Extension MEDIA

Embedded Intel® Solutions

Winter 2016

The Powerhouse Behind the Connected Car

High-End Pharmacy POS System

Implications of the Intel Altera Acquisition

www.embeddedintel.com

www.congatec.us
6262 Ferris Square | San Diego CA 92121
Phone: 858-457-2600| sales-us@congatec.com

congatec

Always at your service!

congatec - technology for 24/7 operation.

conga-MA4

- Intel® Atom™ processor family
- COM Express Mini Type 10 module
- High resolution Intel® Gen 8 graphics

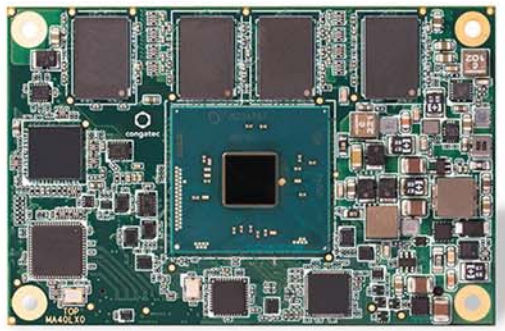

We **simplify** the use of **embedded technology**

intel IoT Solutions Alliance

ADLINK

The Sky is the Limit

Upgrade Now to 6th Generation Intel® Core™ i7/i5/i3 Processor-based Designs With Form, Fit, Function-compatible Boards

Scalability Without Limits - Accelerating Business

COM Express® Type 6 Compact Size Module

COM Express® Type 6 Basic Size Module

Mini-ITX Embedded Boards

cExpress-SL | Express-SL/SLE | AmITX-SL-G

- 4K, H.265/HEVC playback & multi-display support with Intel® Gen9 graphics
- Support up to 32GB dual channel DDR4 (non-ECC or ECC) memory
- Flexible development with cTDP
- Wide voltage range: 8.5 - 20V
- -40°C to +85°C wide temp support (optional)
- Supports device-to-cloud with SEMA* utilities

* Smart Embedded Management Agent, an intelligent middleware to monitor and control your devices.

ADLINK Technology, Inc.

▸ Toll Free: +1-800-966-5200
▸ E-mail: info@adlinktech.com
▸ www.adlinktech.com

embeddedworld2016 Exhibition&Conference ...it's a smarter world

▸ Hall : 1 Stand: 1-538 ▸ Date : 23-25 Feb. 2016

intel IoT Solutions Alliance Premier

Embedded/IoT Solutions

Connecting the Intelligent World from Devices to the Cloud

Long Life Cycle · High-Efficiency · Compact Form Factor · High Performance · Global Services · IoT

IoT Gateway Solutions

E100-8Q

Compact Embedded Server Appliance

SYS-5028A-TN4

Network, Security Appliances

SYS-5018A-FTN4 (Front I/O)

High Performance / IPC Solution

SYS-6018R-TD (Rear I/O)

Cold Storage

SYS-5018A-AR12L

4U Top-Loading 60-Bay Server and 90-Bay Dual Expander JBODs

Front and Rear Views

SC946ED (shown)
SC846S

embeddedworld2016 Exhibition&Conference
it's a smarter world
Nurnberg, Germany
February 23 - 25th
Come visit us @ Hall 1, Booth #1-330

- Low Power Intel® Quark™, Intel® Core™ processor family, and High Performance Intel® Xeon® processors
- Standard Form Factor and High Performance Motherboards
- Optimized Short-Depth Industrial Rackmount Platforms
- Energy Efficient Titanium - Gold Level Power Supplies
- Fully Optimized SuperServers Ready to Deploy Solutions
- Remote Management by IPMI or Intel® AMT
- Worldwide Service with Extended Product Life Cycle Support
- Optimized for Embedded Applications


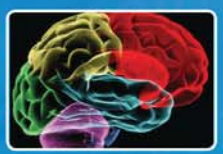



intel XEON inside™

TITANIUM

We Power the Cloud

80 PLUS TITANIUM

Learn more at www.supermicro.com/embedded

Sounding Off About Silence

Welcome to the Embedded World edition of Embedded Intel Solutions.

By Anne Fisher, *Managing Editor*

Sometimes you want silence. Sometimes you don't. Bernd Blume, CEO of autonomous Web hosting service company antsle, is in the "don't" camp when it comes to fan noise. He loves that the Intel® Atom™ Processor C2000 Series product family has just 20W of TDP. That figure has allowed antsle to offer a server with a fanless design producing 0.0 dB. But the TDP feature wasn't the only reason antsle went with Intel®, as Blume's article in this issue explains.

When Communication is Vital

It's not a stretch to associate emergencies with silence's opposite—noise. But while emergencies are not short on noise in general—alarms, cries for help—they can be dangerously low on the "noise" that matters: communication to and among first responders, or among warfighters in a battlefield. Read how the problem of silence when communication is vital is being addressed in the article by Embedded Intel Solutions senior editor Caroline Hayes in this issue: "Communication and Co-operation Create a Tactical LTE Network."

Hayes details how Intel IoT Solutions Alliance Premier Member ADLINK Technology, along with Connect Tech and LCR Embedded, have developed a "complete LTE network in a box." This "wearable" took advantage of off-the-shelf products to trim design time.

Hayes also shares with us in this issue the possibilities Intel and Ford Motor Company see for making our experience as drivers more intuitive. That's one of the topics covered in her article, "Driving Along at Full Speed for Autonomous Vehicles."

With the holidays over, you may feel yourself hurtling along at full speed, trying to keep up with all the exciting developments in embedded. Other developments covered in this issue are the ramifications of Intel's acquisition of Altera--Is an FPGA that behaves like a general-purpose CPU headed our way?--and the IT department's role in new IoT business opportunities.

You'll find more news, videos, white papers and articles to keep you up to speed among the more than 30 EECatalog.com technology channels. And we definitely don't consider silence golden when it comes to our readers. Let us hear from you!

Anne Fisher is managing editor of EECatalog.com. Her role brings her the opportunity to learn about a wide range of embedded solutions for IoT, Mil Aero, Industrial, Sensors and several other fields. She welcomes articles describing innovative approaches to embedded engineering challenges to share with the EECatalog.com readership, as well as opinion and analysis from industry leaders. She can be reached at afisher@extensionmedia.com

Embedded Intel® Solutions

www.embeddedintel.com

Vice President & Publisher
Clair Bright

Editorial

Editor-in-Chief
Chris Ciufo
cciufo@extensionmedia.com

Managing Editor
Anne Fisher
afisher@extensionmedia.com

Senior Editors
Caroline Hayes
David Bursky
Gabe Moretti

Creative/Production

Production Manager
LS Jerrett

Graphic Designers
Nicky Jacobson
Simone Bradley

Media Coordinator
Jon Franke

Senior Web Developers
Slava Dotsenko
Mariam Moattari

Advertising / Reprint Sales

Vice President, Sales
Embedded Electronics Media Group
Clair Bright
cbright@extensionmedia.com
(415) 255-0390 ext. 15

Sales Manager
Elizabeth Thoma
ethoma@extensionmedia.com
(415) 255-0390 ext. 17

Marketing/Circulation
Jenna Johnson

To Subscribe
www.embeddedintel.com

Extension MEDIA

Extension Media, LLC
Corporate Office

President and Publisher
Vince Ridley
vridley@extensionmedia.com
(415) 255-0390 ext. 18

Vice President & Publisher
Clair Bright
cbright@extensionmedia.com

Vice President, Business Development
Melissa Sterling
msterling@extensionmedia.com

Human Resources / Administration
Darla Rovetti

Special Thanks to Our Sponsors

ADLINK TECHNOLOGY INC. · ARTESYN EMBEDDED TECHNOLOGIES · congatec · Portwell · SUPERMICRO · rtd · WinSystems

Embedded Intel® Solutions is sent free to engineers and embedded developers in the U.S. and Canada who design with embedded Intel® processors.

Embedded Intel® Solutions is published by Extension Media LLC, 1786 18th Street, San Francisco, CA 94107.

Intel® Atom™ Processor E3800 Series-based
Single Board Computers
Fanless -40° to +85°C Operation

Intel® Atom™ Processor E3800 Series-based
Industrial Computer Systems
Fanless -40° to +85°C Operation

Intel® Atom™ Processor-Based PC/104
Single Board Computers
Rugged, Stackable Form Factor

Single Board Computers
COM Express Solutions
Power Supplies
I/O Modules
Panel PCs

Accelerate Your Product Development Cycle

Speed up time-to-market with embedded solutions from WinSystems. Our industrial computers include expansion options, so customers can expedite prototyping and integration without the burden of CPU or carrier board design. These proven hardware platforms also provide the building blocks to create customized, application-specific designs. Products are in stock and available for immediate shipment from our Arlington, Texas facility.

Let our factory Application Engineers accelerate your capabilities.

embeddedworld2016

Come visit us in Hall 2 - Booth #644

715 Stadium Drive | Arlington, Texas 76011
Phone: 817-274-7553 | Fax: 817-548-1358
www.winsystems.com | info@winsystems.com

intel IoT Solutions Alliance

ISO 9001:2008 CERTIFIED

WinSystems
The Embedded Systems Authority

Embedded Intel® Solutions

WINTER 2016

DEPARTMENTS

SPECIAL FEATURES

PRODUCT SHOWCASE

On the Cover:
In-Vehicle Infotainment (IVI) applications, such as digital radio, Internet, DVD video and navigation systems, have evolved from novelties to "must-have" options for many car buyers. Consumers are drawn to cutting-edge features, such as hands free phone calls and voice activated navigation systems, which lead to increased safety and convenience. Photo courtesy of Intel.

Innovative Solutions

RTD's Embedded Systems and Enclosures

At RTD, we have developed a full suite of compatible boards and systems that serve transportation, railway, industrial, defense, aerospace, maritime, ground, and research-based applications. We provide high-quality, cutting-edge, concept-to-deployment, rugged, embedded solutions.

Whether you need a stack of modules, or a fully enclosed system, RTD has a solution for you. Call us to leverage our innovative product line to design your own embedded system that is reliable, flexible, expandable, and field-serviceable. Let us show you what we do best. Start here: www.rtd.com/systems

This scalable off-the-shelf system features an Intel® Core™ i7 processor-based SBC and a synchronous RTD power supply. Also available in watertight configurations with locking cylindrical connectors.

AS9100 - ISO 9001 CERTIFIED

DESIGNED AND MADE IN THE UNITED STATES

rtd®

www.rtd.com | **AS9100 and ISO 9001 Certified** GSA Contract Holder | **sales@rtd.com**

Sixth Generation of Intel Core Processors: New Options for Industrial PCs and Embedded

The sixth generation of the Intel Core i CPU family's innovations has benefits for demanding digital signage, medical and industrial automation applications.

By Peter Hoser, Fujitsu

Intel® has launched the sixth generation of its Core™ i CPU family (formerly Skylake) and the corresponding 100 Series chipsets. The first mainboards for traditional desktop PCs have already been available and are now followed by boards for the industrial sector.

Improved hardware features for continuous operation in extended temperature ranges, extensive documentation and software tools, and special lifecycle management distinguishes these latest industrial motherboards. They can be used in many ways, for example in digital signage, industrial automation, POS or in medical technology.

Longevity and Easier Cooling

The DDR4 memory is supported with 2133 MHz. Although the DDR3 memory can still be used with a maximum clock speed of 1600 MHz under the Skylake platform, for industrial customers in particular a change makes sense: DDR4 brings significantly enhanced performance and increases longevity. And DDR4 pricing stands to become more affordable as it establishes itself as a standard technology in the first half of 2016.

Figure 1: Boards developed to target the demands of industrial operations can now take advantage of the sixth generation of the Intel Core i CPU family.

Instead of the LGA 1150 the CPU socket is now an LGA 1151. After the 5th generation Core processors (formerly Broadwell) the Skylake processors belong to the second generation, which uses a fine 14-nanometer structure. The 14nm structure makes it possible to fit more circuits into the same amount of space while reducing switching distances and times. Thus, the CPU can work faster and consume less power.

Industrial customers can use boards with less energy consumption in a higher temperature range. Also, cooling is easier to handle. Furthermore, Intel has revised the functional units on the processor. The cache, for example, can now hold up to 8 megabytes.

Another innovation is the HD 530 graphics. The new graphics engine also promises high performance when playing and editing videos, and it features low energy consumption. For example, the processor uses up only one watt when playing a video in MP4 format.

New Chipsets Specially Designed for Embedded

For embedded systems with a guaranteed availability of seven years Intel offers three chipsets with different functions and price levels. The C236 chipset supports the server processors of the Intel Xeon® E3-1200 V5 processor series. These offer, compared to the previous Intel Xeon E3-1200 V3, a 32-gigabyte Video Memory and DDR4 memory, compared to 16 gigabytes of video memory and DDR3 memory. There are also significantly more USB 3.0 and 2.0 ports and PCIe 3.0 lanes. For the desktop variants of the sixth generation of the Intel Core i7 / i5 / i3 processors the chipset supports four sockets for DDR4 memory with Error Correction Code (ECC). The out of band management is offered by Intel vPro™ Technology and

INTERNET OF THINGS
DEVELOPERS CONFERENCE

May 25-26, 2016
Hyatt Regency Santa Clara, CA

Resolving the Technical and Business Challenges of Getting Connected to the Internet of Things

Two full days of keynotes, training sessions and panel discussions. Topics covered include:

- What you need to know to Secure IoT Devices
- Meeting the Reliability and Security Demands of the Industrial-IoT
- Optimizing Industrial Embedded Implementations
- Build Lower Power and more Secure IoT Devices
- What does Interoperability mean for IoT
- Wide-Area Wireless Technology for IoT and M2M Devices
- The Internet of Things that see: Bringing Visual Intelligence to Embedded Devices
- Network Performance, Reliability and Metrics
- IoT as a Service
- Create Products with Provably High User Value
- IoT Standards
- Embedded Cloud Computing
- Semiconductor Technologies of the Future for IoT

Register now. Come to the **Internet of Things Developers Conference 2016!**

IoT-devcon.com

Intel Active Management technology. Due to its advanced software stability the Intel Stable Image Platform Program (SIPP) guarantees no changes for the chipset and at least 15 months of platform stability for drivers. Up to three independent displays can be used with the Intel C236 chipset.

Figure 2: Fujitsu's D3446-S mainboard. The author notes that industrial customers will find migrating to motherboards with the new Intel chipsets a good move, thanks to robust performance and functionality.

The Intel Q170 chipset is almost identical to the Intel C236 chipset in terms of functionality, but doesn't provide Xeon and ECC support. For lower requirements the Intel chipset H110 is available. It also supports only the sixth generation Intel Core i7 / i5 / i3 processors. Four SATA ports and two DDR4 bases, also without ECC support, can be used. In addition, the PCIe 3.0 is only available for the x16 slot. The Intel H110 chipset supports two independent digital displays.

Limitations

Despite all the changes, there are also certain limitations of Skylake platforms. As with the fourth generation of the Intel Core processor platform PCI support is eliminated. Applications needing PCI support will need to spring for an additional PCI Bridge Controller. Also, in the 6th generation Intel Core processors analog graphics (VGA) is omitted. Accordingly, customers have to resort to adapters or VGA cards.

The Intel 100 Series chipset, unlike the company's fourth generation Core processors, does not allow on all desktop chipsets the use of Intel Xeon processors. The use of Intel Xeon processors requires appropriate server or workstation mainboards with server chipsets such as the Intel C236.

Also, the USB Enhanced Host Controller Interface is omitted. This means that with a Windows 7 installation USB ports no longer work. Thus, the operating system can only be installed on a SATA optical drive and requires a PS / 2 keyboard and mouse. The support of the SATA IDE mode is not continued. Additionally the LPC DMA mode isn't supported for the parallel port.

Future-Proof Investment

For industrial customers, the transition to motherboards with the new Intel chipsets is worthwhile. The new platform is durable and offers state-of-the-art performance and functionality. As soon as the new 6th generation of Intel Core processors and technologies of DDR4 memory has established itself as a standard its price will be cheaper compared to DDR3 memory. The described limitations must be taken into consideration, but can usually be bypassed easily. Therefore the conversion to the new platform as a whole is a future-proof investment for businesses.

Peter Hoser is Director OEM Mainboard Sales, Fujitsu. Hoser has worked for almost 30 years for Fujitsu and all its predecessors such as Siemens, Siemens-Nixdorf and Fujitsu-Siemens Computers. His experience includes five years as developing engineer for various production equipment and another five years as a process engineer in mainboard production. In 1996 he started to develop the sales of mainboards designed and made in Augsburg/Germany. During 20 years of heading the mainboard sales Hoser established Fujitsu as the number 1 vendor of industrial mainboards within EMEA.

Case Study: High-End Pharmacy POS System

Why low-power 4th Gen Intel Core U-Series processors on a Mini-ITX motherboard are the right Rx for a hard-working and versatile pharmacy POS terminal.

By Victor Huang, DFI

Figure 1: High computing capability and rich I/O connectivity are among the requirements for a high-end pharmacy POS that handles tasks such as patient ID validation, prescription printing, keeping track of pharmacy stock, and merchandise promotion in a limited space.

This article describes how a Mini-ITX board with Intel® 4th generation Core™ U-Series processors makes possible a flexible and innovative customizable POS terminal.

The Challenge

Due to the growing importance of smart healthcare applications, improving customer satisfaction, and brand loyalty, a reliable and multifunctional POS system (which optimizes the counter spaces) is becoming an increasingly important element (Figure 1). Thus, in order to stay ahead of the competition and to deliver more valuable and expanded services, pharmacies need a solution to help with business management while improving quality of service for their customers.

From Patient ID Validation to Interactive Videos

The pharmacy POS system requires a high computing capability and a rich I/O connectivity to achieve a wide variety of tasks that covers patient ID validation, prescription printing, stock maintaining and merchandise promotion in a limited space. The all-in-one design of a POS terminal that supports a P-Cap touch screen monitor would be a preferable choice for allowing the integration of options such as: MSR, RFID reader, thermal printer and barcode scanner into one compact machine—one which would make maintenance easy for users. In addition, a secondary display would be needed to broadcast interactive promotional videos and advertisements for waiting customers.

Thin, Energy Efficient and Generous I/O

After evaluating the requirements mentioned above, DFI, which designs and manufactures board- and system-level products for embedded applications requiring strict revision control and long life availability, took on the challenge.

DFI's HU101 Mini-ITX motherboard (Figure 2) meets all the requirements for the pharmacy POS application. This thin and energy-efficient embedded board is fueled by the low-power 4th Gen Intel Core U-Series processors at only 15W TDP, and comes with rich I/O interfaces that include: 2 Gigabit Ethernet ports, 2 USB 2.0 ports, 4 USB 3.0 ports, and 4 COM ports.

The USB ports connect the MSR to the thermal printer to quickly access the customer's information and to immediately print out the prescription. The serial ports support the barcode scanner (for medication scanning) to control the inventory and reduce out-of-stock situations. HU101 is also equipped with 2 HDMI ports and 1 LVDS display port with high resolution for real-time dynamic advertisements that will increase the appeal of product promotions and enhance their impact. Moreover, DFI is able to provide BIOS/OS customization services tailored to customers' unique requirements.

Figure 2: The HU101 Mini-ITX motherboard from DFI, which uses 4th Gen Intel Core U-Series processors, can show pharmacy customers real-time, high-resolution dynamic video.

Victor Huang is a product manager at DFI, which was established in 1981, and is a leading supplier of high-performance computing technology worldwide. DFI focuses on innovative design and manufacture of leading-edge board- and system level products for embedded applications requiring strict revision control and long life availability. The company uses the latest technology platforms and manufacturing techniques to produce cost-effective products for use in medical diagnostic and imaging, ATM/POS, industrial control, kiosk, security and surveillance, digital signage, gaming, and other embedded applications.

Unraveling the Implications of the Intel Altera Acquisition: Designing for Tomorrow, Today

Will FPGA programmability give pause to developers considering GPGPUs?

By Austin Hipes, UNICOM Engineering

> "FPGAs can potentially be used to program any number of specific, specialized tasks into many types of electronics, paving the way for Intel's continued expansion into the Internet of Things market."

The rapidly evolving world of computing requires Integrated Circuits (ICs) that can continue to keep up with an ever-changing list of demands. From virtualization, storage management, analytics and beyond, software developers and OEMs are constantly balancing the needs of the current landscape with the promise of future technologies.

That's why Intel®, one of the world's principal processor manufacturers, recently acquired Altera®, a leading company with a 30-year history of supplying the industry with the latest programmable logic, process technologies, IP cores and development tools. This exciting merger has the potential to bring a new dimension of speed and efficiency into the world of information technology.

Paving the Way for IoT Expansion

As one of the top manufacturers of FPGAs, Altera makes highly programmable chips that allow tasks to be programmed into the IC's body. The result is chips that can quickly and easily be upgraded, reprogrammed or reconfigured. Used everywhere from automotive infotainment systems to high-end network switches, FPGAs can potentially be used to program any number of specific, specialized tasks into many types of electronics, paving the way for Intel's continued expansion into the Internet of Things (IoT) market.

FPGAs are also well suited for many cloud-based and specialized high-performance computing applications that are currently implementing general-purpose graphics processing unit (GPGPU) technology. The programmability of FPGAs offers an increasingly attractive performance advantage over GPGPUs—with the potential to impact everything from traditional IT platforms to scientific computing applications that conduct advanced processes such as genome sequencing and seismic analysis.

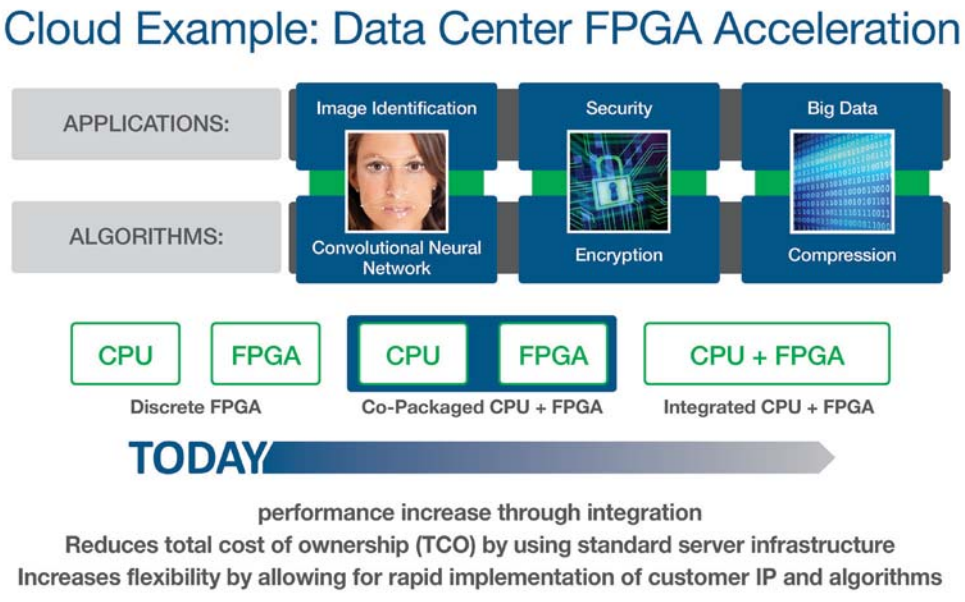

Figure 1: According to the author, Intel's acquisition of Altera could give the information technology sector greater efficiency. Figure courtesy UNICOM Engineering.

In other words, processors with FPGAs can be much faster, have low latency and are more power-efficient when performing specific workloads.

Programmable Workloads and Programmable Fabrics

Intel and Altera have been working together for many years, creating technology that has stood the test of time. This merger comes at a critical time for Intel, as IT professionals demand a processor whose algorithms can be reprogrammed as fast as the technology parameters evolve. As a result of the acquisition, Altera will move from a manufacturing partner to a vital component of the Intel technology family.

By combining the intellectual property of Altera and Intel, the reality of getting an FPGA much closer to a general-purpose CPU becomes more likely. The two main advantages of this include the potential for programmable workloads and programmable fabrics.

Consider, for example, incorporating FPGAs onto the same die as the central processing unit (CPU) itself. In this scenario, a general-purpose Xeon® class CPU could potentially be paired with a programmable FPGA in the same processor package. Sharing workloads makes it possible to share low-latency cache and the same memory bus. Also enabled: high-speed data transfers between the FPGA and the CPU.

Programmable workloads that can benefit from this type of arrangement include optimized search algorithms, image pattern recognition and encryption acceleration (Figure 1). Because FPGAs can be reprogrammed, IT professionals can also adapt their hardware for new workloads over time as their needs or algorithms change.

The combined technology also has a potential impact for IT professionals who require full general-purpose CPU power and want full high-performance FPGAs. Today, Intel connects CPUs and multi-processor systems via a quick path interconnect (QPI), which features a very low latency CPU-to-CPU fabric interface. With FPGAs now included in the family, Intel could potentially make FPGAs that utilize QPI to communicate directly with Xeon® class processors.

For example, let's look at a traditional dual-processor server featuring one general-purpose CPU and a second IC that is an FPGA. The CPU and FPGA in this example can share access to the same memory, while also having their own dedicated memory channels—with the potential of each having their own fabrics to connect to external IO. The fabric channel going into the FPGA could offload much of the initial packet processing workload and pass the remaining data to the general-purpose processor for further work.

This is an interesting scenario not only in dual-processor systems, but also in quad-processor systems where there is more freedom to mix programmable workloads with general purpose CPUs. The possibility of four or eight socket systems with different combinations of CPUs and FPGAs featuring QPI has the potential to substantially reduce latency, as well as address many of the issues that slow down FPGAs in general-purpose systems.

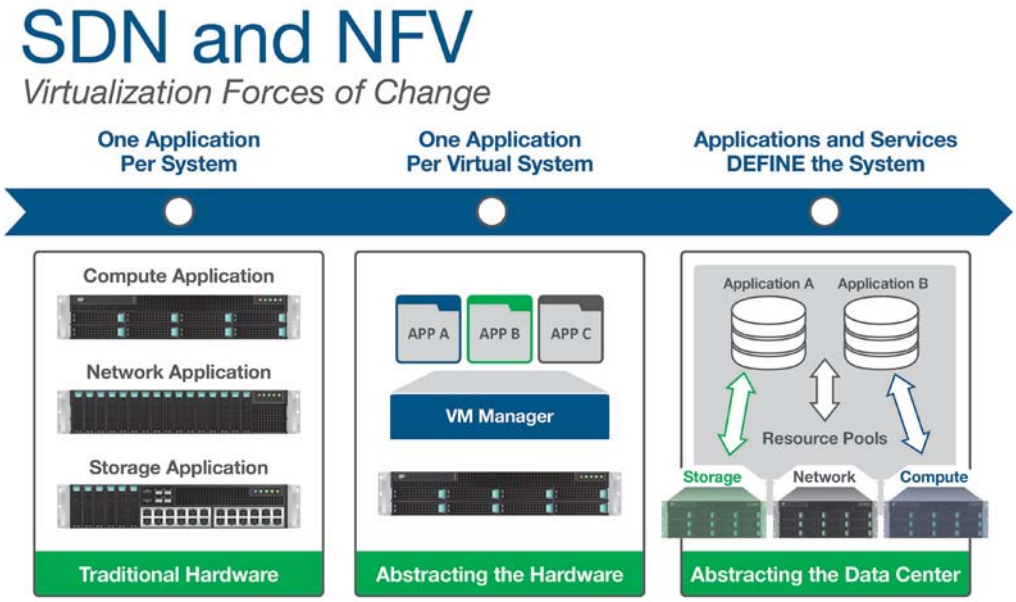

Figure 2: Data centers can decrease the amount of physical equipment they require by having network functions reside on the same physical server platforms as virtual machines (VMs) while logically separating the network functions from their own virtualized workloads. Figure courtesy UNICOM Engineering.

Migration to Network Function Virtualization

For software developers and OEMs serving storage, security and communications markets worldwide, the Intel Altera acquisition has several potential implications. An important area that this may have an impact on is virtualization. Intel multicore processors are a key asset in the transition to virtualization (Figure 2).

From a consumer perspective, virtualization offers well-known benefits such as reduced fixed operational costs and capital expenditures. The problem is that many technology providers are trying to capitalize on virtualization technology without drastically increasing their expenses in the process.

As server platforms continuously support more CPU cores and memory, their ability to support more virtual machines (VMs) per platform also increases. Normally as the number of VMs in a data center increases, so does the need for specialized network equipment, such as firewalls, routers and load balancers, to allow these VMs to communicate with each other and the world outside the data center in a reliable, secure manner. Now, however, we are beginning to see a shift from the use of dedicated network equipment to Network Function Virtualization (NFV).

In simple terms, NFV is a concept where a network node, such as a load balancer or router, is virtualized and runs on traditional IT servers in much the same way that a VM is utilized for a general purpose server workload. By combining many network functions on the same physical server platforms as VMs, but logically separating them from their own virtualized workloads, less physical equipment is needed in the data center. This can save capital expense, space and power, as well as increase network uptime via live migration of services from one NFV VM to another.

Since many proposed NFV functions rely on fixed algorithms for things like switching, routing, compression, load balancing and encryption, using FPGA logic to accelerate these functions makes sense, assuming it is available to the host platforms. By combining FPGA logic with general-purpose CPUs, Intel will be well positioned to take advantage of this market shift by allowing data centers to optimize NFV workloads on servers also able to host more generalized VMs.

Sustainable Solutions

How can software developers design products and business solutions that solve today's deployment challenges and increase business efficiencies, while planning for tomorrow's technological advances? The answer rests in partnering with an application platform and lifecycle support services partner that possesses expertise in solution design, system integration, application management and maintenance services, as well as one that maintains relationships with industry-leading technology providers like Intel. Such a partner understands and can prepare for technological advancements like those promised by the Intel Altera acquisition, and can initiate those changes when the time is right. This partner can provide recommendations for designing sustainable solutions that account for future scenarios, such as processors with FPGA chips, to deliver much faster, low latency and power-efficient IT solutions when the time is right.

Austin Hipes currently serves as the chief technologist and VP Engineering for UNICOM Engineering, a division of UNICOM Global. In this role, he manages UNICOM Engineering's hardware, software, mechanical, and sales engineering groups, as well as supports sales design activities with key customers. Over the last fourteen years, Austin has been focused on designing systems for telecommunications equipment providers, storage solution providers, and network security companies requiring appliance solutions. He was previously VP of Technology at NEI and Director of Technology at Alliance Systems. Austin studied Electrical Engineering at University of Texas at Dallas.

3 Ways Linear Redrivers Differ from ReDrivers

For PCIe 3.0, linear redrivers do the best job walking the line.

By Chris A. Ciufo, Editor-in-Chief, Embedded; Extension Media

Designers know that redrivers are an essential part of solving signal integrity challenges in gigabit systems. But at speeds up to 8 GT/s in PCI Express 3.0 systems, linear redrivers are needed to also comply with PCI-SIG specs and with chipset vendors' datasheets. That is, it's not enough to open the signal eye—the redriver also needs to precisely match PCI Express Link Training waveforms while offering designers the ability to "tweak" channel characteristics (Figure 1).

OPEN EYE

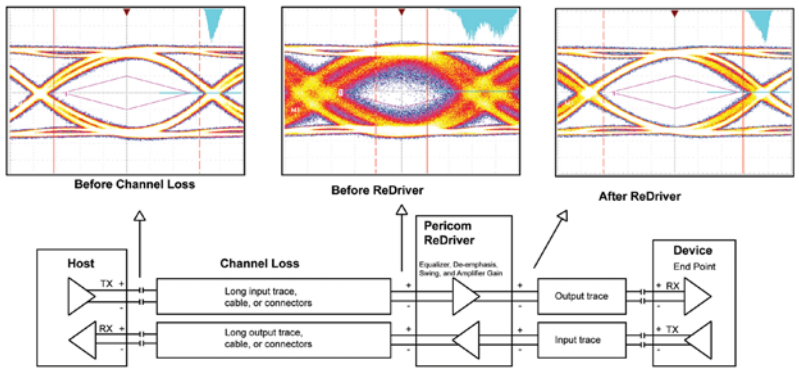

Figure 1: Redrivers amplify and shape high frequency signals to compensate for channel "challenges."

Redrivers Revisited

The continuum of super high-speed, serial communications standards is visually shown in Figure 2, a graphic from integrated circuit (IC) connectivity supplier Pericom Semiconductor. We have previously covered the issues with and mitigation strategies for dealing with signal integrity (SI) in serial standards in the article "Signal Integrity in a GHz World." The reader is encouraged to refer to that article for a refresher.

Here, this article builds on the previous material by arguing that while redrivers absolutely solve the problems as described, as serial standards get faster—and IC chipsets and system design specifications evolve even more—there is a need for a new kind of redriver.

The amplifier portion of the linear redriver is designed to have a "more linear" transfer function (Vout/Vin) than the limiting redriver (Figure 3). While both limiting and linear redriver amplifiers are analog in nature, their gain slopes are very different. This is necessary to comply with Link Training protocols now becoming essential in PCI Express 3.0 and related Gen 3.x standards.

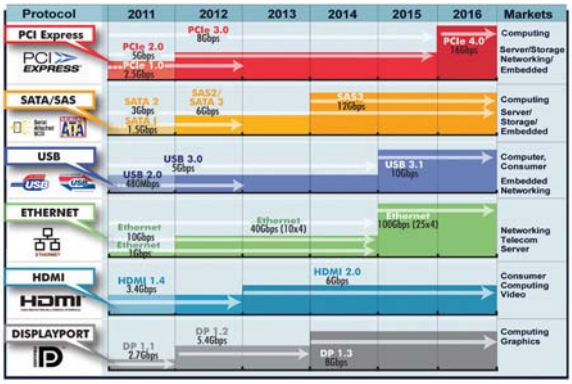

Figure 2: Serial communications standards just keep getting faster. Along with signal integrity challenges also come specification compliance issues such as PCI Express Link Training. (Courtesy: Pericom Semiconductor.)

What is Link Training?

Driven by Intel® and introduced with PCI Express 3.0, Link Training is part of the complicated in-band PCIe enumeration sequence that not only "connects" Source to Destination nodes, but also determines the ideal speed for the channel. Since it's a given that the complete channel from Source to destination will experience signal integrity (SI) issues (refer to "Signal Integrity in a GHz World"), Gen 3 seeks to optimize the transfer speed and minimize bit error rate (BER) based upon the channel's capability.

Note that "channel" refers to the sum of the complete interconnect and boundaries between the Source IC's pins and the Destination IC's pins or leads. For all practical purposes, this includes PCB traces and vias, board connectors, cabling, backplanes and any other electrical path for the 8 Gbps signals. In

fact, the point of Link Training is to verify if the channel is capable of this speed or if the speed should be reduced to PCIe Gen 2 speed—or less.

The Link Training protocol passes a series of waveforms from the transmitter (Source) to the receiver (Destination) across the channel according to a pre-defined state machine sequence. If all of the Link Training protocol waves arrive at the Destination as expected, the channel is deemed capable of full speed. As waves degrade or are not received as expected, this infers information about the channel's speed capability for Gen 3, Gen 2, or even Gen 1 (the slowest) PCIe speed.

Link Training signals use various pre-emphasis and pre-shoot waveforms to represent the different presets. For example, Table 1 shows how the various Link Training Tx signals contain pre-shoot and de-emphasis, which translate to emphasized portions of the input waveform as seen by the linear redriver (Figure 3). Because of these intentional waveform characteristics, the linear redriver must reproduce these various waveforms without any changes—such as limiting (clamping) the signal—which would destroy the Link Training waveform as it passes through the redriver.

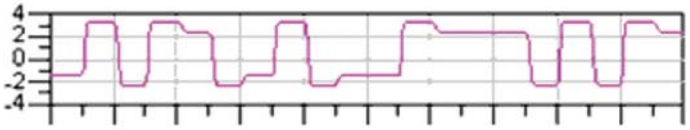

Figure 3: This is an example of a Link Training Tx waveform corresponding to preset P03 with 2.5dB of pre-emphasis.

Preset	Preshoot	Deemphasis
P0	0.0 dB	-6.0 dB
P01	0.0 dB	-3.5 dB
P02	0.0 dB	-4.4 dB
P03	0.0 dB	-2.5 dB
P04	0.0 dB	0.0 dB
P05	1.9 dB	0.0 dB
P06	2.5 dB	0.0 dB
P07	3.5 dB	-6.0 dB
P08	3.5 dB	-3.5 dB
P09	3.5 dB	0.0 dB
P10	0.0 dB	-9.5 dB

Table 1: Examples of Link Training waveforms and associated pre-shoot and de-emphasis. A linear redriver is needed to precisely reproduce the waveform characteristics. Source: PCI-SIG.

Any passive or active impedance, interference, or device in the channel has the potential of adversely modifying these waveforms. While the channel itself "is what it is," the redriver added to the channel to fix signal integrity must at the same time not modify the waveforms in any way!

Therefore, while non-linear redrivers regenerate high-speed signals to open the signal eye to compensate for channel SI effects, linear redrivers precisely reproduce high-speed signals and open the signal eye. With linear redrivers, all of the characteristics of each Link Training waveform are completely reproduced, including amplitude. Hence, the linear redriver compensates for SI challenges without affecting the Link Training protocol waveforms. The difference between linear and non-linear redrivers is shown in Figure 4.

As mentioned above and shown in Figure 4, you can see that the limiting amplifier has much higher gain so that the output signal is driven to its maximum more quickly from the input signal, resulting in cutting off the input waveform. A useful analogy is the compressor/limiter in the audio world, that clamps the output voltage to some limit regardless of how hard the input signal is driven. The linear redriver, on the other hand, has a lower gain, no limiting threshold and therefore allows the output to track the input without clamping the signal to some limit.

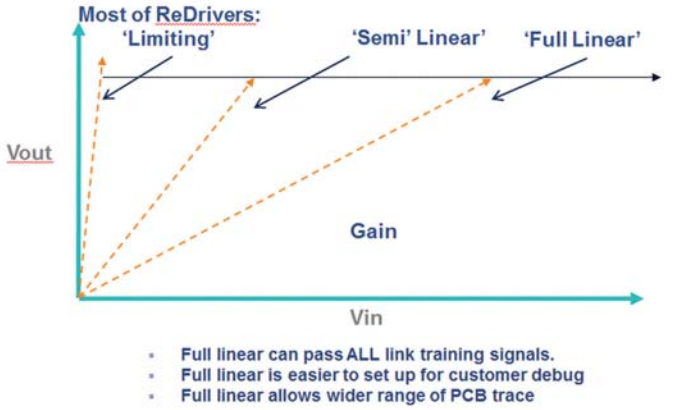

Figure 4: The difference between linear and regular (non-linear) redrivers. Non-linear redrivers compensate for signal integrity problems, but regenerate signals that may not pass PCI Express 3.0 Link Training waveforms.

Waveform Details: Comparison

It's important to understand that while redrivers will "fix" SI problems, they do so by regenerating signals as shown in the Figure 5 graphs provided by Pericom Semiconductor. Notice that the top waveforms in Figure 5 (non-linear redrivers) show some pre-emphasis waveform attenuation. This is just fine for opening the signal eye downstream; however, changing the waveform even just a little bit can cause PCIe Link Training to fail.

The result may be that the channel is deemed insufficient for PCI Express Gen 3 speeds (8 Gbps) and the Source/Destination agree to communicate a slower Gen 2.0 (5 Gbps) speed. Conversely, the bottom waveforms in Figure 5 show how a linear redriver precisely reproduces the input waveform by accurately tracking the source signal. The channel would pass

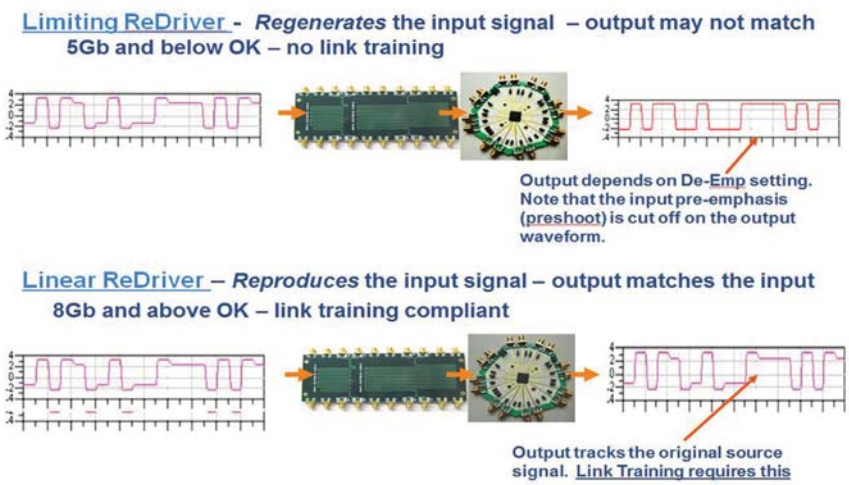

Figure 5: Redriver vs linear redriver waveforms, according to a Pericom Semiconductor lab test.

Link Training, and the devices would communicate at the fastest possible channel speed.

Compatibility and Interoperability Testing

The notional waveforms shown in Figure 5 illustrate the differences between a redriver—sometimes referred to as a limiting redriver—and a linear redriver. Figure 6 shows the logical block diagram differences. In the linear case, the final analog amplifier stage drives the differential channel directly and provides channel equalization, whereas the more common redriver adds a de-emphasis (DE) stage after the amplifier. The DE stage is needed to add back the emphasis that was limited (chopped off) at the input side of the limiting redriver.

What's important with linear redrivers is the accuracy with which they reproduce PCI Express Link Training waveforms. This is measured by plugfests organized by the PCI-SIG and by tests conducted with major PCIe chipset vendors. The details of these tests are often not public; suppliers will only say that their linear redrivers have passed PCI-SIG workshops and that their ICs have been tested by IC vendors.

Pericom Semiconductor, a signal integrity and connectivity IC supplier, has linear redrivers designed not only for PCI Express 3.0 at 8 Gbps, but linear redrivers for 10Gbps Ethernet (KR), USB 3.1, DisplayPort 1.3, Thunderbolt, HDMI 2.0, and SAS 3.0. In the case of PCIe 3.0—the subject of this article—Pericom's redriver family has "passed all the tests for Link Training as measured by a very large PCI Express semiconductor company," says Bill Weir, Sr. Director of Marketing for Pericom.

According to Weir, the company spent many months conducting simulations prior to going into the lab and measuring compliance to Link Training protocol waveforms. The company shares this confidential test data with customers when requested.

The company's PI3EQ8904 and PI3EQ8908A linear redrivers are PCIe 3.0 compliant at 8 Gbps and support four PCIe lanes ('8904) and eight lanes ('8908A), but can also be used at slower PCIe speeds and Gen 1/2/3 protocols when the channel so necessitates. Beyond PCI Express compliance, there are other nice-to-have attributes that designers should consider.

The "Nice-to-Haves"

Beyond the basics—improving SI and compliance with PCI Express Link Training protocol waveforms—eval boards, simulation suites, and programmable features and applications support to assist customers in the debug and validation process are very desirable.

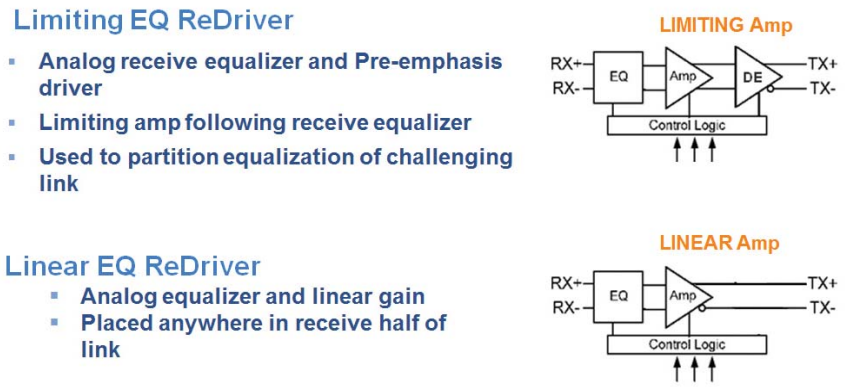

Figure 6: The block diagram difference between regular (limiting) redrivers and linear redrivers. The final analog stage makes all the difference when PCI Express Link Training is required. (Courtesy: Pericom Semiconductor.)

According to Pericom's Weir, the typical customer process is:

1. Simulation of the redriver in the customer channel, typically using S Parameter models. Pericom fully supports this important phase.

2. Use an eval board or customer prototype PCB board with redrivers to do initial lab testing of the actual application.

3. Debug and validation once final customer PCB platform is back and ready to finalize redriver settings for release to production. "Settings" means programmable IC capabilities.

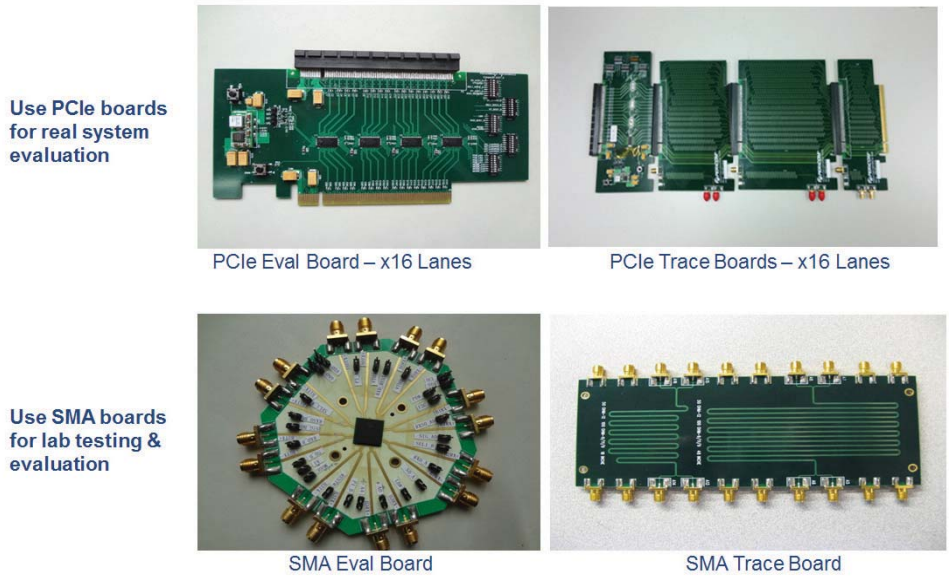

Figure 7: Examples of linear redriver and PCI Express evaluation boards. (Courtesy: Pericom Semiconductor.)

Eval boards allow designers to place a manufacturer's linear redriver into real world scenarios so designers can "tune" their own circuit boards, backplanes, and complete system design. Eval boards can include just the chosen redriver, or might include PCIe traces for "what if" in-system EMI and channel noise tests. The latter might be useful to see if adjacent oscillators or RF transceivers affect a PCIe channel design. Several examples of eval boards from Pericom Semiconductor are shown in Figure 7.

Other nice-to-haves include per-channel programmability of input/output termination, receiver equalization, output swing and overall gain. On the latter, the amplitude of the waveform can be boosted to compensate for channel losses, but the waveform shape remains the same. This assures that the linear redriver doesn't violate Link Training waveforms.

On the subject of programmability, a customer and cost friendly feature offered by Pericom is two-way programming. The linear redriver's attributes can be programmed via local pin-strapping, or via an I2C interface. On the latter, the redriver can be selected for either slave or master mode. In master mode, up to 16 redrivers can be programmed from a low cost EEPROM during boot up. The advantage to this programming method is that the CPU BIOS does not need to be changed in order to program the redrivers, which reduces software complexity and speeds up the design process.

As well, it's often handy to have flow-through pinouts where all inputs are on one side of the device with the outputs on the other to reduce crosstalk and also provide compatibility to double-sided PCIe connectors, such as on a riser card This allows the Rx channels and Tx channels to be routed from the connector on different sides of the PCB. Single voltage supply sources also make for easier designs and lower BOM costs by eliminating multiple DC-DC regulators. Wider temperature ranges from -40 °C to +85 °C are ideal for most consumer and industrial systems. For reference, all of Pericom Semiconductor's linear redrivers include these features, and more. A complete table of redrivers for PCI Express and the other high-speed channels shown in Figure 2 above can be found at https://www.pericom.com/products/redriver-repeater-ic-signal-conditioners/.

This article was sponsored by Pericom Semiconductor.

Driving Along at Full Speed for Autonomous Vehicles

Sensors and processor technologies are already in place for the next generation of assisted driving and self-driving vehicles. Caroline Hayes asks, "How soon can we relinquish control of the car, and will we be happy with technology in the driving seat?"

By Caroline Hayes, Senior Editor

Although analysts cannot seem to agree on a timeframe for autonomous vehicles to be on our roads, there is consensus that they are inevitable. French firm Yole Développement believes that technology is slowly replacing the driver in vehicles. It says that by 2045 over 70% of vehicles sold will integrate autonomous functionalities. More cautiously, Strategy Analytics says that, despite "the considerable market hype," any volume market for fully autonomous vehicles is unlikely before 2025. Its Autonomous Vehicles Service report, Autonomous Vehicle Market Scenarios, offers three possible levels of adoption for automation. The first sees automakers following an evolutionary path to roll-out autonomous vehicles, but such vehicles will be far less in number than conventional car production, up to 2050; a second scenario is the adoption of city cars using autonomous driving to accelerate the market growth; and the third possibility is the increased use of partially autonomous technology, resulting in a slower roll-out.

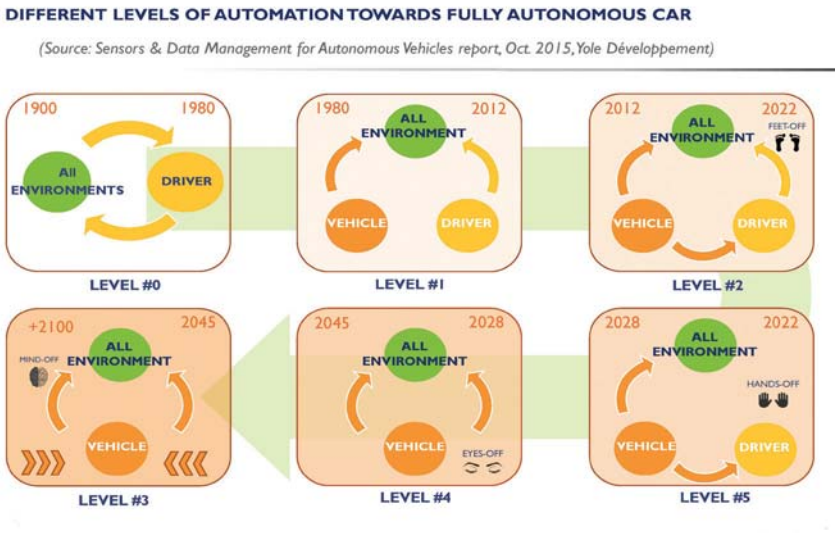

Figure 1: Yole Développement's Sensors & Data Management for Autonomous Vehicles cites the connected car as the first step towards autonomous vehicles.

Falling Sensor and Camera Costs Should Push Adoption

The increased use of sensors to detect the vehicle environment, temperature or seat position will be complemented with short and long range radar sensors, ultra-sonic sensors and Light Detection and Ranging (LIDAR) sensors and cameras. This last sector breaks into sub-groups of Near Infra-Red (NIR) and Long Wavelength Infra-Red (LWIR) cameras. Collectively, these sensors and cameras will add between $10,000 and $15,000 to the cost of the vehicle, predicts the report. This amount is expected to decrease over the next decade, further driving adoption.

The powerhouse behind the connected car is the technology that can transmit and receive data from an internal and external network of sensors. This technology has to analyze and respond to the data scenario in real-time. The level of data throughput is expected to be 1-GByte per second in a vehicle system's RTOS (Real Time Operating System). The levels of computing performance will need to be able to analyze this data quickly enough so that the vehicle can react immediately to any changes, whether that is in temperatures inside the vehicle, or to obstacles/distances changes outside of the vehicle.

Anticipating Centralized Systems' Resurgence

Many in the automotive supply chain are providers of both hardware and software. Intel® believes that supply chain structure has led to the emergence of distributed computing systems within the vehicles to accommodate a growing ecosystem of Embedded Control Units (ECUs). However, this adds to both complexity and the cost of an embedded system. The company has identified a move towards a more centralized

computing system, with its advantages of reducing complexity yet with a breadth of support that can only encourage the adoption of self-driving, and then autonomous, cars.

Another benefit of a homogeneous computing system within the vehicle is that the ever-increasing volumes of data can be transmitted and received safely and securely. A centralized computing system is more secure than one with a selection of discrete technologies. The source of the data is also critical to prevent inaccurate, malicious threats or hacks. Data generated from sources inside the vehicle are more secure, and less vulnerable to attack, than data generated outside of it.

Security Arsenal Grows

Intel's Core™ and Xeon® processors use its Data Protection Technology, with Advanced Encryption Standard—New Instructions (AES-NI). The AES algorithm protects data traffic and, in self-driving vehicles, data will be received not only from sensors and cameras within the vehicle, but also from Web or Cloud services as well as the surrounding infrastructure, such as traffic systems. The integrity of that data is critical, and the inclusion of NI accelerates the AES encryption and decryption without any performance overhead, so processing performance is unaffected.

Another weapon in the security arsenal is Hardware-Assisted Security (HAS), which adds layers of protection to prevent stealthware attacks, where a bug can be introduced and take over a system. Another level of protection is offered with its Platform Protection Technology, which prevents unauthorized software and malware operation. This can be Intel OS Guard to protect deep levels of a system, even if an application has been compromised.

Assisted driving systems use Intel Core i7 processors today. They can be found in adaptive cruise control, emergency braking and lane assistance systems. These are functions that are more compute-intensive than assisted driving functions whose job is to inform the driver of conditions. Examples of this kind of computation are lane departure warning and parking assistance. These require processor performance of less than 100,000 DMIPS and are presently carried out by a processor in the Intel Atom class.

Functions such as adaptive cruise control, emergency braking and lane keeping are categorized as assisted driving and require processor performance above 100,000 DMIPS. These are performed by processors in the Intel Core i7 family. The next level, self-driving, requires even higher DMIPS performance and are the domain of the Xeon processor family (see Figure 2).

Intel's Turbo Boost Technology accelerates the processor performance, enabling the calibre of high performance workloads needed for in-vehicle systems that need to receive, analyze and act upon large volumes of data from various sources, such as information from an external camera, to an ECU and the braking system in less than a second.

The Core i7's Smart Cache sub-system is another way that the processor can be optimized for this and other multi-threading operations that are commonplace in assisted driving systems.

All of this has to be executed in an energy-efficient manner. Intel's integrated memory controller has low latency and high bandwidth (up to 25.6-GByte per second) memory bandwidth to handle data-intensive operations without adversely impacting operation performance.

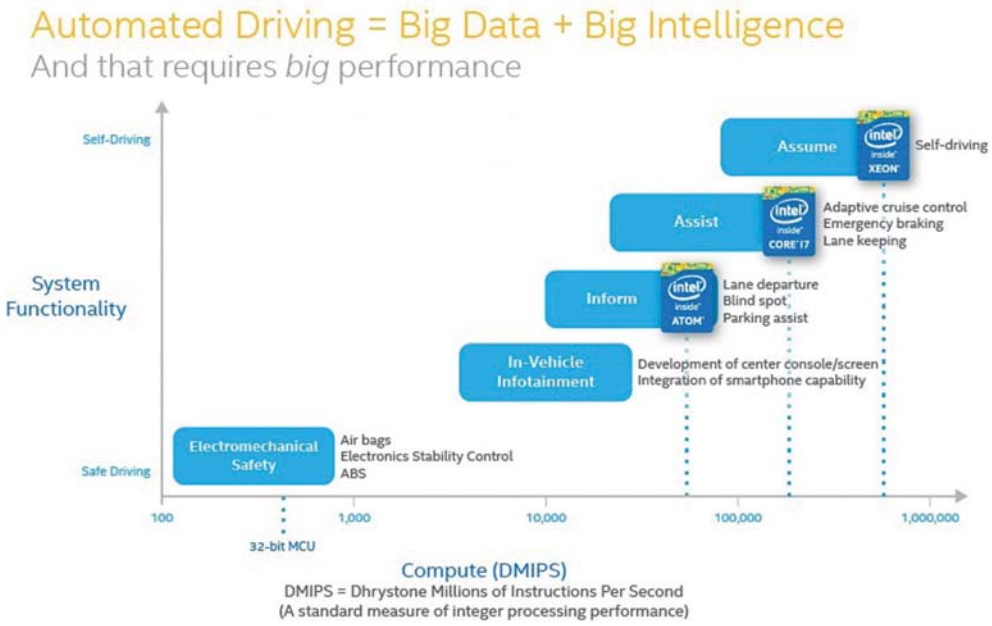

Figure 2: Progressing from IVI (In-Vehicle Infotainment) Intel categorizes driving stages as Inform, Assist and Assume functions, each with their own performance level requirements.

Collaboration

Many automotive manufacturers are working with semiconductor companies to produce specific driver assistance systems.

The Mobile Interior Imaging (Mobii) Project developed by Intel and Ford may be in the prototype phase, but indicates how the smartphone, and consumer electronics, are making their way into the driving experience. It uses interior cameras, integrated with sensor technology and data collected from around the vehicle to personalize the driver's setting and experience.

Figure 3: Project Mobii is a collaborative research project between Intel and Ford to explore how interior cameras and vehicle data can be used to enhance the driving experience.

Intel ethnographers, anthropologists and engineers worked alongside Ford research engineers to develop Project Mobii, described as perceptual computing technology for an intuitive vehicle experience. The aim was to improve the driver experience, using interior cameras with existing car sensor data and the driver's behavioral patterns, explains Tim Plowman, Experience Solutions Architect, Intel Labs.

"Our goal with the Mobii research is to explore how drivers interact with technology in the car and how we can then make that interaction more intuitive and predictive," said Paul Mascarenas, Chief Technical Officer and Vice President, Ford Research and Innovation.

Based around a smartphone app on any Intel phone, one use for Project Mobii is to use the in-vehicle camera to search the interior of the vehicle for, typically, a purse, wallet, or a child's toy, while the owner is panicking inside the house.

Also part of Mobii is face recognition software in a front-facing camera. It automatically identifies a driver and sets in-car console preferences to display information specific to the driver. It can also be used to set the driver preferences, for communications, music and schedule, when the app recognizes the driver is behind the wheel of the car. If the driver is not alone but does not want a passenger to see the selection of guilty-pleasures music tracks, this can be locked and the central display information filtered, so the passenger sees the locked console's navigation map or other choice of graphics.

> "The Mobile Interior Imaging (Mobii) Project developed by Intel and Ford may be in the prototype phase, but indicates how the smartphone, and consumer electronics, are making their way into the driving experience."

The same face recognition software can identify new drivers, which can be approved by the authorized owner of the vehicle. Settings can be adjusted to be appropriate, for example, for a newly qualified driver and with permissions such as driving perimeter distance, speed limit and number of passengers.

Who's There?

Cameras that sense who is reaching for the center screen can be used, speculates Plowman, for safer driving. If the system senses the driver is requesting a destination, it can request a voice command, but if it senses that a passenger is reaching for the center console for navigational directions, then it dis-

plays a keyboard for the destination to be typed in. In this way, drilling down through sub-systems can be eliminated.

The principle of safe driving is to ensure the driver's eyes are on the road ahead. Sensors in Project Mobii can determine if the driver's head is down or turned. It sends an alert to the heads up display and center console to 'remind' the driver of the task in hand and re-direct concentration to the road ahead.

Thinking Ahead

Following connectivity within the car, and connectivity to the immediate infrastructure, the next step is for cars to communicate with each other. Intel Research Scientist Jennifer Healey spoke of vehicles 'gossiping' to each other to create safer roads. In a Technology, Entertainment and Design (TED) Talk in April 2013, she spoke about technology available at the time that would allow cars to exchange data with each other to make roads safer. As she points out, when a driver says a motorcyclist "came out of nowhere," it can't be true; he/she was on the road for perhaps 30 minutes before lane-splitting in front of you. "Cars are three dimensional objects that have a fixed position and velocity," she said. They also travel along published road routes. "So, it is not that hard to make reasonable predictions about where a car is going to be in the near future . . . .as soon as one car sees that motorcyclist, and puts him on the map [noting] position, velocity and estimate he'll continue going 85 miles an hour, you'll know, because....the other car will have whispered [to your car] something like: 'By the way, five minutes, motorcyclist—watch out'. "

Healey proposed that to get to this level of 'chatter' between vehicles, a GPS and a camera in the car are the basic requirements. "Using computer vision, I can estimate where the cars around me are —sort of—and where they are going. What happens if two cars share that data—if they talk to each other?" she asked and answered: "They both improve." Robots, with stereo cameras, GPS, two dimensional laser range finders and short-range communication radio have been demonstrated to talk to each other, precisely track each other's position and to avoid each other. Healey does caution against too much chatter, as this can result in a lot of data packets to process, in which case they need to be prioritized. This can be streamlined by eliminating the ones that are following an expected course, to concentrate on the vehicle that is going off-course. By identifying a road user that could be a problem because he is going off a predicted course, it is also possible to predict the new trajectory. "So you don't only know he is going off-course, but you know how, and you know which drivers to alert to get out of the way," she told the audience.

Assisted driving, in varying degrees, is a reality on our roads today. There are degrees of use in city, freeway and countryside driving that need to be defined, but the sensor and processor technology needed to collect and act on data is available today. Those component parts will not be a roadblock, but reducing costs and gaining public acceptance might be a longer road to travel.

Caroline Hayes has been a journalist, covering the electronics sector for over 20 years. She has edited UK and pan-European titles, covering design and technology for established and emerging applications.

"…more agility at the edge…" Q&A with Eurotech's Chief Marketing Officer

A long-standing and productive relationship with Linux provider Red Hat is just one of the reasons Eurotech finds itself well positioned for the IoT/M2M era.

By Anne Fisher, Managing Editor

Editor's Note: Our thanks to Robert Andres, Chief Marketing Officer, Eurotech, who recently shared his insights on a number of topics.

Embedded Intel Solutions: What are the current challenges and opportunities in the IoT/M2M space?

Robert Andres, Eurotech: Despite the general confusion in the market, there are some visible trends that are making things a lot easier. Many challenges are seen as sound reasons for not implementing IoT /M2M solutions at this time: Complex technology; the lack of specific know-how; an infrastructure that is required; IT/OT integration challenges; security; regulatory issues; certifications; confusing messages from industry and standardization bodies; interoperability of solutions, etc.

The good news is that there are major industry trends that lead to effective and affordable IoT solutions, and to commoditization of distributed systems. Unlike in the past we now have powerful embedded systems, not just general-purpose gateways—often in the form of off-the-shelf, purpose-built HW, and virtualization technology at the cloud datacenter. And on the device side, Enterprise IT technology, best practices and an infrastructure that can be leveraged, along with cloud computing and many well established open and industry standards.

Eurotech has been delivering distributed device solutions (hardware and software) for more than two decades now, and this experience informs our relationship with customers and system integration partners. Stated broadly, our goal is to reduce Total Cost of Ownership (TCO) in M2M/IoT solutions. That involves a dramatic lowering of initial investment, reduced risk in M2M/IoT projects, and much shorter time to market. It is about encapsulating the complexity of distributed systems in the field way beyond the development phase and significantly improving the deployment and operation of IoT infrastructures.

Embedded Intel Solutions: Can you be a bit more specific with regards to Eurotech's IoT offerings?

Andres, Eurotech: Eurotech is one of the leading players in the embedded market, offering a broad product portfolio ranging from CPU boards and SBCs to purpose-built devices to high-performance computing systems. As mentioned before, state of the art M2M/IoT solutions also require well-designed software and communication services layered on top of that hardware.

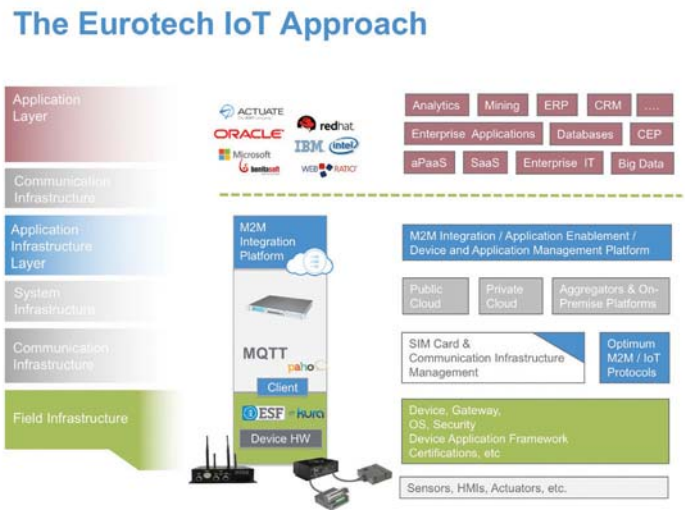

Figure 1: Robert Andres notes that the Everyware Software Framework (ESF) from Eurotech is based on Linux, Java and OSGi and makes possible the implementation of business logic in a very effective way.

One particularly important building block is the software on the device. In order to simplify and optimize the development of device software Eurotech offers an application framework layered on top of the operating system called the Everyware Software Framework or ESF. It is based on Linux, Java and OSGi and enables the implementation of business logic in a very effective way. Not only does it abstract the hardware in order to ensure investment protection with regards to the software, but it also allows new services to be loaded to multi-service M2M/IoT gateways or smart edge nodes dynamically.

Another very important building block we offer our customers today is a cloud-based M2M/IoT platform called Everyware

Cloud that integrates easily into existing enterprise IT infrastructures—offering simple access through standard APIs to real-time and historical data from devices. In addition this M2M/IoT Integration Platform also offers the device life cycle features that ensure a smooth deployment and management of the distributed device infrastructure.

There are other important elements, depending on the nature of the customer's business. One particular aspect I would like to highlight is carrier certifications, which are required in many markets for devices connected to the cellular networks. For many companies that wish to sell their products or services internationally these certifications are a major barrier because of the associated effort and cost. We offer with our ReliaCELL product range a very attractive, pre-certified modular solution.

Embedded Intel Solutions: With regard to effective connection with IT infrastructures, what are the significant ways in which the IT infrastructure itself is changing in response to the growth of the IoT—both on the consumer and on the Industrial (M2M) side?

Andres, Eurotech: Outside the traditional embedded business, in many of the new IoT business opportunities the IT department is playing an increasingly important role. That results in different technology and architecture decisions to allow an effective IT/OT integration. While more is changing on the OT side, IT is also discovering that new technology approaches are required. Interfaces and solutions (including programming of edge devices) have to be done in an IT-centric way (Java instead of C++)—and must also scale.

Embedded Intel Solutions: How is the race to get ahead of those trying to compromise IoT and M2M security affecting Eurotech's plans for being a leader in the IoT/M2M space?

Andres, Eurotech: We see ourselves positioned well when it comes to IoT security aspects. Security is one of the elements that differentiates Eurotech´s IoT offering in the market. Eurotech understood very early on that security has to be approached holistically and in the details of an IoT (or OT) infrastructure—leveraging best practice approaches and technologies from the IT security space where feasible. Of course the OT world is different, and specific aspects differ from the traditional IT world approaches.

But powerful ICT security technology that has been created and is constantly validated, like cryptology solutions and strong authentication methods, certainly can be implemented in IoT solutions. Security in IoT projects is more than the combination of some "perimeter defense" security solutions like white- and blacklisting, firewalls and encrypted communication channels.

While these elements are important, we also believe that proper authentication methods, the use/storing/managing of certificates, and especially an overall architecture and design that eliminates many possible attack scenarios play significant roles. Furthermore it is essential that external security specialists do validate/audit the security measures regularly.

Because of the very fractured world of IoT, Eurotech is one of the very few companies which are able to offer IoT security solutions that span from the IoT gateways or edge devices to the IoT application enablement platform. Examples are an end-to-end security offering based on x.509 Certificate/PKI technology or an architecture that ensures that the "attack surface" of IoT devices is significantly reduced.

Embedded Intel Solutions: What will be some of the overlooked or less obvious effects of cost and other pressures in the IoT/M2M space?

Andres, Eurotech: Distributed device solutions have been created for decades—most of them following a "traditional telemetry" approach where fixed functionality, silo-implementations and a focus on lowest cost per edge node are the primary driving parameters.

But looking at distributed device solutions with the promises made by a modern IoT approach not only requires a lot more agility at the edge in terms of software and functionality (for example, vending) but also aggregation and computational capabilities that require sophisticated and secure device and embedded software management capabilities. For these implementations the primary metrics include a focus on the lowest TCO and the best strategies to protect the investment, especially in software.

It is that software capability at the edge—the software-defined multi-service gateway—that allows many companies to preserve and extend their customer value. (Opposing the thought /fear that IoT technologies level the value of companies deploying such solutions.)

Embedded Intel Solutions: When collaboration with a partner such as Linux provider Red Hat takes place, what are some of the characteristics such a collaboration should have in order to be successful?

Andres, Eurotech: Good partnerships/collaborations build on a common understanding of success factors, technology and an aligned vision.

The Red Hat and Eurotech partnership offers a robust example of such a partnership. Both companies understand the value and necessity of a strong ecosystem to deliver IoT solutions. The companies' respective competencies and products perfectly complement each other (Red Hat on the Enterprise IT side, and Eurotech on the Operational Technology (OT) side with HW and SW). What's more, Open Source and Java are important pillars in both companies' strategies, and all these factors ensure a good "natural" alignment and many positive synergies between the parties.

For Software Developers, Hardware Emulation Rules!

Booting an OS such as Linux gets the designer to the starting line, but how does the real work commence for complex multicore designs destined for gaming, digital signage, and more?

By Lauro Rizzatti

When software developers hear the word emulation, they often think software emulators and not hardware emulation. That's changing, though, as the versatile hardware emulator gains a more widespread reputation for being able to ensure that the embedded system software works correctly with the underlying hardware. It's happening as it becomes a shared resource between hardware and software teams to accelerate hardware/software integration ahead of first silicon. As it should, especially when the ratio of software developers is greater than hardware engineers on a chip project.

Figure 1: Winning the debug process race: Boasting an execution speed up to six orders of magnitude faster than a hardware-description-language (HDL) simulator, hardware emulation can shorten the design debug process without constraining design sizes. [Photo courtesy Official Navy Page from United States of America Candice Villarreal/U.S. Navy; public domain, via Wikimedia Commons.]

Through the years, hardware emulation has been used for hardware verification of ASIC and SoC designs, and its use model is growing. Hardware emulation's speed of execution can be up to six orders of magnitude faster than a hardware-description-language (HDL) simulator, reducing the design debug process without constraining design sizes. Among the largest SoC designs are processor and graphics chips, and Ethernet switch and router chips. All of them approach or even surpass the half-billion ASIC-equivalent gate mark. These are designs that would thwart an HDL simulator.

Executing Embedded Software

Never has hardware/software co-verification been more important as embedded software becomes a large component of an already complex system-on-chip (SoC) design. For designs using an operating system, such as Linux, booting it is a bare minimum. But Linux already exists and it's not code that designers have to write. Although it does have to be customized to the system, with memory maps and drivers that take a simple system call and cause it to do real work on real hardware.

Booting an OS typically gets the designer to the starting line. The real software that does the real work can only start when the OS has been successfully loaded. Attempting to verify such applications would simply stretch the runtime way beyond the OS boot time.

In addition, the designer has to test multi-threaded programs. For some systems, the multiple "threads" are actually turned into multiple processes running independently on different cores, with or without an OS. Packet processing is probably the best example of this: It typically consists of a pipeline of processors, each of which performs a specific piece of the packet processing functionality in turn. All of the threads or processes have to work together seamlessly, passing data around as necessary and seamlessly swapping contexts.

Simulation farms, a popular and well-used resource, are not adequate to execute embedded software. Embedded software needs to be processed in parallel and cannot be split in subsets, all of which takes several billion sequential cycles. Hardware emulation easily can handle a sequential process and is undeterred by design sizes. At speeds of several hundreds of kilohertz or megahertz, it can boot Linux in less than one hour.

Integrating newly designed hardware with newly coded software should be done as early as possible to resolve bugs –– not a problem for hardware emulation. It ensures software and hardware are verified concurrently to confirm they interact correctly. It can be used as soon as the design register transfer level (RTL) code is stable to uncover difficult to find hardware bugs that require many millions of verification cycles. It can trace a software bug into the hardware or a hardware bug in the

software's behavior with the necessary speed, performance, and capacity to handle complicated debugging scenarios, something no other verification tool can do.

The Shortcut Designers Can't Afford

Virtual platforms are becoming increasingly popular as a way of testing software ahead of actual hardware. For general-purpose applications, those not requiring intricate interaction with the hardware, virtual platforms are an effective way of testing application software way before silicon or even RTL availability. However, the strength of a virtual processor is also its weakness: It can execute software quickly because it abstracts out much of the execution detail. Such abstraction is acceptable for higher-level code or applications. For code that touches the hardware, such as drivers, designers can't afford to take that shortcut. Virtual platforms are utterly inadequate for proving that the chip is working because they eliminate the chip's details.

To verify the interaction of embedded software with the underlying hardware, a cycle-accurate representation of the design is needed to be able to trace a bug anywhere in the SoC. A hardware emulator offers an accurate functional representation of the design before the silicon is ready for testing because the emulated design is based on an actual silicon implementation, even though it's not timing accurate.

Functional verification and system validation are hardware emulation strengths. It is used to verify the hardware and software work as intended and validate that the system meets or exceeds the specification. Software programmers are using it to validate embedded software –– applications, diagnostics, drivers, operating systems and software-driven tests, all with a need to process hundreds of millions, even billions of cycles –– extending its use across the entire SoC development cycle.

Hardware emulation gives software developers full, 100 percent visibility into the hardware design, necessary to quickly trace bugs anywhere in the design.

Not to Be Confused

Verifying the combination of hardware and embedded software is estimated to consume about 70 percent of the project cycle. And this verification task is also the most challenging one to perform, so it's no wonder hardware emulation rules. This co-verification process can improve product quality and shave months off a project development schedule.

The makeup of an SoC project team has changed through the years and now is comprised of both software developers and hardware engineers. Combining the two disciplines offers a way to differentiate the product and give it a technological edge. The backdrop hasn't changed much, though. It's still a heavy mix of time-to-market delivery schedules, while balancing cost with design size and low-power considerations, revenue and profitability. Comparatively, debug's a cinch with hardware emulation, as both software developers and hardware engineers benefit. Both groups have come to value hardware emulation's ability to verify hardware and software at the same time. No one will confuse software emulation and hardware emulators again.

Dr. Lauro Rizzatti is a verification consultant and industry expert on hardware emulation. Previously, Dr. Rizzatti held positions in management, product marketing, technical marketing, and engineering.

Communication and Co-operation Create a Tactical LTE Network

Joint development with partner vendors enables ADLINK to develop portable LTE that provides reliable connectivity for first responders when regular networks are non-existent.

By Caroline Hayes, Senior Editor

Communications to coordinate rescue efforts can be the first casualty when they are attempted in emergency situations, or in extreme environments, such as on the battlefield. In disaster areas, such as an earthquake zone, if the communications infrastructure is functioning at all, it is often jammed with harried users, making it hard to access. This means communicators may be unaware of the full scale of the situation and unable to make accurate evaluations for appropriate courses of action.

Figure 1: The LSF-02 is a portable man-pack radio by LCR for first responders and tactical battlefield LTE communications.

As a consequence of disasters such as 9/11 (2001) and Hurricane Katrina (2005), the Federal Communications Commission (FCC) and the International Telecommunication Union (ITU) mandated that first responder communication was to move to Long Term Evolution (LTE) or 4G LTE to allow the different response teams to communicate and coordinate with first responders at or near the scene.

It is also important in emergency or battlefield conditions that communications are as far-reaching as possible. Portable communications units or man-packs can provide mobile and mission-ready tactical LTE networks. These networks can be deployed instantly, allowing communications to resume.

This article examines how system design expertise and some solid relationships brought a varied set of skills together to create an emergency/battlefield communications end-product.

LTE Communications

Chassis, backplane and integrated systems provider LCR has developed the LSF-02, a rugged, instant, tactical LTE network, based on Connect Tech's COM Express Type 6 carrier card, using an SBC from ADLINK and LTE communications software from Radisys. A partnership of multiple vendors can offer development and design support as well as build on proven technologies.

"LTE is a well-defined protocol, not a proprietary one," observes John Long, Vice President, Integrated Systems for LCR. "It can interconnect simply because it is well-defined." In the battle arena, tactical LTE allows communications where there is no functioning infrastructure.

The LSF-02 computing and communications platform, at 20 pounds/ approx. 9 kilograms with batteries, is portable and can be worn on one's person (Figure 1). It has a redundant, field-swappable battery system which uses solid-state SS-2590 rechargeable Li-Ion battery packs, which are external to the sealed case, to allow for continuous operation. Input voltage is 9V to 36V, and the pack has a maximum RF output of 1W (30dBm for each of the two antenna ports).

A 55-pin external I/O, shown in Figure 2, can be used to connect to the Ethernet through another router if required, for TCP/IP communications. Alternatively, it can provide VoIP using the Internet via a backhaul connection.

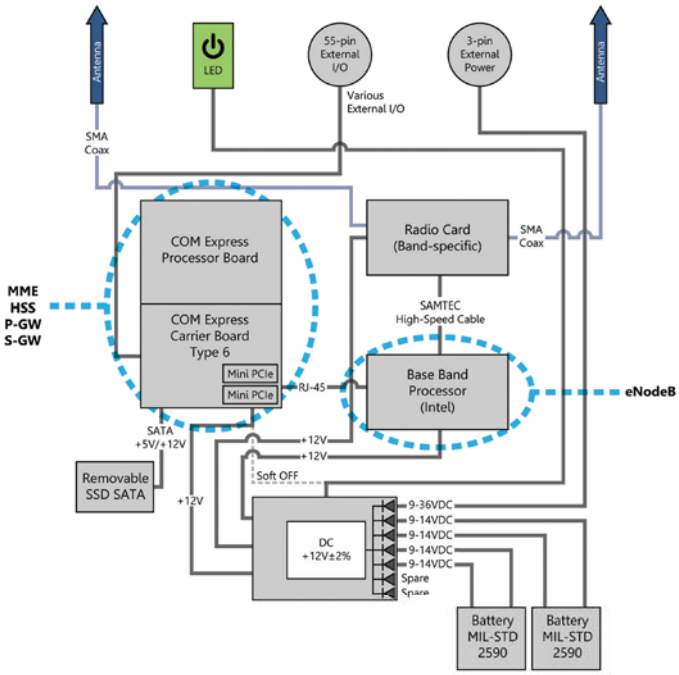

Figure 2: The LSF-02 has two antennae for RF customization, a COM Express Type 6 carrier from Connect Tech and an Express-HL-i7-4700EQ from ADLINK.

The LSF-02 platform is a complete LTE network in a box, says Long, with options for customization on the RF side. It is also modular, supporting two processors and custom RF to meet application-specific needs. The baseband processor puts RF in an Ethernet format for ease of use (see Figure 2).

As well as broadband mobile data connectivity between tablets, smartphones, rugged terminals and displays, LSF-02 allows real-time audio/video calls and multimedia conferencing with a typical range of .62 miles/1 kilometer, depending on terrain and conditions. It is also self-discovering, with the ability to add users or extend the geographical coverage. Up to 16 active LTE users can be supported per unit. This means that powering on multiple LSF-02s in an area creates an automatic "clustered" network, and LTE handsets can immediately begin communicating and sharing data among themselves.

Used as a stationary unit, it is an alternative to the big cell sites designed for thousands of users offered by other providers, such as Nokia and Siemens, says Long. It fills the gap for tactical LTE networks. Long explains that while the LSF-02 can be totally self-contained with no outside communications, it can also be backhauled to other cells by virtue of the two Ethernet ports the CCG011 COM Express Type 6 rugged, ultra-light carrier card from Connect Tech offers. The design requirements for the LSF-02 are shown in Table 1.

Interestingly, Table 1 lists many variations for the internal electronics, especially the RF radio standards. One of the ways this is accomplished in the unit is using open-standard modules such as COM Express and PCI Express (mini). The Mini-PCIe slots provide flexible modularity, while COMe bolts together the primary partnerships used to create the portable LTE base station.

Form Factor

Computer-on-Module (COM) is an ideal choice when a system's I/O functions stay relatively stable over time, but there is a possibility for upgrading the processor as needs evolve. The single-board computer (SBC) processor is usually the mezzanine card that plugs onto the I/O carrier board. COM Express is the most powerful, modular architecture, says Long. "It continues to evolve as it gets deployed. . . It allows a cost-effective solution, as another COM Express module can be chosen for use if needed," he says. Validation is simple, he adds, when using the same type of COM Express module (in this case, Type 6).

Figure 3: ADLINK's SBC Express-HL-i7-4700EQ was selected for the project based on previous collaboration experience, allowing Connect Tech to focus on other parts of the design.

The ADLINK SBC Express-HL-i7-4700EQ (Figure 3) and Connect Tech card run the third-party software that provides the intelligence of the LTE network, in this case supplied by Radisys.

Communicating Relationships

The relationship between the three companies means that they can offer design support. Connect Tech provides LCR with hardware and support, and is also a design partner for ADLINK and provides 3D models for a project's mechanical engineering support.

The carrier card was chosen because it is available off-the-shelf and is a known good fit with the ADLINK SBC Express-HL-i7-4700EQ, says Michele Kasza, VP Sales, Connect Tech. "Around COM Express, there is an assumption that a carrier has to be

custom designed," she says. "Off-the-shelf products can be a perfect fit where the small quantities involved do not make sense for custom design." Another benefit is reducing the development time, she points out. "The custom design process can be very long, whereas Connect Tech delivers an off-the-shelf carrier card in two to five business days." This, concludes Kasza, makes some off-the-shelf cards or boards a perfect fit for those wanting to use COM Express on small production runs.

Deployment Environment:	Ability to survive the rigors of a man-pack environment Ability to operate with direct solar heating
Physical:	Under 20lbs without batteries Operational: • Temperature: -40 °C to 70 °C • Vibration: 7Grms, IEC 6-68-2-63, random, 5-500 Hz, 1 Octave • Shock: 30G/11ms per MIL-STD-516.5 • Immersion: IP67
Power:	Power Switch with visible indicator of power status Support vehicle power Support 2 batteries, where the unit can be operated with only one battery
Minimum External I/O:	1 GbE 1 USB 2.0
Internal Electronics:	RF & Baseband • Support multiple variants of RF Cards (P/A) • Support multiple variants of Base Band processors Compute • Support Various CPU • Support for add in capability such as WiFi or other Storage: • Support minimum of 160G • Must be easily removal for grab-and-go
Others:	Because of dimensions and the deployment environment, the unit must not be significantly damaged if it is used as a step stool. This sounds odd, but it's actually quite important.

Table 1: Design Requirements for LSF-02 by LCR Embedded Systems.

In this case, the rugged CCG011, a COM Express Type 6 card, was chosen over the Type 10 version due to its size, confirms Kasza. At 125mm x 95mm, it is smaller than the Type 10, allowing the end result to keep a small footprint, yet still able to provide the higher processing power required.

In addition to two Gigabit Ethernet ports, the CCG011 has CPU, memory and I/O, including 8-bit GPIO, Mini-PCIe expansion sockets for either two half or full length cards, four SATA ports, two mSATA ports, eight USB ports and two DisplayPort interfaces for DisplayPort, HDMI, DVI or VGA use. Audio is served with an HD audio codec from Cirrus Logic. "[CCG011] is designed to military standards for shock, vibration and temperature range (it operates at -40 °C - +85 °C) and is a proven design for the ADLINK module Express-HL-i7-4700EQ already used by LCR in other projects. It supports the Intel® Core™ i7-4700EQ quad core processor used in the LSF-02, as well as Mobile Intel QM87 Express chipset.

Customers may want to add their own software on top of the box to provide a more full-featured LTE network. Long agrees, pointing out that military customers may want encryption, which involves physically encrypting data before it goes into the Ethernet LAN. This can be achieved by software and run on the ADLINK SBC.

Working with established partners can reduce the development time as products that have been proven to work together can be specified. Selecting off-the-shelf products also accelerates the design time and reduces the cost of customization but allows the desired flexibility for particular end uses and environments. No two emergency or battlefield situations are alike: a reliable, rugged communications pack has to operate in a variety of harsh conditions and must operate 'out-of-the-box' and without fail. This takes system design expertise and some solid relationships that bring a varied set of skills together in one end-product.

This article was sponsored by ADLINK.

Silent Night, Day and Everything in Between

The Intel® Atom® Processor C2000 Series product family enables a completely fanless design, and that's music to the ears of the author, CEO of an autonomous Web hosting company.

By Bernd Blume, antsle

Having just a few companies control the Internet is at odds with your rights. And it's at odds with your wallet. We've seen centralization happening with regard to the Internet since the turn of the century. Everyone rents their Internet server capacity from a bunch of big (some would even say arrogant) companies. These companies charge "taxi" prices. Everyone uses a taxi (to stick with that metaphor for a moment), pays taxi prices and willingly adheres to the rules imposed on them by the taxi lobby because it used to be difficult to own and operate your own server or data center. However, you can now own your "car" (server), be in control and save a mountain of money.

A Change as Virtualization has Evolved

Things have changed. Virtualization software, enabling several virtual servers on a single piece of hardware, has become commonplace and easy to use. Hardware has become more powerful yet less power consuming, enabling fanless designs. Broadband adoption has taken off, giving many of us more bandwidth at home than we used to have in professional data centers just a decade or two ago.

And society has become more aware of privacy concerns, contributing to the trend to host under your own control rather than under Google's or Amazon's. At antsle it's our belief that what is needed is a powerful, 100 percent silent home server to help bring about this control. One that is capable of hosting more than 100 virtual servers, so it can be the user's private data center and replace a lot of Cloud services. We had to make many decisions on how to make that dream of ours a reality.

One decision involved which CPU to choose. Just a few years ago, we probably would have opted for an ARM CPU. We feel our product absolutely has to be silent in order to be suitable to run 24/7 in your living room beside your router. But ARM (let alone other architectures, e.g., MIPS) are just not a good choice for servers. Yes, you can run modern Linux distributions under ARM, but many other OSes are not available on ARM: SmartOS, FreeBSD, (full-fledged) Microsoft Windows, to name a few.

As Versatile as Possible

The CPU that met our requirements best is the Intel Atom Processor C2000 Series product family (formerly Avoton). It's based on the Silvermont architecture and supports an impressive instruction set, including SSE 4.2, CLMUL and AES- I, running only slightly behind the powerful Haswell architecture of desktop PCs. Yet the CPU has just 20W of TDP, enabling a completely fanless design, producing 0.0 dB of noise! In other words, absolute silence.

In the middle of that silence: eight Silvermont cores, each one running at 2.40 GHz. And most importantly for our purposes: Unlike earlier Atom CPUs, the Avoton has Intel VT for x86-based Intel Architecture (VT-x) with Extended Page Tables (EPT). That is necessary to run hardware virtualization using the Linux kernel's KVM feature. We do go with OS-level virtualization in the antsle server for most cases, and we use LXC for that. For use cases where the "guest" inside a virtual server is a flavor of Linux, that will be the way to go for our users. But we want more.

We want our server to be a tool that is as versatile as possible for our clients. So we want KVM virtualization as well. That way, the user can run, for example, 70 OS-virtualized containers running Linux plus six HW-virtualized VMs running other operating systems. Having VT-x including EPT comes in handy for that.

Our system software, antsleOS, manages both container VMs and KVM VMs in a coherent, transparent manner. Also, we wanted to choose a CPU with excellent investment protection for our clients. The future of the Avoton looks bright. A successor codenamed Denverton is on the horizon. It promises 16 cores instead of eight, a larger cache, and the SoC is expected to support 10 GbE, 10 times the speed of today's Avoton. Expected to hit the market some time in 2016, it will enable our users (with an antsle two) to run an extremely powerful, silent data center in their home or office. What a contrast to the times when they had no direct control over their servers whatsoever.

Bernd Blume is CEO of antsle. After a corporate career at IBM and a telecoms carrier, Blume founded software startup Xionet and raised $4M in Venture Capital prior to founding antsle.

ADLIGW-1600 Industrial IoT Gateway

The ADLE3800HDC is based on Intel's first System-on-Chip (SoC) E3800-series Atom processors which use Intel's 22nm 3D Tri-gate process. It offers vastly superior compute performance and energy efficiency and Intel's 7th generation graphics engine for stunning graphics performance. Improved power management capabilities result in standby power measured in milliwatts with days of standby time. Features Include:

Intel E3800 Series Processors
- Intel Atom E3800-3845 Quad
- Intel Atom E3800-3827 Dual

Memory: Up to 8GB DDR3L 1333MHz

I/O Connectivity
- 3x GLAN
- 4x USB2.0
- Optional 1x USB 3.0 (I-PEX)

Network/Cloud Connectivity
- 3x LAN
- Optional 802.11 WiFi

Graphics
- VGA/CRT = 2560x1600@60Hz 24bpp (via DVI conn)
- DVI = 1920x1080@60Hz 24bpp
- DirectX 11
- Open GL 4.0

Mounting
- DIN, VESA, and Bottom Mounting

Come visit the ADL Embedded Solutions booth (#1-554, Hall 1) at the Embedded World Exposition & Conference in Nurnberg, Germany, February 23-25th 2016 to view the latest products and to see multiple live product demos or visit http://ew2016.adl-usa.com. For system pricing, contact ADL Embedded Solutions at 858-490-0597 www.adl-usa.com or email sales@adl-usa.com.

Embedded Solutions

ADL Embedded Solutions Inc.
4411 Morena Blvd., Suite 101
San Diego, CA 92117-4345
(t) 858.490.0597
(f) 858.490.0599 Fax
sales@adl-usa.com
www.adl-usa.com

SharpStreamer™ Pro PCIE-7210 Video Accelerator Card

The Artesyn SharpStreamer™ Pro PCIE-7210 high-performance video accelerator card enables service provider networks to offer AVC & HEVC video transcoding services quickly and dynamically. As an add-on card, the PCIE-7210 offers quick and scalable integration into multiple server architectures to meet the demands of ISPs and MSOs who want to use existing servers and cloud infrastructure to support new video transcoding services.

The PCIE-7210 will enable content owners, broadcasters and service provider networks to lower the cost of handling OTT video and speed the deployment of high density video transcoding and multiscreen delivery. Through its support of Intel® Xeon® processors based on the Skylake microarchitecture, the SharpStreamer Pro card offers a significant performance and functionality boost in a standard PCI Express form factor for compatibility with industry-standard servers.

Key Features
- HEVC transcoding accelerator delivering up to 16 HEVC 1080p30 transcodes
- Provides horsepower to run complete application on each microprocessor
- Designed for lowest in-system latency
- Server-based approach; no dedicated appliances
- Higher H.264/AVC and H.265/HEVC transcoding density than software-only solutions

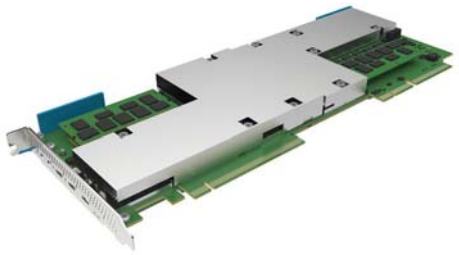

Technical Specs
- (2) Quad Core Intel® Xeon® E3-15xx v5 processors
- Intel® Iris™ Graphics (GT4e)
- 4x DDR4 SO-DIMM; designed for up to 64GB DDR4 memory
- Full height; full length PCI Express form factor

Supported Intel® Products
Intel® Xeon® Processor 3400 Series, Intel® Ethernet 10G Controller, Quad Core Intel® Xeon® E3-15xx v5 processors

ARTESYN EMBEDDED TECHNOLOGIES

Artesyn Embedded Technologies
2900 S. Diablo Way, Suite 190
Tempe, AZ 85282 USA
Phone 1 602 438 5720
TollFree 1 800 759 1107
computingsales@artesyn.com
www.artesyn.com

ATCA-7480 QuadStar™ 40G Packet Processing/Server Blade

The ATCA-7480 blade from Artesyn Embedded Technologies uses two Intel® Xeon® E5-2600 v3 family processors (up to 28 cores per blade), optimized data paths and Artesyn's QuadStar™ interface to deliver the highest performance processing. Optional hardware accelerators can improve the performance of security applications. Scalable memory capacity (up to 512GB) enables faster database access, accelerated pattern matching and helps optimize routing decisions in virtualized network environments. Sixteen memory sockets means cost sensitive applications can still achieve their required capacities using cheaper modules. Artesyn's enabling software for SDN/NFV supports Intel DPDK-accelerated OpenVSwitch, OpenFlow and OpenStack plug-ins, and Wind River's Titanium Server.

Key Features

- Two Intel® Xeon® processors, E5-2600 v3 family
- Scalable performance range with up to 14 cores per processor
- Up to 512GB main memory, DDR4, configurable for highest capacity or cost effective memory configurations
- QuadStar™ 40G fabric interfaces enabling multiple bandwidth and redundancy options
- Hardware off-loading functions for en/decryption and compression (optional)

- Virtualization support including Linux KVM and VMware
- Intel DPDK ready

Supported Intel® Products
Intel Xeon E5-2600 v3

ARTESYN EMBEDDED TECHNOLOGIES

Artesyn Embedded Technologies
2900 S. Diablo Way, Suite 190
Tempe, AZ 85282 USA
Phone 1 602 438 5720
TollFree 1 800 759 1107
computingsales@artesyn.com
www.artesyn.com

Enable Your Intelligent Embedded Solutions

Intel® Xeon® Processor E3-1268L V3

EC532-DL

- Intel® C226 Chipset
- 3 expansion slots: 1 PCI + 2 PCIe x16
- Dual displays: 1 VGA, 1 DVI-I
- 2 GbE, 6 COM, 8 USB, 2 SATA 3.0, 8-bit DIO
- Supports 4 full speed USB 3.0 ports
- 235mm x 153.8mm x 222mm

Ideal for Industrial Automation & Transportation

DFI

Headquarters - DFI Inc.
10F, No.97, Sec.1, Xintai 5th Rd., Xizhi Dist.
New Taipei City 22175, Taiwan, R.O.C.
tel: +886(2)26972986 www.dfi.com

Intel® Pentium®/Celeron® Processor N3000 Family

BW051

- Single channel DDR3L, 1 SODIMM up to 8GB
- 1 Mini PCIe, 1 mSATA
- 1 VGA/HDMI, 1 LVDS
- 1 GbE, 2 COM, 4 USB, 1 SATA 3.0, 8-bit DIO, SMBus
- 100mm x 72mm

Ideal for Mobile & IoT Devices

intel IoT Solutions Alliance

DFI-ITOX, LLC
15 Corporate Place South, Suite 201,Piscataway, NJ 08854, U.S.A.
tel: +1(732)390-2815 www.dfi-itox.com
eStore: www.dfi-itox.com/estore

DFI-ITOX, aka ITOX, is a subsidiary company of DFI Inc. DFI-ITOX

IntervalZero RTX64

RTX64 turns Windows into a real-time operating system (RTOS). RTX64 provides hard real-time and control capabilities to a general purpose operating system that is familiar to both developers and end users. RTX64 consists of a separate real-time subsystem (RTSS) that schedules and controls all Real-time applications independently of Windows. RTX64 is a key component of the IntervalZero RTOS Platform that comprises x86 and x64 multicore multiprocessors, Windows, and real-time Ethernet (e.g. EtherCAT or PROFINET). RTX64 transforms Windows to outperform real-time hardware such as DSPs, allowing you to reduce the development costs for systems that require determinism or hard real-time.

- Reduce BOM costs by 25-50%
- Boost quality and performance
- Scale quickly and shorten product cycle times
- Significantly reduce reliance on proprietary hardware such as DSP

Key Features

- Flexibility to configure how much or little processing capability is used for real-time processes (1 to 63 processors).
- Use Microsoft's Visual Studio IDE and leverage managed code in (C/C++) for Windows, to develop real-time applications.
- SMP aware scheduler uses priority-driven and pre-emptive algorithms to ensure critical thread context switches, and yields to high priority threads occur in the sub-microsecond range.
- Full SDK and runtime with supported headers, libraries and real-time APIs for application development.
- Deterministic memory. Access to Windows file system and registry. Ability to set search paths for process creation and loading of RTDLLs.
- Full control of real-time process threads with 128 thread priorities, and the ability to load balance as needed on dedicated real-time cores while not being impacted by Windows.
- Full integration into Visual Studio, develop build and debug real-time within Visual Studio the same as you would a user level windows application.
- Profile real-time applications, through integrated monitoring functionality, no need to instrument your code.
- Use shared memory to pass data between Windows user level processes and real-time processes.
- Real-time stack and network drivers that provides TCP/UDP/IP networking support for IPv4 and IPv6 (separate purchasable feature in RTX64).

IntervalZero

IntervalZero
400 Fifth Avenue, Fourth Floor
Waltham, MA 02451
USA
(T) 781-996-4481
(F) 781-795-0158
sales@intervalzero.com
www.intervalzero.com

Register now and receive your free admission ticket using the eCode 2ew16P!
embedded-world.de/voucher

Nürnberg, Germany

February 23 – 25, 2016

embeddedworld

Exhibition&Conference

... it´s a smarter world

Feel the pulse of your industry!

Embedded world is THE meeting place for the international embedded community. Get an edge in knowledge today!

embedded-world.de

Trade fair organizer
NürnbergMesse GmbH
T +49 9 11 86 06-49 12
visitorservice@nuernbergmesse.de

Congress organizer
WEKA FACHMEDIEN GmbH
T +49 89 2 55 56-13 49
info@embedded-world.eu

Media cooperation

elektroniknet.de
computer-automation.de
ENERGIE & TECHNIK Solutions for a Smarter World
DESIGN & ELEKTRONIK KNOW-HOW FÜR ENTWICKLER
MEDIZIN-und-elektronik.DE
Elektronik
Elektronik automotive
Markt&Technik
Computer & AUTOMATION
MEDIZIN elektronik

NÜRNBERG MESSE

www.congatec.us
6262 Ferris Square | San Diego CA 92121
Phone: 858-457-2600| sales-us@congatec.com

congatec

Always at your service!

congatec - technology for 24/7 operation.

conga-MA4

- Intel® Atom™ processor family
- COM Express Mini Type 10 module
- High resolution Intel® Gen 8 graphics

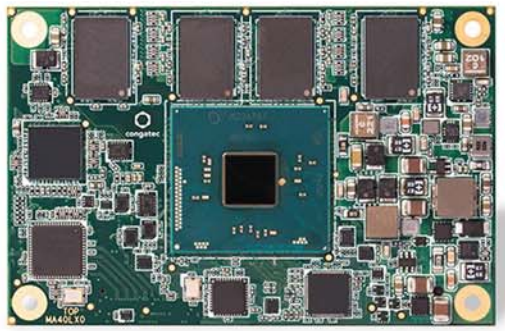

We **simplify** the use of **embedded technology**

intel IoT Solutions Alliance