Extension MEDIA

Embedded Intel® Solutions

Spring 2015

Right Sizing Your Embedded Application

Smart Fabric, Not Wrist Bling, To Lead Wearable Market Growth?

Industrial IoT Looks to SMARC

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

TwinPro™

Up to 4 NVMe • Four or Two Hot-Pluggable Nodes in 2U

Supports up to: 36 Cores and 145W TDP Intel® Xeon® processor E5-2600 v3 Product Family
1TB DDR4-2133MHz in 16 DIMMs • SAS 3.0 (12Gbps) • 12 HDDs/1U • 10GBase-T/ FDR IB
NVMe • Redundant Titanium/Platinum Level Digital Power Supplies

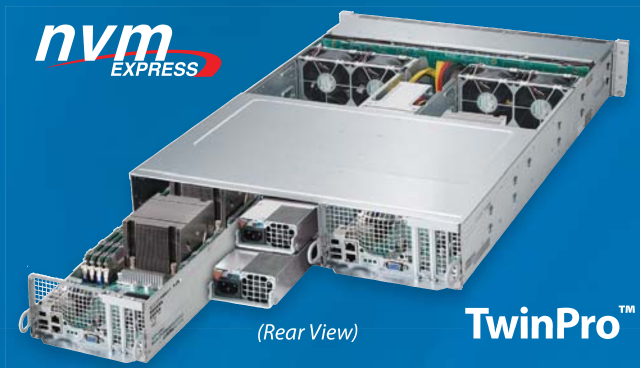

(Rear View)

TwinPro™

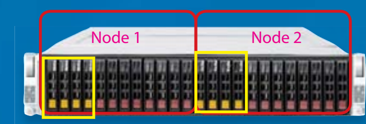

SYS-2028TP-DNC/(T)(F)R
Up to 4x NVMe + 8x SAS 3.0 drives (per node)

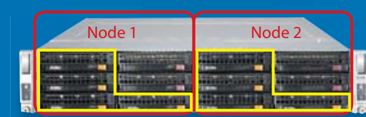

SYS-6028TP-DNC/(T)(F)R
Up to 4x NVMe + 2x SAS 3.0 drives (per node), or
Up to 6x SAS 3.0 support (per node)


(Rear View)

TwinPro²™

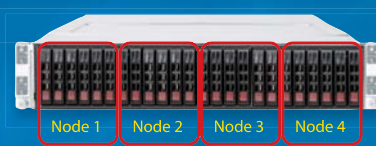

SYS-2028TP-H Series
4 DP Nodes in 2U, 24x 2.5" SAS/SATA drives

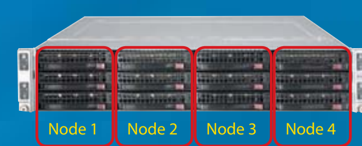

SYS-6028TP-H Series
4 DP Nodes in 2U, 12x 3.5" SAS3/SATA3 drives

Four or Two hot-pluggable Systems (Nodes) in a 2U form factor.
Each Node supports up to:

- 1TB DDR4-2133MHz memory in 24 DIMM slots
- 1 PCI-E 3.0 x16, 1 PCI-E 3.0 x16 "0" slot; and 1 PCI-E 3.0 x8 slot (TwinPro™)
- 4 ports NVMe
- 8 SAS 3.0 (12Gbps) ports with LSI® 3108/3008 controller, with optional SuperCap (CacheVault)
- 8 SATA 3.0 (6Gbps) ports with Intel® C612 controller
- 12 (TwinPro) or 6 (TwinPro²) hot-swap 2.5" HDD drives per Node
- FDR (56Gbps) InfiniBand, Dual 10GBase-T or Dual Gigabit Ethernet LAN options
- Redundant Titanium (96%+) / Platinum (95%+) Level Digital power supplies
- Integrated IPMI 2.0 plus KVM with dedicated LAN
- GPU/Xeon Phi™ option
- SATA-DOM and mSATA support
- Up to 36 Cores per system and 145W TDP dual Intel® Xeon® Processor E5-2600 v3 product family

intel inside™ XEON®

TITANIUM We Power the Cloud X10

Contact us at www.supermicro.com/TwinPro

Coke Machine with a Wii Mote? K-Cup Coffee Machine Marries Oculus Rift?

Talk about some kind of an weird Coke-machine-meets-Minority Report mash-up, but the day is here where vending machines will do virtual reality tricks, access your Facebook page, and be as much fun as Grand Theft Auto. And they'll still dispense stuff.

Chris A. Ciufo, Editor-in-Chief

According to embedded tech experts ADLINK Technology and Intel, the intelligent vending machine is here. Multiple iPad-like screens will entertain you, suggest which product to buy, feed you social media, and take your money.

These vending machines will be IoT-enabled, connected to the cloud and multiple online databases, and equipped with multiple cameras. Onboard signal processing will respond to 3D gesture control or immerse you in a virtual reality scenario with the product you're trying to buy. Facial recognition, voice recognition, and even following your eyes as you roam the screens will be the norm.

For you, the customer, the vending machine experience will be fun, entertaining—and very much like a video game. For the retailer, they're hoping to make more money off of you while using remote IoT monitoring, predictive diagnostics, and "big data" to lower their costs.

The era of the intelligent vending machine is upon us. The machine's already connected to the Internet…and one of many billions of smart IoT nodes coming to a store near you.

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
Darla Rovetti

Special Thanks to Our Sponsors

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.

Embedded Intel® Solutions

SPRING 2015

DEPARTMENTS

FROM THE EDITOR
2 **Coke Machine with a Wii Mote? K-Cup Coffee Machine Marries Oculus Rift?**
By Chris A. Ciufo, Editor-in-Chief

SPECIAL FEATURES

STORAGE & NETWORK SECURITY
6 **Ruggedized Portable Memory Systems: Interview with Paul Plitzuweit, Datakey**
By Anne Fisher, Managing Editor

8 **Right Sizing Your Embedded Application with the OneBank Specification**
By Matthew Henry, ADL Embedded Solutions

10 **Certified Small Form Factor Systems Offer High Performance COTS Building Blocks**
By RJ McLaren, Kontron

13 **Beware These Evil Twins and the Threats to Network Traffic**
By Art Swift, CUPP Computing

ULTRA LOW POWER
14 **Not Enough Just to “Take up the IoT Mantle”: Interview with Eurotech CEO Larry Wall**
By Anne Fisher, Managing Editor

17 **Industrial IoT Looks to SMARC Architecture to Meet Low Power, Performance and Cost Goals**
By Dirk Finstel, ADLINK Technology

MEDICAL & E-HEALTH
21 **Embedded Medical Device Software—How And Why To Use Static Analysis To Update Legacy Code**
Dr. Paul Anderson, Grammatech

STAFF SPECIAL COMPENDIUM
24 **A New Prize in Your Serial Box**
Chris A. Ciufo, Editor-in-Chief, Embedded, Extension Media

28 **Embedded World 2015—No Longer A One-Theme Show**
By Caroline Hayes, Senior Editor

31 **Top 3 Uses for PCI Express Switches**
Chris A. Ciufo, Editor-in-Chief, Embedded, Extension Media

36 **Smart Fabric, Not Wrist Bling, To Lead Wearable Market Growth?**
By Jonah McLeod, Silicon Valley Blogger

37 **DVCon Highlights: Software, Complexity, and Moore’s Law**
By Gabe Moretti, Senior Editor

PRODUCT SHOWCASE

39 **Acnodes Corporation**
39, 40 **Artesyn Embedded Technologies**
40 **DFI Inc.**

cover image courtesy of Intel®

Portwell Empowers Intelligent Solutions

www.portwell.com
info@portwell.com
1-877-278-8899

Portwell

intel
IoT Solutions Alliance
Premier

Ruggedized Portable Memory Systems: Interview with Paul Plitzuweit, Datakey

If your company's experience with ruggedized portable memory systems stretches back to before USB flash drives and SD cards came on the scene, it's well positioned to offer perspective on the milestone developments in this sector.

By Anne Fisher, Managing Editor

Our thanks to Paul Plitzuweit, senior product manager of Datakey, an ATEK Access Technologies brand, who recently offered his take on a number of topics, among them, the question of whether brief life cycles—par for the course for consumer goods—raise problems for medical, military and industrial applications?

EECatalog: What milestones stand out for you and why?

Paul Plitzuweit, Datakey: First off, there are two primary applications for portable memory. One is using the portable memory as a credential for some form of access control; the other device is being used for data transport—moving data from here to there.

For the credential applications, we started out with credentials like badges with bar codes or magnetic stripes. These devices were read-only and relatively easy to copy. They couldn't carry much more than an Identification (ID) or account number. With the introduction of non-volatile, solid-state memory, typically EEPROM, came the ability to store significantly more information and to have data written back to the credential. Our Datakey devices, introduced in the mid 1970's, were one of the first commercial examples of a portable memory device to use EEPROMs. It was literally data on a key. With that jump, a credential could now tell the system who you are, what you have access to, and it can store an audit trail or a record of your last several transactions. Other milestones for credentials would include the use of non-contact credentials like RFID tags or proximity cards, which offered readers that were sealed from the outside environment and provided unlimited cycle life. Yet other enhancements looked to improve the security of the data on the credential, with secure non-volatile memory solutions and smart card devices.

Data transport applications have largely seen a different set of milestones. Before solid-state memory, magnetic media, such as floppy disks were the primary data transport method, especially in PCs. The consumer industry pushed for solutions that had greater memory capacity, faster throughput speeds and ultimately smaller size. Magnetic media gave way to optical discs, which gave way to NAND Flash-based devices such as CompactFlash cards, SD cards and USB flash drives. For years, supporting the use of USB flash drives in microcontroller-based embedded devices was a challenge. This has become easier with USB host solutions and microcontrollers with integrated USB host controllers that include the necessary middleware. Like with credential solutions, data security is often a concern in data transport applications, so the industry has responded with solutions such as encrypted flash drives.

EECatalog: Have there been "hindsight is always 20/20" episodes—for Datakey? For the memory system sector in general?

Paul Plitzuweit, Datakey: One of the "hindsight" events that I have witnessed both with the Datakey line and in the industry in general is how much NAND flash memories are driven by the consumer electronics market. The push for greater capacities, higher performance, smaller geometries, and lower prices has led to short product life cycles with NAND flash memory. This is generally acceptable with consumer electronics products, but can be problematic for products in other industries such as medical, military and industrial that typically have much longer product life cycles.

EECatalog: For storage are the challenges embedded designers face more difficult than in other areas?

Paul Plitzuweit, Datakey: In any design there are always choices and compromises to be made. With the shorter product life cycles of NAND-flash memories, embedded designers are often faced with a choice between the low cost of a consumer memory product, but trading off other aspects such as robustness and long-term availability. There are certainly opportunities for products that can eliminate some of the shortcomings. One such example that comes to mind is Spansion Long-Term Memory, where it has committed to extended product life cycles with its NAND and NOR Flash products for embedded customers who operate in market verticals that demand long-term availability.

EECatalog: What is the role secure, reliable transport of data is playing and will play for the IoT?

Paul Plitzuweit, Datakey: I think the biggest challenge in the IoT space is security, especially when you consider military, medical and industrial applications. For major adoption of IoT applications in these spaces the security issue must be addressed. Authentication of the endpoints, the sensors, and encryption of the data both in transit and at rest in the cloud are vital.

EECatalog: What new developments are you keeping an eye on and why?

Paul Plitzuweit, Datakey: Like many, I am very interested in the Internet of Things and specifically how that will impact the markets we serve, such as military, medical and industrial. Other developments that have been taking place for a bit longer would be the proliferation of wireless communications in embedded devices. In the medical industry, for example, Datakey products have been used to positively pair a medical transmitter and receiver(s). As you can imagine, in an operating room setting it is critical to make sure the data you are receiving is coming from the source you expect.

Other developments that I am anxiously watching are all of the new non-volatile memory technologies, such as the various non-volatile RAM devices, like FRAM and MRAM and potential NAND-flash successors such as RRAM.

EECatalog: A question you raised in your blog was: "How can designers of embedded systems eliminate or limit their exposure to […] attacks that come through USB ports?" Could you expand on that a bit in order to advise designers of embedded systems on avoiding additional risks besides USB port attacks?

Paul Plitzuweit, Datakey: The main issue with security in these situations is to give thought to the connector itself. On a personal computer, a USB port makes sense because the manufacturer doesn't know which USB devices a user might want to plug into that port. With an embedded device, chances are much more likely that a port might be added for a single purpose, to support the use of a USB flash drive for example. If that is the only purpose of that USB port, designers can reduce their risk by simply moving to a proprietary connector, rather than the universal Type A USB receptacle. Many of our RUGGEDrive™ customers selected the product (that delivers USB flash drive or SD card functionality, but uses our proprietary SlimLine™ contact system), simply because it eliminates these types of security risks.

Anne Fisher is managing editor of EECatalog.com. Her experience has included opportunities to cover a wide range of embedded solutions in the PICMG ecosystem as well as other technologies. Anne enjoys bringing embedded designers and developers solutions to technology challenges as described by their peers as well as insight and analysis from industry leaders. She can be reached at afisher@extensionmedia.com

Right Sizing Your Embedded Application with the OneBank Specification

An evolution that's in step with the ascent of speedier and more sophisticated bus architectures is one reason PC/104 effectively serves embedded applications across a wide range—encompassing machine vision, security, high-speed data collection, and more.

By Matthew Henry, ADL Embedded Solutions

Most embedded products will provide adequate processing power when appropriately sized for the job at hand. It is I/O bandwidth and product expandability that become critical data points when defining your platform and application. Data storage I/O can bottleneck your system and keep it from performing at an optimal level. Is it very hot or cold in your area of operation? Is your application mobile or do you have limited power available? It is these factors that help differentiate one form factor from another.

Intel ® Architecture Raises the Bar

Perhaps you need a lot of horsepower for data processing, or you need a Windows-based platform. The PC/104 footprint allows enough PCB for the implementation of some of the highest performers available from the Intel® Core™ processor product lines. Currently, PCIe/104 is able to support the Intel 4th Generation Intel® QM87 chipset with the Intel® Core™ i7-4700EQ processor, giving it a robust set of features with near-server class performance employing six SATA /mSATA ports up to SATA 6 Gbs, USB 3.0 and advanced video like DisplayPort 1.2, HDMI 1.4 and eDP.

Common legacy interfaces are still included for those looking to retrofit an existing platform, like VGA video, RS-232 COM ports, dual Gigabit Ethernet and the ability to connect to a host of peripheral I/O boards to get to the interface needed via the PCIe/104 bus. 3rd and 4th generation Intel Core architecture raises the level of performance in the PC/104 market space. This architecture has opened the door to emerging markets including machine vision, high-speed data collection, security/surveillance and mil-aero, to name a few.

The PC/104 Embedded Consortium has released the OneBank™ specification, opening another door for peripheral designers. General consensus for peripheral designers has been that the x16 lane of the PCIe/104 connector is mostly wasted real estate, since most peripherals typically require a x1 PCIe lane only. The four x1 lanes in the PCIe/104 connector are carried in the first of three banks. The other two banks can vary, depending on whether you have a Type 1 or Type 2 PCIe/104 implementation, with the first bank remaining constant for both PCIe bus implementation types. The OneBank connector contains four x1 lanes or can be utilized as a single x4 lane. Care has been taken to ensure that the OneBank connector is compatible with the three-bank PCIe/104 bus architecture and can be mated directly to a full-sized three-bank connector. See Figure 1.

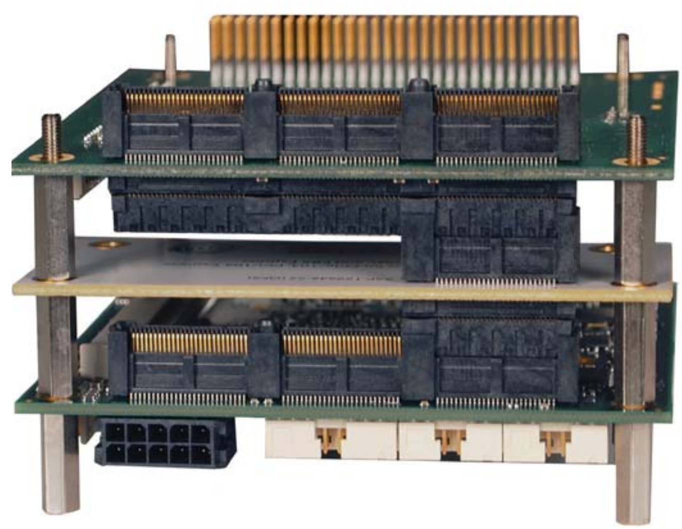

Figure 1: ADLE3800PC (Type 2) with OneBank™ Intermediate Peripheral Card.

The Greatest Threat?

There are several environmental factors that can put your hard work into an early grave when defining an embedded computer system. The top three system killers, in no particular order, are temperature, shock and vibration, and power.

Thermal management may be the largest threat to the lifetime of an embedded design. Although your system may be operating in a benign environment, thermal pathway issues can shorten product life significantly. Thermal issues can occur in a transmissive (radiative or convective) thermal solution if the thermal interfaces are not properly aligned for optimal heat transference. Misalignment when assembling or an extreme shock or vibration event could cause misalignment…especially if the thermal solution has not been adequately secured during system integration. Convection cooling systems are susceptible to thermal anomalies if the supplied airflow is impeded by dust buildup on fan blades, heatsinks, components and filters.

Extreme heat will not only affect the life of an embedded system, but also the performance. Most Intel architecture based embedded board designs include signaling for thermal management purposes. Onboard sensors monitor the processor die temperature as well as the board temperature at various locations and provide the necessary data for thermal management via the cooling fan (usually via BIOS and/or O/S based tool) by controlling fan rotation speed. This same technology is used to prevent an over-temperature condition within the processor by dynamically throttling back the CPU clock multiplier to try to keep it from reaching over temperature shutdown. If thermal pathway issues exist, your system could be perpetually running in a “throttled” condition, giving you less than 100 percent performance.

Headache Removal

Input power can also cause an array of unforeseen and sometimes unidentifiable problems for a system designer. While many COTS embedded products run on a single +5VDC rail, the quality of power supplied can vary greatly depending on the application. PC/104 boards are often integrated into larger assemblies such as environments where only mobile power is available. Power sources in planes, trains and automobiles vary greatly. Remote or “wearable” applications could use a renewable energy source, run on batteries, or a gas powered generator. Most ships use power from an onboard power generator to supply the onboard equipment. Most embedded boards simply do not have the necessary PCB real estate to include the power conditioning circuits that are found on desktop boards. Therefore, choosing a quality power supply with a high switching frequency can go a long way to removing embedded CPU design headaches.

Source power can pose unique problems that can shorten the life of a board. The severe power instability that can occur during a brown-out make this situation especially dangerous. This condition can send power spikes into connected equipment, often causing catastrophic component failure. Electrostatic discharge (ESD) concerns are also ever-present. ESD is always a concern when handling electronic devices, but ESD damage can also occur to a board during installation and a system where transient voltages may exist. Inadequate input filtering, hot-plugging of user interface devices, or poor grounding techniques can all lead to ESD-type system failures. Transient voltage damage can occur slowly over time, causing latent failures that are hard to identify.

Shock and vibration failures are much less subtle. Vehicles provide excellent examples of both shock and vibration. Provided that the necessary enclosure space is available, what are the chances of an ITX, mini-ITX, COM or even an Epic board surviving a vehicle crash, extreme airplane turbulence, or the concussion blast of artillery on the battlefield? One thing is for certain, PC/104 has a high rate of survivability in extreme conditions. Larger PCBs are more susceptible to flex conditions that fall outside of acceptable norms. This can result in a higher rate of failure in a high shock and vibration environment. Smaller form factors like PC/104 have a higher survivability rate. While COM is a small form factor, it employs a modular approach that adds to the risk of a shock failure because of the added connector interface. Vibration can create resonant frequencies within a PCB and can cause solder failures. Small form factor boards survive a wider range of vibration as they are less affected by lower vibration frequencies than larger form factors.

Customization

PC/104 form factor boards offer the highest level of customization of any standards-based COTS embedded form factors. There is no carrier board to design, and generally, no onboard consumer grade connectors that limit the product implementation. I/O’s can be custom cabled to meet specific needs, allowing the board to be deeply embedded within a system. The PC/104 specification has evolved with the rise of faster and more advanced bus architectures. Slower ISA and PCI use bus interfaces (PC/104 and PC/104-Plus) similar to VME style connectors, providing a deep and secure connection to added peripheral boards, and still maintains a presence in the field on platforms that have survived for decades. The current generation of PCIe/104-Express includes support for the fastest bus architectures and is well suited to support Intel technology advances for the foreseeable future. It remains the form factor of choice for rugged, extreme temperature embedded applications.

Matthew Henry is the verification engineer at ADL Embedded Solutions, based in San Diego, California. He is U.S. Navy educated in Electrical Engineering and has worked in the embedded computing market segment for the past 24 years as an FAE/OEM customer engineer and product manager.

Certified Small Form Factor Systems Offer High Performance COTS Building Blocks

Data gathering missions such as surveillance and reconnaissance are more often delegated to unmanned or portable systems—increasing the importance of small form factor design.

By RJ McLaren, Kontron

Military embedded designers are consistently tasked with increasing functionality and performance in smaller packages. Mobile deployments can include unmanned aerial vehicles (UAVs), vetronics and avionics systems—environments which allow no option to expand system footprint in order to make room for increased power or performance. SWaP (Size, Weight and Performance) concerns elevate in priority, along with the need for reliability and a strategic design path to accommodate inevitable future upgrades in the same small space. Pre-integrated systems meet a number of these needs, offering developers a highly effective starting point for their designs.

Validated high performance systems based on small form factor Computer-on-Modules (COMs) create a 'trusted COTS' design environment, relying on standard-based building blocks to simplify development. Using these certified systems, developers can quickly prototype designs by adding their own application-specific software, and testing functionality without extensive costs or development timelines. This provides ideal support to the military's agile acquisition initiatives that require Proof-of-Concept (PoC); design and deployment can be significantly accelerated, and developers can focus on the I/O customization aspects of their system, including support for both Linux and Windows operating systems.

Designing with Certified COMs-Based Systems

By packaging COM Express-based COMs in a rugged housing, designers have access to an application-ready platforms ideal for rugged, high-performance deployments. Operating as a nearly complete computer mounted on a carrier board, COM Express offers one of the smallest form factors available for military systems; I/O is customizable and designers have ready access to a range of module sizes that optimize SWaP considerations. Integrating mezzanine options adds tangible value, allowing developers to create new systems without significant modification to an original base design.

Now available to meet specific application requirements are complete systems which include rugged baseboard, power module, housing and appropriate I/O connectors. Systems on the market include. fanless, fully enclosed systems which weigh less than six pounds and offer efficient thermal management in a small 8.5 (W) x 5.5 (D) x 3.5 (H)-inch form factor. Designers can access a range of processor options to scale computing performance as needed for specific application requirements; for example, systems are available which can perform as a very low power Intel® Atom™ processor-based implementation or as a more powerful dual Intel® Core™ i7 processor system. The pre-validation of shock and vibration environments to the broad spectrum of UAV, tracked vehicle, shipboard and avionics environments makes performance possible in temperatures ranging from -40° to +71°C in order to reliably handle the physical rigors of military applications.

Flexible and Rugged, By Design

These enclosed systems offer designers access to COMs' extended thermal characteristics "by design." This process incorporates COMs that have been specifically re-engineered and validated for reliable performance in extended temperature applications, such as mil/aero and industrial environments that may range from -40° to +71°C. In a by design COM, complete systems are tested by suppliers, with individual components designed-in to ensure mission-critical performance and reliability under specific environmental conditions. In some cases, the processor may not be fully temperature rated by design and would then be 100% temperature screened for each production module. As another example, Kontron implements a special Rapid Shutdown circuit design on selected by design modules; this provides an onboard mechanism for the system to survive a high-energy pulse such as a nuclear event or high energy electromagnetic pulse (EMP). These types of sealed IP67 systems are geared to support high-rugged applications like vehicle- or helicopter-based computing needs.

Rugged COMs-based systems capitalize on the overall design value and flexibility of the platform, which provides the chipset I/O to the carrier board via rugged board-to-board connectors. Flexible I/O combinations are incorporated into the design via the application-specific customization of the carrier board, and could include LAN, SATA, video, audio and multiple USB or PCI Express ports. Video processing and display features are integrated, creating a notable advantage for graphics-heavy imaging and sensor data processing applications common to the modern military.

System Advantages

Systems such as COBALT leverage the COM Express Type 6 pin-out and its future design options. Type 6 reallocates legacy PCI pins to enable the digital display interface and additional PCI Express lanes. The importance for developers is that some of the extra PCI Express lanes can be routed to serial-based mezzanine card slots such as mPCIe and XMC to create expansion options. This also results in a performance jump contrasted to devices incorporating earlier pin-out options, as well as an enhanced fourth generation graphics architecture. Designers are enabled with a critical advantage for advanced video applications such as surveillance for situational awareness.

Because small form factor systems based on COM Express modules don't require a backplane (relying instead rely on a specific baseboard to host the module and I/O), there is space for only a small number of mezzanine expansion slots. System engineers must determine in advance the ideal configuration based on the necessary performance for the particular application or system. For example, the mPCIe mezzanine card offers a very small form factor in an off-the-shelf card, based on a standard specification maintained and developed by the PCI-SIG (PCI Special Interest Group); mPCIe can deliver specialized I/O such as video encoding, ARINC 429 or MIL-STD-1553 or more common wireless specifications like Wi-Fi, GSM or LTE. Embedding this type of functionality in the mezzanine rather than in a custom-designed board enables a longer and more flexible product lifecycle for deployed systems. Reusing these proven designs in smaller systems allows OEMs to extend functionality, even while avoiding customization and its related development costs and resources.

Figure 1: ADLE3800PC (Type 2) with OneBank™ Intermediate Peripheral Card.

Application Examples

Using Kontron's Computer Brick ALTernative (COBALT), we can illustrate a sample system based on a COM Express solution and integrating the industry's most current x86 modules. In this design scenario, system performance can evolve by swapping out modules to access processor advancements and avoid requalification of the design. Ruggedized, thermal performance is validated at the board level from -40° to +85°C, and this system's baseboard stack provides all necessary interconnects for the COM Express board and XMC and mPCIe (mini PCIe) interfaces. Based on rigorous environmental requirements, all I/O from the baseboard relies on a proven rugged connector, while all external I/O employs a 38999 type MIL circular connector. For OEMs, the XMC and mPCIe modules leverage industry standard boards that provide additional I/O functionality, capitalizing on the system's basic design to easily build a specific system profiles without significant modification.

A set of standard profiles features performance attributes that are applicable to a variety of mil/aero deployment – further simplifying design, reducing costs and accelerating development timelines. Performance may include increased storage capabilities with either fixed or removable solid state drives, wireless connectivity requirements including Wi-Fi, WiMAX or a 3G/4G modem, or an L2/L3 GbE switch for additional network port connections.

Removable Storage Profile

This is a common application requirement. A storage device may need to be removed and replaced in the field; security concerns may warrant the storage device itself is physically removed rather than having its data transmitted over a network interface. By using a certified small form factor system as a foundation for the application design, developers can have easy access to a variety of storage types and capacities. System developers can choose ideal options that meet current needs, updating or reconfiguring systems easily as requirements change. Extensive system changes or replacement is not required as higher capacity or better storage options become available, or as new program requirements are defined by military leadership.

Situational Awareness Profile

There is a growing need for video capture, compression, storage and analytics within rugged small form factor systems. The COM Express standard itself addresses long-life video support as a native feature within the chipset; standard I/O access is currently enabled for VGA, LVDS, SDVO and now Displayport, DVI and HDMI. Video cards are not always required (can be added via the XMC module), and neither is borrowing access from CPUs already space-constrained for processor, chipset and memory.

With the addition of a video encoder profile, enclosed COMs-based systems can perform video capture and compression, and support up to 4x composite video feeds as a dedicated video or data acquisition system. Data storage is flexible; video can be recorded on fixed or removable solid state drives, or off-loaded via wired GbE network connectivity or wirelessly via the Wi-Fi or modem option. Onboard 2x HDMI display outputs readily accommodate real-time display and playback.

Server Application Profile

High-speed signaling or image detection is frequently deployed in extreme rugged settings such as onboard vehicles or helicopters; applications onboard may include a full sensor array, including biometric, thermal, motion or radar. All of these sensors return data back to the networked server onboard the vehicle or aircraft; this can be reliably handled by a certified COMs-based system incorporating a rugged switch option to create a switch-stacked profile. The switch-stacked profile enables sensors and devices such as IP cameras to connect to a central server solution. Further, certified COTS-based systems may add an application on top of the network server capability, evaluating data, communicating information or delivering command alerts.

Determining which profile offers ideal performance may have designers balancing latencies and overall signal processing speed. For example, systems require some compression to transmit analog data from a camera to the onboard network. In certain applications such as situational awareness, even milliseconds to accommodate this process may not be acceptable, potentially requiring a choice between a dedicated hardware solution or network-based system.

When low latency is required, a dedicated video or data acquisition profile may provide ideal performance. Such systems are optimized to efficiently handle a specific task through special algorithms versus a server profile which is designed for more generic processing performance. While high performance and suited for extensive data collection and storage, the server profile's performance is directly related to the limitations of its connected devices.

Servers can be used effectively for offloading, using two channels to route data for handling. One channel may execute certain functions related to certain data types, for example delivering critical data to a command and control application, while the second channel routes less critical information back to the server for offloading or data analysis. Multiple certified systems can also work in concert, for example the COBALT situational awareness system can be connected to a system based on the server profile. Systems can also be configured to act as sensors, for instance adding an XMC signaling card and then networking the system back to the server profile.

Certified Platforms Add Design Value

Gathering and sharing data has become an essential component of the modern military, relying on high performance processing to fuel leadership decisions, enable comprehensive situational awareness, and ensure the overall safety and security of military personnel. Data gathering missions such as surveillance and reconnaissance are more and more often delegated to unmanned or portable systems – increasing the importance of small form factor design. Flexible, cost-effective design is the primary purpose of certified small form factor systems, leveraging a trusted COTS platform as a generic building block solution for these types of rugged, high performance applications.

By capitalizing on COMs-based certified systems that couple mezzanine modules with carrier boards, military designers can readily use and reuse low power, reliable designs to develop smaller high performance systems. Optimized as pre-validated military design platforms, rugged small form factor systems reduce development cycles, enable fast PoC and act as flexible, standards-based building blocks for the range of mil/aero deployments. Maximizing flexibility with upgradable processors and perfect fit design by virtue of carrier board and mezzanine card options, developers have a fully tested, cost-effective design path that meets mil/aero requirements and gets to market quickly.

R.J. McLaren is portfolio manager, Commercial Avionics & Military Products, Kontron.

Beware These Evil Twins and the Threats to Network Traffic

What we don't know about miscreants glomming onto our Internet activities can hurt us.

By Art Swift, CUPP Computing

NPR and Ars Technica, working with Pwnie Express, cooperated on an interesting experiment in what can be learned from snooping an individual's network traffic. Some of it was to be expected. Ongoing revelations about the activities of the NSA and various hacker groups make it harder to be surprised by the information that we reveal through our Internet traffic.

In the span of a red light, a substantial snapshot of what's going on in your phone—possibly adding up to enough information to identify who you are—can be collected. Courtesy Wikipedia Commons.

On the other hand, insights about how bad actors can latch onto our Internet traffic were eye opening. Particularly interesting was the burst of traffic that occurred when a mobile phone connected to a Wi-Fi network. It's possible to imagine this leading to some pretty scary scenarios. Consider an innocuous-looking utility truck parked at a major intersection. What's actually inside the truck is a Wi-Fi hotspot, broadcasting the Service Set Identifier (SSID) "attwifi." The phone of any AT&T customer who has not changed the default connection settings will automatically connect to the network. Once connected, phone apps running in the background—particularly those that only update when connected via Wi-Fi—may choose to "check in."

The browser in the phone might refresh some, or all, open pages. The email app might download new messages. And various social media apps can check in for new info. In the span of a red light, a substantial snapshot of what's going on in your phone—possibly adding up to enough information to identify who you are—can be collected. Needless to say, it also includes where you've been, and when you were there.

Leaky Bucket

You might thing that you're safe when you get to your favorite coffee shop, which happens to be a Comcast hotspot but, alas, you're not. The "xfinitywifi" that your phone automatically connects to could be the actual hotspot, or it could be an evil twin network from the guy at the table next to you, broadcasting his own "xfinitywifi." If your phone connects to it, then you've become an unwitting participant in a classic man-in-the-middle attack, and everything you do will be snooped, giving your attacker a foot in the door of your online persona, which can easily be leveraged for various purposes.

So your standard off-the-shelf phone arrives as a very, very leaky bucket when it comes to information security. It's possible—advisable, even—to address these leaks piecemeal, service-by-service, e.g., with encrypted messaging or email. But it's critical to try to address the problem at a system level as well. One way to tackle the problem, for example, is with an inexpensive dedicated-hardware security solution packaged in a microSD card suitable for a phone or tablet.

And the bottom line? Businesses today need a more secure, container to shore up your leaky bucket, and help you ensure that your private data remains private.

Art Swift is CEO of CUPP Computing, which provides security solutions for mobile systems, such as tablets, smartphones, remote service devices and the Internet of Things (IoT). CUPP Computing is headquartered in Oslo, Norway, and has operations in Netanya, Israel and Palo Alto, CA

Not Enough Just to "Take up the IoT Mantle": Interview with Eurotech CEO Larry Wall

As embedded systems designers and chip makers turn away from a past that showed a bias toward performance and form factor, with power consumption concerns on the back burner, and as the IoT matures, what's the best way forward for industrial, mil-aero, medical, automotive and other applications?

By Anne Fisher, Managing Editor

Larry Wall
Eurotech Inc.

"So many companies have taken up the IoT mantle," Eurotech CEO Larry Wall told EECatalog recently, explaining that the reason some enterprises are jumping on the IoT bandwagon is to "be part of the hype." Rather than seeing embedded developers taken in by "big brands and big promises," Wall advocates, "working with companies who have proven they know how to build, connect and manage devices." Other must-haves according to Wall: standards-based building blocks—from embedded boards and programmable multi-service gateways to connectivity and cloud-based processing.

Wall also told us, when asked how he would like to see the Industrial, Medical and Military IoTs mature, that wider adoption of open source, standard protocols such as the Eclipse Kura Project, the OSGi container, and the Linux operating system would help. "Open source software with a broad reach gives developers flexibility to accelerate IoT solutions more smoothly and efficiently," he said. And this is true, according to Wall, no matter where an application lives in the Internet of Things.

He explained that designers and customers need to connect easily to the Cloud, and they need to be able to manage devices remotely to accommodate large and varied networks of devices. Wall names privacy and security as crucial demands that are ballooning in importance as the IoT evolves. "These requirements are never-ending issues, of course, and require designers and suppliers to evolve as capabilities grow and change, and frankly, as hackers and their capabilities grow and change," he says.

Read on for Wall's answers to additional questions on ultra-low-power designs, small form factors and the IoT.

Figure 1. Lured in? Eurotech CEO Larry Wall notes the problem of embedded developers being lured by "firms that want to take up the IoT mantle, but which may not have a track record for building, connecting and managing devices." [Image By Ankara (Own work) [CC BY-SA 3.0 http://creativecommons.org/licenses/by-sa/3.0)

EECatalog: What philosophy should a manufacturer of low-power Computer-on-Modules (CoMs) follow to assure that its offerings are more than just the sum of their parts?

Larry Wall, Eurotech: The best thing is to provide flexibility and give customers options. Every application is different, every industry is different, and the CoM must allow the designer to optimize the hardware to meet their needs. For instance, Eurotech's Catalyst BT Ultra Low Power CoM offers the Intel® Atom™ processor E3800 Series quad-core, dual-core and single-core options depending on the need. (Figure 2). It also supports both Windows and Linux and has the OS installed for faster integration with the device.

EECatalog: As you consider what has led to low-power (or other) milestones and breakthroughs, are there principles and practices to be recognized that could apply to challenges the industry is experiencing today, and to anticipating future needs?

Wall, Eurotech: Future-proofing is a key in low-power design. Designers must start with a clear understanding of the customers' application and the problem they are trying to solve, and then apply standard or customized products to meet that need. Most importantly, they have to allow for change or expansion. The closer you match with the customer, the better you match their needs.

Software also plays a huge role in helping developers be prepared for future needs. With applications built on the right framework, development work doesn't have to be lost then repeated when moving to a new hardware platform. The right framework allows an application to be ported to a new hardware platform in a straightforward manner.

Additionally, full systems—processors, software, I/O, communications—must be designed with a focus on power consumption to enable new applications operating in remote and battery-powered systems

EECatalog: In what areas are you finding designers struggling most as they strive to optimize and simplify the design process, particularly when working on ultra-low power designs?

Wall, Eurotech: One of the difficulties we see designers face in low-power design is understanding the application well. Designers must understand the application use cases fully, including the wake up and sleep needs and possible future needs as the application evolves over time. The best way to design a power-efficient system is to have the device turn off or go to sleep and then start when necessary. Understanding the balance for the application is essential to decide which mode works best.

Connectivity is another challenge we see in the market. It is now common for devices to contain multiple radios and connectivity options that are software- or environment-defined. When considering the wake up cycle, if the device doesn't find the network or takes a while to connect, the system is drawing off battery power during that delay. Both static and mobile devices have challenges with connectivity, so the right modes must be employed—RF, satellite or cellular depending on the application. No device is an island—in today's world of M2M systems and devices connected to the Internet of Things, connectivity is key for all designs.

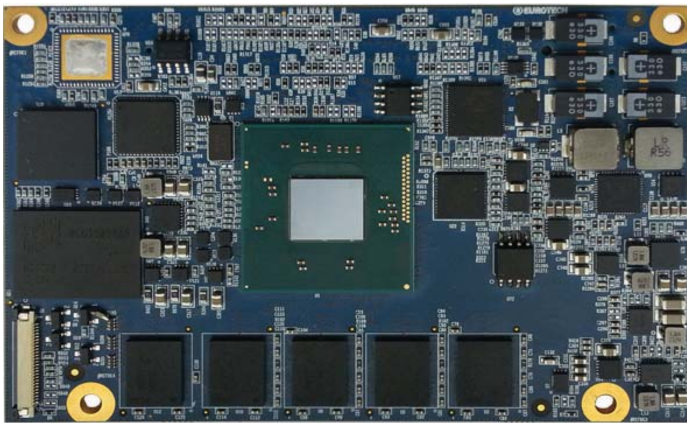

Figure 2. The Catalyst BT is an Ultra Low Power CoM based on the Intel® Atom™ SoC. Photo courtesy Eurotech.

Last, some designers struggle with selecting the right software and development tools to optimize and simplify the design process, especially for future proofing designs. Open source versus commercial systems, coding in Java or C++ and what software frameworks to use are common questions.

EECatalog: What are some approaches to addressing the challenges you've just named?

Wall, Eurotech: As designers struggle with understanding their application and how to optimize for low power, embedded system providers like Eurotech can enable several power modes so customers can make these choices and optimize a low-power design. In addition Eurotech employs an extensible embedded controller to manage many diverse functions such as battery monitoring, A/D conversion, or even connectivity status monitoring including wireless connections and GPS interfaces. The controller can handle any low-performance function that requires some periodic servicing.

Connectivity options are improving with designers able to choose wired, WiFi, cellular, Bluetooth Low Energy, and others. However, challenges are also present. For example, one of the challenges of cellular technology in particular is carrier certification. This process can cost thousands of dollars and take many months to complete. Designers have to understand the cellular landscape and then navigate the testing process, which requires extensive expertise most designers do not have.

Eurotech offers a pre-certified cellular adapter, the IP67 Certified ReliaCELL, to solve this problem, with a suite of ReliaCELL units certified on different networks. Designers can simply plug the ReliaCELL in and mount it, and the expensive and time-consuming cellular certifications are already taken care of. This flexibility gives organizations options to deploy devices and services in different regions.

From a software perspective, integrated hardware and software offerings simplify the design process so designers and developers experience a reduced amount of subsystem integration. Open source software for device developers is becoming more prevalent.

EECatalog: What, if anything, are you seeing as problems or challenges that are not getting their fair share of attention in the area of ultra-low-power design and in the embedded ecosystem more generally?

Wall, Eurotech: We see the embedded ecosystem by default evolving into the IoT ecosystem. With all the hype around IoT, the challenge for designers is to sift through all the marketing promises to determine what is real and deployable now. On a related note, it is expected that innovation will continue to occur at a rapid pace for all IoT tiers. During this period, embedded system suppliers must provide technical and business assurances to customers to reduce the real and perceived risk of solution adoption.

EECatalog: What is the role management from a remote location is playing and will play to help the Military Internet of Things, the Medical Internet of Things, and the Industrial Internet of Things mature?

Wall, Eurotech: Designers and customers need to connect easily to the Cloud, and they need to be able to manage devices remotely to accommodate large and varied networks of devices.

In these three vertical markets in particular, remote management saves time, money, and even lives. In the military, it can be dangerous to manage a sensor on-site. Remote management can save lives by enabling updates, configuration and troubleshooting without physically touching the device. In the medical field, remote management means a network of thousands of devices can be updated all at one time remotely, so patients can get the most up-to-date care available. And in the industrial market, managing devices remotely saves money when technicians need not be sent to service devices in the field.

One example of a solution that enables a complete set of remote device management operations, over MQTT, either on demand or on schedule is Eurotech's Everyware™ Software Framework (ESF), a Java OSGi software framework for M2M/IoT multiservice gateways, smart devices and IoT applications, and which interacts with Eurotech's Everyware Cloud Web Console.

ESF enables the delivery of the latest generation Pervasive/Embedded Computer hardware platforms and all the fundamental components needed to quickly develop complex applications as part of the device software as well.

In summary, the military and medical markets require higher and broader security capabilities and predictable reliability in wide operating environment range. All markets benefit from protocol standardization, but the industrial markets, where there is large volume and proliferation of sensors, could benefit the most in the early phases of broad IoT adoption.

EECatalog: What are the top 3 technologies, capabilities or features Eurotech wishes were available?

Wall, Eurotech: Eurotech would like to see more acceptance of open source software as a starting point for development. Developers can begin their development with open source software, then move to a fully supported version for extra confidence and assurance for deployments.

We would like to see ultra-low-power sleep modes available in chipsets and connectivity options accessible in chipsets. We are adding these options separately, but look forward to seeing WiFi or Bluetooth integrated directly into chips.

We'd especially like to see more deployment of connected devices taking advantage of remote management capabilities, instead of relying on inefficient processes based on decades-old technology.

EECatalog: What trends should system developers keep an eye on?

Wall, Eurotech: There is a trend towards having more sensors close to where data is measured. Intelligence is getting pushed in both directions, with big data analysis on one end, and collecting more data on the other end. System and software designers must develop the skills to understand the tradeoffs and pros and cons of varying models of computing distribution. These design decisions cannot be made at the subsystem level, but require comprehensive understanding of the system-level requirements and objectives.

System developers must learn more about data compression and setting data up for cloud services for analysis. The big data trend will continue to accelerate. We will continue to see data structured for use in traditional and big-data analytics systems, but we will also see a demand for IoT data that's structured for easy integration into vast numbers and types of enterprise application beyond analytics, CRM and ERP.

Anne Fisher is managing editor of EECatalog.com. Her experience has included opportunities to cover a wide range of embedded solutions in the PICMG ecosystem as well as other technologies. Anne enjoys bringing embedded designers and developers solutions to technology challenges as described by their peers as well as insight and analysis from industry leaders. She can be reached at afisher@extensionmedia.com

Industrial IoT Looks to SMARC Architecture to Meet Low Power, Performance and Cost Goals

As industrial, digital signage/POS, medical, automotive, smart energy and other IoT strategies unfold, SMARC is emerging as key to a new generation of embedded computing applications.

By Dirk Finstel, ADLINK Technology

Isolated systems symbolize the past. A network of smaller form factor sophisticated devices with power-saving thermal characteristics signals the future. When this network makes Industrial Internet of Things (IIOT) applications possible, it does so by achieving both low power consumption and healthy computing performance from control devices. Designed to play an essential role in a network that enables intelligent industrial applications is the SMARC™ computer-on-module form factor, which powers IoT devices using Intel® x86 System-on-Chip (SoC) and ARM technologies.

Redefining the Low-Power Landscape

Embedded computing today encompasses markets where users expect energy-efficient systems with a generous helping of intuitive user interfaces based on touch screens and voice recognition. These interfaces promise revolutionized medical treatment, more sophisticated industrial equipment control and Internet access in far-flung areas once thought off limits.

As x86 architecture has evolved, Intel has addressed low-power requirements with smaller, more efficient Intel® Atom™ and Intel® Quark™ processors. The latest Intel Atom processor E3800 SoC is optimized for systems that require efficient imaging workflows, delivery of secure content and interactive, real-time processes. The Intel Quark SoC X1000 targets applications in automotive and wearables, where lower power, lower cost and smaller size take priority over higher performance.

Optimized for the mobile phone and tablet markets, the ARM architecture extends the power envelope into low-energy applications that have previously found it difficult to adopt standard form factors, and so had to absorb high up-front custom board and module design costs. At the same time, ARM offers scalability into upcoming high-performance computing and 64-bit platforms that make it possible to build advanced, highly energy-efficient server platforms.

Intel has been instrumental in defining the core microprocessor and instruction-set architecture and the architecture of peripherals. Proprietary or open-standard embedded-computing products based on the x86 architecture leverage Intel's chip-level expertise and employ a common I/O interface. Through the use of common connector pinouts, it is possible for customers to select from a wide range of hardware- and software-compatible peripherals with which they can customize their end products. The ARM environment is more complex and differentiated. In contrast to the PC environment, in which the core module comprises a processor and Northbridge and Southbridge device, the ARM market focuses on SoC products, each usually optimized for a particular application. Historically, there has been far less focus on building standard I/O definitions; each SoC would be used on a custom board design. There is also a wider range of I/O options that ARM platforms provide, depending on their target market, with less emphasis on standard buses such as PCI Express.

The result has been the introduction of a number of proprietary form factors and connector definitions that lock the customer into a vendor's offerings and which may not have support for more than a generation of silicon as they move to different SoCs. Some vendors claim the use of a standard form factor—sometimes piggybacking ARM support on an existing x86-focused specification—but with additional custom connectors to support I/O lines that cannot be supported by the primary connectors.

SMARC: A Unifying Architecture

Supported by a number of embedded computing module vendors and held by the vendor-independent Standardization Group for Embedded Technologies (SGeT), the Smart Mobility ARChitecture (SMARC) provides an open-standard definition for ARM-based and x86-based SoC embedded computing solutions, optimized for low power, cost efficiency and high performance. SMARC also supports systems that need more

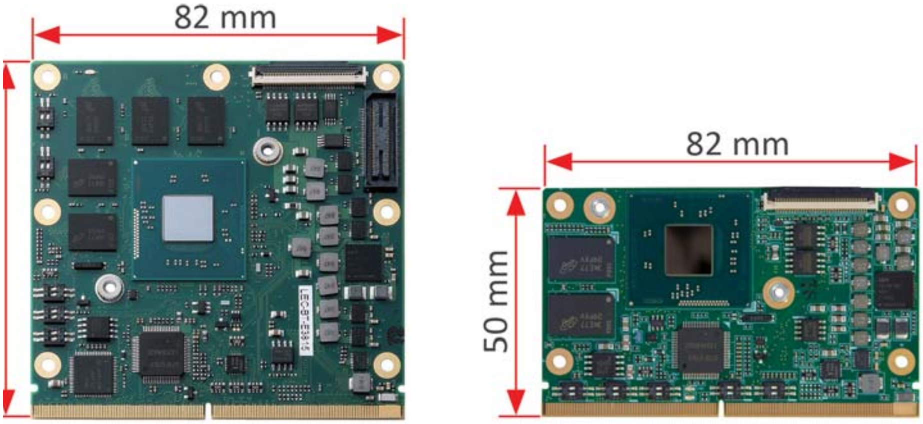

Figure 1: SMARC defines two sizes of module: a full-size module that measures 82mm x 80mm, and a short module for more compact systems that measures 82mm x 50mm.

The SMARC module is designed to support a combined height above the carrier of less than 7mm (See Figure 1). The PC heritage of most computer-on-modules (COMs) has led to the assumption that all COM boards will be used with a heatspreader, which adds to overall package height. The typical combined height of the processor board and heatspreader alone exceeds the height of a package that includes both the SMARC COM and carrier board. Many SoCs do not require a dedicated heatspreader because of their lower overall power consumption. The SMARC format allows for this, making it more suitable for systems where space is at a premium.

compact solutions than are offered by the PC-oriented form factors.

As SoCs do not need the support chips of a PC platform and draw less power, the amount of board space that needs to be reserved for power converters and power supply lines is greatly reduced. This allows the use of a smaller form factor, facilitating use of SMARC-based modules in low-power portable equipment. SMARC CPU modules are expected to have an actual power intake between 2W to 6W, allowing for passive cooling and thus further reducing subsequent design effort and overall cost. The standard allows for up to 9W continuous power draw for more demanding scenarios.

Based on the proven connector as it is employed by Mobile PCI Express Module (MXM) video modules, SMARC defines two sizes of module: a full-size module that measures 82mm x 80mm, and a short module for more compact systems that measures 82mm x 50mm. The edge connector supports 314 electrical contacts. For systems that are to be used in harsh environments, shock- and vibration-proof versions of the connector are readily available. The temperature range of the connector extends from –55°C to +85°C.

The SMARC MXM connector guarantees a high degree of signal integrity, required by high-frequency serial interfaces. For example, on 2.5GHz signals as employed by PCI Express Gen2, the insertion loss of the connector is just 0.5dB. In comparison, the insertion loss encountered on the connection scheme used by previous generation MXM connectors is significantly higher at 3dB. SMARC also supports a wide input voltage range, reducing the need to use additional DC/DC converters on the core module and overall power dissipation. A SMARC module can support input voltages from 3V to 5.25V. Originally designed to support PC-class hardware, the many other formats are restricted to a nominal 5V input.

Case Study: SMARC-based IoT Gateway Powers Machine Failure Prediction Application

Industrial machinery is subject to nearly constant shock and vibration, which generates fatigue and wear on materials and components. Forecasting potential problems in order to implement preventive measures and maintenance or equipment replacement is critical to sustain performance and avoid costly downtime and damage.

Machines or tools using a rotary axis, for example, frequently exhibit vibration and noise from unbalanced rotation. This imbalance reduces the life of shaft bearings. To avoid unacceptable vibration and noise, a process called dynamic balance is employed, using accelerometer and tachometer readings to calculate and adjust the unbalanced mass and deviation angle. Temperature and voltage readings can also be used to diagnose machine health.

Real-Time Simultaneous Monitoring

Conventionally, data acquisition modules collect measurements from sensors located on individual machines. Based on the data, portable vibration detectors are used to check potential problem machines one by one, a time-consuming and labor-intensive practice, with unavoidable errors and slow reaction times.

In environments where multiple machines are monitored, development of secure and reliable networking of devices from the edge to the cloud is crucial to realizing the much needed real-time simultaneous monitoring. However, current manufacturing environments utilize a wide variety of fieldbus communication protocols such as Modbus, PROFIBUS, DeviceNet, CANopen, PROFINET, EtherNet/IP and EtherCAT. The disparity among them presents a significant communication barrier impeding delivery of field data to the upper layers of the cloud structure. Overcoming this barrier to enable data flow to the cloud

Figure 2. Combined with cloud services, SMARC-based IoT gateways can push data to the cloud, deliver remote monitoring and control capabilities, and offer rich libraries and tools allowing OEM customers to easily configure and design intuitive GUIs viewable on any browser-based device for 24/7 data access.

without replacement of legacy equipment is a major challenge for IIoT applications.

In addition, a reliable, secure data connection to the cloud must be in place to protect sensitive data. Security is critical to IoT-based operations at both the device and network level. Secure boot at the device level, as well as access control and authentication, application whitelisting and firewall and intrusion prevention systems can all help combat security threats while allowing the connection of legacy equipment to the cloud.

Finally, factory environments present demanding physical conditions. Industrial machines are frequently located in areas experiencing regular shock and vibration, extreme temperature and humidity and the presence of liquid and solid contaminants. IIoT devices require dedicated design and construction to withstand these conditions and ensure reliable operation.

Reliable Connectivity and Secure Data Transfer with Intel ® x86-based SMARC boards

The advent of cloud-based IIoT solutions has vastly improved the field of machine failure prediction, with real-time data transmission, remote monitoring and control, and enhanced accuracy, efficiency and economy. IIoT solutions require a vertical integration of networked field devices such as sensors and data loggers, IoT gateways (controllers) and cloud servers. The IoT gateways collect data from the field, implement primary data analysis ("fog computing") and push the analyzed data to a cloud server for more advanced analytics. At each point, network connectivity is critical.

IIoT gateways powered by Intel x86-based SMARC boards include characteristics that make reliable connectivity and secure data transfer possible. These gateways can offer a fanless, compact footprint that delivers both low power and high performance. When built on the Intel® IoT Gateway platform, a portfolio of pre-integrated and pre-validated software and hardware bundles developed to address the challenges facing IIoT, these gateways enable the connection of legacy industrial devices and other systems to the cloud.

Designed in collaboration with McAfee and Wind River, Intel IoT Gateway platform solutions are intelligent systems integrating technologies and communication protocols for networking and embedded control with reinforced security. McAfee Embedded Control maximizes security by dynamically monitoring and managing whitelists that allow only pre-authorized programs to run in the system. Wind River Intelligent Device Platform XT and Wind River Workbench provide a proven development environment.

Combined with cloud services, SMARC-based IoT gateways can push data to the cloud, deliver remote monitoring and control capabilities, and offer rich libraries and tools allowing OEM customers to easily configure and design intuitive GUIs viewable on any browser-based device for 24/7 data access.

Application Scenario: Deployment of Dynamic Signal Acquisition Modules

A possible application scenario includes deployment of dynamic signal acquisition modules and a thermocouple input module to acquire and collect sensor data and transmit it to an IoT gateway controller for primary computing. The gateway then sends only processed meaningful data to the cloud server for further analysis and utilization. Computing performed at the gateway, or fog computing, alleviates cloud computing loading and reduces data traffic volume over the network, improving connectivity and reducing networking costs.

For peak performance of local analytics and other applications, ADLINK's own SMARC-based IoT gateway, the MXE-200i, features the dual-core performance of the Intel Atom processor E3826 (see Figure 2). This processor supports Streaming SIMD Extensions (SSE, SSE2, SSE3, SSSE3) for fast, efficient digital signal processing of sensor and instrumentation data such as the FFT algorithm. To accommodate a wide range of indoor and outdoor environments, the MXE-200i features a compact footprint—120 (W) x 60 (H) x 100 (D) mm, and rugged, fanless design with an environmental temperature tolerance of –20°C to +70°C. The platform offers rich I/O for sensors and data acquisition modules, as well as optional wireless modules for Wi-Fi and 3G/4G/LTE.

Conclusion

The application of IoT technology helps to implement real-time and remote monitoring over industrial machinery, allowing more precise forecasting of possible equipment malfunction and responsive advance measures, minimizing unexpected failure and consequent damage and reducing repair, maintenance and human resource costs, for improved precision, yield, and quality. Most importantly, the intelligence gained from analyses of accumulated big data aids development of more advanced insight into equipment management and purchasing.

Industry Impact of Low Power Improvements

In this new era of connected computing technology, intelligent systems add global value as standalone systems evolve from their foundation into connected networks that communicate with each other and the cloud. OEMs and developers can anticipate a convergence of increasingly connected devices, answering demand for real-time data gathering and sharing, nonstop communication, new services, enhanced productivity and more. Systems will solve business problems by being smart and connected, which is becoming a priority, adding business value such as cloud economics for compute and data operations.

As IoT strategies unfold—for example in healthcare, smart metering, digital signage/POS and retail banking, factory floor systems and connected buildings—the business case for intelligent services increases. Minimal power consumption is a primary driver in this renaissance. Low-power designs powered by SMARC-based building blocks support mobile, portable or embedded devices. And they capitalize on thermal characteristics to manage fully enclosed, passively cooled designs as the key to enabling connectivity anywhere. Thanks to its focused support for low-power architecture and backing from multiple vendors, SMARC is the key form factor for a new generation of embedded computing applications.

Dirk Finstel is CEO, EMEA, & EVP, Global Module Computer Product Segment at ADLINK Technology. He has 20+ years of in-depth experience in leading embedded computer technology, with a proven track record in embedded modules. Finstel has held executive-level positions at embedded computing companies since he founded Dr. Berghaus GmbH & Co. KG in 1991, and has been responsible for global technology, as well as research and development and setting technological strategy. Finstel holds a BS in Computer Engineering & Science.

Embedded Medical Device Software—How And Why To Use Static Analysis To Update Legacy Code

Is the ability to bring high-quality embedded medical applications to market endangered?

Dr. Paul Anderson, Grammatech

In common with most of the rest of the embedded software development industry, the code in medical devices is becoming more complex. Unsurprisingly the frequency and cost of failures of medical devices continues to rise. In fact, a recent FDA report stated medical device recalls have jumped 97% between 2003 and 2012 —and one of the four primary causes of recalls is software. To understand how costly these recalls can be, a recent Reuters article reported on a single medical device recall that cost $221 million.

In their research, the FDA identified software design failures as the single most common cause for device recalls, with an average of 15% of all recalls caused by software. Based on the 1,190 recalls in 2012, this means about 178 medical device recalls in 2012 were due to software defects. It is also interesting to note that 80% of all recalls were attributed to devices manufactured in the U.S.

Which prompts the question...why is this still happening when we have better developers with better processes and better tools? The short answer is that there is no single reason. The data from the FDA tells us that despite increased sophistication in processes and analysis technology, medical device manufacturers "might" actually be getting worse at delivering high-quality embedded applications, not better. So what are all the potential reasons for this?

Networking—As an element of medical devices is now exposing software to new quality and security issues.

Third-party code—Increasingly, be it open source or commercial code, medical device manufacturers are relying on third-party suppliers, so deep familiarity with the code base is being sacrificed.

The fundamental challenge medical device manufacturers face is how to build software that takes advantage of the latest hardware, yet can still survive and thrive in the wild where it increasingly needs to be network aware (and secure). For many companies, this means the modernization of legacy code bases. When modernizing any mission- or life-critical software system, static analysis is an excellent tool because the complexity of the process is simply too great in most cases to entrust to code reviews.

Figure 1. Modernizing legacy code bases can be a boon to the task of developing medical device software that is network aware, secure and capitalizes on hardware innovations.

Migration to Modernity

I have personally had a hand in a number of high-visibility efforts to address software safety—including working directly with the FDA and assisting with forensic investigation of major software failures other government agencies conducted—which is why I have been so impressed with the work done by Micrel Medical Devices. A leading European medical device manufacturer, Micrel Medical Devices has taken an advanced, process-driven approach to migrating key legacy code bases to a more modern language and structure. To accomplish this, the company began with the following assumptions:

Micrel Medical Devices' fast growing installed base of medical devices that use both new and legacy code has made delivering safe and reliable devices more complicated.

The best way to ensure consumer safety and satisfaction was to modernize legacy code, and moving aggressively to modernize would reinforce the company's commitment to patient care.

However, modernizing legacy code is challenging and would require significant changes to development processes as well as the use of new analysis tools.

A Careful Look at NASA

The legacy code bases that needed modernization were written in assembly language. Micrel Medical Devices' development leadership decided to move to modern dialects of C and C++. The company also adopted a safety-critical coding standard. Originally, the company's internal coding standards had been created by summarizing findings from code reviews and test results, then integrating these findings with factors such as safety, expandability, modularity, readability and efficiency. Unconstrained use of C and C++ can be highly risky for several reasons. For example, it is easy to write programs that have undefined behavior, and the runtime provides few safeguards to prevent such behavior from silently corrupting the program's state. Also, aspects of the semantics such as how values are implicitly converted from one type to another are non-obvious and are frequently misunderstood by even the most experienced programmers. To avoid such pitfalls, Micrel Medical Devices' leadership understood that new standards would need to be created or adopted.

The company decided the best course would be to adopt an existing, proven external standard for C/C++ , and after evaluating multiple standards, elected to adopt the NASA/JPL coding standards and the Power of 10 Rules, which were developed by Gerard J. Holzmann of the NASA/JPL Laboratory for Reliable Software.

The development team at Micrel Medical Devices found these standards most closely matched its existing internal coding standards, which would make it easier for developers to adhere to them. Additionally, the team believed these standards closely matched the rigor of the company's existing internal code safety standards. Using this respected standard would also make possible a tangible reference point for compliance for regulatory bodies that are critical to the company's continued growth.

Extreme Testing

Driven by its commitment to continuous process improvement, Micrel Medical Devices searched for a static analysis tool that would deliver the required compliance to internal coding standards and also expose the code to a deeper level of analysis. To accomplish this, the firm's R&D team devised an "extreme testing" regimen to identify the most advanced static analysis tool, capable of delivering the precision and depth of analysis necessary to support its product line, classified by the FDA as Class II B.

The winning static analysis tool would need the flexibility to support the company's existing internal as well as its newly adopted external coding standards. During the evaluation, Micrel Medical Devices tested static analysis tools from multiple vendors against the same code base. The evaluators judged the tools on many factors, including ease-of-use, compatibility with the company's coding standards, depth of analysis and overall fit with the development process.

Performance—defined as the ability to assure the highest quality software—would be the most critical decision making point. Ultimately, the team broke the key components of "performance" into three key factors that the company would base its static analysis investment upon:

- Depth of Analysis—Due to the life-critical nature of the company's products, Micrel Medical Devices developers needed an automated tool that would excel at pinpointing hard-to-find defects that may be deeply hidden.

- Flexible Reporting—Beyond rudimentary graphical compliance metric representations, the firm needed text outputs that could be compared easily with previous outputs from earlier code versions.

- Native JPL Standard Support—This would empower developers to modify rule sets with ease.

- After months of rigorous testing, Micrel Medial Devices chose GrammaTech's CodeSonar.

Conclusion

Software quality for medical devices can hold lives in balance, as can the software used in other safety-critical products, so there is simply no way to overstate the importance of software quality in this industry. That said, selecting the right tools can be challenging because the technology is so sophisticated, and the problems they are being used to solve are very complex. The selection process can also be difficult sometimes because static analysis tools are often strategic investments for companies, which can pose budgetary challenges.

In a 2013 McKinsey & Company study, it was reported that "using medical device company financials, proprietary benchmarks, and expert insights, we (McKinsey) estimated the total cost of quality for the medical device industry—including day-to-day quality costs as well as cost and revenue loss from non-routine events—as $17 billion to $26 billion per year, or 12 to 18 percent of industry revenue."

With a price tag of this magnitude, it is vitally important to build a software development process from the ground up to deliver sustainable quality. Patching old processes while racing to deliver new products in this competitive marketplace clearly leads to failure. Micrel Medical Devices' approach to the modern challenges posed by rapid advances in technology is, in my view, a textbook example of how to do this the correct way.

Dr. Paul Anderson is VP of Engineering at GrammaTech.

A New Prize in Your Serial Box

Advanced, programmable I2C and SPI T&M equipment opens up flexible options.

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

Two popular low-speed serial buses continue to live on, but changes are afoot in the test and debug equipment used by embedded systems engineers. Philips developed the serial Inter-IC bus (I2C) in the 1980s, and Motorola years ago developed the Serial Peripheral Interface (SPI) bus as an equally simple serial data link between processors and slave devices. I2C uses two wires; SPI uses four (Figures 1 and 2). I2C leaves a smaller footprint and can connect multiple masters to multiple slaves using two wires, but has lower speed performance. On the other hand, SPI offers much faster data rates, but individual lines are required for multiple slaves. Both protocols are simple, work well and are ubiquitous in all kinds of embedded systems—from processors and peripherals to non-volatile memories like EEPROM and Flash to sensors and slow I/O devices.

But ubiquity has its downside. There's a common belief that debugging systems with these standards is equally simple and straightforward. Not always true. And since they're so common, designers may think that low-cost rudimentary test and measurement (T&M) tools are "good enough" for these modest buses.

The truth is, evolution of both SPI and I2C demands better tools to debug newer features; remote- and network-equipped T&M equipment is needed in modern engineering and production environments; and "if this, then that" automated style data/event management on both I2C and SPI buses demands more sophisticated equipment to isolate data/events and to debug these complex embedded multi-slave systems. And despite the prevalence of high-speed serial I/O like PCI Express on the latest CPUs and peripherals, SPI and I2C are still appearing on new SoCs, MCUs and peripherals.

To keep up with change, it's about time for a better T&M tool.

Speed Demons?

In a digital environment, no one is ever going to consider 1 kHz "fast." Yet the bottom end of I2C clocks at a mere 1 kHz, while SPI starts out at 31 kHz. These low speeds constrain both the Master sending signals on the bus and the speed at which the Slave responds and/or remits data. The faster the clock, the more information can be transmitted per cycle. Test equipment designed to work with SPI or I2C should meet the latest specs—or come close. Trying to debug a system whose speed exceeds the T&M equipment is problematic.

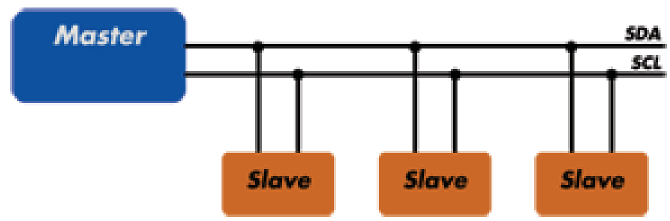

Figure 1: I2C uses a two-wire interface where the lines are shared by the master and multiple slaves and typically maxes out at 1 MHz. (Courtesy: Total Phase.)

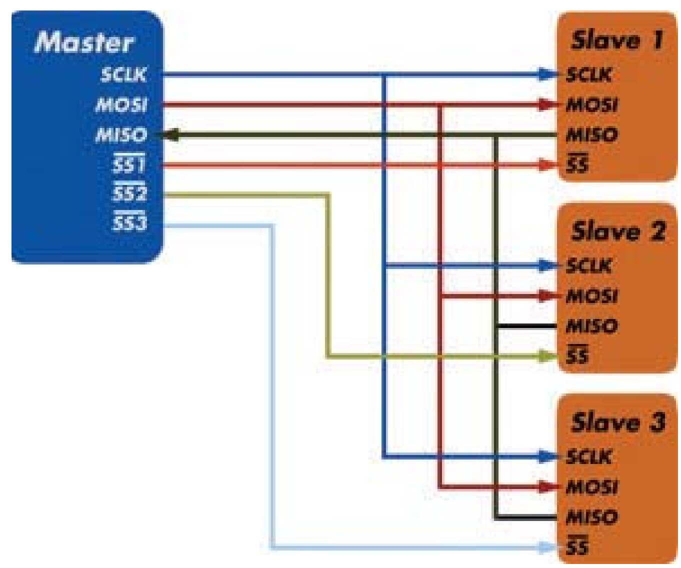

Figure 2: SPI uses a four-wire interface that requires a separate line for each slave and typically clocks up to 8 MHz. (Courtesy: Total Phase.)

An example of the current state-of-the-art is the Aardvark™ I2C/SPI Host Adapter from Total Phase that supports an I2C Master and Slave at 800 kHz, or an SPI Master at 8 MHz and an SPI Slave at 4 MHz (Figure 3). Total Phase also offers the Cheetah™ SPI Host Adapter, a master-only device capable of programming speeds up to 50 MHz. Yet more speed is on the horizon.

Figure 3: Low-cost USB-connected Aardvark I2C/SPI and Cheetah SPI host adapter T&M equipment from Total Phase. (Courtesy: Total Phase.)

Evolutionary Improvements to SPI

Since its introduction over 30 years ago, SPI has been used at progressively higher clock rates and lower voltages. Today's state-of-the-art memory devices operate at speeds up to 133 MHz and operate at voltages as low as 0.9 Volt. In addition to the evolution of voltage and clock rate for standard SPI, Intel is driving the Enhanced Serial Peripheral Interface bus or eSPI as a replacement for the company's Low Pin Count (LPC) bus. eSPI saves pins compared to LPC (always a plus on an IC) and will operate at up to 66 MHz but at a reduced voltage of 1.8VDC. eSPI also differs from SPI by the definition of an Alert# signal that is used by an eSPI slave to request service from the master.

Like SPI, eSPI can be used as a peripheral bus to communicate with sensors in the growing Internet of Things (IoT) movement. SPI and its variants offer a much faster way to communicate with non-volatile memory such as EEPROM or the more common Flash memory. Popular T&M equipment on the market does not currently support eSPI, so engineers will need to upgrade their hardware when eSPI gains traction. "Upgrade" means set aside the old and buy new.

According to Wikipedia, I2C, which has been popular and stable since its introduction in 1982, has also seen an evolution since its inception :

- In 1982, the original 100 kHz I2C system was created as a simple internal bus system for building control electronics with various Philips chips.

- In 1992, Version 1 added 400 kHz Fast-mode (Fm) and a 10-bit addressing mode to increase capacity to 1008 nodes. This was the first standardized version.

- In 1998, Version 2 added 3.4 MHz High-speed mode (Hs) with power-saving requirements for electric voltage and current.

- In 2007, Version 3 added 1 MHz Fast-mode plus (Fm+), and a device ID mechanism.

- In 2012, Version 4 added 5 MHz Ultra Fast-mode (UFm) for new USDA and USCL lines, using push-pull logic without pull-up resistors and added assigned manufacturer ID table.

Although now a part of the I2C specification maintained by NXP (which bought the original Philips I2C product line), not all I2C host adapters and T&M equipment support FM+, UFm, or Hs and may be limited to slower (called "standard") I2C speeds.

One other important point about I2C speeds: the simple two-wire interface and Data/Ack protocol lends itself well to faster speeds. It's not uncommon for proprietary I2C versions to have faster speeds, limited only by the signal integrity of the wires and connections. T&M equipment that can scale up and accommodate these non-standard faster speeds is highly desirable and becomes obsolete much more slowly, if ever.

Flash Programming

Besides watching the SPI or I2C bus for debug purposes, SPI has become the de facto interface for programming non-volatile memories. SPI-enabled Flash is available from most manufacturers such as Micron and Spansion, and is also used at the Flash module level or on-chip microcontroller interface for easy programming. SPI T&M equipment is the most popular and economical way to program these memories, even in a production (manufacturing) application. But imagine using single-user bench tester equipment in a manufacturing operation.

The downside of using most low-cost SPI programmers for production programming is their need for a host PC and a slow USB 1.1 connection. The USB speed can be alleviated with an upgraded USB 2.0 (480 Mbps) connection, but every SPI T&M box requires an instance of programming software running on its own PC. Although multiple instances can run on a single PC, it's still a lot of PCs to control locally (and manually) on the manufacturing floor. Even if networked, each PC must be running a virtual machine remotely controlled by the manufacturing engineer via the LAN. Additionally, long distance runs of USB may suffer signal degradation issues, often complicating the production environment by requiring control PCs in close proximity to programming stations.

There's another downside of using a "simple" SPI programmer in this scenario: each T&M box requires its own USB drivers. Brand "A" will have its own software running on the PC, which will certainly be different from Brand "B." Maintaining these programmers—and the occasional firmware updates needed to program new memory types and densities—can lead to a challenging configuration management problem. Plus, there's the PC disk images needed to assure compatibility. These are not insurmountable problems, since this is a common set-up in many manufacturing situations, but it's a system that is ripe for a better way.

Is Your T&M Equipment on the Level?

As previously mentioned, SPI and I2C are mature technologies that have found their way into countless Master and Slave devices. Their simplicity has made them popular, but the lack of rigorous standards and interoperability certification has left device manufacturers plenty of leeway to create their own protocol variations and voltage levels. This latter issue—voltage—is partially a function of design flexibility and partly due to the digital evolution from 5VDC TTL down to 3.3VDC and lower. For example, Intel's proposed eSPI goes down to 1.8VDC; some devices accept voltages as low as 0.9VDC.

The trouble is that not all SPI or I2C T&M equipment handles masters or slaves of all voltages. If the target (slave) device operates at 3.3VDC I/O, but the T&M host adapter works at 5VDC, it's up to the test engineer to create a level-shifter board that converts signals to the appropriate levels. While straightforward, level shifting adds complexity to what should be a simple debug and test environment. Level shifting circuitry adds additional timing delays, can degrade signal integrity and can possibly introduce data or I/O errors due to termination impedance, phase changes and other injected unknowns. It would be ideal if the T&M equipment supported multiple voltage levels, or even programmable ones that match tomorrow's slave devices.

Future-Proof Features

As we've seen so far, despite the simple nature of both SPI and I2C and their decades-long history, both buses continue to evolve (speed, voltage, protocols, proprietary features, etc.). T&M equipment needs to remain flexible enough to "fit" all sizes lest engineers be forced to use multiple host adapters depending upon the device-under-test (DUT). And it's likely both I2C and SPI will continue to evolve as component vendors consider new protocols and bus commands.

For example, queued serial peripheral interface (QSPI) is an intelligent SPI controller (Master) that allows some data transfers without CPU intervention—it's sort of like RDMA from a local queue. According to Wikipedia , "peripherals appear to the CPU as memory-mapped parallel devices", and QSPI is particularly useful in real-world control of A/D converters. One can imagine this feature being particularly useful in IoT sensors or low-speed factory automation I/O.

And then there are devices that require interrupts, a feature not covered at all by SPI's Master-controlled Slave-Select (SS) line. In regular SPI, the Master exerts all the control and only communicates with a Slave when it so desires. But SPI offers an elegant, protocol-agnostic mechanism for streaming data from sensors, audio devices, A/D converters and myriad other proprietary I/O.

It would be particularly handy to allow the Slave to interrupt the Master or otherwise offer some feedback (flow control)

Figure 4: Promira Serial Platform USB- and Ethernet-based T&M tool for SPI and I2C. (Courtesy: Total Phase.)

mechanism via interrupt. T&M equipment designed for "standard" SPI would have no way of accommodating any kind of custom SPI protocol or extra interrupt line. Yet the common headphone jack on many smartphones is one example of how an interrupt line is used to stop the audio CODEC when the headphones are removed from the jack. What kind of T&M SPI host adapter can test and debug this behavior?

FPGA-based Flexibility

So far we've seen that while most I2C and SPI host adapter T&M tools can get the basics done—debugging and programming sensors, MCUs and non-volatile memories—there is room for improvement. In most scenarios, whether a single-user bench set-up, a LAN-based manufacturing line, or faster or non-standard versions of the serial buses, there's room for additional features in many T&M tools.

Total Phase, already a leading supplier of I2C, SPI, USB and CAN development tools, is changing the T&M paradigm by creating an FPGA-based, extremely programmable I2C/SPI solution. Called the "Promira (Latin for "forward" and "miraculous") Serial Platform," the device addresses all the challenges and current device deficiencies described above.

To improve speeds between the host adapter and the PC host, both USB 2.0 and 10/100/1000 Ethernet are present. As mentioned above, USB's 480 Mbps (and double-sized 512 B data packet) reduce timing delays and substantially improve packet transfer between the T&M tool and the PC, allowing faster debugging and data downloads to/from the Promira platform's greater-than-1 Gbit onboard storage. But it's the Ethernet LAN connection that really gets noticed. Beyond the faster connection speed—which allows faster non-volatile memory programming—there's no need for a local PC host for configuration or bottlenecking the throughput.

This makes the device ideal in performance, speed and convenience for factory floor memory programming applications. An optional DIN rail mounting kit allows installation of multiple programmers into industrial racks

Figure 5: User-programmable level shifting is done via drop-down menu. (Courtesy: Total Phase.)

instead of literally using Velcro to affix programmers to shop fixtures. SPI programming speed has been increased to 12.5 MHz (and there's headroom above that for future upgrades to support >80 MHz SPI, dual and quad I/O, and eSPI).

LAN connectivity also eliminates the need for local PC images and USB drivers, and facilitates firmware updates. The Promira platform's software is revised over time and creates the opportunity for remote automation using a standard Total Phase API that is backward compatible with previous generation non-programmable T&M tools. In short: the onboard FPGA with LAN access can realize countless current and future functions and be easily re-architected.

The software programmability of the tool provides additional flexibility for voltage level-shifting, too. Voltage conversion boards aren't needed because a user drop-down menu in a Promira platform GUI selects from 0.9VD to 3.3VDC (Figure 5). Besides the convenience, engineers get a more accurate view of SPI or I2C data transactions by eliminating in-line level shifting, facilitated instead by logic in the T&M tool.

Even Further into the Future

Savvy readers will note that the primary hardware difference between Promira platform and previous generation I2C and SPI T&M tools—including those by Total Phase—is the fully programmable FPGA-based architecture. FPGAs can realize numerous logic, processors, peripherals and buffer/memory combinations. The device also includes additional digital and analog front-end capabilities, making future changes as simple as an FPGA recompile via a LAN-based firmware upgrade.

This means that any future versions of SPI or I2C, such as eSPI, QSPI, MDIO, SDIO, new protocols or even the addition of non-standard interrupt lines as previously described, can be implemented by a software update. In effect, Promira might be the last host adapter embedded systems engineers or Flash programmers ever need. Updates and new features are planned as technology evolves.

A particularly compelling feature of an "infinitely" programmable SPI/I2C controller is the idea of "concurrency": allowing the device to be both master and slave, or contain two, independent masters...simultaneously. Under software control via Ethernet or USB, engineers can simulate various combinations of master/master or master/slave functionality with a single T&M device. This means software can be debugged before either the Master or Slave hardware is ever available.

Total Phase's existing Control Center Software and Flash Center Software are now Promira-enabled. These API-based and dashboard-displayed GUIs allow rapid programming and master/slave emulation, and will offer sniffing plus intelligent filtering of data packets matched against user-selected events in the near future. Additionally, Total Phase provides Rosetta™ Language Bindings along with the API, enabling engineers to create their own custom applications in C, C++, C#, Python, VB6 and VB.Net. With some programming skills, "if this, then that" scenarios can be created to facilitate more sophisticated debugging and emulation. As well, this kind of programming can enable Promira platform as an IoT node that intelligently sends serial I2C or SPI data to the Internet based upon predefined events. In this case, the device functions more like an intelligent shadow RAM with remote diagnostic capability.

We can easily envision more layers of sophistication as Total Phase rolls out future application upgrades, making the Promira platform a real prize in your serial (tool) box.

This article was sponsored by Total Phase.

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 World 2015—No Longer A One-Theme Show

The acronym anyone following the embedded industry expected to see in Nuremberg's Messe halls while exploring Embedded World 2015, was IoT (Internet of Things).

By Caroline Hayes, Senior Editor

In fact, the IoT was one of the themes of the conference program that accompanied the industry show ('We are the Internet of Things'). Certainly, the booths in the halls did not disappoint, with IoT mentioned somewhere on most stands, but there was another, emerging movement that began to take root at this year's show, as the call for heterogeneous computing gathered momentum.

If you would have had a bratwurst sausage only at stands that had no mention of the IoT, you would have been hungry all week. As Scott Aylor, Corporate Vice President and General Manager, AMD, put it in his keynote presentation at the 'We are the Internet of Things' conference: "Who would have thought three little letters would generate so much interest?".

ARM of course had a large presence at this year's Embedded World. With over 50 billion processors sold worldwide, ARM drew lots of mentions in other companies' literature, and its mbed IoT Device Platform won the 'embedded AWARD 2015' in the software category (Figure 1).

At the show, the company also launched the mbed IoT Starter Kit – Ethernet Edition. Designed to accelerate the prototyping of connected devices, the Starter Kit is to ready these devices for 'smart' products that use cloud services. Enabling Internet-connected devices to communicate directly with IBM's Bluemix cloud platform, the kit contains ARM mbed OS, ARM Cortex-M4-based development board from Freescale and a sensor I/O application shield. Kits to be released later will run the ARM mbed OS and use ARM mbed Device Server software.

Automotive and Security

Naturally, for a show based in Germany, automotive design figured, and Freescale used the opportunity to showcase its i.MX 6SoloX applications processor, which is a heterogeneous multicore device for the connected vehicle, home and the ubiquitous IoT (Figure 2).

Figure 1: ARM took home the Software 'embedded AWARD 2015' in the software category for its ARM mbed IoT Device Platform.

The heterogeneous multicore i.MX 6SoloX applications processor from Freescale addresses secure connectivity in vehicles, as well as in the home

The SoC integrates ARM Cortex-M4 and Cortex-A9 cores with cryptographic cipher engines and a configurable resource domain controller that allows peripherals to be locked or shared by the CPU cores. A secure messaging semaphore unit enables cooperative, multi-OS software to access shared peripherals in safety. Another feature is the secure boot and protected data storage. These hardware capabilities enable users to architect custom security

Figure 2: The heterogeneous multicore i.MX 6SoloX applications processor from Freescale addresses secure connectivity in vehicles, as well as in the home

solutions. The company's security theme anticipates the challenges of having 25 billion connected things by 2020 (Gartner figures.)

Freescale announced several initiatives in Nuremberg, including teaming with the Embedded Microprocessor Benchmarking Consortium (EEMBC) to identify critical embedded security gaps, and, with other consortium members, to establish guidelines to secure IoT transactions and endpoints. It also announced the Freescale Security Labs to enhance IoT security technologies from the cloud to the end-node and an education program for start-ups on IoT security best practices and to provide access to its partner ecosystem. [Editor's Note: At the time of going to press, NXP has announced its intention to buy Freescale Semiconductor, with the acquisition expected to be finalized in the second half of 2015.]

Video Drivers

The human machine interface (HMI) figured heavily at the show, with companies vying for attention. At AMD, a plethora of displays showed the virtues of a heterogeneous architecture with a graphics processor unit coupled with the main processor. The one that caught my eye was a 3D synthetic vision cockpit display demonstration. The CoreAVI H.264/MPEG2 video decode driver suite (Figure 3) with integrated OpenGL SC graphics driver uses the AMD Embedded Radeon E8860 graphics processor and Universal Video Decode engine on Wind River's VxWorks 7 RTOS and Curtiss-Wright's VPX3-133 single board computer, based on the Freescale T2080 and VPS3-716 graphics module.

The CoreAVI thread safe video decode driver architecture allows unmanned aerial vehicles (UAVs), ground control stations, 360° geographical information systems, moving maps and other real-time, safety-critical video applications to decode multiple, independent, simultaneous compressed HD video streams, which are then rendered through the video decoder API application programming interface (API).

Figure 3: AMD showcased CoreAVI DO-178C certifiable H.264/ MPEG2 video decode driver

Another video example was to be found at the FTDI booth, where the company introduced four Embedded Video Engine (EVE) ICs, making up the FT81xQ series. They support higher screen resolutions at 800 x 600 pixels, up from 512 x 512 in the earlier FT800Q and FT801Q series. The intention is to enable seven-inch (and above) screen sizes for point of sale and information kiosks, with smooth video playback, due to the enhanced algorithms. The quartet is made up of the FT810Q, with 18-bit RGB interfacing and resistive touch functionality, the FT811Q, with 18-bit RGB interfacing and capacitive touchscreen functionality (with five-point touch detection), the FT812Q (resistive) and FT813Q (capacitive touchscreens) with 24-bit RGB.

Working in partnerships

Many companies announced collaboration, with Green Hills Software having a foot in both ARM and MIPS core camps, with an announcement with Imagination Technologies and ARM. The first announcement was that Green Hills' tools and compilers now support more of Imagination's MIPS CPU IP. This includes support for the microMIPS code compression instruction set architecture, available now and with support for MIPS 64-bit Warrior M-class and I-class CPUs, with hardware virtualization, available in the second half of this year.

Figure 4: The congatec IoT kit was announced at Embedded World 2015

In the ARM camp, Green Hills also jointly announced with ARM, collaboration on a compiler optimized for the ARM Cortex-R5 processor for automotive design. The latest version of the compiler scored 1.01EEMBC Automarks/ MHz using the Cortex-R5 Traveo automotive MCUs from Spansion—a 30% increase on previous performance scores.

The beauty of Embedded World is that all voices are heard. Another core provider, Cortus, based in the south of France, announced three new partners porting to its architecture, highlighted its two latest cores (the 32-bit APS23 and APS25, based on a Cortus v2 instruction set, and the collaboration with Blunk Microsystems to offer Target Tools IDE for Cortus software development. The Eclipse-alternative embedded software CrossStep IDE and TargetOS RTOS are available now for the APS processor cores. It is, says the company, fast, small, and pre-emptive RTOS ported to the APS architecture.

In addition to Blunk, the company also announced a dual IP v4/v6 stack from Oryx Embedded and secure point-to-point connection software from Nabto.

CycloneTCP, from Oryx Embedded, is a dual IPv4/IPv6 stack and offers seamless interoperability with existing TCP/IP systems, ready for the next generation of Internet connected designs, using the IPv6 protocol.

Nabto's device software offers Cortus licensees "Skype-like" secure point-to-point connectivity from devices to smartphones, PCs and big data systems. The device software has a small firmware footprint for minimalist processor cores.

Finally, if longevity of supply is an issue, congatec announced that it is extending availability of Qseven, COM Express, XTX and ETX COM (computer on modules) based on AMD's Embedded G-Series SoCs. Now, customers have a secure supply, promises the company, for 10 years, three years more than previously available—or an extension in Q1 2024, bringing with it the security of long-term availability and reduced non-recurring engineering costs.

True to the theme of IoT, it also announced a Qseven IoT kit to simplify IoT application development (Figure 4). It contains a Qseven COM based on the Intel® Atom™ processor, a compact IoT carrier board and a seven-inch LVDS single touch display with LED backlight and accessories, including AC power supply and 802.11 WLAN antenna.

The congatec IoT kit was announced at Embedded World 2015

Caroline Hayes has been a journalist, covering the electronics sector for over 20 years. She has worked on many titles, most recently the pan-European magazine, EPN. Now a freelance journalist, she contributes news, features, interviews and profiles for electronics journals in Europe and the US.

Top 3 Uses for PCI Express Switches

Listen up: here's where you're going to use a PCIe switch.

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

I'll bet on every digital designer's desk you'll find a stapler, coffee mug (usually with a weeks' worth of crust to add flavor), paper clips and a PCI Express switch IC. And if it's not an actual IC, it's a design that either has or needs one.

My point: PCIe Gen2 (Gen3 is still ramping up) is about as common on digital PCBs as is a terminal block or pin header. Nearly every CPU, MCU, fast peripheral, bus or nonvolatile (NV) memory sports PCIe. To connect all of these "dots" together requires a PCI Express switch or one of the creative variations thereof.

Here are three really common uses for the lowly PCIe switch.

Reason #1: Port Expansion and Fanout

This has nothing to do with dredging the harbor to make room for luxury condos. Fanout is needing more PCIe connections in your design than are available on your CPU, SoC, MCU, Southbridge or multi-I/O controller.

CPU vendor Intel is pretty good about having two or three PCIe ports on its x86 Intel® Core™ i7 and Intel® Atom™ processors.. Not so from Freescale on i.MX ARM-based SoCs: there's only one on some devices. The answer is to use a PCIe switch to increase fanout as shown in Figure 1. Here, there's one upstream PCIe Gen2 port and five downstream ports using Pericom Semiconductor's PI7C9X2G612GP packet switch. Myriad permutations offer all kinds of ports and features.

Practical examples of this kind of fanout improvement include host bus adapters, where the CPU/SoC uses the switch to talk to local resources or bus/backplane-based resources. As well, a PCIe switch can be used between a CPU and two memory controllers (RAID, SSD, etc.) to switch between different memory arrays or banks and increase density or improve speed through memory striping (by ping-ponging memory banks). Figure 2 shows a multi-function printer that switches between the SoC, scan and print engines, and four PCIe endpoints.

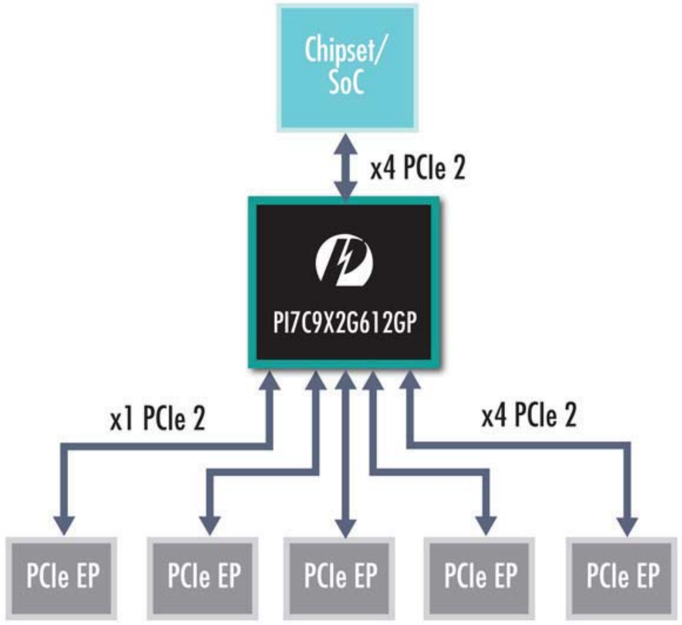

Figure 1: A six-port (12 lane) PCI Express switch increases fanout from one x4 to five additional multiplexed/switched devices. (Courtesy: Pericom.)

Reason #2: Lane Swapping

Not going to make the obvious joke with this one. Suffice it to say, this is sort of like Endian conversion or the old-style connector gender changers: converting from one type of PCIe configuration to another type.

This one's easy to understand. It's like having a drawer full of USB cables but none of them have the right end connector to fit your device. PCI Express host adapters such as CPUs come configured only one way: say, one x8. But suppose you need two sets of x4? Eight sets of x1 (one lane per device)? A PCIe switch can do this "lane swapping" for you.

Figure 3 shows a six-port, eight-lane switch that lane-swaps from a single x4 to four or five x1 ports, where port #5 is available under specific conditions. Although shown on a PC add-in card, the architecture could easily be embedded and the decoders could be peripherals or

Part Number	Ports	Lanes	Clock Buffer	Pd Typ. (W)	Temp (C°)	Pkg. Size (mm)	Package Type	Pins	Configuration Upstream	Configuration Downstream
PI7C9X2G608GP	6	8	Yes	0.9	-40 ° +85 °	15 x 15	LBGA	196	x1 or x2 or x4	4 or 5 x1
PI7C9X2G608EL	6	8	Yes	0.9	-40 ° +85 °	10 x 10	aQFN	136	x1 or x2 or x4	4 or 5 x1
PI7C9X2G312GP	3	12	Yes	1.1	-40 ° +85 °	15 x 15	LBGA	196	x4	2 x4
PI7C9X2G308GP	3	8	Yes	1.1	-40 ° +85 °	15 x 15	LBGA	196	x4	2 x2
PI7C9X2G612GP	6	12	Yes	1.1	-40 ° +85 °	15 x 15	LBGA	196	x4	1 x 4 + 4 x1
PI7C9X2G606PR	6	6	No	0.5	-40 ° +85 °	15 x 15	LBGA	196	x1 or x2	4 or 5 x1

Table 1. Lane-swapping options made possible by the use of PCIe switches from Pericom.

coprocessors. Table 1 shows the huge range of lane-swapping options on several PCIe switches from Pericom.

Reason #3: Bridging, Conversion and Advanced Functions

Beyond the available port permutations and lane-swapping capabilities, switches offer other benefits to designers. The switch shown in Figure 1, for example, includes device power management and a built-in clock buffer to save external components. It supports isochronous traffic to some of the ports for guaranteed latency to time-critical data devices like network data planes. These kinds of advanced functions make up the final reason for choosing a PCIe switch.

Most PCIe switches merely move data from one port to another, but some include packet routing functions. As the name implies, a routing criteria supplied by the local host tells the switch that only certain types of packets can be routed between ports—making the switch more than just a fanout improvement; rather, the switch functions as an intelligent local controller. Another variation on this theme is a bridge/switch, called a “SWIDGE” by supplier Pericom. One example of a SWIDGE, the PI7C9X442SL, bridges between USB 2.0 and PCIe 2.0. This device turns an I/O- or resource-limited system into a multi-port multiplexer that moves data between dissimilar bus types.

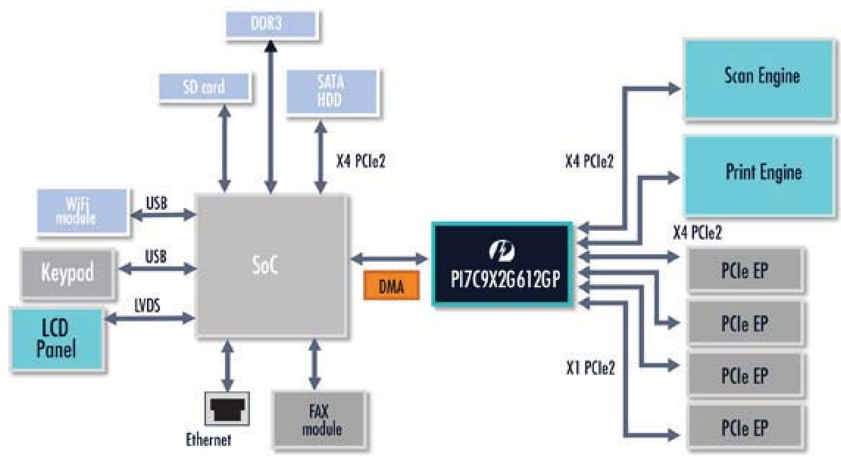

Figure 2: PCIe Gen2 switch expands one port on the SoC to six peripherals. (Courtesy: Pericom.)

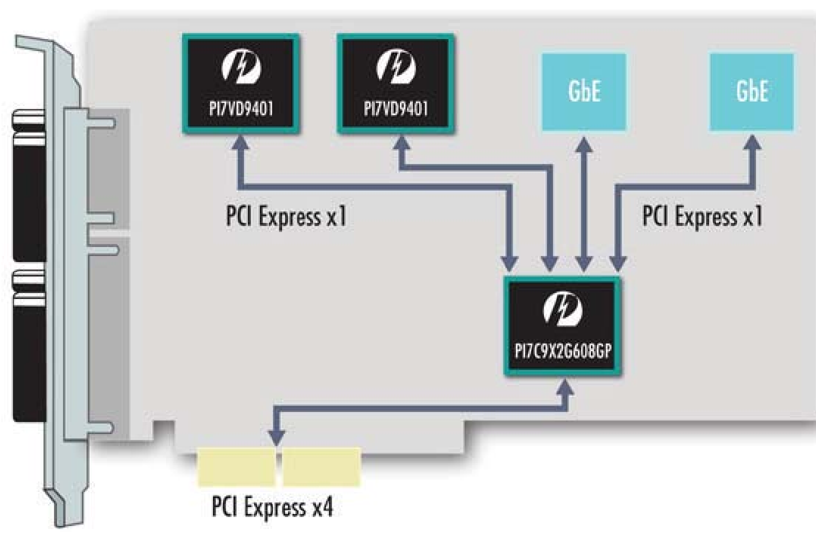

Figure 3: An example of lane swapping, from one x4 to four x1’s. (Courtesy: Pericom.)

The device shown in Figure 4, for example, functions as a fanout improver with one upstream and two downstream PCIe 2.0 ports. Additionally, there are four USB 2.0 host controller ports, a feature that essentially turns a single PCIe x1 lane CPU into a one that has two PCIe and four USB 2.0 ports. No USB 2.0 drivers or software are needed when the switch/bridge implements USB 2.0.

Intelligent routing functions are also added with non-blocking bridging and “store and forward” packet switching—all with minimal CPU intervention. Collectively, the switch/bridge becomes a flexible I/O controller with DMA boasting a full 16 Gbps when all ports are running flat out.

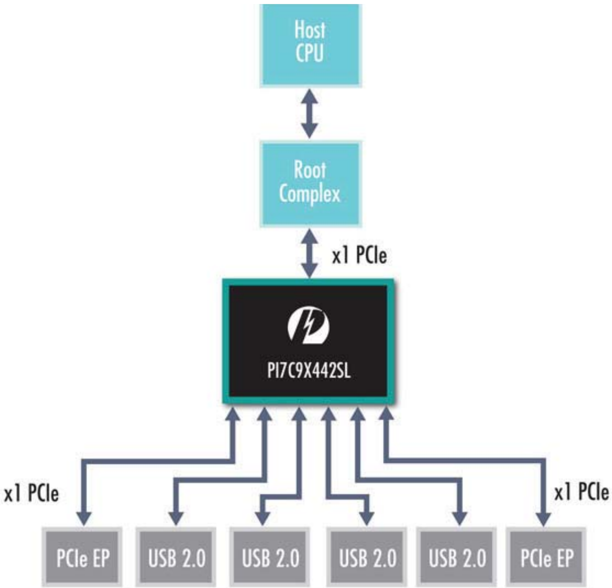

Figure 4: A bridge/switch combination called a "SWIDGE" that increases fanout, adds USB 2.0 I/O and has local intelligent packet forwarding. (Courtesy: Pericom).

Besides USB 2.0, other I/O conversions are possible in a PCIe switch such as between PCIe and RS-232/RS-422/and RS-485. Pericom's PI7C9X7952 converts between PCIe 1.1 and serial ports via on-board dual UART blocks. Truly, the permutations are limitless.

Switch Your Thinking

We've presented a short list of three reasons to consider PCIe switches, and each of them should convince you that if there's not a PCIe switch on your desk right now—or in your design—there should be.

This article was sponsored by Pericom Semiconductor.

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.

Extension MEDIA
嵌入式英特尔®
解决方案
Embedded Intel® Solutions
Summer 2014
Mobile and IoT Drive Embedded Small Form Factor Evolution
Development Options Evolve with Advances in Intel® Processor Performance-Per-Watt
Advanced Image Stabilization Techniques for Tablet Camera Performance
Meeting the Power Neutrality Challenge in Mobile Handsets
Fall 2014
Next Generation of Xeon Processors Rolls Out at IDF
Once Earmarked for FPGA Always Earmarked for FPGA?
From SuperServers to Smart Beer Kegs
Winter 2014
SoC FPGA Relies on Many Cores and 14 nm Tri-Gate Process
The Internet of Things Defines the Future RTOS
Automotive Ethernet: No Simple Answers
Smart Antennas Drive Evolution in M2M Applications
www.embeddedintel.com
Gold Sponsors
Winter 2015
Securing Embedded Devices in the IoT Era
Automotive Applications Pursue Higher Performance
Test Equipment Tackles PCIe Receiver Challenges
Spring 2014
Intel SoC Enables Next-Gen Networking Devices
The Power Challenge: How Low Can You Go?
Meet the Challenge of Securing the Smart Grid
Modules Take Portable Medical Equipment Design to the Next Level

Smart Fabric, Not Wrist Bling, To Lead Wearable Market Growth?

Coming soon: Smart Fabric as Coach?

By Jonah McLeod, Silicon Valley Blogger

There's incredible hype surrounding the wearable market and most of it is aimed at the wrist where activity monitors continue to move from crude step counter to something much more tripped out. The ultimate example is certainly the $350-ish Apple Watch slated for release in the spring if rumors are to be believed. In his article, "Wearable market set to explode," author Dan Cook declares, "the evolution of products designed to measure heart rate, blood pressure, weight and body fat, and to track workouts, will move toward smartwatches." But, there are other opportunities being overshadowed by the designer bling and it centers around smart fabric.

Consider the measurement opportunities for smart fabric. In the insoles of shoes, smart fabric can sense running style, pronation, gate, contact order, and fit. Covering the head in a skull cap, smart fabric can not only tell the force a player is hit but where on the head the force occurred. Inside a boxer's gloves, smart fabric can tell how effectively the punch connected to an opponent and, inside a batter's glove, the behavior of his grip throughout a swing. Sewn onto the arm of clothing, smart fabric can provide the same kinds of controls as buttons on a smart phone: answer call, volume control, advance a track or go back a track.

Of these applications, running is the one with the greatest market potential for smart fabrics. According to the report issued June 15, 2014 entitled "2014 State of the Sport – Part II: Running Industry Report," published by Running USA, "Since 2004, total running/jogging participation (run/jog 6+ days/yr) has increased 70% to a record of nearly 42,000,000, according to the NSGA. Females in the 25-34 age group category lead participation totals with more than 5.6 million in 2013, and since 2012, according to NSGA, more women run than men in the USA (both genders are at record highs)." Furthermore, a typical running shoe is retired after 300 to 500 miles of use, at 5 miles a day, that's roughly once or twice a year, more frequently than that $300 wristwatch.

One company poised to capitalize on this performance-crazed market is Berkeley, Calif.-based, Bebop Sensors. BeBop's proprietary Monolithic Fabric Sensor Technology integrates all of the sensors, traces and electronics into a single piece of fabric, thus enabling increased sensitivity, resolution, range of deployment and robustness in a practical size. The sensors in the fabric accurately detect pressure, bend, location, rotation, angle, and torsion, to enable the creation of a 3D representation of these forces. This produces the meaningful feedback for athletes seeking to eek out the milliseconds of performance improvement they hunger for. The watch will tell you how fast your heart is beating, your temperature and oxygen level of our blood, great if your concerned about your health, less useful for getting ahead of the game.

Company founder Keith McMillen developed the Bebop fabric technology for his musical instruments company Keith McMillen instruments. The fabric performs the touch function in fabric as the touch sensors in a smart phone or tablet. The uniqueness in Bebop's smart fabric is its ability to provide a 3D profile of the applied force. For example, the smart fabric beneath the keys of the QuNexus keyboard produces a unique MIDI CC number for every key that is struck. In addition, the fabric sensor detects a key being held after being struck and the orientation of the force of the key in 3D space, a unique sound resulting from the force and orientation of the force.

All of this capability applies equally well if the smart fabric is worn in a running shoe. The 3D force profile can be used to compute ground reaction force (GRF), the shock wave or force that occurs when a runner's foot strikes the ground. GRF comprises a vertical, horizontal, anterior and posterior component. Researchers speculate that the part of a runner's foot—front, rear, or middle—that contacts the ground in relation to his/her body's centre of mass is key to performance. Having an insole sensor that pings the runner when he/she hits the sweet spot on each stride could provide a real competitive advantage—in effect, a coach providing continuous advice in real time during a race.

The fabric is still looking for the OEM that will add the hardware and software to turn the potential into a final product. As for cost, in a typical wearable application such as the athletic shoe, if Fitbit is an example, the hardware and software runs around $17.36. When included in the bill-of-material, BeBop's insole fabric solution will be competitive the company asserts. Expect this to be flying off the shelf.

DVCon Highlights: Software, Complexity, and Moore's Law

By Gabe Moretti, Senior Editor

The first DVCon United States was a success. It was the 27th Conference of the series and the first one with this name to separate it from DVCon Europe and DVCon India. The last two saw their first event last year and following their success will be held this year as well.

Overall attendance, including exhibit-only and technical conference attendees, was 932.

If we count, as DAC does, exhibitors personnel then the total number of attendees is 1213. The conference attracted 36 exhibitors, including 10 exhibiting for the first time and 6 of them headquartered outside of the US. The technical presentations were very well attended, almost always with standing room only, thus averaging around 175 attendees per session. One cannot fit more in the conference rooms that the DoubleTree has. The other thing I observed was that there was almost no attendees traffic during the presentations. People took a seat and stayed for the entire presentation. Almost no one came in, listened for a few minutes and then left. In my experience this is not typical and points out that the goal of DVCon, to present topics of contemporary importance, was met.

Process Technology and Software Growth

The keynote address this year was delivered by Aart de Geus, chairman and co-CEO of Synopsys. His speeches are always both unique and quite interesting. This year he chose as a topic "Smart Design from Silicon to Software". As one could have expected Aart's major points had to do with process technology, something he is extremely knowledgeable about. He thinks that Moore's law as an instrument to predict semiconductor process advances has about ten years of usable life. After that the industry will have to find another tool, assuming one will be required, I would add. Since, as Aart correctly points out, we are still using a 193 nm crayon to implement 10 nm features, clearly progress is significantly impaired. Personally I do not understand the reason for continuing to use ultraviolet light in lithography, aside for the huge costs of moving to x-ray lithography. The industry has resisted the move for so long that I think even x-ray has a short life span which at this point would not justify the investment. So, before the ten years are up, we might see some very unusual and creative approaches to building features on some new material. After all whatever we will use will have to understand atoms and their structure.

For now, says Aart, most system companies are "camping" at 28 nm while evaluating "the big leap" to more advanced lithography process. I think it will be along time, if ever, when 10 nm processes will be popular. Obviously the 28 nm process supports the area and power requirements of the vast majority of advanced consumers products. Aart did not say it but it is a fact that there are still a very large number of wafers produced using a 90 nm process. Dr. de Geus pointed out that the major factor in determining investments in product development is now economics, not available EDA technology. Of course one can observe that economics is only a second order decision making tool, since economics is determined in part by complexity. But Aart stopped at economics, a point he has made in previous presentations in the last twelve months. His point is well taken since ROI is greatly dependent on hitting the market window.

A very interesting point made during the presentation is that the length of development schedules has not changed in the last ten years, content has. Development of proprietary hardware has gotten shorter, thanks to improved EDA tools, but IP integration and software integration and co-verification has used up all the time savings in the schedule.

What Dr. De Geus slides show is that software is and will grow at about ten times the rate of hardware. Thus investment in software tools by EDA companies makes sense now. Approximately ten years ago, during a DATE conference in Paris I had asked Aart about the opportunity of EDA companies, Synopsys in particular, to invest in software tools. At that time Aart was emphatic that EDA companies did not belong in the software space. Compilers are either cheap or free, he told me, and debuggers do not offer the correct economic value to be of interest. Well without much fanfare about the topic of "investment in software" Synopsys is now in the software business in a big way. Virtual prototyping and software co-verification are market segments Synopsys is very active in, and making a nice profit I may add. So, it is either a matter of definition or new market availability, but EDA companies are in the software business.

When Aart talks I always get reasons to think. Here are my conclusions. On the manufacturing side, we are tinkering with what we have had for years, afraid to make the leap to a more suitable technology. From the software side, we are just as conservative.

That software would grow at a much faster pace than hardware is not news to me. In all the years that I worked as a software developer or managers of software development, I always found that software grows to utilize all the available hardware environment and is the major reason for hardware development, whether is memory size and management or speed of execution. My conclusion is that nothing is new: the software industry has never put efficiency as the top goal, it is always how easier can we make the life of a programmer. Higher level languages are more powerful because programmers can implement functions with minimal efforts, not because the underlying hardware is used optimally. And the result is that when it comes to software quality and security the users are playing too large a part as the verification team.

Art or Science

The Wednesday proceedings were opened early in the morning by a panel with the provocative title of Art or Science. The panelists were Janick Bergeron from Synopsys, Harry Foster from Mentor, JL Gray from Cadence, Ken Knowlson from Intel, and Bernard Murphy from Atrenta. The purpose of the panel was to figure out whether a developer is better served by using his or her own creativity in developing either hardware or software, or follow a defined and "proven" methodology without deviation.

After some introductory remarks which seem to show a mild support for the Science approach, I pointed out that the title of the panel was wrong. It should have been titled Art and Science, since both must play a part in any good development process. That changed the nature of the panel. To begin with there had to be a definition of what art and science meant. Here is my definition. Art is a problem specific solution achieved through creativity. Science is the use of a repeatable recipe encompassing both tools and methods that insures validated quality of results.

Harry Foster pointed out that is difficult to teach creativity. This is true, but it is not impossible I maintain, especially if we changed our approach to education. We must move from teaching the ability to repeat memorized answers that are easy to grade on a test tests, and switched to problem solving, a system better for the student but more difficult to grade. Our present educational system is focused on teachers, not students.

The panel spent a significant amount of time discussing the issue of hardware/software co-verification. We really do not have a complete scientific approach, but we are also limited by the schedule in using creative solutions that themselves require verification.

I really liked what Ken Knowlson said at one point. There is a significant difference between a complicated and a complex problem. A complicated problem is understood but it is difficult to solve while a complex problem is something we do not understand a priori. This insight may be difficult to understand without an example, so here is mine. Relativity is complicated, black matter is complex.

Conclusion

Discussing all of the technical sessions would be too long and would interest only portions of the readership, so I am leaving such matters to those who have access to the conference proceedings. But I think that both the keynote speech and the panel provided enough understanding as well as thought material to amply justify attending the conference. Too often I have heard that DVCon is a verification conference: it is not just for verification as both the keynote and the panel prove. It is for all those who care about development and verification, in short for those who know that a well developed product is easier to verify, manufacture and maintain than otherwise. So whether in India, Europe or in the US, see you at the next DVCon.

Gabe Moretti has been in EDA for 45 years. First as an individual contributor with TRW Systems and Compucorp. Then as a manager with Intel and Signetics. He has been a member of the executive management team with EIS Modeling (a company he founded), HDL Systems, and Intergraph/Veribest. From 2000 to 2005 he was technical editor for EDA at EDN. Since then Gabe has run his own consulting company, GABEonEDA. He has a B.A. in Business Administration and a Master in Computer Sciences.

PCV5104 - MIL-SPEC Panel PC Features 1.83GHz 4-Core Intel® Celeron® Processor N2930

Acnodes, a quality provider of industrial computers, announces the release of the PCV5104, an ultra rugged military grade Panel PC. This Panel PC is engineered and certified to withstand extremely harsh environments, which is ideal for military mission critical applications. With its MIL-STD-810G compliance, it can survive wild temperature ranging from -13 °F to 140 °F. It also boasts of anti shock and vibration resistance, fully IP65 rated anti corrosion panels and fanless design that increases durability and provides longer life while reducing down time.

Based on the Fintek F81866 chipset, PCV5104 is powered by the latest low power consumption high performance 1.83GHz 4-Core processor Intel® Celeron® processor N2930. This rugged Panel PC comes standard with one on board DDR3 SO-DIMM socket with a maximum of 8GB system memory and one 2.5 in. SATA Hard Disk Drive along with dual Ethernet, one COM port, two USB ports, and a HDMI port to accommodate a wide range of connectivity requirements. To facilitate access to wireless network, the PCV5104 supports on board WiFi and 3G. Moreover, the PCV5104 has proven to provide critical information on and off the battle field. With a 10.4 inch high image quality 5-wire resistive type touch screen monitor, the PCV5104 supports up to 1024 x 768 pixels resolution, offers 16.2 million colors, and sunlight readable 800 nits high brightness level with a viewing angle of 176 degrees.

In addition to panel mounting, PC5104 is compliant with VESA mounting standards. Standard VESA 75mm spaced mounting holes are available on the rear for easy mounting to any fixture or wall which enables users to run demanding applications in the field, floor or office environment.

Additional information on Acnodes' PCV5104, please visit http://www.acnodes.com/pcv5104.htm

ACNODES

Acnodes Corporation
14628 Central Ave.
Chino, CA 91710 USA
909-597-7588
909-597-1939
sales@acnodes.com
www.acnodes.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.

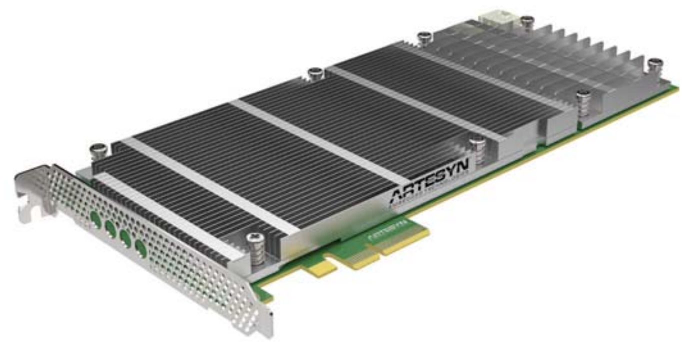

The SharpStreamer PCIE-7207 is equipped with the Intel® Media Software Development Kit featuring Intel® HD Graphics with fixed-function hardware acceleration, monitoring, and processor subsystem O/S and management tools for easy integration with server host processing environments.

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