is making a big entrance in the Electronic Design Automation (EDA) scene, highlighting top companies to watch in 2012-2013 plus the industry's top innovators, whose designs serve as the foundation of EDA today and are shaping design automation for tomorrow's systems. Check out the EDA Digest Resource Guide for columns, articles, product profiles, and more at: http://opsy.st/SLME4X. And find extended coverage including videos and social media updates at the EDA Digest TechChannel at: http://tech.opensystemsmedia.com/eda.

2012 Top Innovators in EDA: Serial innovation, collaboration, and patience

By Mike Demler

One of the first observations I had when reviewing our 2012 Top Innovators in EDA was that there are a number of individuals who are "serial innovators," having contributed to several of the most significant developments in EDA over a period of many years. While an innovation may most often be thought of as a new idea, I hope you will agree that the true measure of an innovation is the impact it has on changing how things are done. That takes time, and innovators who have done that more than once stood out in the group of nominees.

Read more: http://opsy.st/TPVCTE

AMD
EMBEDDED
SOLUTIONS
R-SERIES APU

The Next Generation of Embedded APUs

Innovate Now !

Introducing the AMD Embedded R-Series Accelerated Processing Unit

- Exceptional graphics performance (up to 576 GFLOPs), enabling reliable and low power applications with amazing user experiences and smooth, vivid HD video.
- Supports AMD Radeon™ Dual Graphics technology[1] for more than double graphics performance compared to using discrete graphics alone.
- Innovative CPU architecture integrates dedicated resources that deliver exceptional performance, with shared resources that reduce power consumption and die space.
- Drive up to 4 displays[2] directly or up to 10 independent displays by pairing an AMD R-Series APU with an AMD Radeon™ Embedded discrete graphics processor or card.
- Supports a wide range of parallel compute capabilities to suit the requirements of many embedded applications.

 [1]AMD Radeon Dual Graphics technology combines the processing power of select AMD APUs and select AMD Radeon GPUs and can support displays connected to either the APU or the discrete GPU. Windows Vista® or Windows® 7 operating system required. [2]Support for the 4th display requires the use of DisplayPort 1.2 multi-streaming technologies with compatible monitors and/or hubs. The number and types of supported displays may vary by board design.

AMD