Extension MEDIA

Embedded Intel® Solutions

Winter 2015

Securing Embedded Devices in the IoT Era

Automotive Applications Pursue Higher Performance

Test Equipment Tackles PCIe Receiver Challenges

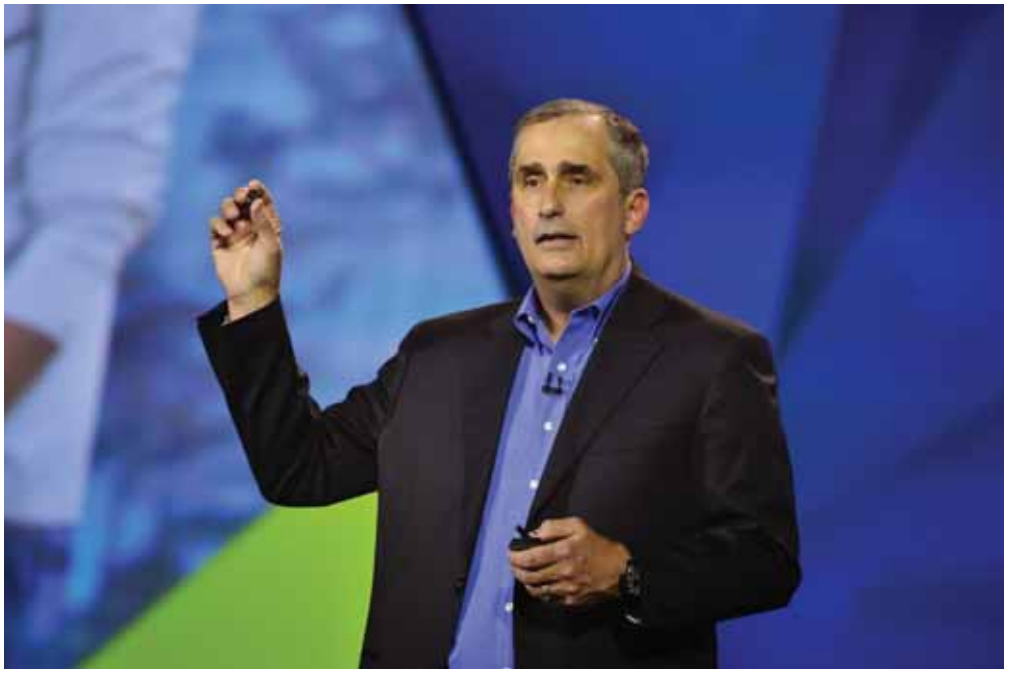

www.embeddedintel.com

We simplify the use of embedded technology

FAST AND COMPACT
conga-TC97

- COM Express Compact Type 6 module
- Dual-core 5th generation Intel® Core™ i7 processors
- 3x DisplayPort 1.2, up to 4k resolution
- Intel® Advanced Encryption Standard New Instructions (Intel® AES-NI)

HIGH END PERFORMANCE
conga-TS87

- COM Express Basic Type 6 module
- Quad-core 4th generation Intel® Core™ i7 processors
- 3x DisplayPort 1.2, up to 4k resolution
- Intel® Advanced Encryption Standard New Instructions (Intel® AES-NI)

POWERFUL AND SMALL
conga-MA3/conga-MA3E

- COM Express Mini Type 10 module
- Intel® Atom™ and Intel® Celeron® processors
- Gen 7 Intel® HD graphics
- Extended temperature range option

HIGH PERFORMANCE QSEVEN
conga-QA3

- Qseven module
- Intel® Atom™ and Intel® Celeron® processors
- Gen 7 Intel® HD graphics
- Extended temperature range option

Find more details at: **www.congatec.us**

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

intel IoT Solutions Alliance

congatec

MicroCloud

24/12/8 UP Server Nodes in 3U

High-Density · High-Performance · High-Efficiency · Easy-to-Service · Cost-Effective
Space Saving · Optimized for Cloud Computing and Web Hosting

SYS-5038MA-H24TRF
Up to 24x 2.5" SATA3 HDDs,
or 48x 2.5" slim SDDs
Intel® Atom™ C2750 series

24 UP Nodes / 12 Hot-pluggable sleds in 3U
SYS-5038ML-H24TRF
Intel® Xeon® E3-1200 v3 product family

12 Hot-pluggable Nodes in 3U
SYS-5038ML-H12TRF / SYS-5037MC-H12TRF
Intel® Xeon® E3-1200 v3 product family

8 Hot-pluggable Nodes in 3U
SYS-5037-MC-H8TRF / SYS-5037MC-H86RF / SYS-5037MR-H8TRF
Intel® Xeon® E5-2600 v2 / E3-1200 v3 / E3-1200 v5 product family

- High density with 24/12/8 UP nodes in 3U
- Excellent expansion capability with 1 PCI-E 3.0 x8 LP slot per node
- Up to 2x 3.5" hot-swap or 4x 2.5" fixed SATA3.0 (6Gbps) HDDs per node
- High efficiency Platinum Level (94%) 1+1 redundant power supplies
- IPMI 2.0 remote management plus KVM with dedicated LAN per node
- Intel® Atom™ C2750 series, Intel® Xeon® E5-2600 v2 and E3-1200 v2 / 1200 v3 product families
- Up to 130W high-performance CPU support available

intel inside™ ATOM™

9 Digital! We Power the Cloud

SUPERMICRO X10 SERVER PLATFORMS

www.supermicro.com/MicroCloud

SMCI-20140616 - A

Intel is the Silent Standards Bearer

Intel works behind the scenes to create more standards than it gets credit for.

Chris A. Ciufo, Editor-in-Chief

One of the great things about Intel—the company that got us to think about the microprocessor "inside" our PCs—is how it works behind the scenes to make technology actually work. Much has been written over the last few months about the production delays of the Intel® Core™ and Intel® Xeon® processor (codename Haswell) 14nm Tri-Gate technology. Yet it's easy to forget that this is the company that accelerated production of Tri-Gate (FinFET) technology at 22nm and beat its own estimates. As a result, we have 3rd Generation Intel Core i7 processors and server-class Intel Xeon processors that no one can match by ASP, performance, or $/watt.

Taken-for-granted technologies like Wi-Fi, MiraCast, and even Android were catalyzed by Intel's engineers. Wi-Fi, for example, only became prevalent when the Intel® Centrino® processor technology initiative (remember that one?) coupled the CPU with Internet connectivity. MiraCast is now the technology that lets Android-based tablets, phones and the ChromeCast stick stream wireless videos to a local screen. This started out as Intel® Wireless Display (Intel® Wi-Di). And finally, Intel has more developers working on Android than anyone besides Google (maybe more). That's because the company is winning smartphone designs with Bay Trail and SoC-based Intel® Atom™ processors, plus setting its sights on any GUI-enabled doodad such as automotive infotainment.

And don't forget that if running embedded Linux on your favorite SoC is in your plan—even if it's not an Intel® x86-based SoC—the Yocto standard for creating that Linux distro is the brainchild of some really smart Intel engineers and developers. What about PCI Express? You get the idea: Intel incubated that, too, from an idea called "NGIO" in response to competing LVDS proposals.

At the Intel® Developer Forum 2014, Intel took the wraps off even more technologies designed to disrupt our world. The tiny Edison Compute Platform (Quark SoC/Wi-Fi/Bluetooth) is positioned to compete against other popular low-power architectures while bringing in Intel's humongous x86 ecosystem of code and connectivity. The recently evolved Intel® Internet of Things Solutions Alliance puts Intel's arms around an entire ecosystem of embedded products and software—including intelligent gateways—to bring reality to smart sensors and previously "dumb" endpoint nodes.

The company recently announced the Intel® IoT Platform, designed to facilitate end-to-end IoT connectivity. Standards are sure to follow as various Intel IoT Solutions Alliance partners work this comprehensive initiative, including Taiwan's ADLINK, along with cloud providers like Booz Allen Hamilton.

Bringing in Wind River for device configuration and McAfee for security (embedded and enterprise), Intel is also delivering a roadmap of integrated hardware and software products that spans from edge devices out to the cloud. Ever the standards bearer, expect Intel to create specs for API management and service creation software, edge-to-cloud connectivity and analytics, intelligent gateways and a full line of scalable x86 processors.

Chris A. Ciufo is editor-in-chief for embedded content at Extension Media, which includes the EECatalog print and digital publications and website, Embedded Intel® Solutions, and other related blogs and embedded channels. He has 29 years of embedded technology experience, and has degrees in electrical engineering, and in materials science, emphasizing solid state physics. He can be reached at cciufo@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

Managing Editor - China
Jane Lin-Li

Contributing Editors
Dave Bursky

Creative/Production

Production Manager
Spryte Heithecker

Graphic Designers
Nicky Jacobson
Caldin Seides

Media Coordinator
Yishian Yao

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
Michael Cloward
mcloward@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
Rachael Evans

Special Thanks to Our Sponsors

ADLINK TECHNOLOGY INC. ARTESYN EMBEDDED TECHNOLOGIES congatec SUPERMICRO WinSystems An Employee-Owned Company

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

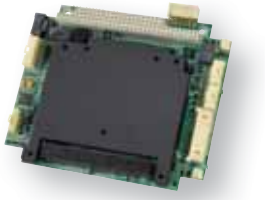

Intel® Atom™ Processor-Based PC/104
Single Board Computers
Rugged, Stackable Form Factor

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.

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

Call 817-274-7553 or visit www.winsystems.com.
Ask about our product evaluation!

WinSystems®
An Employee-Owned Company

Embedded Intel® Solutions

WINTER 2015

DEPARTMENTS

SPECIAL FEATURES

LAST WORD

PRODUCT SHOWCASE

ON THE COVER: Brian Krzanich, Intel's Chief Executive Officer, gave a keynote speech at the 2015 International CES in Las Vegas, Nevada. While not all the technologies described in this issue would come to mind when one thinks "consumer electronics," a surprising number of articles describe needs that are common across the embedded world--consumer or not--security, connectivity, and future-proofing. All photos courtesy Intel®.

Securing Embedded Devices in the IoT Era

In this corner, a deeply embedded device with limited resources. In the other, pervasive determination to exploit connectivity for the purpose of getting up to no good.

By Daniela Previtali, Wibu-Systems and Michael Weinstein, Wind River

In a recent study, Spanish security researchers reported that smart meters installed by a utility in Spain to meet government energy efficiency goals lacked basic safeguards, leaving room for hackers to carry out billing fraud or even cause blackouts. Weak encryption used in these smart meters allowed the researchers to get hold of the encryption keys used to scramble all the information that the smart meter shares with "nodes" sitting higher in the power distribution system. Using the keys and the unique identifier associated with each meter the researchers were able to spoof messages being sent from the power-watching device to a utility company and make the smart meter under-report the energy use. Shared IDs, poor protection against tampering and data formats that would be easy to fake have been identified as problems for smart meters deployed in other countries, such as the US and the UK, too[1].

With Secure Boot, the foundation of trust has been established, but the device still needs protection from various run-time threats and malicious intentions.

Just in 2014, multiple data breaches at JPMorgan Chase, Home Depot, Albertsons and others compromised in excess of 150 million accounts in the US alone. Piracy and reverse engineering of embedded devices and software also remain a big issue that costs embedded device vendor billions in lost revenues. A German Engineering Federation (VDMA) study indicates that 9 in 10 companies with over 500 employees are affected by piracy that caused €7.9 billion in losses for the German economy in 2013 alone. In 51 percent of the cases, the complete machine was subject to plagiarism[2].

Clearly, the importance of connected embedded systems being impermeable to cyber-attacks, acts of industrial sabotage and data theft has become paramount. But how can one safeguard deeply embedded endpoint devices that usually have a very specific, defined mission with limited resources available to accomplish it? Embedded devices are designed for low power consumption, with a small silicon form factor, and often have limited connectivity options. They typically have only as much processing capacity and memory as needed for their tasks. And they are often "headless"—that is, there isn't a human being operating them who can input authentication credentials or decide whether an application should be trusted; they must make their own judgments and decisions about whether to accept a command or execute a task. For example:

- In factory floor automation, deeply embedded programmable logic controllers (PLCs) that operate robotic systems are typically integrated with the enterprise IT infrastructure. How can those PLCs be shielded from human interference while at the same time protecting the investment in the IT infrastructure and leveraging the security controls available?

- Similarly, control systems for nuclear reactors are attached to infrastructure. How can they receive software updates or security patches in a timely manner without impairing functional safety

1. Smart meters can be hacked to cut power bills (BBC news) http://www.bbc.com/news/technology-29643276
2. Wind River Blog Hackers, Crackers, and Pirates: How to Protect Embedded Devices in the Internet of Things http://blogs.windriver.com/wind_river_blog/2014/10/hackers-crackers-and-pirates-how-to-protect-embedded-devices-in-the-internet-of-things.html

or incurring significant recertification costs every time a patch is rolled out?

- IoT sensor hubs aggregate a representative data set from numerous packets of sensed data. How can these real-time operating system (RTOS)-based devices open those packets, validate their integrity, analyze their contents and verify that these actions have taken place securely without compromising the speed and performance?

The answer is in designing systems for security from the start and incorporating a comprehensive set of security features to efficiently and effectively protect devices and data throughout their lifecycle.

Designing for Security

Security cannot be thought of as an add-on to a device, but rather as integral to the device's reliable functioning. Software security controls need to be introduced at the operating system level, take advantage of the hardware security capabilities now entering the market, and extend up through the device stack to continuously maintain the trusted computing base.

Building security in at the OS level is critical, since adding it at the user or application level is ineffective, expensive and risky. Enabling security at the OS level can also take the onus off device designers and developers to configure systems to mitigate threats and ensure their platforms are safe.

Protecting Devices at Every Stage

Security must be addressed at every stage—from boot-up to operation to data transmission to powering down (Figure 1). Being able to add hardware-based security to software-only features can help significantly harden device security overall.

Secure Boot

When power is first introduced to the device, the authenticity and integrity of the software on the device must be verified using cryptographically generated digital signatures to prevent the injection and execution of malicious code. In much the same way that a person signs a check or a legal document, a digital signature attached to the software image and verified by the device ensures that only the software that has been authorized to run on that device, and signed by the entity that authorized it, will be loaded. Binaries must be verified at every stage of the boot-up process. If a component fails to pass signature verification, boot must stop.

Runtime Security

With Secure Boot, the foundation of trust has been established, but the device still needs protection from various run-time threats and malicious intentions. Preventing unauthorized execution and other forms of tampering with system code is a critical component for securing devices in operation. A solution that can decrypt (using AES or other encryption) and verify digital signatures (using Elliptic Curve Cryptography (ECC), for example) of downloadable kernel modules and real-time processes can effectively protect the integrity of the system and safeguard intellectual property from piracy and code from reverse engineering.

Access Control

User management features are required to safeguard devices from unauthorized access and enable the definition and enforcement of user-based policies and permissions, implementing restrictions and controlling access to the device based on user credentials.

Network Security

It is critical for a connected device to incorporate features to effectively secure network communications using technologies such as SSL (Secure Sockets Layer protocol), SSH (Secure Shell protocol), IPsec and IKE.

Data Protection

Technologies such as encrypted containers can help safeguard data when the device is powered down, as data in

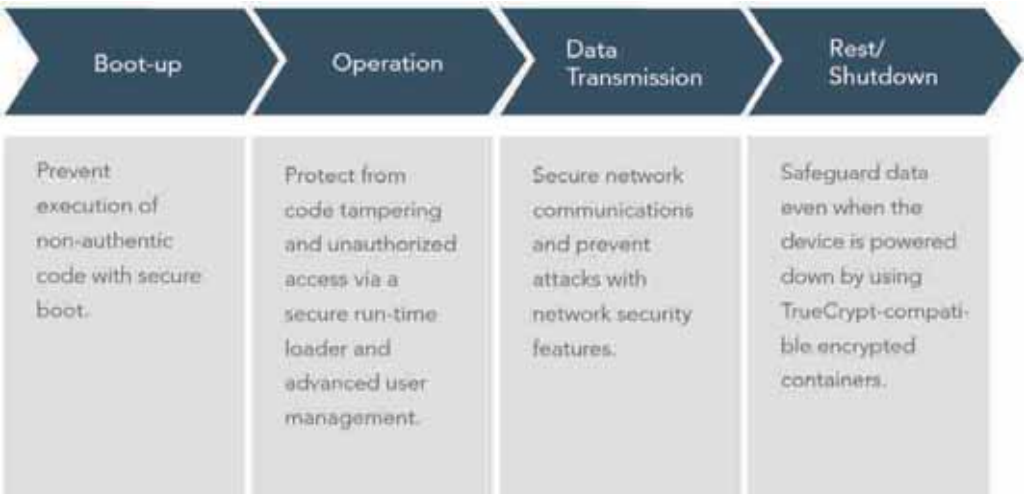

Figure 1: Security throughout the embedded devices' lifecycle.

containers remain encrypted even when the device is idle or powered off.

Security and Performance In Balance

In today's demanding market, a controller must not only deliver maximum performance, but also provide seamlessly integrated security. Strict security policies and potent firewalls prevent unauthorized intrusions. Communications and data exchange are subject to additional scrutiny through separate processes. However, the control of a smart plant needs to go beyond this traditional paradigm and offer new security features well suited to harsh environments, and fully reliable for industrial processes.

Let's imagine a European power grid vendor that develops its own PPC-based controller running VxWorks. The production of such hardware, including the download and testing of the related firmware, would be carried out in China. It becomes mission critical to transfer the license from the vendor headquarters to the production facility through a secure channel, while maintaining full control over the workflow.

The vendor can easily attain protection against know-how piracy, reverse engineering and tampering, by introducing Security Profile and CodeMeter Security. In particular, IP protection would be achieved through the combined use of CodeMeter and the Secure ELF Loader from Wind River and Wibu-Systems. Reverse engineering protection would be ensured by CodeMeter high encryption standards, which would make it impossible to analyze the ciphered firmware. Copy protection would be the result of CmActLicense, the soft license container bound to a secure element on the embedded system. Tamper protection would be reached through code signing operated by authorized team members only, secure boot and signature verification performed by the Secure ELF Loader in VxWorks. The staff would in fact own CmDongles, in the form of CmStick for USB ports and CmCards for SD slots; the private key would then be securely stored in the smart card chip, and the dongles would be configured for use with a password, set to expire after a pre-determined time.

The return in investment stems not just from ensuring optimal security standards, but also from redesigning the licensing blueprint, and introducing scalable business models based on logistic efficiency and feature on-demand dynamics, which could be realized with the Secure ELF Loader and CmActLicense.

A Safe and Secure RTOS for IoT

Powering billions of embedded devices, VxWorks® is the world's most widely deployed real-time operating system. Enhanced by Security Profile for VxWorks, the RTOS provides a comprehensive set of software-based security features that enable manufacturers of intelligent embedded devices deliver security in their products. The expandable, upgradable architecture of VxWorks separates the core kernel from middleware, applications, and other packages, enabling bug fixes, upgrades, and new feature additions to be performed as frequently as necessary and without disrupting other technologies in an installation.

Security Profile for VxWorks is readily expandable and can be enhanced with Wibu-Systems' CodeMeter® hardware-based security to enable a comprehensive solution for security-sensitive applications. With software and hardware components as well as activation-based licensing, the joint solution delivers an optimal way to protect devices, data and IP in the Internet of Things.

Daniela Previtali is a global marketing manager at Wibu-Systems, responsible for both corporate and channel marketing strategy and activities.

Michael Weinstein is senior product marketing manager at Wind River driving the product marketing efforts pertaining to the VxWorks real-time operating system and the global automotive business.

Video Analytics Alter the Digital Signage Landscape

The intersection of the Internet of Things (IoT) and the embedded ecosystem is just one reason video analytics is energizing the digital signage sector.

By Anne Fisher, Managing Editor

congatec AG is an associate member of the Intel® Internet of Things Solutions Alliance and collaborates closely with Intel.

Editor's note: Our thanks to Dan Demers, director of marketing -Americas, congatec, who recently explained how a number of developments are helping digital signage systems capitalize on video analytics.

Dan Demers, director of marketing—Americas, congatec

EECatalog: Digital signage is certainly a growing market segment. Are there any technologies or market demands that are contributing to this growth?

Dan Demers, congatec: I think it is fair to say that by now we have all noticed the vast increase in LCD screens that provide us information—from those in grocery stores to roadside signs. The digital approach offers retailers and suppliers of information more options and increased flexibility when it comes to attracting and informing us. Adding to these options today is the field of video analytics, which is opening the next realm of possibilities for those companies trying to reach and understand consumers. Video analytics has a role to play as well in enabling companies to predict what it is we want to see and learn more about. An organization can also use video analytics to learn more about itself. As video analytics become more sophisticated, they will become a major vehicle for true communication between information provider and information receiver. Today's embedded computer platforms are poised to provide the necessary backbone to make video analytics-based systems a reality on a large scale.

EECatalog: What are some of the key requirements for the technology needed by digital signage companies that participate in video analytics?

Demers, congatec: In the not so distant past, systems with the type of horsepower needed to execute video analytics software were quite expensive and often not ideal in size for many installs and applications. In addition, many platforms lacked adequate security features. However, today many silicon providers offer low-power, multicore processors that are much lower in cost and substantially easier to package in smaller, lighter weight systems. What's more, many of these systems are much easier to package for extended temperature ranges, which opens up more opportunities for location installs of the signage. It is also common to see security features embedded in the silicon these days. When you step back and look at what is available in the embedded market for multicore, security rich, small form factor solutions, the number is quite high. For instance, COM Express modules that measure a mere 95mm by 95mm are housing the latest Intel® Core™ processors. That is a very small form factor that packs a lot of power, the kind of power needed for video analytics.

EECatalog: What other areas of concern can come into play for video analytics-based digital signage systems?

Demers, congatec: As mentioned previously, one area of concern can often be environmental. Not all digital signage systems are in temperature-controlled environments like shopping malls or doctors' offices. In these cases consumer-derived systems are not going to work, at least not for long. Environmental concerns make a case for sourcing a truly embedded and ruggedized solution. In addition to the physical environment that the system resides in, the operation cycle is something that must be considered. If the organization is looking to install a system that runs 24/7, again a consumer-derived system will not be adequate. In both of these cases using a system that has a higher level of design concern is necessary.

EECatalog: How do embedded systems benefit video analytics based digital signage systems?

Demers, congatec: Unlike basic digital signage systems that often utilize consumer-oriented hardware, video analytics systems need robust and long-life product. Embedded systems,

such as those based on Computer-On-Modules, offer flexibility and faster time to market. Obviously there is more than one company in the video analytics space, so like most industries, beating your competition to market with a product is always a desirable outcome. Embedded systems offer a nice base platform to build from in order to create that perfect fit solution. As I mentioned earlier, today's embedded systems often take advantage of security features native to the silicon. So, this layer of concern is also addressed in embedded systems for the digital signage supplier.

EECatalog: Where do you see digital signage moving to in the future?

Demers, congatec: Video analytics really do open the door for a lot of additional information to be obtained. Based on that, I expect to see even more sophisticated video analytics features being integrated into digital signage. At some point, I would expect nearly all digital signage systems to have some form of video analytics capability. As the costs of components and embedded solutions fall and the capabilities increase we will see more opportunity for the proliferation of video analytics. Because of the sheer number of potential installs, I also believe silicon providers and embedded solution providers will focus more and more on this area. I once heard that the business goal for Coca-Cola was to have a Coke within arms reach of everyone on the planet. Not sure that was actually a goal of Coca-Cola, however, I got the point, instant availability. I believe the digital signage market has a similar goal. I am sure they would like to have digital signage in as many places as possible and be able to reach consumers and people in need of information as quickly as possible. As the cost of goods decreases, this opens up opportunities to put digital signage in places it does not exist today. We hear a lot about IoT these days. When you break down a digital signage system with video analytics capabilities it is obvious that this is a system that can heavily benefit from IoT features found in the embedded space. We are talking about low-power, long-life, secure and flexible solutions. All based on platforms that have proven IoT track records and don't mandate that the digital signage supplier recreate the wheel.

Dan Demers is Director of Marketing for the Americas at congatec Inc. He holds a B.B.A. degree in International Business from Grand Valley State University, Grand Rapids, Michigan and an M.B.A. from Ashford University, Clinton, Iowa. Demers has over 13 years of experience in embedded computing having worked with Fortune 500 companies in the medical, military, and communications markets.

After the Operating Room Surgeons Have Left…Medical, Industrial and Harsh Panel PCs Remain

From mobile to medical, ADLINK and its German PENTA division get ruggedly specialized—such as panel PCs that can be decontaminated after a messy surgery.

By Chris A. Ciufo, Editor-in-Chief

ADLINK has changed over the past couple of years from a "fast follower" offering primarily Intel® technology-based, open-standard single-board computers (SBCs) to a company rife with intellectual property (IP), unique products and a mission to bring new products to market segments such as Industrial and Medical.

The objective of hospital containment protocols is to minimize the spread of deadly diseases. Certainly a key part of those protocols is extensive decontamination of ER/OR equipment, including instruments and the essential medical PCs. Not just any piece of electronic equipment can sustain a literal hose-down, chemical flush, or aggressive wipe-down with FDA sanitizer. This is the territory for seriously rugged medical, industrial and food safety PCs by PENTA ADLINK.

"PENTA" comes from the Greek word meaning "five." Let's briefly take a look at five key characteristics that strike me as noteworthy compared to a "regular" consumer-grade panel PC:

1. Extended temperature and fanless
2. Rugged by design, but dying for your touch
3. IP rated for exceptional liquid ingress protection
4. Applicability to specialized markets like medical, food processing and harsh industrial
5. Mil-Spec "ilities" at the corporate level, similar to a DoD defense contractor

#1: Fanless, not Headless

PENTA ADLINK's panel PCs are used in some pretty tricky applications. From their literature, examples abound in factories such as Geobra Brandstätter, a German manufacturer, of 2.6 billion Playmobil figurines per year.

Figure 1. Rugged, sealed panel PC used in manufacturing.

Installation in "unheated rooms, cool warehouses…heated rooms and the spraying department" implies operation beyond 0 – 25 °C. In fact, PENTA ADLINK's Giant Series 17- to 19-inch PCs operate from 0 – 50 °C by careful component selection to minimize power consumption—in this case, an Intel® Atom™ processor replaces an Intel® Core™ processor.

At cold ambient temps, component self-heating and strategic PCB layouts can warm PLLs after a few moments to facilitate timing lock and system boot. Note that these kinds of PCs typically operate 24/7 and are rarely booted—high MTBFs are a given. In the case of the Geobra factory, the Giant Series PCs boast over 50,000 hours MTBF.

At high T's, heat must be dissipated without ventilation slots or fans. This includes the internally mounted power supply (110/220VAC or 24VDC): it's not possible to use "wall wart" bricks in industrial installations. Careful placement of the warmest components to minimize the thermal path to the panel's case facilitates component conduction cooling. Another neat trick is using passive

Figure 2. Cart-mounted medical PC is fully sealed and able to withstand disinfectants and "substantial" liquids.

custom-designed heat pipes that conduct energy away from components with a longer thermal path to the case edges.

This technique is employed in the company's HPERC rugged "shoebox" family and also in PENTA ADLINK's PCs. Flat heatspreaders with a plain external surface make the system easy to clean without any cooling-ribs where dirt can gather.

As well, buried thermally conductive layers embedded in PCBs plus top-side conduction ribs can move heat to the case. LED-backlit LCD panels (rather than fluorescent lighting) also minimize heat. Collectively, these mechanical and other electrical design techniques keep heat within specs without necessitating costly extended temperature ICs.

#2: Rugged Enough to Touch

Panel PCs mounted on trollies, trucks, forklifts, or just rolling carts take a beating. Sure, there's a smooth concrete or linoleum floor right up until the equipment is banged into a wall, expansion joint, elevator threshold, or "carefully" jockeyed down the stairs. Stiffened cases prevent flex while enhancing thermal conductivity, memory components are soldered instead of being on fragile DIMMs, and items like connectors and standoffs are carefully selected to maintain rigidity and conductivity.

Ruggedness extends beyond just mechanical constraints. To assure reliability in the face of possible environmental adversity, dual BIOS PROMs afford boot redundancy in the event of a BIOS corruption (or to roll back to a known good revision).

Human machine interfaces (HMI) in factories or in hospital isolation wards rarely rely on mechanical keyboards because they can't be sufficiently cleaned or disinfected. While smooth membrane keyboards are occasionally employed, the tablet revolution has spawned on-screen virtual keyboards and buttons using rugged LCD overlays. The trouble is that neither latex nor other kinds of gloves work with traditional capacitive touchscreens. Instead, PENTA ADLINK equips panel PCs with glove-friendly resistive touchscreens. Precise inputs, such as making parts list entries, are done using a stylus or even a retracted ball point pen.

Figure 3. Close-up of wipe-down disinfecting process. No moisture will penetrate the IP65 interior.

Figure 4. Examples of food processing PCs, also fully sealed and supporting high-temp liquid washdown.

#3: Cry Me a River, my PC Stays Dry

Cooling a medical panel PC without a vent is impressive enough, but the real reason for sealing the enclosure is to prevent dust, dirt, and liquid ingress. The "Ingress Protection Marking" (IP) part of IEC 60529 has a two-digit code that grades solids (first digit) and liquids (second digit). In wet, harsh environments like oil exploration or a hospital ER, IP65 means a panel PC is "dust tight" (6) and protected against "water jets" (5).

This latter describes a three-minute test with 12.5 liters of water per minute at 80 – 100 kPa of pressure at a distance of 3 m with a nozzle having a diameter of 6.3 mm. Water can be sprayed from any direction with no harmful effects. PENTA ADLINK medical PCs such as model Medical i7 T withstand IP65 and are listed by the manufacturer as being "easy to clean; disinfectable." The company counts surgical equipment supplier Zeiss as a customer in the Opmi Lumera 700/Calisto Eye product used by ophthalmologist surgeons in OR applications.

For extraordinarily harsh industrial applications such as food preparation or in meatpacking factories, certain panels and panel PCs can withstand the top-of-the-line IP69k rating. Wikipedia.org describes this as "powerful high temperature water jets" which include high temperature spray downs. PENTA ADLINK indicates these products survive six side jets using 80 °C water at 1,450 psi. These products are impressive in applications requiring hygienic food safety preparation such as meat rendering plants.

#4 and #5: Application-Specific, with Mil-Spec Credibility

By now it should be obvious that these panel PCs are anything but ordinary. In fact, they remind this author of products designed for Mil-Spec, harsh environment applications like avionics or space craft. In many ways, they are very similar. Some models don't use plastic housings; instead, stainless steel (class 1.430 AISI 304) is used in select food and pharmaceutical applications.

In operating rooms, for instance, specialized medical PCs must protect patients and OR personnel from ESD, leakage currents caused by faulty earth grounds, or even ΔV between the equipment or I/O peripherals such as the hospital's Ethernet LAN. These kinds of medical PCs are certified under DIN EN60601-1 and have galvanically separated serial I/O and NICs. Even HMI interfaces like case-mounted keys or barcode scanners are isolated and certified.

Carrying the Mil-Spec analogy further, mother ship ADLINK uses military-like Highly Accelerated Life Test (HALT) to wring out new designs. Consisting of shock/vibration plus voltage and temperature margin testing, the gauntlet of tests is based upon MIL-STD-202 and other "four corner" stress tests routinely used on DoD/MoD contracts. In addition, HumiSeal 1B31 acrylic conformal coating is available should equipment need beyond-IP65 moisture protection.

These kinds of "ilities" (reliability, survivability, traceability, etc.) demonstrate that ADLINK and PENTA ADLINK equipment will still be running long after the surgeons—or oilfield roughnecks—have gone home.

Chris A. Ciufo is editor-in-chief for embedded content at Extension Media, which includes the EECatalog print and digital publications and website, Embedded Intel® Solutions, and other related blogs and embedded channels. He has 29 years of embedded technology experience, and has degrees in electrical engineering, and in materials science, emphasizing solid state physics. He can be reached at cciufo@extensionmedia.com.

This article is sponsored by ADLINK.

Trends and Drivers in Fail-Safe Architectures for Rail Systems

No veering off the tracks: The means to head straight for differentiation from the other guys is here now for rail application developers.

By Shlomo Pri-Tal, Artesyn Embedded Technologies

The market for embedded computing technologies in rail applications is following a similar trend as has been seen in other embedded market spaces. A layer of the technology value chain becomes 'table stakes'—delivering limited competitive advantage to a point that it makes sense for application providers to reallocate R&D resources to differentiating elements of the end product and buy the base technology from companies who are dedicated to that technology. We are witnessing this transition in the rail market for embedded computers that are certified to safety integrity level four (SIL4), the highest level. These embedded computers offer a certified, commercial off the shelf (COTS) generic fail-safe platform allowing rail application developers to focus their R&D resources on differentiating applications.

This trend is driven by a number of emerging trends in the global rail industry.

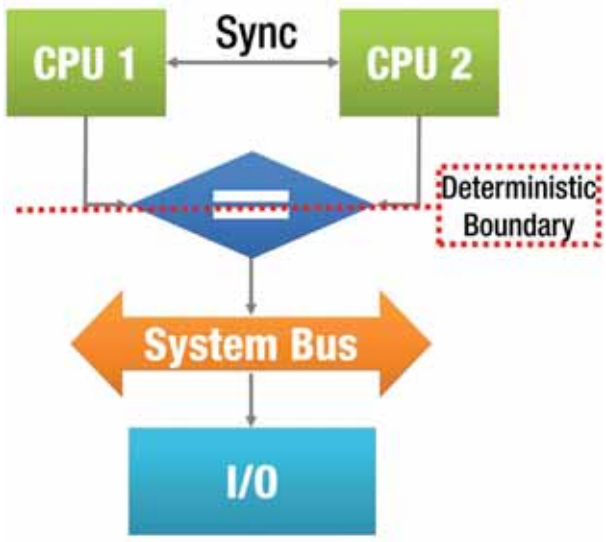

Figure 1. When operating in hard lockstep, the processors' clocks are synchronized and, before allowing a transaction to drive external equipment, all data and address bits driven by the two processors are compared.

In the past few years we have witnessed an explosive growth in global investments in public rail transportation, in particular high-speed rail and metro, caused mainly by the effort to reduce a nation's carbon footprint by replacing inefficient automobile-based transport with efficient mass transportation. This is particularly evident in emerging economies such as China and India, as well as established economies in the Far East, Africa and South America. While less so in Europe and North America, we do witness growth in these markets due to other factors such as pan-European rail standardization as well as modernization of the rail infrastructure to enhance safety.

The system runs Wind River's VxWorks 653 operating system, which has been deployed in many fail-safe avionics-certified applications....

However, a growing market, while creating an attractive target for COTS products, will not on its own cause an outsourcing trend. Additional safety, technical, and commercial factors come into play.

As train speed increases to 300 kilometers per hour and above, reliance on computers that control the rail infrastructure and the trains themselves increases exponentially. As an example, stopping a train that travels at 300 Km/h will only take 2 minutes or so, but during those two minutes the train will travel 10 kilometers, requiring real-time and continuous monitoring of the rail network to provide early alerts of potentially hazardous events. High-speed, high-availability, and fail-safe computer-based control equipment must be deployed to guarantee safe operation under

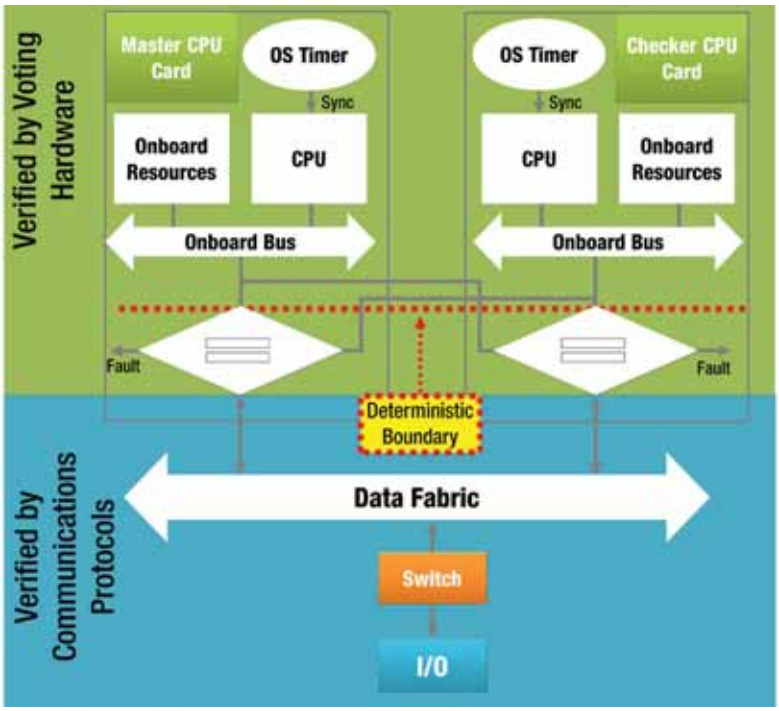

Figure 2. The deterministic boundary is not at the processor itself but rather at the edge of the processor and before packets are placed on the data fabric.

all conditions. High-performance and high-availability computing expertise is relatively widespread, however fail-safe computing has been the domain of a few expert companies, located mostly in Europe (Alstom, Bombardier, Siemens, etc.) for SIL4 certified systems, and Japan (Nippon Signal, Hitachi, etc.) for certification to Japanese safety standards and deployed locally. Fail-safe know-how has not been prevalent in other markets that are investing in rail networks, relying on mostly European vendors for acquiring the fail-safe systems (e.g., India, Africa, South Korea) or for forming joint ventures with these same European vendors to develop fail-safe systems for the local market (e.g., China).

The demand for SIL4 certified equipment has been further fuelled by safety incidents that have driven governmental bodies to make it mandatory for all new installations to be SIL4 certified, and that non-SIL4 certified equipment in use today must be upgraded to SIL4 certified equipment. For example, the South Korean government mandated that rail equipment be upgraded to SIL4, and the Indonesian rail authorities have recently issued an RFP to upgrade their infrastructure to SIL4 certified equipment.

Another interesting trend in the global rail market is the aspirations of Asian application providers and rail integrators to expand their reach and penetrate overseas markets. Witness Hitachi's establishment of a design center in London, recent announcements from Chinese vendors of wins in the US and Africa, as well as efforts by South Korean vendors to expand into former Soviet Union countries. Almost without exception, SIL4 certified equipment is a mandatory requirement.

A few major and factors emerge from these trends that are the root cause for the emerging trend to outsource SIL4 certified application platforms:

1. The lack of SIL4 development expertise by Asian rail application providers and the barrier that poses to aspirations to expand into overseas markets.

2. The threat to western vendors posed by the entry of Asian vendors into the global rail market and the price erosion that would likely bring (witness the impact Huawei had on the global telecom market).

3. The prevalent architecture implemented by existing fail-safe computers is no longer capable of handling the required performance, requiring an expensive development effort in 'table stakes' base technology.

Lockstep Architectures

Most rail systems today use an architecture called hard lockstep, whereby two processors execute the same instruction at the same time and drive their respective address and data buses in synchronization.

When operating in hard lockstep, the processors' clocks are synchronized and, before allowing a transaction to drive external equipment, all data and address bits driven by the two processors are compared. If the bits are exactly the same, then the address

Figure 3. The implementation of a dual 2oo2 architecture to deliver high availability.

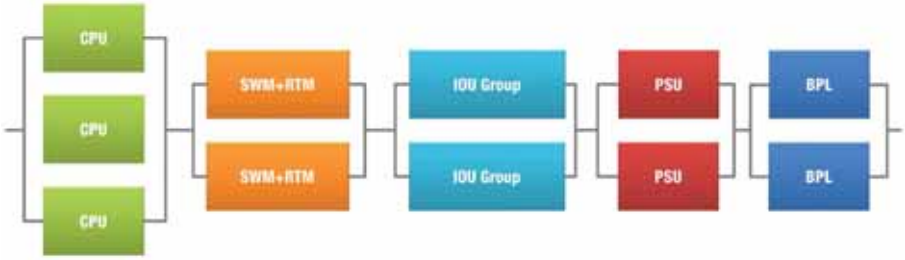

Figure 4. In 2oo3 voting, three computing elements execute the application, and if the three don't agree then the system determines which one is at fault, disables it, and continues running with two.

Dual Redundant 2oo2 System	Feature	Single Redundant 2oo3 System
2oo2 Only	Voting Logic	Must switch from 2oo3 to 2oo2 and back to 2oo3
Box Fail-over	HA Model	Module Fail-over
Not Required	On-line Fault Isolation	Required
Not Required	Hot Swap	Required
Not Required	On-line Module Reintegration	Required and performed by User Application
Failing CSC is diagnosed off-line; low risk of total outage	Human Factors	Risk of total outage due to hot-swapping wrong module
Simplex	Backplane	Redundant

Figure 5. Voting method comparison.

and data information are allowed to change the state of external equipment. If they do not compare, then a failure is declared and the system is brought to a safe state and is prevented from driving external equipment.

Since, in hard lockstep, comparison is performed at the address and data bits of the processors, a primary and mandatory requirement is that the two processors must execute the same instruction, at the same time, to the same external resources (memory, cache, I/O, etc.). To do so, the processors themselves must be deterministic. We call the boundary created by the comparators the deterministic boundary (Figure 1).

Unfortunately, hard lockstep cannot be implemented using modern processors. The first problem is that modern processors do not guarantee deterministic behavior.

Multi-threading creates multiple paths for the execution of the program. Responses to soft errors in memory and I/O will cause divergent execution paths and timing. For example, errors that are caused by cosmic rays and change a bit in the register are not synchronized and not deterministic. This is more prevalent in current technologies because of the geometries of the transistors, which are so small that cosmic rays can flip bits. Also, other CPU features such as power management and cache operations introduce non-determinism.

The second problem is that it's practically impossible to synchronize the data pairs of two different modern CPUs. The use of on-chip devices to multiply clocks prevents synchronized operation, multiple memory channels and serial peripheral interfaces also make it impossible and it's not practical to synchronize buses operating in excess of 1Ghz.

Another problem with a hard lockstep system is that it is fundamentally a closed system. Everything is tuned to work together, and it has to be all synchronized such that it's very difficult to upgrade technologies without affecting the total system. So the bottom line is that hard lockstep is just not possible any more with advanced processors.

Artesyn has developed an alternative approach we call data lockstep architecture, whereby a deterministic boundary is created at the output stage of the processor board to the system data fabric that connects the processors to external devices. Before the processor boards are allowed to change the state of external equipment by driving packets on the data fabric their packets are compared to ensure that they are the same. If they are the same, then the transaction is forwarded to external equipment; if the packets do not compare, then a failure is declared, and the system fails safe; i.e., it is prevented from changing the state of external equipment.

As shown in Figure 2, the deterministic boundary is not at the processor itself but rather at the edge of the processor and before packets are placed on the data fabric.

The benefit of data lockstep is that it makes it possible to use modern processors and deliver the performance required by modern rail applications.

2oo2 or 2oo3

There are two methodologies for voting in a fail-safe system. They are called two-out-of-two (2oo2) and two-out-of-three (2oo3).

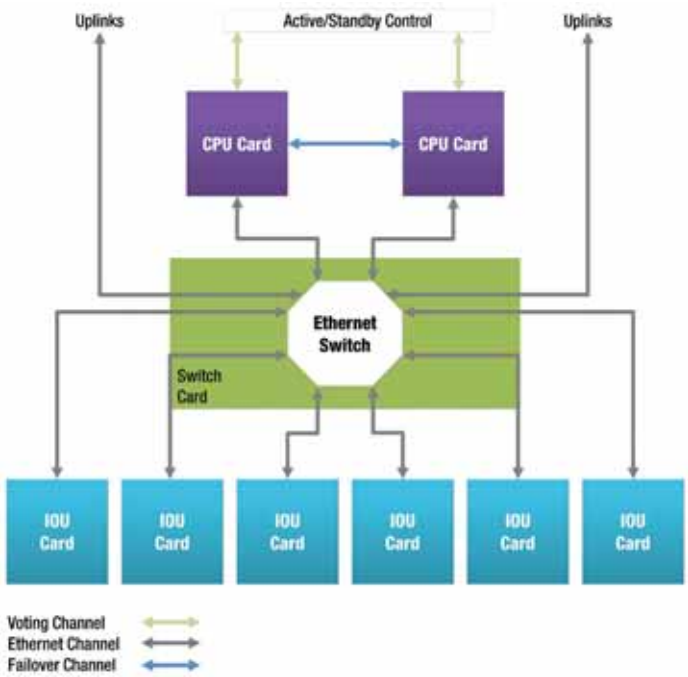

Figure 6. All I/O modules are connected via Ethernet such that expanding the system from local to remote or expansions in the I/O environment is straightforward and scalable.

Figure 7. The ControlSafe Platform uses an OS, VxWorks 653 from Wind River, with a track record in numerous fail-safe applications, including avionics.

In 2oo2 voting, two computer elements compare the results of their computation and, if they compare, the transaction is driven to external equipment. If they don't compare, a fail-safe state is entered. As shown in Figure 3, Artesyn's ControlSafe™ Platform implements a dual 2oo2 architecture to deliver high availability. In case the first ControlSafe Computer fails, the second redundant one takes over and continues running the application.

In 2oo3 voting, three computing elements execute the application, and if the three don't agree then the system determines which one is at fault, disables it, and continues running with two.

If the two disagree, then the system enters its fail-safe state and is prevented from changing the state of external equipment.

While both of these voting methods deliver the required safety and availability, the 2oo3 method is more complex to implement than the 2oo2 method. In the 2oo2 method, in case of a mismatch, the failed CSC enters its fail-safe state and the second CSC is enabled to run the application. No failure analysis, or fault isolation, hot-swap or re-integration is required.

On the other hand, in case there is a mismatch in a 2oo3 voting, failure analysis, fault isolation, switching to 2oo2 voting mode, module hot-swap, module reintegration, and re-enabling 2oo3 voting are all required. This is complex, and complexity leads to design errors.

For this reason, Artesyn's ControlSafe Platform chose the 2oo2 voting method. A simple design is a safe design.

ControlSafe Architecture Highlights

Artesyn's ControlSafe Platform employs data lockstep synchronization and 2oo2 voting. The system runs Wind River's VxWorks 653 operating system, which has been deployed in many fail-safe avionics-certified applications, including extensions to assure the task level synchronizations needed to implement data lockstep.

All voting is implemented by hardware using proprietary FPGAs, making it transparent to application software, and easing porting of existing applications.

The architecture is flexible and expandable. All intra system communications are over the data fabric and are based on Ethernet. All I/O modules are connected via Ethernet such that expanding the system from local to remote or expansions in the I/O environment is straightforward and scalable (Figure 6).

In conclusion, the ControlSafe Platform is a cost-effective, modular and a scalable system that is based on open industry standards. The system is future-proof and provides protection for the customer's investment because the architecture enables upgrades to both the CPUs and the I/O modules independently of each other.

It is designed to offer a COTS SIL4 certified platform bringing to customers all the benefits of outsourcing table-stake technology – accelerated time to market, significant savings in R&D and certification costs, and the ability to focus their effort and their R&D on differentiations from their competitors.

Shlomo Pri-Tal is the Vice President of ControlSafe Platforms at Artesyn Embedded Technologies. Artesyn Embedded Technologies is the new name for the former Embedded Computing & Power business of Emerson Network Power. In his present role, Shlomo directs all activities required to develop and bring to market a SIL4 certified fault tolerant computing and communications platform targeted at safety critical applications.

Artesyn Embedded Technologies is an Associate Member of the Intel® Intelligent Systems Alliance.

PCI Express Receiver Testing Responds to New Challenges

The Internet of Things (IoT) encompasses a lot, but it has no room for slackers in the data rate increase department. One outcome: test equipment with advanced features to handle PCIe Receiver complexity.

by Thorsten Götzelmann, Keysight Technologies, Inc. (formerly Agilent Technologies electronic measurement business)

The receiver (RX) specification for PCI Express evolved with specification revisions and data rate increases. For instance the reference point for receiver parameters is for 2.5GT/s and 5GT/s outside the chipset hosting the PCI Express RX, while for 8GT/s and 16GT/s the reference point is within the chipset. The requirements for a stress signal used to test a receiver are developed to a higher detail and complexity level. Three major PCI Express specification layers and different device under test (DUT) categories as well as different operation modes result in setup and test differences. Mastering this complexity to set up a receiver test bench for a specific test can be challenging.

PCI Express Specifications

Base Specification

The PCI Express Base Specification is the foundation for the PCI Express specification framework. From a physical layer perspective it specifies transmitter, channel and receiver parameters as well as possible clocking architectures and the logical sub block. Every PCI Express use model refers to the base specification. Most relevant for chipset testing, the base specification supports synchronous operation as well as asynchronous operation. Three different types of clocking architectures are possible: common reference clock (CC), which is synchronous; data clocked (DC), which can be synchronous or asynchronous, or independent reference clock (IR), which is asynchronous. Originally, asynchronous operation was allowed in the absence of spread spectrum clocking (SSC) only. Asynchronous operation in the presence of SSC was introduced early in 2013. Separate reference clock no SSC (SRNS) is used to describe asynchronous operation without SSC and separate reference clock independent SSC (SRIS) is used to describe asynchronous operation with SSC. Different test requirements are defined for 8GT/s and 16GT/s RX testing for synchronous and asynchronous operation.

CEM Specification

The largest PCI Express use model is most likely the PCI Express extension slot, which is defined by the PCI Express Card Electromechanical (CEM) specification. Two different device types need to be considered for receiver testing: add-In cards (AIC) and mainboards (system). CEM uses synchronous operation only, and it is the only PCI Express ecosystem providing a mandatory compliance certification program, including physical layer tests. Every device found on the PCI-SIG integrators list had to pass compliance testing at one of the PCI-SIG compliance workshops. The PCI Express Architecture PHY Test Specification (CTS) defines the required physical layer compliance tests. Because CTS tests are designed to be manageable in a workshop environment, receiver testing according to the CTS can be less stringent and less complete than receiver testing according to the base specification.

M-PHY Test Specification

M-PCIe replaces the physical layer of PCI Express with the PHY layer defined for M-PHY. Receiver testing therefore has to be performed according to the M-PHY specification and not the PCI Express specification. M-PHY is a physical layer defined by the MIPI® Alliance.

PCI Express Receiver Test Requirements

Test requirements and calibration methodologies are not the same for the different transfer rates. With revision 3.0 the specification reference point moved into the chip, and composition of the stress signal became more complex. In addition, the methodology describing the inter symbol interference (ISI) channel to be used for receiver testing differs for 2.5GT/s, 5GT/s and 8GT/s / 16GT/s.

The test requirements also specify backward compatibility. Therefore a device capable of higher data rates must also be compliant at the lower data rates.

PCI Express 2.5GT/s

Base Specification	CEM Specification	PHY Test Specification
Yes	Yes	No

Receiver specifications are defined at the receiver pins. Specifications are identical for different clock architectures and synchronous or asynchronous operation modes. A simple receiver mask is defined only. In the absence of a random jitter (RJ) specification, solutions today usually use the RJ defined for 5GT/s. Base specification testing does not require de-emphasis, but testing according to CEM does.

Stressor mix:

- ISI via an external channel. ISI should be the major deterministic jitter (DJ) component. CEM testing requires the PCI-SIG Compliance Base Board (CBB) and Compliance Load Board (CLB). CBB for gen1 and gen2 needs to be modified for receiver testing.
- RJ
- Sinusoidal jitter (SJ) to supplement ISI for necessary eye closure
- Common mode sinusoidal interference (CM-SI), base specification only

PCI Express 5.0GT/s

Base Specification	CEM Specification	PHY Test Specification
Yes	Yes	No

Receiver specifications are defined at the receiver pins. The base specification defines different parameters for CC and DC based receiver designs. The CEM specification does not apply CM-SI but adds a secondary high frequency jitter tone. Residual SSC (rSSC) is introduced for CC use cases; rSSC is a triangular phase modulation which is applied to the stressed data signal only but not to the reference clock. It represents the worst case delta a receiver can experience between SSC on the reference clock and SSC on the incoming data signal.

Stressor mix:

- ISI via an external channel. ISI should be the major DJ component. CEM testing requires the PCI-SIG Compliance Base Board (CBB) and Compliance Load Board (CLB). CBB for Gen1 and Gen2 needs to be modified for receiver testing.
- Spectral filtered RJ (sRJ) with higher RJ amplitude for frequency spectrum up to 1.5 MHz and lower RJ amplitude for frequency spectrum between 1.5 MHz and 100 MHz.
- SJ to supplement ISI for necessary eye closure
- SSC:
  - o rSSC is used for CC based implementation except for CEM based system tests, since SSC is defined by the system's reference clock.
  - o SSC is used for DC based implementations
- CM-SI, base specification only
- Secondary high frequency SJ tone for CEM specification only

PCI Express 8.0GT/s

Base Specification	CEM Specification	PHY Test Specification
Yes	Yes	Yes

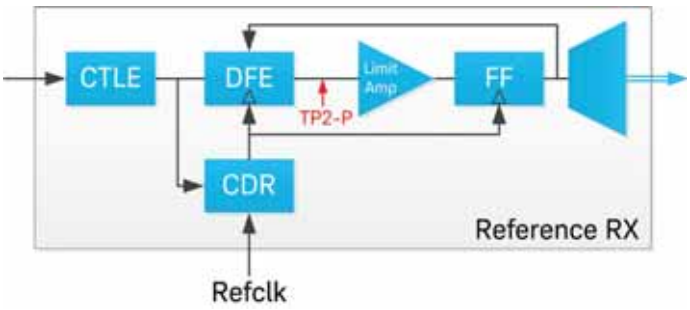

Figure 1. 8GT/s RX specification reference point TP2-P.

The increased transmission rate via basically the same channel makes RX equalization necessary, and, consequently, testing of receiver gain more important. The receiver specifications are more thorough and are defined within the receiver after CTLE and DFE. This reference point is called TP2-P. As a consequence of this definition point shift, embedding of a behavioral RX package as well as simulation of the equalizer stages and clock recovery are required for the stress signal calibration.

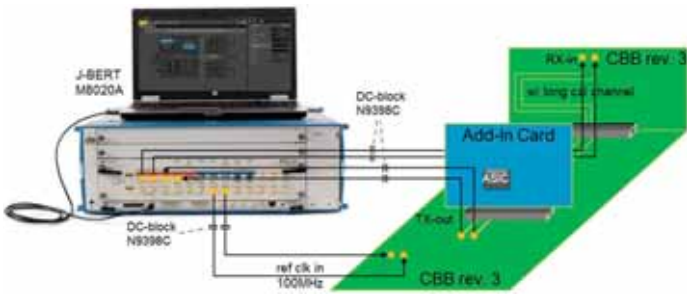

Figure 2. Example setup for PCIe 8GT/s Add-In Card (AIC) receiver test.

Stressor mix:

- ISI via an external channel. CEM testing requires PCI-SIG Compliance Base Board (CBB), riser and Compliance Load Board (CLB) for Gen3 for long channel and CBBGen2 for short channel test.
- RJ with 10 MHz high-pass filter applied

- SJ, different jitter tolerance masks for CC and SRNS/SRIS
- SSC, SRIS only:
  - o Triangular down spread @ 33 kHz for stressed voltage test
  - o Sinusoidal down spread @ 33 kHz for stressed jitter test
- DM-SI
- CM-SI, base specification only

Next to improvements of physical capabilities of the PCIe receiver, a TX equalization optimization procedure, which takes current channel, transmitter and receiver characteristics into account was added to the link training.

PCI Express 16GT/s Outlook

Base Specification	CEM Specification	PHY Test Specification
Yes	Yes	PHY Test Specification is pending

PCI Express revision 4.0 will include 16GT/s. The specification is not released and is a work in progress at this time. However, PCI-SIG workgroups have started work on revision 4.0, and the 16GT/s receiver specification will likely follow the methods of the 8GT/s receiver calibration with improvements targeting better comparability between different test setups. The adjustment of eye width and eye height based on RJ and DM-SI in the PCIe 8GT/s receiver test calibration in combination with +/-2 dB tolerance band for the channel turned out to be problematic. Setups on the lower loss side can require significantly more RJ to close the eye to target values compared to setups that are on the higher loss side. But most receivers can deal with ISI caused by channel loss better than RJ, and thus the two setups can lead to non-comparable test conditions. To fix this situation the standard will tighten the tolerance bands for the test channels, which will require some sort of ISI adjustment; e.g. through a selection of different ISI channel boards.

Preliminary stressor mix:

- ISI via an adjustable or selectable external channel. CEM testing will most likely require test fixtures developed and provided by PCI-SIG.
- RJ with 10 MHz high-pass filter applied
- SJ, different jitter tolerance masks for CC and SRNS/SRIS
- SSC, SRIS only:
  - o Triangular down spread @ 33 kHz for stressed voltage test
  - o Sinusoidal down spread @ 33 kHz for stressed jitter test
- DM-SI
- CM-SI, base specification only

Link Equalization Tests

A new category of tests had to be created to test the new link equalization procedure introduced in PCIe 3.0. This new category of link equalization tests requires link layer capabilities within test equipment since the respective link training status state machine (LTSSM) stages have to be performed and can no longer be bypassed for those tests. BERTs had to be extended with a LTSSM rather than using a clever pattern sequence which looked like a handshake through the different training steps performed. Classical BERT architectures, where pattern generators, de-emphasis signal converters and error detectors / analyzers are entirely separate building blocks or even entirely separate instruments, will struggle with fast enough response times if it is possible at all. Thus newer BERT generations combine all three functions within one instrument.

The link equalization tests can be differentiated into transmitter and receiver tests. The receiver tests are very similar to the classical receiver tests with the exception that the DUT's receiver negotiates with the BERT transmitter the de-emphasis and pre-shoot settings used for the actual receiver test. The calibration of the stress signal is identical to the classical 8GT/s receiver test. Link equalization transmitter tests focus on two things:

- Physical waveform based on link partner requests. This includes physical waveform changes as well as ensuring that the physical waveform is within target specification.

- Response time to link partner change requests. This includes logical response time as well as physical response time. The logical response time is the time between change request and acknowledgment to the link partner while the physical response time is duration from change request until the waveform actually changes.

Conclusion

From revision to revision the PCI Express standard stretched the limits of data transmission through FR4, and thus receivers became more complex. More sophisticated test setups and test equipment with higher capability integration are the result. Test equipment manufacturers had to respond to the new requirements. A good example is the development in Keysight's BERT systems. The introduction of the J-BERT M8020A BERT System is a showcase for an integration of classical BERT building blocks into one instrument to enable new capabilities like a LTSSM targeted to the new test challenges.

Thorsten Götzelmann is a BERT application engineer in the Digital Photonic Test Division of the Electronic Measurements Group at Keysight Technologies.

Linux and Android Q&A with Wind River

Answers on a set of questions covering topics from automotive IVI to wearables to writing for multicore processors.

By Chris A. Ciufo, Editor-in-Chief

Editor's note: Our thanks to Dinyar Dastoor, Wind River's Vice President of Product Management, who recently offered his insights on a number of questions.

EECatalog: Automotive software has been described as a boom area for Linux, but will uncertainty on the legal landscape hinder this? Why or why not? What specific steps can be taken to keep the boom from being a bubble?

Dinyar Dastoor, Vice President of Product Management, Wind River

Dinyar Dastoor: As automotive systems become more connected and sophisticated, the role of software is indeed becoming more important than ever. However, with increased capabilities and connectivity comes increased complexity and exposure to greater threats.

There is a greater need for security not only around personal data but also, ultimately, around safety in the car as well. Unlike a smartphone or PC, a hacked automobile system can result in physical harm. These new demands are creating new hurdles in automotive certifications and regulations in the areas of safety as well as security.

All of these trends result in a rise in software complexity in the car. However, automotive companies don't traditionally come from a history of software expertise. Take for example, the handling of software licenses and copyright. Given all of these new complexities, it is now more critical than ever for auto companies to partner with appropriate experts. This is a sweet spot for Wind River. Wind River is well positioned with its experience in embedded, open source and deep vertical industry expertise (across markets such as auto, A&D, industrial and networking). In addition to the need for software expertise, partners with experience working with complex automotive projects including certification challenges, will be highly valuable. Wind River provides world-class professional services and global support to help customers better manage the highly detailed requirements, complex software integration and intricate project planning associated with the automotive industry.

EECatalog: Google has announced its intention to target the automotive IVI space either natively, or via smartphone connectivity. How well is Android suited for in-vehicle (headed) applications? What are the pro's and con's of Android in automotive?

Dinyar Dastoor: Android has been gaining tremendous momentum in automotive. Not surprisingly, open source technologies have risen in popularity in the car as they help accelerate the rate of innovation in automotive systems. Android lends itself well to the rich multimedia demands of IVI systems with its rich framework of middleware, applications and development tools. That said, with increased capabilities and connectivity comes increased complexity and exposure to greater threats. Auto companies will need to be more diligent to build in security as they work through the design process.

With Google "Projected Mode," a lot of automotive requirements are addressed and this makes it easier to address low and mid end IVI systems with a connected smartphone as an enhanced "amplifier." Android is the market leader in mobile devices and this is the biggest add-on that it brings to automotive. IVI systems get enlarged with consumer features. On the other hand, software components like Bluetooth phoning, audio management or power management have to be adapted to be able to solve automotive use cases and requirements.

EECatalog: What would be the most notable differences today, as compared to one year ago, for checklists comparing Android and Linux on such topics as GUI support,

availability of code packages, foothold in traditional enterprise ecosystems, the wearables market, security and other criteria of your choice?

Dinyar Dastoor: One of the key value proposition differentiators between the two operating systems remains Android's included touch user interface for direct human interaction and wireless connectivity, especially in the consumer device markets. Outside of the mobile phone and tablet markets, use in wearables and automotive continues to grow for Android. For the wider enterprise markets, Linux, Real Time Operating System and other operating systems remain the top candidates based on their existing ecosystem of developers and applications that are more specific to needs of non-phone/tablet markets, as well as the increased need for safety, security, certification and long-term support. As mobile phones and tablets become more disposable devices with less than two years of product life, the need for support and updating is less critical. However in other areas such as industrial, medical, networking, and aerospace and defense, where embedded devices have expected lifecycles ranging from 5-20+ years, it's critical to have solutions where implementing bug fixes, updates, and security patches to the originally selected operating system version is possible. It makes a difference in the decision process as well as re-certification efforts.

EECatalog: What's new with the latest Linux kernel? Is this significant?

Dinyar Dastoor: Some significant changes with recent Linux kernel releases are:

Improvements on enabling dedicated cores and isolated applications. This allows enhanced performance and predictability in user space applications.

Better integration of soft real-time capabilities to the Linux kernel. Finer grained locking, improved synchronization techniques, Refuelling Control Unit (RCU) enhancements and the Earliest Deadline First (EDF) scheduler are a few examples.

Continued scalability improvements like faster suspend resume, VFS, VMA caching, etc.

KVM, container and namespace evolution to support sandboxed applications while leveraging hardware capabilities to offer improved performance.

Thousands of minor bug fixes, security improvements and maintenance updates.

EECatalog: As recently as two years ago, we were still hearing about the gyrations of writing Linux code for multicore processors. Is this now straightforward...or still a challenge?

Dinyar Dastoor: Designing for multicore remains a challenge as developers must take into account more complex designs. While the industry has long been talking about multicore, many developers are still fairly new to putting this type of design into action and need all the help they can get. Also, teams must seek and get familiar with newer "multicore-aware" development tools and methodologies in order to maximize the potential of multicore design.

Synchronization within the kernel can also be challenging, but there are more constructs and advanced techniques available in newer kernels. This includes RCU improvements and modes, new barriers and synchronization primitives, per CPU variables, developing communication channels such as kernel debus and tools that can detect and warn of locking / synchronization issues. These primitives are leveraged by the latest generation of programming language that takes care of multicore and synchronization issues on behalf of the developer.

Chris A. Ciufo is editor-in-chief for embedded content at Extension Media, which includes the EECatalog print and digital publications and website, Embedded Intel® Solutions, and other related blogs and embedded channels. He has 29 years of embedded technology experience, and has degrees in electrical engineering, and in materials science, emphasizing solid state physics. He can be reached at cciufo@extensionmedia.com.

Introducing Embedded Hardware Abstraction Through Virtualization

Networking, industrial, automotive, medical and other applications are using integrated hypervisors to pursue improved performance and more effective multiprocessor consolidation.

By Jens Wiegand, Kontron

Embedded developers no longer see virtualization's natural habitat as limited to IT. Instead, they're viewing virtualization as a way to save on the resources that would otherwise be committed to developing and maintaining non-differentiating software. Embracing a new approach to software optimization, OEMs are consolidating multiprocessor systems to reduce Bill-of-Material (BOM) costs and deliver more innovative capabilities within a device.

By making efficient use of multicore processors through multi-OS configurations, developers can introduce "software appliances," while retaining fine-grained control over the assignment of hardware resources. Real-time performance and determinism can be maintained for the real-time operating system (RTOS), by removing the barriers to performance and consolidation. Key to knocking down these barriers are integrated hypervisors that accelerate and simplify virtualization in embedded platforms.

To the Next Level of Functionality

Development teams needing to deliver the next level of functionality (or performance) with their devices are considering multicore options. And we're seeing systems evolve from single core to low-core-count multicore as well as from low-core-count multicore to "manycore." Virtualization "carves up" the processor to match each application's need for performance, robustness and flexibility.

Consolidation is a factor as well, where devices that currently utilize multiple single-core processors can, for example, consolidate two single-core processors into one dual-core processor. This frees up board space, saves power, reduces heat generation and manages BOM costs. Consolidation to a dual-core processor is of course not the only approach to take; in certain cases, the device can be consolidated on a single core with virtualization, thereby saving even more power.

Addressing Legacy Technology

Reasons vary as to why teams develop a particular device using multiple processors. For instance, the application could demand multiple operating systems, one for real-time processing and one for the user interface. Possibly, one or more of its processors were dedicated ASICs or

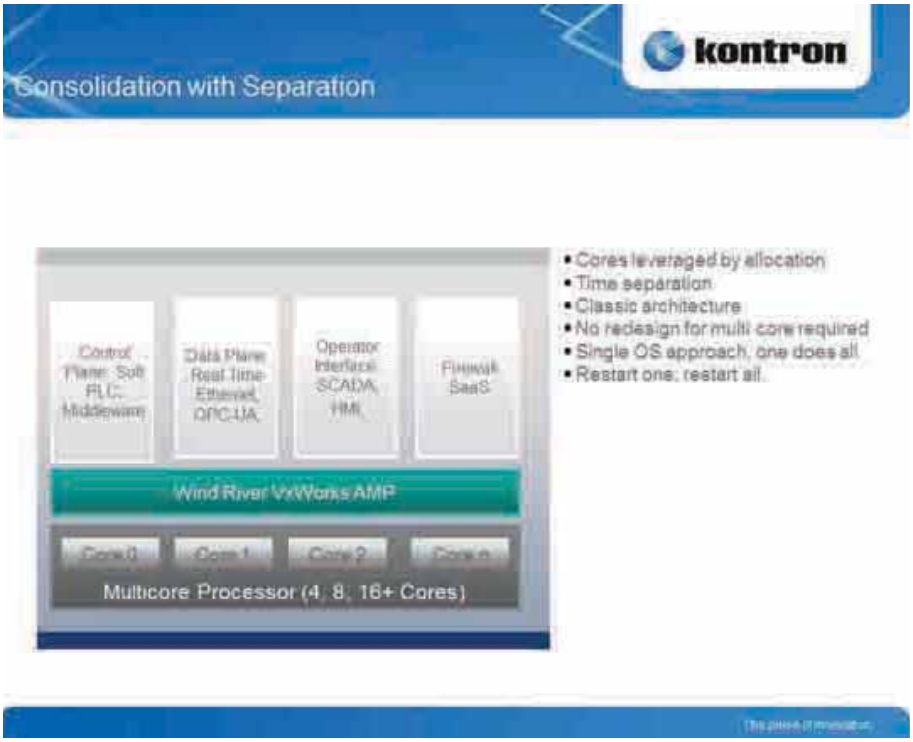

FPGAs or had a defined safety or security requirement. Or, devices running important intellectual property in kernel mode on an RTOS could use multiple processors to keep data separated from a GPL OS on the second core. Now, using virtualization on single- or multicore processors can efficiently remove these barriers.

When migrating to multicore and virtualization, developers must keep in mind that very few projects actually start with green field development. Often a certain amount of legacy code needs to be reused, and it's even the possible there is a legacy operating system in the picture. The reliability of systems must remain intact, as processors will be running much more code, which must be kept separate to avoid software defects that can be very difficult to debug. Virtualization supports these issues by creating virtual boards that can execute legacy code, while keeping different virtual boards strictly separated.

Integrated Hypervisors Add Design Value

It is essential that hypervisor technology—the system at the heart of virtualization—minimizes impact on performance of the system itself. System architects must select a set of hypervisor features and capabilities that improve the overall architecture, and avoid problems with integrating various runtime components and tools from different sources. However some hypervisors designed for desktops or servers add too much performance overhead and lack sufficient availability, safety, security and performance.

Embedded manufacturers are addressing this with hypervisors optimized for specific embedded performance requirements, which is illustrated by Kontron's Size, Weight and Power (SWaP) Platform initiative (Figure 1). The initiative comprises a broad spectrum of computing platforms that deliver high-performance embedded hypervisors based on integrated runtime components and tools. Using this approach, partitioned and safety-critical systems are accessible to more easily support developers and integrators of workload-consolidated systems with a single-vendor solution.

These designers may be building differentiated devices that require embedded operating systems as well as support for Windows or Linux. Embedded hypervisors allow the execution of Microsoft Windows seamlessly, adjacent to other virtual boards in the system. Virtual boards are protected from one another, and applications on Windows and other virtual boards can smoothly communicate and collaborate.

Virtualization in Networking—Use Case

In a typical networking use case, a multicore chip (8+ cores) may be configured by a Supervisor to have one control plane virtual board running Linux and/or an RTOS, with all other cores used to run small specific tasks such as fast forwarders. These small tasks can run at maximum speed as they typically involve only a small code set and largely operate directly from cache. In older systems, they tasks would have been dedicated to FPGAs and ASICs in order to maximize performance; however, the small tasks would have been difficult to program and often impossible to update during runtime.

Multicore processors often still have dedicated hardware to queue packets, for example, coming in from Ethernet cards and then queuing them for the processing cores. Consider a router, for example. Its primary task is to evaluate incoming packets and determine which port should receive them. By having dedicated hardware for smaller tasks such as buffer management and crypto operations, multiple cores can run in parallel and achieve very high throughput. The benefit of virtualization in this case is that legacy code can be reused, mostly in the control plane. Virtualization allows the system to scale to very high core counts, using multiple operating systems—such as Linux for control or VxWorks® for data. Virtualization further enhances reliability by making possible separation between the control and data plane cores.

Virtualization in Industrial Embedded—Use Case

Industrial control, medical or automotive applications have single core processors or perhaps a low dual core processor, and ideally need to run a different OS on each processor. They are typically looking to consolidate an existing dual processor design in order to reduce cost, or may be looking to add new functionality to an existing design.

Using a hypervisor to provide core virtualization allows systems to run multiple virtual boards that are strongly separated. Each virtual board can contain a different OS, for example VxWorks, which boots up quickly to control the device and establish and maintain a safe state of performance, while the Human Machine Interface (HMI) OS (i.e., Linux or occasionally Windows) boots up slowly.

In industrial automation processes, using a multicore design strategy accelerates and drives convergence between back-office or gateway functionality (e.g., OPC Gateways/Servers) and graphic (HMI) functionality, SCADA middleware, or .Net-based applications that currently require Windows XP and potentially Windows V7 in the future. A combination of both back office performance from VxWorks or Linux and graphics performance through Windows is highly attractive for customers to reduce cost of goods sold (COGS), and gain advantages on product lifecycle management.

In applications featuring machine or robotic controllers, convergence is between the Programmable Logic Controller (PLC) and graphic (HMI) functionality. The control function is usually real-time critical and runs softPLC functions on a real-time OS such as VxWorks, while the machine user interface is often based on Windows XP, Windows CE and increasingly Linux. Some end users have opted to replace Windows CE/XP with Linux for cost reasons and lifecycle management issues surrounding Windows, such as daily or weekly patches that cause unpredictability and a lack of transparency. In comparison to many virtualization approaches coming from the IT/server space, determinism is a key requirement in these use cases and requires real-time performance, scalability and proper support by tools.

Advancing Embedded Virtualization

Virtualization is playing an important role in embedded design, offering unique advantages to support complex, connected systems. Developers can more easily consolidate devices such as industrial or medical equipment that requires real-time capabilities combined with a rich, human-machine interface. It also allows Windows to be used as the GUI interface, freeing embedded operating systems to interface with other components and systems. Hardware consolidation reduces power consumption and costs—enabling developers to compete by offering more fully featured, innovative devices. Scalable, reliable multicore systems speed time-to-market by reusing existing legacy applications, and capitalizing on readily available commercial and open-source software.

Embedded platforms pre-integrated with powerful hypervisor technology are furthering virtualization advancements by supporting hardware consolidation, enhanced device functionality and improved adoption of multicore processors. Risk is reduced, and developers can access virtualization from the same trusted manufacturing resources that provide their embedded operating systems. These factors are working together to continue the shift in perception where virtualization is not limited to IT/server markets—it is poised to herald new advancements in embedded arenas as part of a well-defined approach to multicore design.

Jens Wiegand is CTO at Kontron.

What's In A Name: When Is The IoT Industrie 4.0?

Caroline Hayes noted some of the booths at electronica 2014 were branded IoT, but as the industrial manufacturing market is one of the cornerstones of the Germany economy, it was not really a surprise to see that the German phrase Industrie 4.0 was equally used.

By Caroline Hayes, Senior Editor

I had a little taster of Industrie 4.0 the week before at NI Day in London. At his keynote, Eric Starkloff, Vice President Global Sales aligned Metcalfe's Law (Robert Metcalfe proposed that the value of a network is proportional to the square of the number of connected users. In other words, as more users become connected, the information value within those networks also increase) to the smart factory. As devices become connected, sensory technology is used for the networked nodes to sense and control the world around them, using a combination of touch, motion and pressure, using cameras, gyros and sensors. This creates the smart grid, the smart factory and smart machines.

"There are opportunities for disruption, and new applications," he told attendees, "with intelligence monitoring for freight and intelligence in alternative energy sources, for example wind turbines communicating with the grid".

An alternative phrase is Cyber Physical systems, although this, explained Starkloff, is usually used in research and academia to identify computation, communication and control.

Suddenly, I was still left unsure of where industrie 4.0 ends and IoT begins. I had the opportunity to speak to Rahman Jamal, Global Marketing Director, National Instruments, later that day. He is based in Germany, so he was the person to ask: What is the difference between IoT and Industrie 4.0?

"It is the smart factory for tomorrow," he began. "Computer integrated manufacturing, a simulated factory with computerized and automated manufacture." It is an initiative of the German government to create computerized manufacturing (the smart factory). It follows water and steam power to mechanize production, electric power, a revolution that saw mass production and the third, digital revolution. "It is more than 'smart factory', it is a convergence of four or five technologies, driven by processing power (Moore's Law); connectivity (Metcalfe's Law); wireless needs for bandwidth and power and the role of 5G, which could be MIMO, GDMF, or micrometer wave-based." Jamal notes that while all four factors are driving Industrie 4.0, true productivity arises when these factors meet and work together.

Another definition I came across that day is Internet of Big Things. This is systems such as smart cities, smart factories, where large amounts of data are being sent and received on a network infrastructure, rather than the regular Internet.

NI's interest is in the Big Analog Data, i.e. natural sources such as light and speed. Its role is to acquire, analyze and abstract the data, says Jamal. When anomalies are found, they can be analysed; in contrast sensors produce raw data but are connected; adding the intelligence is the distinction, says Jamal.

The company provides a foundation for platform-based smart systems, says Jamal, allowing companies to integrate IP on a single platform. Abstracting from the platform and application with LabView produces a focus on the embedded design, rather than the underlying hardware and software, he explains.

He offers the example of Xilinx's Zynq processor, which with software design instruments to reconfigure the processor, and USB 3.0 technologies, the user can simply acquire and abstract date. "There is no need to worry about details of Zynq in Compact Rio or PCI bus, PXI platforms. . . embedding the IP in target platforms allows the user to focus on the domain expertise – for example, the smart grid".

"Custom hardware is dead," declares Jamal. "No one redesigns a phone, they write a new app: software is the instrument."

Although Jamal was making the point that different members of a team can work on different parts, with a common approach using LabView, I am not sure that hardware engineers need to review their retirement plans just yet. NI works hard to ensure that the data for the latest Intel® Atom™ processor, USB protocol or FieldBus/Ethernet version is in the spec, there is still the expectation that new hardware will be introduced, even called for, to meet the demands of Industrie 4.0 or the IoT for that matter.

Caroline Hayes has been a journalist, covering the electronics sector for over 20 years.

The Importance Of Estimating Power Losses In Consumer Power Supply Magnetic Components

Examining an LED lighting string application to discover it's not worth it to be in the dark on inductor performance at high switching frequencies.

By Cathal Sheehan, Bourns, Inc.

Using inductors can significantly diminish power losses in consumer electronics at high switching frequencies. Yet relatively little information exists on inductor performance in these circumstances. The mystery exists despite abundant tools to help engineers design and simulate power supplies. For instance, one important data point that may not be readily available about inductors is the accurate AC resistance of a particular magnetic component. Another area where it's important to shed some light involves the role core losses play in surface mount power inductors. Yes, core losses are small compared to AC losses in many surface mount power inductors, but this article will show that it is nevertheless a critical evaluation point. That is why it is a recommended design step to estimate power losses before selecting an inductor for a particular power supply.

Estimating Power Losses

There are multiple factors to consider when estimating power losses. A main component to evaluate is the forward converter, which is the industry-standard name for an isolated buck converter. The elements to consider in the design phase are the input and output voltages, the transformer and the switch current. The block diagram illustrates the basic elements of the open loop section of a single-switch forward converter:

Using an LED lighting string application as an example, Table 1 lists the typical design requirements for a power supply. For simplicity, the control circuit and feedback loop are omitted in the example.

The peak voltage at D1 is calculated as

$$VD1 = \frac{8\,V + 0.5\,V(diode) + 0.1\,V(pcb\ conduction\ loss)}{D} = 13.23\ V.$$

The voltage measured across the inductor ΔV during the first interval is, therefore, 5.23 V.

The subsequent inductance value is:

$$L = \Delta V x D x \left(\frac{1}{f}\right) / \Delta I = 21.58\ \mu H.$$

However, since 21.58 µH is not a standard inductor value, the LED lighting power supply example would need to select a 22 µH inductor, which is the closest standard inductor available.

Based on the requirements for this application, designers might select a shielded power inductor in a surface mount package that has a resistance of 62 mΩ and a RMS current of 2.3 A. An educated first look would indicate that the calculated losses are due to the DC resistance of the winding and the DC current of 2 A. This calculation gives a power loss of 0.248 W. The typical information provided on this sample power inductor's datasheet shows that it can conduct up to 2.3 A_{rms} giving a total power dissipation of 0.327 W, and may have a temperature rise of 40 °C at full load. If the ambient temperature is 40 °C, then the inductor is expected to rise in temperature to less than 80 °C at the required current of 2 A. These calculations would take a designer in one direction, but to find the optimal power inductor solution, however, requires taking the following additional factors into account.

Thermal: Copper has a very high temperature coefficient. The resistance at 43 °C ambient will increase to 67 mΩ. The DC loss in the winding will, therefore, be 0.268 W at a current of 2 A.

AC Resistance: The distribution of current in each layer of a multi-layered wound inductive component will be unevenly distributed depending on the frequency of the current. The magnetic fields generated by the multiple layers are responsible for this effect, which is also known as "proximity." There is a one-dimensional solution to the complex differential equations describing the ratio of the AC resistance to the DC resistance, which is called

the Dowell equation. This equation enables us to calculate losses in the inductive component due to proximity.

The AC current in the inductor is a saw tooth waveform, and can be written as Irms = I1/3. For this device, the RMS values are calculated at 0.22 A given that the current in the inductor is 0.75 A peak to peak.

Taking the AC resistance at the switching frequency is 2.4 Ω would generate a loss due to proximity of 0.112 W in the inductor.

Core Losses: The ferrite core also generates some losses due to the eddy currents generated by the swings in flux as current rises and falls in the coils of the inductor. These are known as core losses where Faradays Law is used to calculate the flux density "B" in the wound core.

$$B = \frac{1}{NA} * \int edt : \text{Faraday's Law}$$

Using this calculation, the peak flux in this application is 26 mT. By convention, the change in flux density or ΔB is taken as one half the peak flux or 13 mT. By checking the core data, it is known that the loss at 210 KHz at a ΔB of 13 mT is 50 mW.

The total dissipation in the inductor is now determined to be 0.438 W (0.268 W + 0.12 W+0.05 W), which is far different than the original calculation 0.248 W proved now to be erroneous. The updated measurement indicates that the device is actually above the rated full power of the inductor at 0.325 W.

This information can significantly save designers' time in the component selection process. It is very important when selecting the right magnetic component to calculate AC resistance curves along with the understanding that AC losses are more significant than core losses.

Determining the Optimal Inductor Solution

Taking into consideration the series power loss estimates provided, the optimal solution for this particular application is for designers is to select an inductor that has the same dimensions as the example component but with a lower DC resistance. The most optimal inductor for the LED lighting application would be one that has the same core as well as the exact footprint (12.5x 12.5 x 6 mm) with the identical number of winding layers, which means that the ratio of AC to DC resistance is unchanged. Applying Dowell's equation to determine AC resistance on a different power inductor device such as one that is shielded with a current range of 1.70 A to 9.8 A. Using this inductor at 210 KHz, the AC resistance is 1.7 Ohms, giving a total power loss, including core losses, of 0.3 W (0.170 W + 0.082 W + 0.05 W), which is well below the rated maximum DC power of 0.688 W for this particular device.

Designing for Improved Power Supply Reliability

To achieve the highest power supply operational reliability, it is crucial to avoid selecting an inductor purely based on the rated DC current as written in the datasheet. The selection process must also include evaluating data to help prevent power losses that can lead to other complications, such as overheating and premature failures in the field. While core losses are often mentioned as being a problematic source, AC resistance can be much more detrimental to the application and is frequently overlooked. If AC resistance information is not available, designers are typically forced to "over-specify" the inductor as to allow enough margin to accommodate any additional losses. The curve charts are a great resource for design engineers in estimating power losses to determine the optimal inductor for their next power supply design.

Cathal Sheehan is Market Manager of the Consumer Market Segment at Bourns, Inc. He has held several roles in Bourns in product management and application engineering. Sheehan holds a Masters Degree in Electronic Engineering from University College Cork and a Masters in Business Administration from Open University.

NEW ADLE3800PC – Intel® Atom™ Processor E3800 Series PCIe/104 SBC

The ADLE3800PC is based on Intel's first System-on-Chip (SoC) Intel® Atom™ processor E3800 product family, which is built using Intel's 22nm 3D Tri-gate process. It offers vastly superior compute performance and energy efficiency including 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.

Key Features

- Intel Atom processor E3800 Series SoC, DC/Quad
- Junction Temperature Rated at -40C to +110C
- Up to 8 GB DDR3L-1333, 1.35V SoDIMM204 Socket
- Onboard mini PCIe/mSATA Socket
- 4x USB 2.0, 1x USB 3.0, 2x Serial COM, 2x SATA 3 Gb/s, 2x GLAN Ethernet

Applications Areas:

- Military & Defense Rugged SFF
- Rugged Mobile Computing
- Portable Medical Devices
- Mobile Autonomous Systems for Civil, Commercial and Defense Applications Including:

– Unmanned Ground Vehicles – Robotic Subs – Unmanned Avionics – Unmanned Buoys and Other Surface Vehicles

Embedded Solutions
Embedded. Works Anywhere.™

ADL Embedded Solutions
4411 Morena Blvd Suite 101
San Diego, CA 92117 USA
858-490-0597
858-490-0599
sales@adl-usa.com
www.adl-usa.com

ADLMES-8200 High-Ingress Protection (IP) Modular Enclosure Systems

The ADLMES-8200 is a highly innovative embedded enclosure design. Its highly configurable modularity makes it possible to expand or reduce a system without replacing the entire enclosure. Side wall modules may be added or removed as system requirements evolve. Three standard profiles provide quick turn inventory availability. A broad portfolio of PC/104 SBC options ranging from low-power Intel® Atom™ processors to high performance 4th Generation Intel® Core™ i7 processors are available.

Key Features

- Modular Sidewall Design Supports Variable PC/104 Stack Heights (2 - 6 Cards) or Expanded 3.5" SBC Intelligent Systems
- High and Low IP (Ingress Protection) Systems Possible via High IP, Modular Chassis Design Coupled with Full Custom, Quick-Turn I/O Panels
- Broad Portfolio of PC/104 SBC Options Ranging from Low-Power Intel Atom processors E3800 to High Performance 4th Generation Intel Core i7 processors
- Fully Supported by ADL Embedded Solutions' Team of Solidworks Engineers for Model and or Design Support
- Options for MIL-STD 461, and MIL-STD 704/1275
- Designed for MIL-STD 810 Shock & Vibration

Potential Applications Include:

- Military and Aerospace – Rugged SFF
- Rugged Industrial Applications
- Communications Applications
- Mobile Routers and Other Network Appliances
- Railway Train Control
- Transportation
- Imaging Applications

Embedded Solutions
Embedded. Works Anywhere.™

ADL Embedded Solutions
4411 Morena Blvd Suite 101
San Diego, CA 92117 USA
858-490-0597
858-490-0599
sales@adl-usa.com
www.adl-usa.com

ATCA-7480 Packet Processing Blade

The Artesyn ATCA-7480 blade is based on two members of the Intel® Xeon® processor E5-2600 v3 product family with up to 28 processing cores per blade. Data paths to main memory and I/O are highly optimized and, combined with the scalable memory capacity of up to 512GB, allows for shorter response times of database accesses, accelerated pattern matching and it can help optimize routing decisions in virtualized network environments.

Cost sensitive applications will benefit from the possibility to use up to sixteen memory DIMMs with lower capacity while still achieving attractive memory capacities. The blade can be combined with optional hardware accelerators directly connected to the CPUs. The accelerators are optimized for assisting encryption/decryption algorithms and can greatly enhance throughput of encrypted data in security applications.

The ATCA-7480 features a QuadStar™ interface consisting of four 40G Ethernet channels that connect to the shelf's backplane. System integrators have the choice of combining channels to different kinds of redundant and non-redundant topologies depending on the applications bandwidth and availability requirements.

With enabling software for Artesyn's SDN/NFV solutions, the ATCA-7480 blade includes support for the Intel® Data Plane Development Kit with accelerated OpenVSwitch,

OpenFlow and OpenStack plug-ins for managing virtualization services on the computing platform. The blade will also support Wind River's Titanium Cloud.

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

SharpStreamer Video Accelerator Card

The Artesyn SharpStreamer PCIE-7207 high-density video accelerator enables service provider networks to offer video transcoding services quickly and dynamically. As an add-on card, the SharpStreamer PCIE-7207 offers quick and scalable integration with existing and standard 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.

With a focus on the high-density and low power demands of video streaming applications such as OTT streaming servers, mobile network optimization, video conferencing and broadcast equipment, Artesyn employs multiple Intel® Core™ i7 processor and GPU accelerated devices in a small and scalable PCI Express card footprint that is easily deployable in off-the-shelf servers.

Compared to dedicated appliances the SharpStreamer solution is more easily deployable, portable, and does not constrain operators to only one type of equipment to monetize OTT streaming content. It also offers network scalability for increased subscribers to pay as you go adding more cards and density from small to large servers as needed. Compared to software-only solutions, the SharpStreamer PCIE-7207 requires far fewer servers and much less operational cost to power video transcoding services.

The SharpStreamer PCIE-7207 is equipped with the Intel® Media Software Development Kit featuring Intel® HD Graphics with fixed-function hardware acceleration, monitoring, and processor subsystem O/S and management tools for easy integration with server host processing environments.

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

EC200-BT Series - Intel® Atom™ Processor E3800 SoC Fanless Rugged Embedded Systems

High Flexibility, High Scalability, Rich I/O Connectivity

Modular-concept construction offers unlimited configurations
1 Vertical USB port for securing sensitive data from unauthorized access
Low-power and fanless solution for industrial automation

Fanless

COMPUTEX d&i awards 2014

EC200-BT
1 Mini PCIe
275 x 59 x 203 mm (W x H x D)

EC210-BT
1 Mini PCIe + 1 PCI
275 x 94 x 203 mm (W x H x D)

EC220/EC221-BT
1 Mini PCIe + 2 PCI/1 PCIe x16, 1 PCI
275 x 114 x 203 mm (W x H x D)

Replaceable I/O modules

1 DVI, 1 VGA, 4 USB, 6 COM, 2 LAN
For POS System Application

1 VGA, 2 USB, **6 COM,** 2 LAN
For Machine Control Application

Customize Your Own I/O Module
Offers Efficient Customized Service

EC200-BT Front View

Model Name	EC2xx-BTA881	EC2xx-BTAG60	EC2xx-BT6060
Processor	Intel® Atom™ Processor E3800 Series		
Memory	2 DDR3L SODIMM up to 8GB / 1 DDR3L SODIMM up to 4GB		
Storage	1 2.5" SATA drive bay 1 mSATA via the mini PCIe socket		
Rear I/O Ports			
DVI	1	0	0
VGA: DB-15 D-sub	1	1	1
LAN: RJ45	2	2	2
USB 2.0 Type A	4	2	2
COM: DB-9 D-Sub	6	6	2
GPIO	8-bit	16-bit	0
Power Input	Wide Voltage 9~36V DC-in		
Front I/O Ports			
USB 2.0 Type A	4	4	4
COM: DB-9 D-Sub	4	4	4
Wi-Fi module antenna hole	2	2	2

DFI®

DFI Inc. (Headquarters)
100, Huanhe St., Xizhi Dist., New Taipei City 22154, Taiwan (R.O.C.)
Tel: +886(2)2694-2986 Fax: +886(2)2694-5815
E-Mail: inquiry@dfi.com.tw http: //www.dfi.com

intel IoT Solutions Alliance

Designing with Intel® Embedded Processors?

Embedded Intel® Solutions delivers in-depth product, technology and design information to engineers and embedded developers who design with Intel® Embedded processors

Visit
www.embeddedintel.com

Subscribe Today at
www.embeddedintel.com
Free!

Tighter Budgets Are Better For This Systems Supplier

General Micro Systems (GMS) has evolved from VME SBCs to rugged bespoke "shoebox" systems for military vehicles.

By Chris A. Ciufo, Editor-in-Chief

Ben Sharfi, founder of General Micro Systems in Southern California, is a well-known colorful personality in the VME industry. Known for his "That's ridiculous!" outbursts at VITA trade meetings and not-so-subtle advertisements poking his competitors, he started building VME SBCs in 1979—practically coincident with the introduction of the 6U form factor. Today, GMS still makes SBCs but is finding more traction in building very integrated rugged small form-factor systems.

Edited excerpts follow.

Chris Ciufo, EECatalog: Ben, we go back a long way. Give us a very brief history of GMS and explain where the company is today.

Ben Sharfi, GMS: We incorporated in 1979 and shipped our first VME board in 1980, right at the same time both Force Computer and Motorola introduced theirs, making us the third in the new market. Our part numbers were "V01" then; now we're shipping "V399" and up. So we're a few part numbers later than those early boards. Today, we're doing lots of rugged systems and many of them are for the Army dealing with some real issues and problems they're having.

EECatalog: What kinds of issues is the Army facing?

Sharfi: Cross-domain, multi-domain and security are big issues. The Army assigns certain [computer and data access] privileges based upon domains such as SIPRNet and NIPRNet [Secret Internet Protocol Router Network; Nonsecure IP Router Network] and others, depending upon the mission. The issue is that every domain requires a separate, segregated system, which can get big and needlessly redundant in vehicles and mobile platforms, given the state of today's technology and processing power.

There's a big movement to address the size and weight problem using multi-domain, which is one system that can address two or more domains. Additionally, there's cross-domain where one system can talk to another domain within the same environment. This is the traditional partitioned "red/black" NSA architecture that the military has used for a while.

We were hired by the Army to create a multi-domain system that runs red/black within the same box. The Army plans to install this on a variety of vehicles to save space and weight, and our box will be running SIPR and NIPR simultaneously in one small box. Also, we're offering a cross-domain architecture that provides a secure networking switch. This is possible because of our RuggedCool technology and the great capabilities in some of the latest processors from Intel and others.

EECatalog: Before we dive into ICs, let's talk about VME where it all started for GMS. Your website shows lots of systems—what about VME and VPX?

Sharfi: Our commitment to VME is second to none, and we've also got a whole bunch of VPX boards, too. We still do a lot of community work through VITA and are committed to these specs because we have to be. But one of the problems with [VPX] is that it's crazy expensive and crazy big. And even worse, despite everyone's best efforts, there's really no standard.

In the old days, you could plug VME boards from two vendors into the backplane and they would interoperate. The same was true of CompactPCI. These boards would 'play with' other boards. The OpenVPX standard has in some ways actually made this worse, even with the predefined profiles. Now anyone can add any pins they want and create any kind of format.

EECatalog: Ben, you're going to have to explain that because this is totally different from my understanding of VPX and OpenVPX. You're saying that a VITA 65 OpenVPX board from Curtiss-Wright might not be interoperable with a GE Intelligent Platforms board?

Sharfi: Absolutely not! The [VITA] specs opened up the pinouts to allow vendors a great deal of flexibility…to the point of not interoperability. The issue [is often] PCI Express, and I have to match my PCI Express lanes to them. This is the fault of the

standard, because if you look at the [personal computer] PCIe cards on the market you'll only see two choices: x1 and x16. The x8 is a version of the x16 format. And these cards will work in any PC.

But that's not true in OpenVPX: you can have two x2, two x4, eight x1, three x4, and any way that I want to, because the pinout is open. I've talked to [VITA] about this and the problem is that this is run by committee and we no longer have the VME plugfests like we used to. They were designed to see who runs with each other's board, and we'd write up the issues and they'd get fixed to assure interoperability. Any vendor could work [interoperate] with anybody. Right now if we did the same test, nobody's going to work with anybody.

EECatalog: But VPX—and OpenVPX—is a growing market with many designs and the DoD is definitely deploying systems. How is this issue being solved?

Sharfi: From VITA's perspective, there are multiple configurations that specify how the boards are designed to interoperate, and it's up to the user to decide which configuration to choose and assure that the chosen vendors adhere to it. Here's why this doesn't work in practice. The graphics board, for example, may be a x16 [Editor's note: this is common for graphics boards.] But the Intel® Xeon® processors (codename Haswell) CPUs that we're using are not going to support x16 without a PCI Express fanout switch.

The solution is like what Apple does: dynamic, configurable PCI Express that changes the lanes around when cards are plugged into the desktop Mac. This uses a PCIe switch that configures to whatever the host processor wants and gives the user more options. We've done this very thing on our 3U and 6U VPX cards using a PLX PCI Express switch based upon user input.

EECatalog: What's new in the area of rugged shoeboxes that don't use VME or VPX?

Sharfi: The Army's WIN-T program is alive and well, and GMS is the exclusive supplier of multi-domain boxes in Bradley, MRAP, Stryker, and all of the program's [6] ground vehicles. We've been doing this for a while, supplying four boxes per vehicle. Based upon that, we're seeing a need for night vision in MRAPs in a program named—you guessed it—Night Vision.

Eventually replacing the HMMWV, MRAP can carry equipment and troops in all terrain conditions including total darkness. The current configuration calls for using six "Rover" L3 day/night camweras to provide the driver and support personnel with a 360-degree view around the vehicle. As you can guess, six real-time cameras fusing data to a single screen with under 1 frame of latency [30 ms] is a massive computational task.

The Army looked at many different solutions, including VPX, rackmounted equipment, and regular servers. The problems with these were cost, space, and most importantly, heat. Don't forget that this is a very rugged environment. Tough challenge.

EECatalog: I'll bet this most recent PR you sent me for the 4:1 "Tarantula" shoebox was your answer, right?

Sharfi: You got it. It's literally four systems combined into one system. It's smaller than a shoebox and replaces what would have taken four 2U chassis in the past. It's based upon the the Intel® Xeon® processor (EP) with 10 cores running at 2.4 GHz each, with 12 cores shipping sometime soon. It's literally the fastest embedded server class product out there. We add 512-bit wide memory to get extreme memory speed to talk to one of up to six hardware secure virtual machines—do you see the relationship to the multi-domain product I described earlier?

We add an analog to digital camera feed converter, plus a hardware video switch matrix to pipe in the data. The camera data is converted a lossless open standard called "GigE Vision", then fed to a switch. This intelligent video switch is based upon Vitesse 24 Gbit Ethernet plus two 10 Gbit Ethernet ports. This is a Cisco-like switch but it's embedded inside this rugged shoebox.

The multiplexed video is then sent to the processors—all in under a frame cycle. Finally, we created removable canisters that can store up to 32 TB of SSD storage to keep the data for post analysis. Our storage solution replaced separate RAID controllers and storage—and it's all in our rugged shoebox!

EECatalog: The creativity of these boxes is impressive. What's the secret?

Sharfi: There are several things we do that make this possible. First, our systems are very highly and tightly integrated. You couldn't stuff this much in without some serious engineering. And secondly, the competition for our products is often OpenVPX, backplanes, and modular chassis and power supplies. But our customers sometimes don't care about a modular, open standards approach. They're dealing with a box-level LRU [line replacement unit]. This allows us to bring our best technologists to solving program requirements.

The tighter the budgets are, the more successful we seem to be.

Chris A. Ciufo is editor-in-chief for embedded content at Extension Media, which includes the EECatalog print and digital publications and website, Embedded Intel® Solutions, and other related blogs and embedded channels. He has 29 years of embedded technology experience, and has degrees in electrical engineering, and in materials science, emphasizing solid state physics. He can be reached at cciufo@extensionmedia.com.

Register and
get your tickets now!
embedded-world.de

Nürnberg, Germany
February 24 – 26, 2015

embedded world 2015
Exhibition&Conference
... it´s a smarter world

THE gathering for the embedded community!

The world's largest event for embedded technologies brings together all the players in the embedded sector for a productive conversation.

Make sure you are there when contacts and networks are being cemented on the international level and new trends are being set.

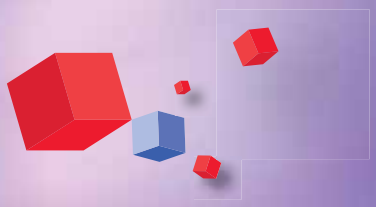

Media partners

elektroniknet.de
computer-automation.de
energie-und-technik.de
MEDIZIN-UND-elektronik.DE
Markt&Technik
DESIGN & ELEKTRONIK
KNOW-HOW FÜR ENTWICKLER

Elektronik
Elektronik automotive
ENERGIE & TECHNIK
Computer & AUTOMATION
MEDIZIN elektronik

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

Congress organizer
WEKA FACHMEDIEN GmbH
Tel +49 (0) 89. 2 55 56-13 49
info@embedded-world.eu

NÜRNBERG MESSE

ADLINK TECHNOLOGY INC.

Smart IoT Gateway Solutions

Connectivity, Manageability, and Security

Powerful Yet Compact Embedded Platform

MXE-200i

- Intel® Atom™ processor E3845/E3826
- Intel® IoT Gateway
- 2x mPCIe for ZigBee/WiFi/BT/3G/LTE
- Ultra-Compact 120 x 100 x 55 mm housing

Embedded IoT Gateway Platform

MXE-100i

- Intel® Quark™ SoC X1021
- Intel® IoT Gateway
- Optional onboard ZigBee/WiFi for connectivity
- Ultra-Compact 120 x 100 x 55 mm housing

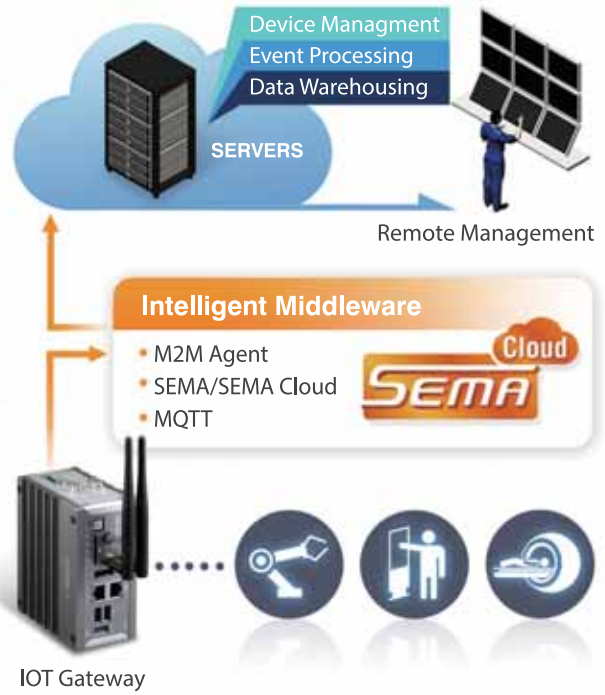

ADLINK Technology, Inc.

Toll Free: +1-800-966-5200 ▶ info@adlinktech.com ▶ www.adlinktech.com

embedded world 2015 Exhibition&Conference ... it's a smarter world

Booth 1-538 in Hall 1

intel IoT Solutions Alliance Premier