Extension
MEDIA

Embedded Intel® Solutions

Winter 2014

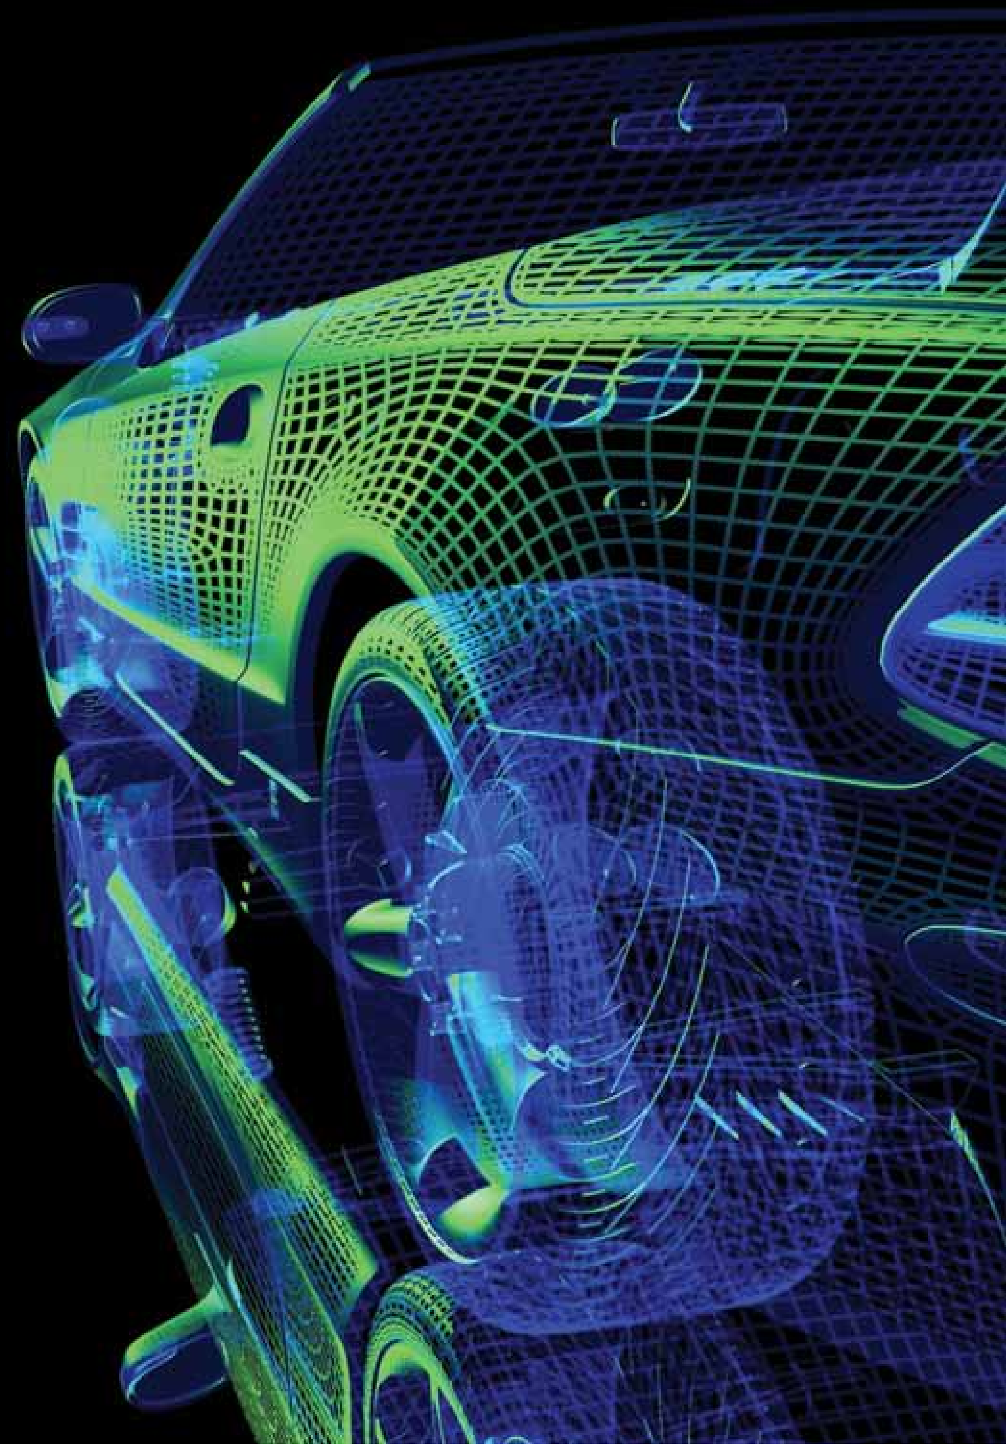

SoC FPGA Relies on Many Cores and 14 nm Tri-Gate Process

The Internet of Things Defines the Future RTOS

Automotive Ethernet: No Simple Answers

Smart Antennas Drive Evolution in M2M Applications

www.embeddedintel.com

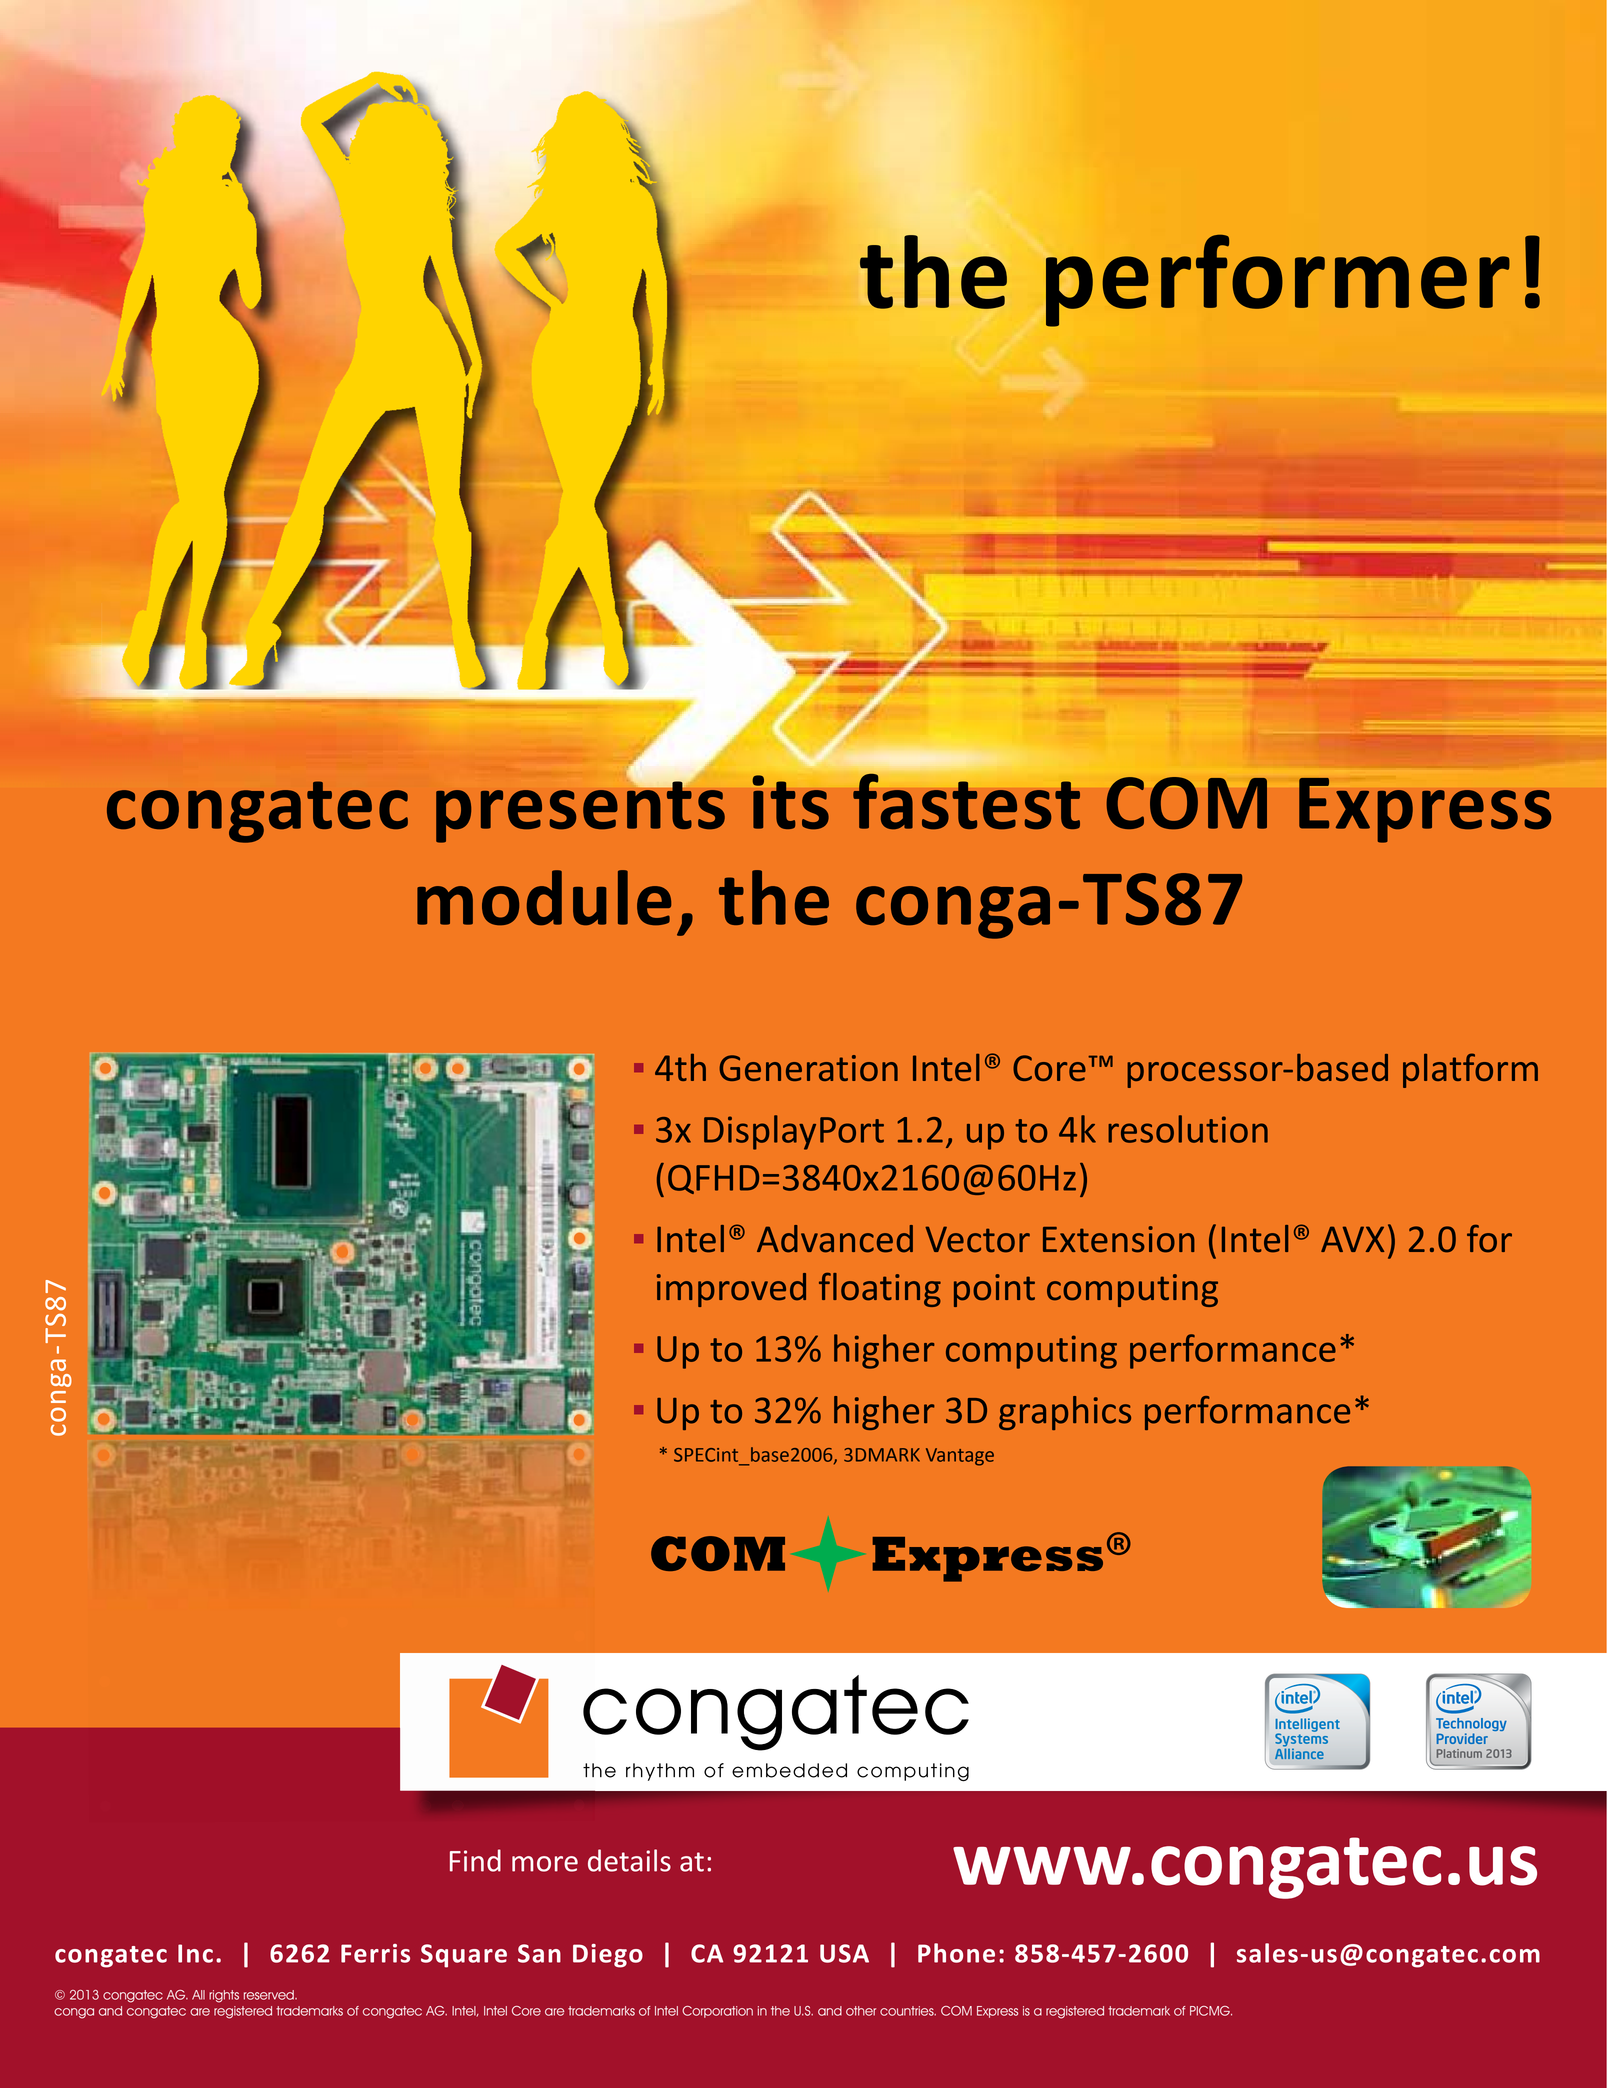

the performer!
congatec presents its fastest COM Express module, the conga-TS87
conga-TS87
4th Generation Intel® Core™ processor-based platform
3x DisplayPort 1.2, up to 4k resolution (QFHD=3840x2160@60Hz)
Intel® Advanced Vector Extension (Intel® AVX) 2.0 for improved floating point computing
Up to 13% higher computing performance*
Up to 32% higher 3D graphics performance*
* SPECint_base2006, 3DMARK Vantage
COM Express®
congatec
the rhythm of embedded computing
intel Intelligent Systems Alliance
intel Technology Provider Platinum 2013
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
© 2013 congatec AG. All rights reserved.
conga and congatec are registered trademarks of congatec AG. Intel, Intel Core are trademarks of Intel Corporation in the U.S. and other countries. COM Express is a registered trademark of PICMG.

Intel® Atom™ Processor Small, Fanless Embedded PCs

For your next design consider WinSystems' single board computers powered with a single- or dual-core Intel® Atom™ processor. Our Industry Standards-based SBCs have a wealth of onboard I/O, plus PC/104, SUMIT, and MiniPCI(e) expansion connectors.

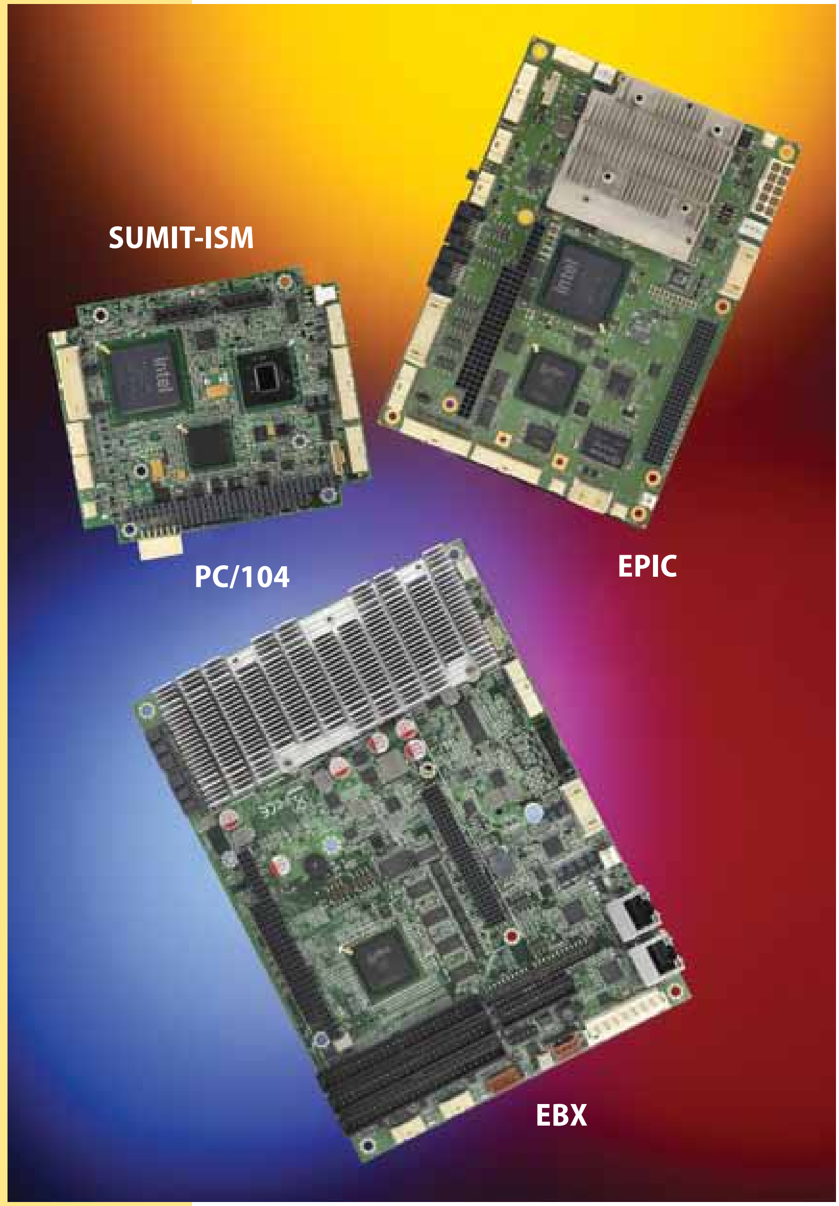

- Industry Standard Platforms
  - EPIC – 4.5 x 6.5 inches
  - EBX – 5.75 x 8.00 inches
  - PC/104 – 3.6 x 3.8 inches
  - SUMIT-ISM™ – 3.6 x 3.8 inches
- Software Support
  - Windows® CE, XPe, and WES7
  - Linux
  - x86-compatible RTOS
  - Quick Start Development Kits
- Stackable Input/Output Modules
  - A/D and D/A
  - Digital and Serial
- Accessories Include
  - Cables and Adapters
  - Power Supplies and Memory
- Communication Expansion
  - 802.11 a/b/g Wireless
  - 10/100/1000 Mbps Ethernet
  - Supports MiniPCI and MiniPCI(e)
- Long-life Product Availability
- Extended Temperature Operation

Our SBCs are the right choice for industrial, pipeline, transportation, instrumentation, medical, and MIL/COTS applications. Call us today.

Call 817-274-7553 or
Visit WinSystems.com/AtomEIS

Ask about our 30-day product evaluation

WinSystems®
An Employee-Owned Company

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

Complexity Spreading vs. Intelligent Embedded in the IoT

Which phrase is more indicative of the Internet-of-Things depends upon your viewpoint—namely, from the process or the products.

By John Blyler, Vice President & Chief Content Officer

The IoT brings many challenges to the traditional design of embedded systems, including the integration of analog/mixed-signal and RF technologies into microcontroller-based sensor-fusion devices.

Some may think that IoT is just the latest marketing ploy to rehash the old computation-communication convergence "wars" from the last decade. (See "Are Mobile Communication and Computer RF Technologies on a Collision Course?," June 2000) http://www.chipdesignmag.com/blyler/2013/10/18/flashback-to-june-2000-are-mobile-communication-and-computer-rf-technologies-on-a-collision-course/

Certainly, there are similarities between these two concepts. For example, both require the ongoing convergence of analog (mixed-signal) and digital systems. But today's IoT adds the element of sensor systems to the mix. Many engineers may view this as a trivial addition to the overall design equation, as sensors, instrumentation, and data-acquisition technologies have been understood since the '50s.

What is new is the scale of integration, which adds a new level of complexity. This is not the deep complexity of a new transistor device structure like a FinFET or the associated 3D manufacturing challenges. Instead, the IoT requires a spreading out of reasonably well-understood technologies, which is a different kind of complexity. The end result is a complexity spreading or even smearing-out process. (From a device or product perspective, this is known as smart sensors or intelligent embedded.)

This is not a new phrase. In the world of cellular code-division-multiple-access (CDMA) design, the term "complexity spreading" often refers to various algorithmic designs related to spread-spectrum systems. However, in terms of the smart sensor technology (sensors with some processor intelligence), I would argue that complexity spreading refers to the pushing out of this intelligence as far as possible—in this case, to the interface with the analog physical world experienced by the sensor. Whenever "intelligence" is pushed outward, so too is the technology needed to enable that intelligence—namely, digital microcontrollers, processor, memory, and interface (wired or wireless) systems.

By nature, complexity spreading requires low-power systems. Sensors are often located in remote regions or, if not remote, in very restricted areas. This necessitates extreme power conservation, an essential element in the mobile wireless space.

The other result of complexity spreading in IoT is the data explosion caused by such intense connectivity. The sensor, data and cellular networks must all adapt to handle the predicted growth. Companies are responding by developing a wide range of solutions to meet multiple needs, thus creating a very diverse network. To be successful, such approaches will require much collaboration, a major increase in bandwidth, lower network latency, and lower-cost, low-power, sensor-rich connected devices.

Whatever your view of the IoT market, rather from process or product, the important issue is to deal with the complexity spreading that results from the coming world of connected objects.

John Blyler covers today's latest high-tech, R&D and even science fiction in blogs, magazine articles, books and videos. He is an experienced physicist, engineer, journalist, author and professor who continues to speak at major conferences and before the camera on. John is the Vice President, Chief Content Office for Extension Media, which includes the brands Chip Design, Solid State Technology, Embedded Intel® Solutions and others. He holds a BS in Engineering Physics and a MSEE. John plays the piano and holds a black belt in TKD.

Embedded Intel® Solutions

www.embeddedintel.com

Vice President & Publisher
Clair Bright

Editorial

Vice President & Chief Content Manager
John Blyler
jblyler@extensionmedia.com
(503) 614-1082

Managing Editor
Cheryl Coupé
ccoupe@extensionmedia.com

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

Managing Editor - China
Jane Lin-Li

Contributing Editors
Cheryl Ajluni, Mark LaPedus, Craig Szydlowski, Nicole Freeman,

Creative/Production

Production Manager
Spryte Heithecker
Graphic Designers
Nicky Jacobson
Jacob Ewing
Media Coordinator
Yishian Yao
Senior Web Developer
Slava Dotsenko
Mariam Moattari

Advertising / Reprint Sales

Vice President, Sales
Embedded Electronics Media Group
Clair Bright
cbright@extensionmedia.com
(415) 255-0390 ext. 15
Sales Manager
Michael Cloward
mcloward@extensionmedia.com
(415) 255-0390 ext. 17

Marketing/Circulation
Jenna Johnson

To Subscribe
www.embeddedintel.com

Extension MEDIA

Extension Media, LLC
Corporate Office

President and Publisher
Vince Ridley
vridley@extensionmedia.com
(415) 255-0390 ext. 18
Vice President & Publisher
Clair Bright
cbright@extensionmedia.com

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

Human Resources / Administration
Rachael Evans

Special Thanks to Our Sponsors

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

WINTER 2014

DEPARTMENTS

SPECIAL FEATURES

TECHNOLOGY APPLICATIONS

LAST WORD

PRODUCT SHOWCASE

congatec
the rhythm of embedded computing

COM Express®

conga-TS87

- 4th Generation Intel® Core™ processor-based platform
- 3x DisplayPort 1.2, up to 4k resolution (QFHD=3840x2160@60Hz)
- Intel® Advanced Vector Extension (Intel® AVX) 2.0 for improved floating point computing
- Up to 13% higher computing performance*
- Up to 32% higher 3D graphics performance*

* SPECint_base2006, 3DMARK Vantage

Formfactor	**Formfactor COM Express™ Basic, (95 x 125 mm), Type 6 Connector Layout**
CPU	Intel® Core™ i7-4700EQ processor (4x 2.4 GHz, TDP 47W) Intel® Core™ i5 processor planned Intel® Core™ i3 processor planned Intel® Celeron® processor planned Intel® Turbo Boost Technology, Intel® Hyper-Threading Technology (Intel® HT Technology), Intel® Advanced Vector Extensions 2.0 (Intel® AVX2), Intel® Advanced Encryption Standard New Instructions (Intel® AES-NI), Integrated dual channel memory controller, up to 25.6 GByte/sec. memory bandwidth, Intel® HD graphics with dynamic frequency up to 1GHz, Intel ® Clear Video HD Technology Intel® Virtualization Technology (Intel® VT), Intel® Trusted Execution Technology (Intel® TXT), Intel® Streaming SIMD Extensions 4.2 (Intel® SSE4.2), PCLMULQDQ Instruction, Intel® Secure Key, Intel® Transactional Synchronization Extensions (Intel® TSX)
DRAM	2 Sockets, SO-DIMM DDR3L@1.35V up to 1600MT/s and 16GByte
Chipset	Mobile Intel® 8 Series Chipset: Intel® QM87 / HM86 chipset @ planned Intel® Celeron® processor variants
Ethernet	Intel® Ethernet Connection I217-LM GbE LAN Controller with Intel® Active Management Technology (Intel® AMT) 9.0 support
I/O Interfaces	7x PCI Express™ GEN. 2.0 lanes, 1x PEG GEN 3.0 (8GT/s), 4x Serial ATA® with 6 Gb/s, 4x Serial ATA® with 3 Gb/s (AHCI) RAID 0/1/5/10 support, 2x ExpressCard®, 4x USB 3.0 (XHCI), 8x USB 2.0 (EHCI), LPC bus, I²C bus (fast mode, 400 kHz, multi-master)
Sound	Digital High Definition Audio Interface with support for multiple audio codecs
Graphics	Next Generation Intel® HD Graphics with OpenCL 1.2, OpenGL 4.0 and DirectX11.1 support; up to three independent displays: HDMI 1.4, DVI, DP, VGA; High performance hardware MPEG-2 decoding, WMV9 (VC-1) and H.264 (AVC) support Blu-ray support @ 40 MBit/s
LVDS	Dual channel LVDS transmitter, Supports flat panels 2x24 Bit interface, VESA mappings, resolutions up to 1920x1200, Automatic Panel Detection via EDID/EPI
Digital Display Interface (DDI)	3x DisplayPort 1.2a / TMDS (DVI, HDMI)
CRT Interface	350 MHz RAMDAC, resolutions up to QXGA (2048x1536 @75Hz)
congatec Board Controller	Multi Stage Watchdog, non-volatile User Data Storage, Manufacturing and Board Information, Board Statistics, BIOS Setup Data Backup, I²C bus (fast mode, 400 kHz, multi-master), Power Loss Control
Embedded BIOS Features	AMI Aptio® UEFI 2.x firmware, 8 MByte serial SPI firmware flash
Security	The conga-TS87 can be optionally equipped with a discrete "Trusted Platform Module" (TPM). It is capable of calculating efficient hash and RSA algorithms with key lengths up to 2,048 bits and includes a real random number generator. Security sensitive applications such as gaming and e commerce will benefit also with improved authentication, integrity and confidence levels.
Power Management	ACPI 4.0 with battery support
Operating Systems	Microsoft® Windows 8, Microsoft® Windows 7, Linux, Microsoft® Windows® embedded Standard
Power Consumption	Typ. application: tbd., see manual for full details, CMOS Battery Backup
Temperature:	Operating: 0 .. +60°C Storage: -20 .. +80°C
Humidity Operating:	Operating: 10 - 90% r. H. non cond. Storage: 5 - 95% r. H. non cond.
Size	95 x 125 mm (3.74" x 4.92")

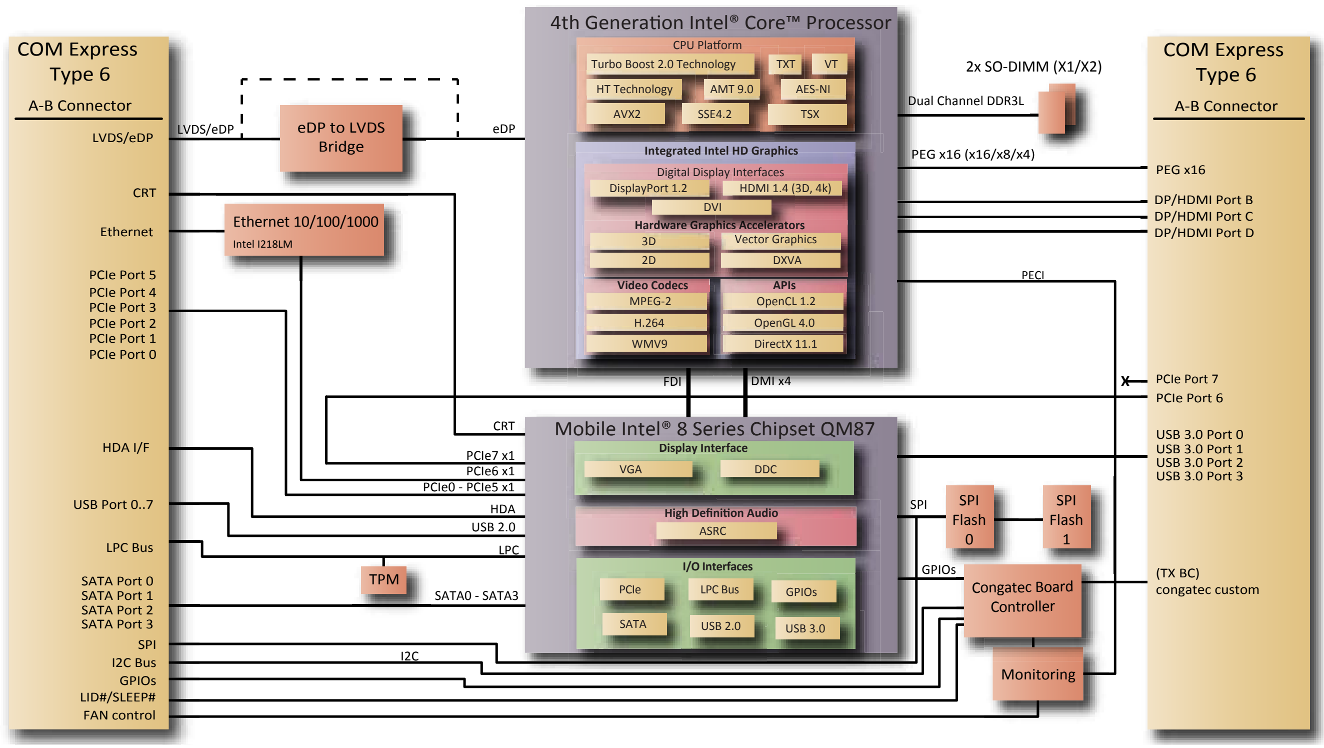

Article	PN	Description
conga-TS87/i7-4700EQ	046804	COM Express Type 6 Basic module with quad-core Intel® Core™ i7-4700EQ processor with 2.4GHz, 6MB L2 cache and 1600MT/s dual channel DDR3 memory interface
DDR3-SODIMM-1600 (2GB)	068767	DDR3 SODIMM memory module with 1600 MT/s (PC3-12800) and 2GB RAM
DDR3-SODIMM-1600 (4GB)	068770	DDR3 SODIMM memory module with 1600 MT/s (PC3-12800) and 4GB RAM
DDR3-SODIMM-1600 (8GB)	068780	DDR3 SODIMM memory module with 1600 MT/s (PC3-12800) and 8GB RAM

Engineering Tools / Accessories		
conga-TEVAL	065800	Evaluation carrier board for Type 6 COM-Express-modules
conga-LDVI/EPI	011115	LVDS to DVI converter board for digital flat panels with onboard EEPROM
COM-Express-carrierboard-Socket-5	400007	Connector for COM-Express carrier boards, height 5mm, packing unit 4 pieces
COM-Express-carrierboard-Socket-8	400004	Connector for COM-Express carrier boards, height 8mm, packing unit 4 pieces

Contact Information

congatec Inc.
6262 Ferris Square
San Diego,
CA 92121 USA
+1 858-457-2600 Phone
+1 858-457-2602 Fax
www.congatec.us

SoC FPGA Relies on Many Cores and 14 nm Tri-Gate Process

Altera's multicore FPGA uses Intel's tri-gate (FinFET) 14 nm 3D process technology to combine logic, four ARM A53 cores and OpenCL as a heterogeneous SoC replacement.

By Chris A. Ciufo, Editor-in-Chief

At the core of Altera's recent Stratix 10 announcement is… multiple cores. Four ARM Cortex-A53s to be exact, plus a bunch of DSP blocks, more logic than ever before done by the company, double the on-chip clock rate and software designed to create a multicore system instead of just a dense FPGA. In effect, it's a system-on-chip (SoC) that bests the company's recent Arria 10 SoC.

But also at the core of the 2014 next-gen FPGA family is Intel's tri-gate (FinFET) 14 nm 3D process technology, to which Altera has the exclusive FPGA rights for 12 years. It's Intel's tri-gate transistors that give Altera the edge—at least on paper—in the war of high density FPGAs. Details on Altera's Stratix 10 (the follow-on FPGA to the Stratix V) are still scarce, though the company's second pre-announcement added more details and we can guess at others.

The new multicore Stratix 10 FPGA SoC will tape out in 2014 and may possibly ship by late 2014. Let's take a look at some of the elements that make up the Stratix 10's performance goals.

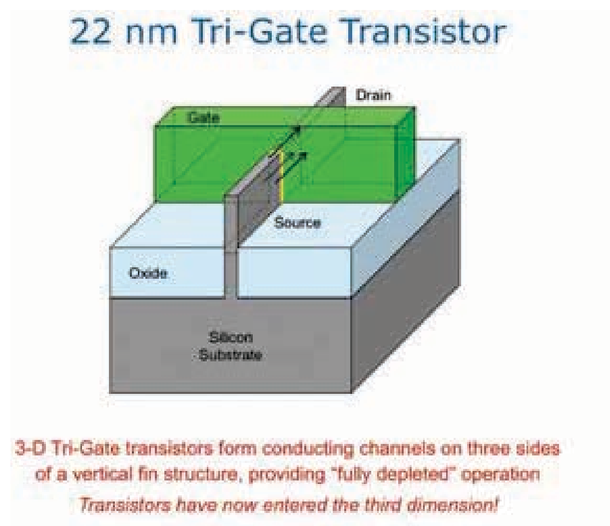

Figure 1: Intel's tri-gate transistors propagate Moore's Law, increase silicon density, and reduce power.

Tri-Gate's Secret Sauce

Altera is relying heavily on four factors to achieve the aggressive specs for Stratix 10: a refined architecture, the ARM Cortex A53, software tools such as OpenCL and Intel's tri-gate 14 nm process.

Intel has consistently battled Moore's Law through new process innovations. Following two-dimensional high K (dielectric) metal process improvements at 45 nm in 2007, Intel unveiled 3D transistors in 2011 in 22 nm feature size. As oxide thickness and transistor dimensions approach mere Angstroms, the company essentially made transistors "smaller" by stacking them in layers and wrapping the silicon channel with three gates instead of two. This allows more transistors per square nanometer by using the Z axis, and they're smaller and more power efficient. The concept of a non-planar double-gate transistor with a fin wrapped around the channel is called a "finFET," although Intel calls theirs "tri-gate" (Figure 1).

In early 2013, Intel and Altera inked a deal that allowed Altera to become the only FPGA company with access to

Figure 2: Altera's Stratix 10 SoC FPGA relies on four elements to achieve what the company calls "unimaginable" performance. Intel's 14 nm tri-gate process is the root of the performance discontinuity from current 28 nm FPGAs. (Source: Altera.)

Intel's 14 nm tri-gate process. Altera's CEO John Daane said at the time that he believed Intel was two to four years ahead of the competition. (This is debatable as TSMC moves from 20 nm to 16 nm.) Intel is already producing 22 nm 4th generation Intel® Core™ and Intel® Xeon® processors (codename

Industry's First Gigahertz FPGAs and SoCs

- New ultra-high performance FPGA architecture
- 2x the core performance of prior generation highend FPGAs
- >10 TFLOPs of single-precision floating-point DSP performance
- >4x processor data throughput of prior-generation SoCs

Break the Bandwidth Barrier with Unimaginable High-Speed Interface Rates

- 4x serial transceiver bandwidth from previous generation FPGAs for high port count design
- 28 Gbps backplane capability for versatile data switching applications
- 56 Gbps chip-to-chip/module capability for leading edge interface standards
- Over 2.5 Tbps bandwidth for serial memory with support for Hybrid Memory Cube
- Over 1.3 Tbps bandwidth for parallel memory interfaces with support for DDR4 at 3200 Mbps

Table 1: Stratix 10 performance targets. (Source: Altera.)

Intel® Core™ processor (codename Broadwell) is scheduled for production in 2014.

Despite a recent delay that necessitated design rule changes to Intel's 14 nm process, Altera as an Intel foundry customer will certainly follow shortly after Intel's production of the Broadwell microarchitecture-based processors.. According to Chris Balough, Altera's senior director of SoC product marketing, design software for the Stratix V will be available Q1 2014 and first silicon will tape out sometime in 2014.

"Unimaginable" Performance

In June and in November, Altera began predicting some pretty big numbers for Stratix 10…so much so that the nomenclature leaped from Stratix V (available now) to Stratix 10. The company is making what Balough called a "hyperbolic claim, something we never guessed was possible," he told me. There will be a 2x increase (from 28 nm) in logic fabric speed to over 1 GHz. Even though density increases, the SoC device will also achieve a 70 percent power savings with a target 4x. Other technology targets are shown in Table 1.

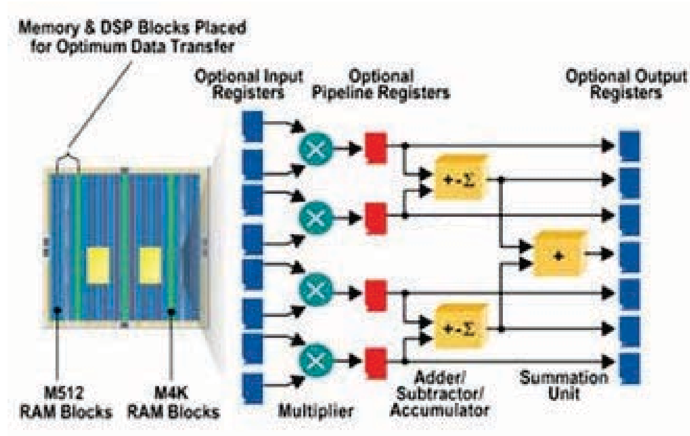

Figure 3: Stratix V (current generation) DSP block.

There is credibility in these claims, albeit based upon scaling factors. The new, in-production Arria 10 SoC is based upon TSMC's 20 nm process and is 15 percent faster than Arria V, consumes 40 percent less power and boasts a 50 percent processor system improvement. For Stratix 10, the magic again is Intel's 14 nm tri-gate process as described in the whitepaper "The Breakthrough Advantage for FPGAs with Tri-Gate Technology" (http://www.altera.com/literature/wp/wp-01201-fpga-tri-gate-technology.pdf).

Altera has provided no details yet on architecture other than the notional marketing chart shown in Figure 2. Fundamentally, Stratix 10 processing will take place in FPGA gates, four ARM Cortex-A53 processors and a series of DSP blocks. Software tools meld it all together into a powerful one-chip system.

DSP Blocks and ARM Cores

Current-generation Stratix V devices contain two columns of DSP blocks surrounded by constellations of logic array blocks. Each DSP block (Figure 3) can be configured for up to eight 9 x 9 bit multipliers, four 18 x 18 bit multipliers and one 36 x 36 bit multiplier. These DSP blocks run at 333 MHz and provide data throughput performance of 2.67 giga-MACs per block. The largest Stratix V device (EP1S80) has 22 DSP blocks.

With the forecasted 1 GHz fabric on Stratix 10 over Stratix V's 333 MHz, one might expect a 3x performance increase with no logic changes. However, Altera's Balough told me they're expecting a 6x throughput

Figure 4: Stratix 10 contains four Cortex-A53 cores, ARM's most impressive multi-threaded 64-bit core.

increase for Stratix 10 to "greater than 10 TFLOPs." It's entirely likely the number of DSP blocks will double, RAM block sizes will increase and there will be some fine tuning within the blocks and routing fabric. However, it's also possible the 6x claim is taking into account other FPGA resources for data movement and manipulation such as the ARM processors.

An FPGA's DSP sub-systems are a major reason FPGAs are chosen for high-performance algorithms like video CODECs or data-plane processing. On the Stratix 10, there are also the four ARM Cortex-A53 processors. Altera press materials cite an 8x throughput improvement over 28 nm FPGAs, and Altera says that ARM claims it's the highest power efficiency of any 64-bit processor. That's probably not hard to argue considering ARM's leadership in mobile and low-power devices...and Intel's struggle to catch up.

The four Cortex-A53 cores shown in Figure 4 are members of ARM's Cortex-A50 series announced in 2012 and are based on the ARMv8 64-bit architecture. The A53 chosen for Stratix 10 is the "little brother" to the A57, ARM's single-threaded, deep pipeline monster targeting servers. Both devices are designed for gigahertz performance in heterogeneous SoCs. For example, Stratix 10's A53s will target applications that offload x86 host processors, while AMD's HieroFalcon server accelerator 64-bit SoC (2H 2014) uses the Cortex-A57 to complement AMD's on-chip x86.

Stratix 10's A53s are software compatible with previous generation 32-bit ARM Cortex-A9s in Altera Cyclone, Arria and Stratix devices. This allows design sockets to upgrade to the latest FPGAs while migrating forward operating systems, unmodified software and IP cores (Figure 5). The A53 can access up to 256 TB of memory with ECC for on-core L1/L2 caches; both features pinpoint data center and high-end heterogeneous computing applications. To fully capitalize on Stratix 10's logic, DSP and A53 resource elements, the company is planning a suite of heterogeneous-focused tools.

Sweet Tool Suites

Altera's Chris Balough asserts that "all of Stratix 10's processing elements are compelling enough by themselves," but the "unimaginable performance" comes into play when they work together as a system. Altera is counting on OpenCL and the SoC Embedded Design Suite (EDS) to capitalize on the FPGA's heterogeneous elements.

The company has endorsed C-based design entry using OpenCL for several years via the SDK for OpenCL product. First demoed by Altera at SuperComputing 2012, the OpenCL design flow allows designers to work in C and easily mix logic and on-chip resources into heterogeneous multicore architectures. The high-level design flow alleviates the typical RTL coding required for most FPGA tasks. It also mimics the way embedded developers mix and match processing resources at the board or system level.

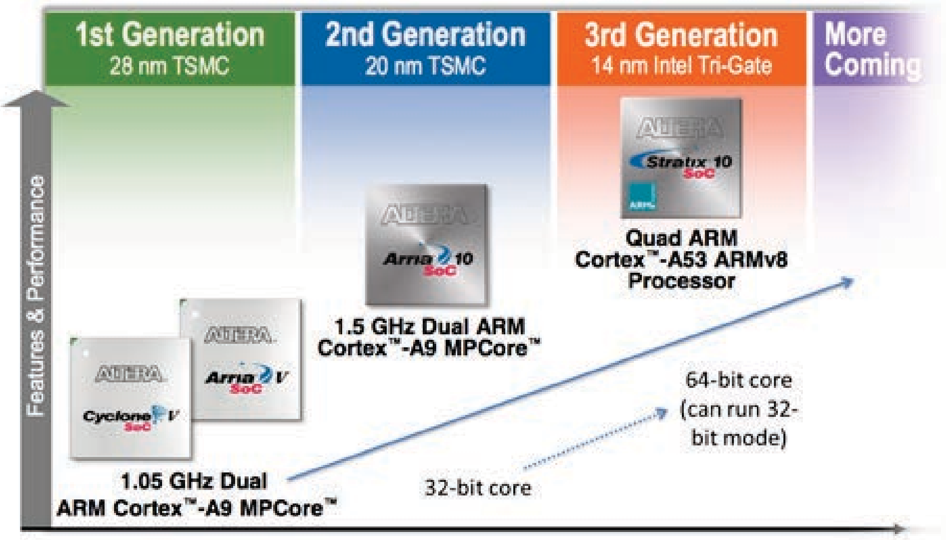

Figure 5: Stratix 10's quad Cortex-A53 cores can run in 32-bit mode for upgradeability from previous generation devices.

OpenCL code is also portable and lets designers load-level or performance-tune applications to take advantage of the massively parallel nature of FPGAs and the Cortex-A53 CPUs. In October of this year, Altera announced the SDK for OpenCL is conformant to the OpenCL 1.0 standard and is listed on the Khronos Group list of conformant products. This means that Altera can "provide a validated cross-platform programming environment" designed to accelerate algorithms at significantly lower power versus other cross compiler alternatives.

Practically speaking, OpenCL allows Stratix 10 designers to port or create hardware accelerators to take advantage of the FPGA's parallelism. This has always been possible when coding in RTL but not necessarily practical and usually not easily portable. The OpenCL SDK is complemented by EDS which provides FPGA-adaptive debug. Altera's Chris Balough describes EDS as "a native multicore debug environment with intrinsic FPGA debug capabilities built in." EDS supports real-time, in-system, whole-chip debug and visualization, including the ARM Cortex-A53 cores. That's because EDS includes Altera's existing DS-5 Altera Edition software, ARM's own development tool that's been tailored for Altera FPGA devices.

Both the OpenCL SDK and EDS software suites exist today for production devices, and will soon be released supporting the Stratix 10 so lead customers and early adopters can start their designs.

Intel® QuickAssist Technology Capitalizes on FPGA Resources

Altera was the first FPGA company on record to support Intel®'s QuickAssist Technology, a set of instructions and APIs that seek to offload host CPU instructions to a co-processor like an FPGA. Xtreme-Data was the board vendor who demoed it at the 2008 Intel® Developers Forum n Intel® Xeon® processor and Mathworks' Simulink graphical block diagramming software.

Altera's Stratix 10, with faster transceivers, DSP blocks and 1 GHz fabric, is likely to be a favorite co-processor for Intel host CPUs like the latest 4th Generation Intel Core devices or Intel's next-generation 14 nm tri-gate Broadwell microarchitecture-based CPUs. According to it's website, Intel, QuickAssist Technology (http://www.intel.com/content/www/us/en/io/quickassist-technology/quickassist-technology-developer.html) is ideal for computational workloads including cryptography, data compression and pattern matching—all applications that are algorithmically "heavy" and can take advantage of an FPGA's parallelism via mixed resources.

In the case of Altera's Stratix 10 next-gen devices, CPU offload via Intel QuickAssist Technology could also utilize high-level language OpenCL to construct heterogeneous "hardware macros" which balance workloads between the FPGA fabric, DSP blocks and the decision-making capability of the ARM Cortex-A53. All of these resources, of course, work as an accelerator sub-system toIntel Xeon and Intel Core processors, and now the newest Intel® Atom processors C2xxx—all of which contain Intel QuickAssist Technology.

Conclusions: Altera's Design Targets

We've written quite a bit about Altera's burgeoning roadmap over the past twelve months. During an interview in July 2012 with Brad Howe, Altera's senior VP of R&D, he talked about the need for FPGAs to evolve past their growing complexity and focus instead on "silicon convergence" (http://eecatalog.com/fpga/2012/07/17/fpgas-the-best-of-both-worlds-until-theyre-not/). He was of course referring not only to heterogeneous FPGA SoCs, but also on reducing design complexity with OpenCL.

Stratix 10, with its "unimaginable performance" made possible by Intel's 14 nm tri-gate process, epitomizes the future of big density FPGAs. "It was an explicit choice to make [Stratix 10] a heterogeneous computing platform," said Altera's Chris Balough.

We're anxious to see the company release details on the architecture, cite power estimates based upon notional reference designs and detail more of the specs. As we went to press, Altera revealed some data on Stratix 10's transceivers (4x bandwidth and 28 Gbps backplane switching), and indicated that the device is 3D-capable for packaging with ASSPs and memories.

Also, it will be interesting to see if Altera targets Big Data applications with Stratix 10 as I had predicted in June 2013 ("Does Altera have 'Big Data' Communications on the Brain?" http://eecatalog.com/caciufo/2013/06/05/a-slew-of-recent-altera-high-performance-announcements-over-the-last-three-months-can-only-mean-one-thing-the-company-is-targeting-big-data-in-a-big-way/#comment-20464). The company's string of acquisitions and posturing sure points that way.

We'll have to wait until 2014 for concrete information on the Stratix 10 heterogeneous SoC FPGA and its heavy lift capabilities. Until then, all this "unimaginable" information remains a technology announcement. But a compelling one, indeed.

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.

The Internet of Things Defines the Future RTOS

The RTOS of the future will give embedded systems manufacturers a competitive edge in the IoT by helping them bring industry-leading devices to market faster while reducing risks and development and maintenance costs.

By Prashant Dubal , WindRiver

Driven by the convergence of cloud technology, rapidly growing data volumes and increasingly connected devices, the Internet of Things (IoT) poses new challenges and presents a host of new opportunities that businesses of all sizes and industries can seize right now. This system-of-systems is fundamental to realizing business value—unlocking the insight hidden in data, identifying and creating new services, enhancing productivity and e¬fficiency, improving real-time decision making, solving critical problems, and developing new and innovative user experiences. Billions of intelligent devices and systems make up the Internet of Things. The majority of these "things" are embedded systems, many of which are running a real-time operating system (RTOS).

To fully take advantage of the opportunity offered by the Internet of Things, manufacturers of embedded systems must meet multiple challenges:

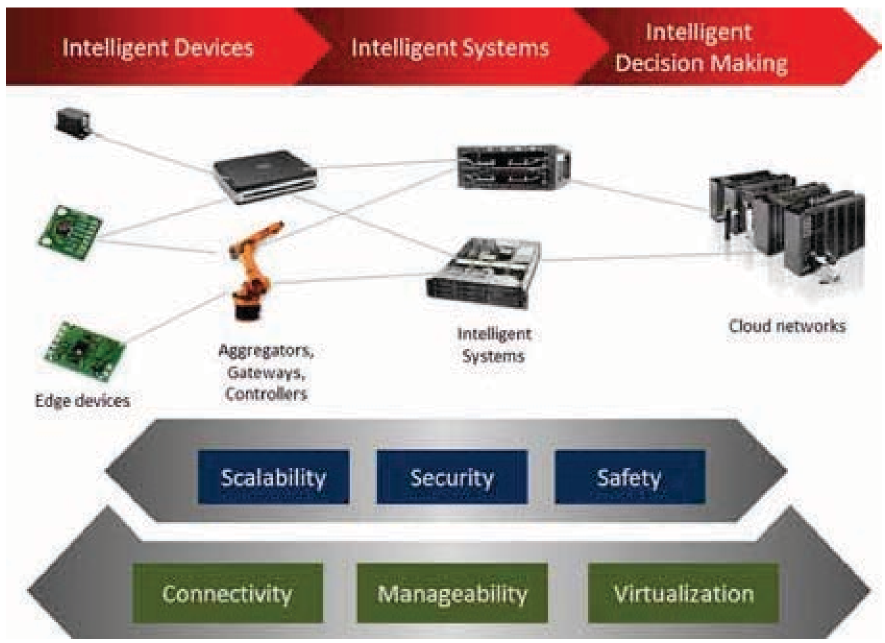

Figure 1: Core attributes and platform features of an RTOS for the IoT

- Bring connected devices to market faster
- Differentiate products with leading-edge features and capabilities
- Address security risks that pervasive connectivity of the Internet of Things entails
- Build flexibility into existing products so as to be able to tap new market opportunities as they emerge
- Ensure the product offering remains relevant and competitive as markets evolve
- Reduce system development costs and risks

To help manufacturers of embedded devices meet these challenges, an RTOS must evolve to deliver the scalability, modularity, connectivity, security, safety and a cutting-edge feature set that are demanded by the new, highly connected, security-conscious, remotely managed world of machine to machine (M2M) networks and the IoT (see Figure 1).

Scalability

The Internet of Things can create an incentive for manufacturers of embedded devices to maintain a broader product portfolio that includes different classes of devices ranging from small form factor, simple, single-application devices to large-scale, complex, multi-application systems. A single RTOS that can scale to meet the unique memory footprint, functionality and processing power requirements of multiple product classes can help manufacturers of embedded systems increase the return on their operating system investment, cut development costs by leveraging the economies of scope and reduce time to market (see Figure 2).

Modularity

The IoT and M2M landscape is evolving faster than the release cycles for the traditional RTOS, which means the design and deployment of the RTOS need to adapt. Traditionally monolithic in nature, an RTOS has been delivered all at once as a large bundle of software, board support packages (BSPs), middleware, operating system and tools. Updates to this baseline have been mostly for bug and security fixes rather than to add new features due to the prohibitive amount of coding and testing required to implement them.

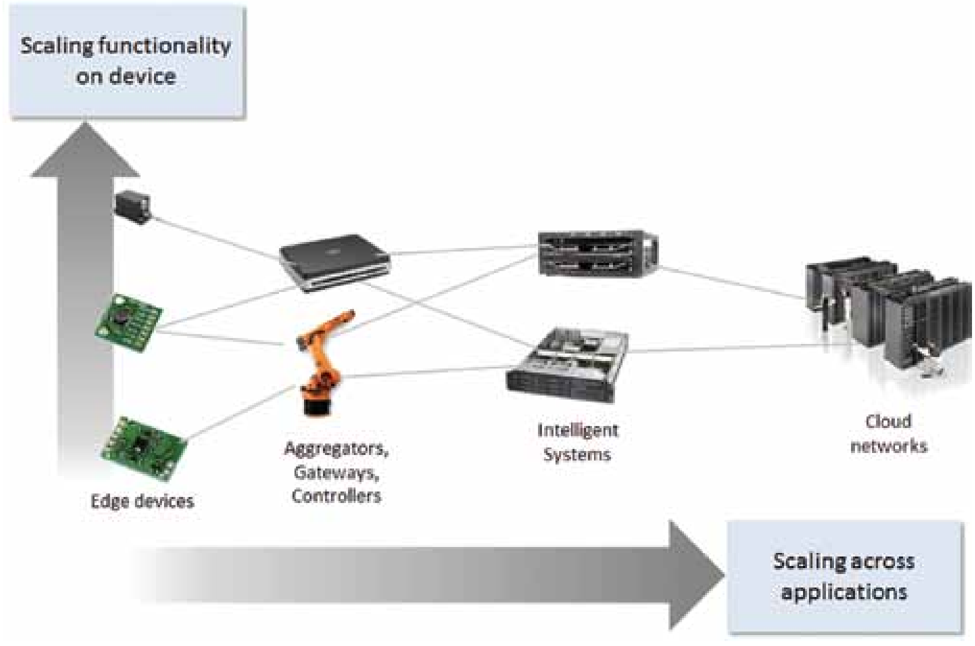

Figure 2: A modern RTOS must support these axes of scalability in order to deliver the most value in the IoT

The days of dedicated functions with little or no updates or expansion are over. Intelligent devices need to adapt to changing needs in the network. The reinvented RTOS must be built on a modular, upgradeable, future-proof architecture that separates the core kernel from middleware, protocols, applications and other packages. The RTOS of the future will provide a stable core so that add-on components can rely on this stability for a relatively extended period of time; for example, three years. Middleware, new protocols and other packages can be added or upgraded without changing the core. Components for all aspects of the RTOS, above the base kernel, can be provisioned by an application store model.

A modular architecture of an RTOS will help manufacturers of embedded devices better differentiate their products and maintain them competitively over longer periods of time by enriching them with new features and capabilities without changing the system core as standards and market requirements evolve. The new RTOS will also allow manufacturers to extend the useful life of the system core to several generations of products, which increases the return on their investment in the operating system.

Connectivity

While traditionally isolated, embedded devices are increasingly connected to corporate or public networks for a wide range of applications that are forming the Internet of Things. Small standalone sensor devices are being connected together using low-power wireless technology. Industrial control systems are interconnected and controlled remotely. Medical devices used at the home send diagnostic data back to a hospital.

A reinvented RTOS for IoT needs to support industry-leading communications standards and protocols such as CAN, Bluetooth, Continua, ZigBee, Wi-Fi and Ethernet, and deliver high-performance networking capabilities out of the box. In addition, a modular nature of the new RTOS can help retrofit existing devices with the required connectivity options so that many of the previously disconnected devices can be brought online without reworking the core of their embedded software.

Security

A critical aspect of IoT is security, and next-generation embedded systems must be designed with security in mind as their pervasive connectivity results in a substantially larger exposure to threats. A winning RTOS for IoT would give customers the flexibility to design their embedded system to the necessary level of security by leveraging a comprehensive set of built-in features covering (see Figure 3):

A good RTOS needs to support security features not only to protect against malware and unwanted or rogue applications, but also to deliver secure data storage and transmission and tamper-proof designs. Operating system-level support for these features is critical since adding them at the user or application level is ineffective, expensive and risky. Take, for example, sensor hubs that aggregate a representative data set from numerous packets of sensed data. These RTOS-based devices will require the logic to open those packets, validate their integrity, analyse their contents and verify that these actions have taken place securely. Security threats and vulnerabilities are ever-changing. An RTOS needs to support the secure upgrade, download and authentication of applications to help keep devices secure going forward.

Safety

Safety is paramount in many embedded operating systems because they control machines that can endanger life, or their malfunction can cause injury or death. Although well-established in aerospace, medical and industrial markets, safety standards are being applied by regulators to new mar-

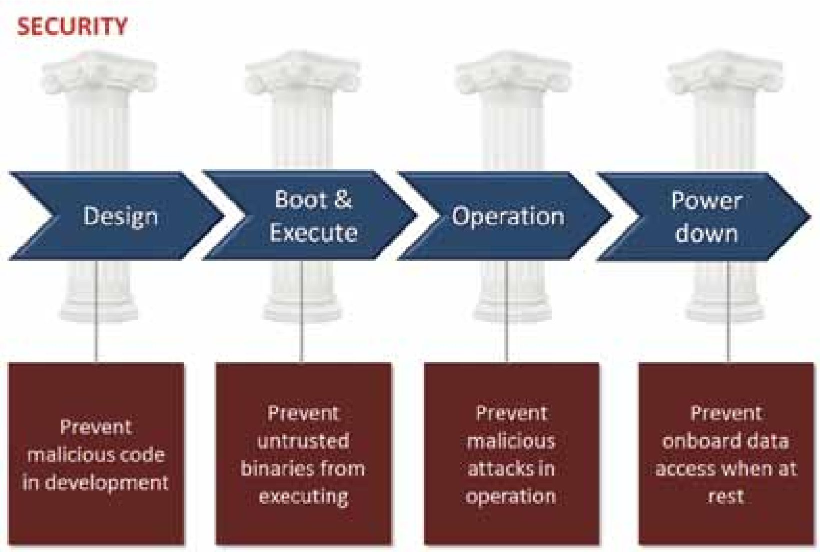

Figure 3: Four pillars of RTOS security

kets. Also, better applications of existing standards to such systems as smart grid meters or medical devices are sought. As standards evolve, manufacturers increasingly look to RTOS vendors to deliver the appropriate safety and security capabilities and certifications, so as to make it easier for them to obtain required safety and security certifications for their end products.

The reinvented RTOS must be built on a modular, upgradeable, future-proof architecture that separates the core kernel from middleware, protocols, applications and other packages.

Cutting-edge Feature Set

A broad feature set delivered by the modern RTOS and its ecosystem of compatible third-party applications is essential to enabling manufacturers of embedded systems to create a differentiated product offering and secure a sustainable competitive advantage.

- Rich user interface. With customer experience and user interface becoming key differentiating features for products ranging from mobile phones to medical devices to industrial control systems, powerful human-machine interaction capabilities are becoming a must for an RTOS for IoT. This includes quality 2D and 3D graphics engines, support for multiple monitors and touch screens, as well as rich graphics designer tools.

- Custom-tailored RTOS. Embedded systems manufacturers who are early adopters of IoT in industries such as networking, industrial and medical can enjoy even faster times to market by leveraging an RTOS that has been purposely customized and packaged to address the needs of their industry out of the box. For example, an RTOS for the industrial vertical would provide industrial device manufacturers with essential multimedia and connectivity middleware, including drivers and protocols for connected devices on the factory floor, wireless peripherals and other devices within the network infrastructure. An RTOS customized for medical devices would incorporate technology solutions designed to meet the unique needs of medical device manufacturers related to getting approvals from the U.S. Food and Drug Administration. A platform for network equipment manufacturers would enable them to rapidly create, test, deploy, maintain and manage high-quality wired and wireless infrastructure devices. Such a platform would also offer an extensive suite of security protocols to protect network data.

Compatible Software and Hardware Ecosystem

In addition to delivering rock-solid real-time performance and other cutting-edge features, an RTOS of the IoT era must support a broad ecosystem of tested and verified complementary hardware and software solutions. This would allow device manufacturers to differentiate their product offering with leading-edge features and capabilities, accelerate time-to-market through rapid, lower-risk integration of best-in-class third-party technology and cut costs by deploying systems integrated and validated out-of-the-box.

Summary

The era of the Internet of Things requires a modular, configurable and expandable RTOS.

The reinvented RTOS will add improved scalability, connectivity, security, safety, and an extended feature set to the solid real-time performance, low latency and multi-core processor support of the RTOS of today. The RTOS of the future will give manufacturers of embedded systems a competitive edge in the world of IoT by enabling them to bring industry-leading devices to market faster while reducing risks and development and maintenance costs.

Prashant Dubal is a product owner at Wind River managing VxWorks, development tools and the infrastructure product management team. He has held several roles at Wind River, including technical account manager and solution architect. Dubal has 13+ years of experience in the embedded industry, and holds a bachelor's degree in electronics engineering from Mumbai University, India.

Automotive Ethernet: No Simple Answers

Industry leaders speak out on the status of automotive Ethernet and where the technology needs to go.

By Cheryl Coupé, Managing Editor

As Microchip Technology's Henry Muyshondt states, "Ethernet means different things to different people." And that makes a discussion of the adoption of Ethernet in automotive applications tricky. In our roundtable discussion, I've pulled Muyshondt's explanation into a sidebar, which deserves consideration even outside the specifics of automotive applications. Along with Muyshondt, senior marketing manager for the Automotive Information Systems Division of Microchip Technology, Inc., we have Armin Lichtblau, business development director for the Automotive Network Design at Mentor Graphics (Deutschland); Joel Hoffman, automotive strategist for Intel's expanding Automotive Solutions Division; and Nick DiFiore, director of the Automotive Segment for Xilinx. Many thanks to our panel for their thought-provoking responses!

EECatalog: What's your view of the status of Ethernet adoption in automotive applications?

Henry Muyshondt, Microchip: For streaming data within the vehicle, the automotive industry has already pretty much settled on the MOST standard (Media Oriented Systems Transport). Over the last 13 years, it has come to be used in more than 150 automobile models already on the road, from most of the major car makers, in all regions of the world. There are more than 100 million MOST devices in use, as multiple devices are used in each vehicle. This technology can handle both Ethernet frames, as well as streaming information, all running in parallel. The higher-layer protocols used in the IT industry can communicate over the MOST standard without needing any changes, other than at the low-level link layer. Ethernet frames are sent unmodified, and other channels are available to transport audio and video data without overhead.

In terms of the actual Ethernet physical layer, it is not well suited for automotive applications due to electromagnetic compatibility issues as well as the challenges associated with running standard Ethernet cabling within the vehicle. There are proprietary technologies, such as Broadcom's BroadR-Reach that could be applied, but those are not really standard Ethernet and at this point have not received the wide adoption among carmakers that some in the industry press would suggest. There are many challenges that would need to be solved, in order to implement a low-cost, standards-based Ethernet solution in automobiles, and it will be several years before the current IEEE efforts toward defining a Gigabit physical layer for automobiles become a reality.

In short, Ethernet-style frames are definitely seeing increased use in automotive applications. The Ethernet physical layer, not so much.

Armin Lichtblau, Mentor Graphics: Particularly for new generation E/E designs and often AUTOSAR systems, Ethernet in the car is a requirement. The Ethernet communication is required as diagnostics over Ethernet to achieve high-speed, end-of-line programming in one use case. The other use case is the adoption of high-speed applications in the car environment, e. g., for video signal processing on safety and active safety systems (collision camera, etc.).

Joel Hoffman, Intel: Ethernet adoption is occurring slowly in automotive. This is partially due to the entrenched nature of existing closed solutions (such as MOST, CAN, Flex Ray and others) that have been developed specifically and exclusively for automotive, along with differing ideas on how to deploy the technology (such as the proprietary "OPEN" protocol by Broadcom).

These issues existed in the tech and enterprise segments as well, when Token-Ring, Asynchronous Transfer Mode (ATM) and other complex designs claimed technical advantages until more advanced silicon and software were created for the broader Ethernet market. If these issues had lingered, we would not have the cost-effective enterprise cloud and connectivity that we have today.

Nick DiFiore, Xilinx: Ethernet for automotive applications is rapidly growing among automotive OEMs to meet the growing bandwidth demands for new in-car audio and video (A/V) features. Many automobile vendors are planning to roll out Ethernet-based infotainment and advanced driver assist systems (ADAS) products in their 2015-2017 models. As a result, several high-profile automotive OEMs including GM, BMW and Hyundai have joined the AVnu Alliance—an industry forum dedicated to the advancement of professional-quality audio video in markets that include automotive, professional AV and consumer electronics. Xilinx is one of the founders of the AVnu Alliance.

EECatalog: Where does the technology need to go to cement Ethernet's role in the car?

Muyshondt, Microchip: If by Ethernet you mean packet-based communication technology, that role is pretty well cemented in the car for future applications, as mentioned above. The cement is still hardening, but there is no doubt that for inter-domain communication, both inside and outside the vehicle, packet-based communication is very relevant.

Much remains to be seen about what is implemented at the physical level, to transport these packets of information. So far, there are more than 150 vehicle models that have standardized on MOST technology for their high-speed network. MOST, in turn, has been enhanced to include packet communications and a dedicated Ethernet packet channel that can move unmodified Ethernet frames. We expect Ethernet to be used for some types of communication within the vehicle, using the existing MOST infrastructure. Additionally, a next-generation, multi-gigabit version of the MOST standard is already being developed for future applications. Even if packet communication is used for all in-vehicle applications, the MOST technology's full bandwidth can be allocated to Ethernet-style communication, providing an automotive-proven physical layer that also has the flexibility to allow streaming communication in parallel with the packet-based channels.

The MOST standard has already cemented itself with companies such as General Motors, Volkswagen and Toyota, along with most of the other major car makers of the world. These are large-volume manufacturers that have already proven the MOST technology's cost-effectiveness for even midrange to low-end vehicles. Additionally, the MOST Ethernet Packet channel is helping car makers implement Ethernet-style communications in vehicles.

Ethernet-style communications are already in place, while the Ethernet physical layer really isn't being used in vehicles.

Lichtblau, Mentor Graphics: Ethernet stacks need to be implemented in the Embedded Basic software stack which is used as the central operating platform in the car ECUs. The tooling provided for the ECU development needs to be set up to configure and test the Ethernet components. The IP communication protocol needs to be adjusted with higher level protocols to accommodate the automotive requirements.

Hoffman, Intel: OEMs need to take a leadership position in solving any remaining technical issues and pulling their supply chain into the ecosystem. This begins with audio systems implementing audio-video bridging (AVB), and extends to other deterministic control systems in the vehicle. While there are risks with implementing any new technology, automakers have great influence in stimulating the development of related software that will be needed. For the first cycle or two there may not be significant savings; however, the first adopter will gain the most in the end through reduced development costs, reduced equipment costs and long-term vehicle fuel savings due to vehicle weight reduction.

DiFiore, Xilinx: There are two main issues:

1. 1Gbps physical transport in automotive environments. There's a need for a low-cost, single, unshielded twisted-pair (UTP) solution that meets automotive EMC requirements. This need is currently met by Broadcom's Broad-R-Reach technology for 100Mbps Ethernet networks. The industry still needs an affordable 1Gbps UTP solution to really cement use of Ethernet in automotive applications, especially for A/V uses. Broadcom and other Ethernet vendors are working on technology to meet the 1Gbps requirement.

2. Guaranteed quality of service (QOS) over Ethernet for reliable, deterministic, real-time data delivery and audio/video streaming over Ethernet in a car. This need is not fully satisfied by current IEEE Ethernet standards. However standards dealing with bandwidth reservation, time synchronization and packet prioritization are emerging and the AVnu Alliance is developing ways of adapting these standards for automotive needs. Many automotive OEMs are currently evaluating Ethernet extensions against their list of requirements. At this point, it's not clear if OEMs will agree upon a complete, standardized Ethernet stack or if different OEMs will simply adopt variants of Ethernet AVB.

EECatalog: What are the challenges you expect to address?

Muyshondt, Microchip: The biggest challenge is to have the right system, at the right point in time, at the right cost point. The networking infrastructure inside vehicles has much different requirements than the typical IT infrastructure, in terms of robustness and surviving in a harsh environment. Microchip is a leader in both consumer and automotive applications, and I can tell you that it is not easy to simply take a consumer device and harden it for automotive applications.

The device has to start out with the design goal of being used in automotive applications.

Bandwidth costs money. You can either spend that money on the physical connection, or you spend it on bigger processors with more computing capabilities. The challenge is to balance all sides of the equation to obtain the optimum system. This requires participation by the people designing the system. Carmakers have to be involved from the ground up in specifying the various parameters, including how the system will be supported after it rolls off the assembly line, and for many years after it rolls away from the dealership. Only carmakers and their suppliers can properly assess the correct trade-offs among the various competing objectives of a networking and communication system. Their timeframes for choosing a technology have to look five years or even a decade into the future, in order to select the appropriate technology to use. Carmakers have to be technology decision-makers, and not just the users of what other industries may provide.

Infrastructure, such as a networking system, doesn't provide a lot of differentiation between cars. As such, it is in everyone's best interest to use common "plumbing" to move the data behind the dashboard. Carmakers need to cooperate with each other on a technology they can truly influence, rather than relying on devices primarily built for other industries that have much higher volumes, and therefore much higher influence, than the automotive industry. Their products have very long life cycles, measured in dozens of years rather than a few months, and therefore their choices need to have built-in reliability and robustness. They also have to have the confidence that the actual devices they

Ethernet vs. IP – Henry Muyshondt, Microchip Technologies, Inc.

The term Ethernet means different things to different people. Ethernet really refers to the IEEE 802.3 standard that defines the physical and data-link layers used to connect computers. It specifies the four twisted-pair CAT5 wires and the electronics they connect to on a typical office or home computer, along with a particular format for a packet of information. It does NOT refer to the higher communication protocols above the physical layer, such as the various Internet protocols and other communication mechanisms used in the IT communications world. It also does NOT refer to things like Wi-Fi, Bluetooth, or other networking physical connections, even when those connections use packets of information similar to those used in a proper Ethernet system. Many people say Ethernet when what they really mean is Internet Protocol (IP) or another packet-based communication protocol, and not the physical layer. This has led to some confusion in the industry when referring to Ethernet.

IP communication is certainly the wave of the future, especially as it relates to communicating between different domains within the automobile, or to the world outside the vehicle. This is already starting to happen, with many telematics applications. We expect these types of communications to continue to expand their presence in the automobile. We also expect packet-based communication to coexist with other more efficient streaming technologies, as engineers exploit the advantages that each technology brings to the table.

For example, when independent systems need to communicate—particularly with off-board systems, such as the Internet or OEM diagnostics systems—they can take advantage of packet communication's characteristics, such as being able to adapt to changing interconnections between the data's source and destination. They can also take advantage of standardized protocols that are designed to work over unreliable connections that might require retransmission of packets, and where the path between source and destination may include links that disappear, while others appear along the way, as computers are turned off or links are otherwise interrupted. These mechanisms might also be needed to communicate between domains inside the vehicle that are developed separately from each other, such as engine control and infotainment. This type of communication provides a common language for disparate groups to talk to each other. Ethernet-style frames and packets are a useful tool for this purpose.

Other applications, such as when there is a continuous flow of information between devices, can use more efficient mechanisms that don't require the overhead of packet communications. If you have a continuously flowing stream of information—say video going from a camera in the driver-assist system to an in-cockpit display, or audio going to an amplifier—there is no need to break up the stream into separate packets, which add a significant amount of addressing and error-correction information. These packets would then have to interrupt a processor that would need to process each one, discard up to 3/4 of the transmitted overhead bits, and again assemble the data into a continuous stream to be fed to a digital-to-analog converter of some sort. Ethernet also doesn't supply the higher-level protocols that are needed to manage communication channels, transmission errors, control of various devices, etc. Such higher-level protocols do exist, but they are no longer part of the Ethernet realm. In fact, they could be used over many other physical layers that are not called Ethernet. Packet communication of streaming data wastes a significant amount of bandwidth and very significantly increases processor performance requirements to handle the increased interrupt load along with the software stacks needed to process the packets. These stacks make determinism difficult and introduce varying amounts of latency that affects audio and video presentation.

select to implement a technology will still be manufactured more than a decade later, when the consumer and IT devices of today are but a long-forgotten memory.

Ethernet-style communications can help address some of these challenges, alongside streaming technologies such as the MOST standard. The Ethernet physical layer, on the other hand, just isn't well suited for automobiles.

Lichtblau, Mentor Graphics: Mentor is committed to provide an Ethernet stack as part of their AUTOSAR solution and will fully support Ethernet in cars and diagnostics over Ethernet with their Embedded SW AUTOSAR stack as well as the AUTOSAR tooling.

Hoffman, Intel: Intel believes in standards and has benefited most when standards are widely adopted. Our participation in automotive forums including AVnu Alliance (avnu.org), GENIVI Alliance (genivi.org), Wi-Fi Alliance (wi-fi.org) and others bring the best of the ecosystem together to agree on common methods for solving common problems. Collaboration with these groups allows Intel to produce silicon solutions such as the Intel® Ethernet Controller I21X family with support for AVB. Intel also develops proof point technologies to jump start the industry including reference implementations such as Open AVB (https://github.com/intel-ethernet/Open-AVB).

DiFiore, Xilinx: Xilinx solutions already address Ethernet QOS challenges and the company has worked with Digital Design Corporation (DDC, a Xilinx Alliance Program member) to develop a complete Ethernet AVB solution called EAVB for automotive applications for use with Xilinx FPGAs and the Xilinx Zynq All Programmable SoC. Ethernet AVB implementations on Xilinx All Programmable platforms allow our customers to develop and field system designs well ahead of their competitors and to react quickly as the automotive uses of Ethernet and the related standards evolve and grow.

Cheryl Berglund Coupé is managing editor of EECatalog.com. Her articles have appeared in EE Times, Electronic Business, Microsoft Embedded Review and Windows Developer's Journal and she has developed presentations for the Embedded Systems Conference and ICSPAT. She has held a variety of production, technical marketing and writing positions within technology companies and agencies in the Northwest.

The Rise and Challenges of M2M Applications

A "how to" list of practical considerations for adding wireless to machines.

By Herbert Blaser and Carl Fenger, u-blox

2013 could be the year when machine-to-machine (M2M) communications exceed human-to-human communications for the first time in history, with even more machines connected to the Internet than people. Mobile resource management systems, meters, robots, vending machines, security systems, asset trackers, vehicles and emergency call systems all belong to this growing population of chatting machines.

There are many considerations to think about when designing inter-connected M2M products. Many new standards, both wireless and positioning, are in transition. It is important to consider the long term anticipated lifetime of products, and in which markets those products will serve. As well, designers must consider whether it is important to include support for next-generation performance and network coverage, or instead to design for easy product upgradeability.

Increasingly Inter-connected Machines

The M2M growth is occurring at the same time that we are running out of IP addresses. But the future of the Internet relies on IP version 6, which supports 2 to the power of 128 addresses, more than enough for every grain of sand on Earth to have its very own address. It is thus no surprise that LTE, the fourth generation of mobile networks (4G), is designed to deliver services such as data, voice, and video all over IP version 6.

The motivation for this networking revolution is simple—all devices and applications that can profit from being connected to the Internet eventually will be connected. While consumer devices are the most visible applications of mobile connectivity, humans aren't the only ones using the internet. It is the invisible applications that are growing the fastest…the silent conversations between millions of machines exchanging data 24 hours a day, 7 days a week, with no human intervention.

All that's needed to join the network is to embed any device with a small, low-cost (wireless) modem. Applications reporting on location, speed or navigation information also require a GPS or GNSS (Global Navigation Satellite System) receiver. Both components, with an antenna, can fit easily in a device much smaller than a mobile phone.

This is happening across all sectors of the electronics industry at this very moment.

Equipping devices with M2M communications capability, however, has special requirements, depending on the application. It is important to consider these requirements when thinking not only of the initial design, but also about product longevity (how long the device should operate before needing replacement), geographical coverage (it was initially designed to work in only one region, but now needs to work in another one), or compatibility with unavoidable wireless network upgrades, 2G to 3G to 4G.

Technical Considerations

Below are some important technical features to consider when designing M2M applications, and how they can affect the design of specific types of devices.

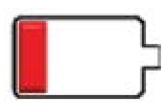

1) Power consumption

For portable tracking, security or personal safety devices, time between battery charging is one of the most important features. For example, a container-mounted tracking device that has to recharge once a day is too frequent, as a typical trip could take several days by air or road, and up to several weeks if shipped by sea.

For consumer devices such as personal tracking or health monitoring devices, mobile phones have set the standard for expectations—battery life should last a minimum of 3 days. When comparing modem and GNSS receiver specifications of such applications, not only operating and standby current consumption are relevant, but also power saving modes. These include auto-wakeup features and intelligent power-saving modes such as the ability to log data autonomously without waking up the host processor. Ideally these components should be in a minimum-power mode most of the time, waking up only when absolutely necessary.

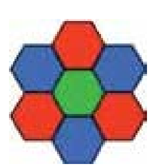

2) Cellular network compliance

In which regions should the device work? With global mobility increasing for both people and goods, it is important to consider that a modem that works in one region, may not work in another (GSM is supported by two main frequency bands worldwide, UMTS by six and LTE over 30.) For these types of applications it is important to identify where the device should work geographically, and to anticipate that this area could expand in the future.

A resource management system that must monitor shipments in all regions of the world should have either a quad-band GSM modem, or 6-band UMTS modem. For a device that you don't expect to move, e.g., a residential electricity meter, only a single frequency band is necessary. Other applications may need some additional consideration. A vending machine whose location is often forgotten, can always "phone-home", but it must then be equipped with a modem that operates in regions where it is, or could be located.

3) Operator approval

Any wireless device that communicates via GSM, UMTS or LTE will require operator certification before it is allowed to access their network. To significantly ease the certification process, the modem embedded in the device should also be operator certified. Check the list of modem certifications against the regions where the tracking device should operate.

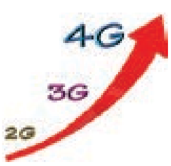

4) Wireless modem upgradeability

Although it is tempting to rely on GSM/GPRS for remote metering applications where only small amounts of data are communicated, the frequency bands for GSM are already being considered for re-allocation to 3G and 4G services. In the case of Automatic Meter Reading systems, retro-fitting hundreds of thousands of remote utility meters is expensive. It may therefore be wise to design with the technology of the future in mind. This means either already designing with UMTS/HSPA or LTE modems, or at least future-proof your hardware design such that modem upgrade is as cost-effective as possible.

5) Nested modem design

Your M2M device today may need to adapt to a new mobile or GNSS standard tomorrow, or address a new customer demand coming from a region that uses a different frequency band or satellite receiver standard. Ideally, you can react to this market requirement by simply making variants of the firmware, antenna and modem or GNSS receiver of your existing design. In reality, this is a nightmare unless nested design is an inherent property of your vendor's products. In particular, PCB layout issues can generate a long list of expensive design and logistics problems.

The best way to avoid this issue is simply to use components that have layout compatibility across the entire range of wireless modem (GSM, UMTS, CDMA and LTE) modules or GNSS receiver (GPS, GLONASS, Galileo, and BeiDou) modules. With this solution, a single PCB layout can be designed for use by all end-product variations. This issue can be addressed with these questions: Does your component vendor support a nested design philosophy? Does their next generation modem fit comfortably on the same PCB footprint of their current modem product? Do they provide documented technical support to help you make a successful nested design?

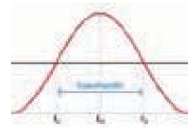

6) Bandwidth requirements

For many of today's tracking applications, only a low bandwidth connection is required to support tracking and text messaging. If only data is needed, then simple GPRS is sufficient. If a voice-channel option is required, then at least GSM/GPRS. If a video stream to support visual surveillance is desired, then UMTS/HSPA is the better choice. For applications requiring high-definition video or low latency such as telehealth terminals, LTE will be the technology of choice. One thing that is certain is that the tracking applications of tomorrow will always require increasing bandwidth. Select your modem based not only on the requirements of today, but also 3-5 years from now, or make a selection where modem upgradeability comes with minimal costs (refer to nested design above).

Position Requirements

7) Automotive requirements

Especially for vehicle- mounted systems where temperature, humidity and vibration conditions can be extreme, look for components that are automotive qualified, conforming to AEC-Q100 and manufactured in ISO/TS 16949 certified sites. Qualification tests for each component should conform to the ISO16750 standard: "Road vehicles – Environmental conditions and testing for electrical and electronic equipment". This is important not only for vehicle- mounted devices, but also for industrial devices that must operate outside, in ships or in railcars.

8) Support of emergency call systems

There is increasing global legislation to equip new cars with automated systems that can report an accident, as well as aid recovery in the case of theft. The US, Europe, Russia and Brazil have all established nationwide initiatives supporting these systems and that will increasingly be required by government mandate. For these systems, a specific modem feature, the "In-band modem", is often required. It enables data to be sent over the modem voice channel, similar to how a fax machine sends data over the telephone lines. This is required due to the higher prioritization of the voice channel over data in mobile networks. In an accident scenario, voice channel availability is higher than data channels such as GPRS or HSPA which may not even be available in remote areas. The voice channel is therefore a crucial link for transmitting data to an emergency response center.

Questions to ask your modem vendor concerning emergency call support: Do your modems support in-band modem? Is it supported on both 2G and 3G modems? Are your modems suitable for automotive applications? Do your satellite receivers support Dead Reckoning? Do you provide both GPS and GLONASS receivers? (See related points above.)

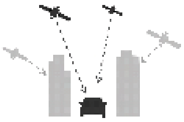

9) Assisted positioning

For M2M applications requiring reliable position information in urban environments, the availability of an assisted positioning system should be considered. Especially in cities where satellite visibility is often blocked by tall buildings, drop-out of positional overview can be overcome by calling up a remote A-GPS server. This is a simple process of downloading a few bytes of satellite orbital data from the Internet using a wireless modem. With this aiding data, visible satellites need only be visible for a few seconds to calculate a position, and not the full 30 seconds it takes to receive an entire 1500 bit satellite frame.

For certain applications, such as low-end navigation systems, a momentary lapse of position can be tolerated. For other applications such as vehicle emergency call or road-pricing systems, even temporary loss of position can have unacceptable consequences, making assisted positioning an attractive feature. What to look for when considering assisted positioning: Does the positioning (GPS) receiver vendor support an online assistance service? How reliable is the service, for example, does it provide guaranteed availability? Which regions of the earth does the service support? Does the vendor include client software that supports the service transparently? Do the positioning receiver and the wireless modem have an interface to support the service? Is the service available for GPS and GLONASS?

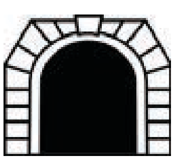

10) Dead reckoning support

For vehicle-based telematics systems, such as insurance tracking systems, the ability to accurately record position, heading and velocity information is crucial. In tunnels, however, the absence of satellite signals means this information must be temporarily generated by a parallel system. An important technology to supplement satellite signals is Dead Reckoning, which extrapolates location and speed based on input from vehicle sensors.

Questions to ask about your positioning receiver: Does it support Dead Reckoning? Can it be plugged directly into the vehicle CAN bus to acquire the data? Can it directly interface to vehicle sensors such as Gyro and odometer? Does the vendor offer a complete, proven system with evaluation environment? Are the components automotive grade (see point below)?

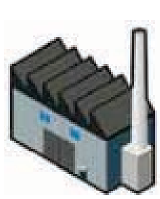

11) Indoor positioning

Unfortunately, GPS does not work deep indoors, nor do any other satellite-based navigation systems. The extremely weak signals are easily blocked by walls, metal, or even a thin sheet of water. For applications where an approximate position indoors is required, combining a satellite receiver with a wireless modem can overcome this problem via a hybrid solution that exploits the visibility of 2G or 3G cells. As GSM or UMTS signals easily penetrate walls, if the boundaries of visible mobile cells are known, an approximate position can be calculated based on where the cells overlap. This solution requires wireless connection to an external service, similar to the assisted positioning solution mentioned above. Questions to ask about your source of positioning receiver and wireless modems: Do they support such a solution? Is it proven or only in theory? Do they provide an online service, and is it in operation? Can your chosen satellite receiver and wireless modem support the service? How accurate is it?

12) Positioning system compatibility

Until recently, GPS was the only system you needed to consider. Now, with Russia's GLONASS and Japan's QZSS systems online, plus China's BeiDou and Europe's Galileo on the horizon, compatibility with GPS plus at least one other satellite system will be required available to increase system reliability and accuracy as well as to fulfill regional government mandates for compatibility with their own systems.

Often, parallel operation that uses two systems simultaneously will be part of the specification. An example is Russia's new ERA-GLONASS vehicle emergency call system that requires GLONASS compatibility. Questions to ask about your GPS/GNSS receiver: Does it provide multi-GNSS support? Does it provide parallel GPS/GLONASS or GPS/BeiDou?

Wiring It All Together

These are just some of the considerations you may want to think about when designing your M2M products. Remember that many new standards, both wireless and positioning, are in transition. It is important to consider the long term anticipated lifetime of your product on the market, and which markets your products will serve. Also consider whether it is important to include in the design support for next-generation performance and network coverage, or instead to design for easy upgradeability of your products in the field.

Carl has 25 years experience working in the semiconductor, software, and telecoms industries based in the USA and Switzerland. He is currently Product Communications Manager at u-blox. Carl holds a Bachelors of Science in Electrical Engineering from the University of California.

Herbert has 20 years international experience in the semiconductor industry. Before joing u-blox, he spent 14 years at Philips Semiconductors/NXP working in several areas of embedded cellular technologies. He is currently VP of Business Marketing at u-blox specializing in strategy, business planning and leading the communications team. Herbert holds Masters Degrees in Electrical Engineering and Industrial Management from the Swiss Federal Institute of Technology in Zurich.

SHIM Spec Unlocks Power of Hardware Accelerators, Speeds Software Development

New SHIM working group by the Multicore Association seeks to define ways to describe complicated IC hardware via XML, creating better software tools and faster time to market.

By Chris A. Ciufo, Editor-in-Chief

Moore's Law allows hardware improvements to run ahead faster than developers can write code for new, sophisticated accelerators like CODECs, encryptors or multi-/many-core processors. Just look at Intel's roadmap demonstrating exponential process shrinks almost yearly, doubling transistors and constantly adding new hardware functions.

Or consider the rapid evolution of I/O standards like USB 2.0 to USB 3.0, PCI Express Gen 2 to Gen 3, Serial RapidIO, DDR4, 802.11ac Wi-Fi and so on. These all use specialty silicon hardware. In fact, according to Markus Levy, co-founder of the Multicore Association: "semiconductor manufacturer documentation for some specialty accelerators—never mind the entire SoC device—can run in the hundreds if not thousands of pages." Using all these cores, accelerators, and I/O interfaces optimally requires sophisticated tools to make the programmers' job manageable. But how can the tools vendors keep up without huge personnel investments, or slowing down market momentum while they spin-up on new functions?

An answer might be SHIM, the "Software-Hardware Interace for Multi-Many-Core," now an active working group under the Multicore Association. The WG is headed by Masaki Gondo, CTO of software provider eSOL, and includes primary contributing member companies: Cavium, CriticalBlue, eSOL, Freescale, Nagoya University, PolyCore Software, Renesas, Texas Instruments, TOPS Systems, Vector Fabrics and Wind River (as of December 2013).

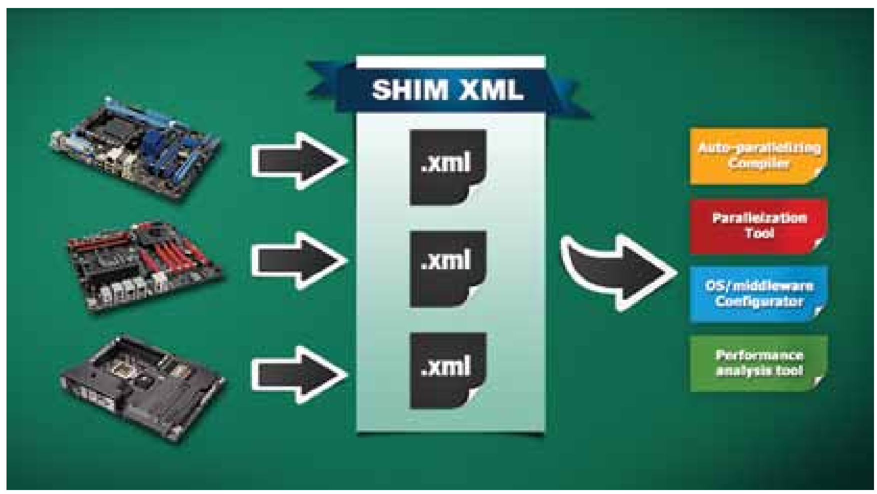

Figure 1: The SHIM specification defines ways for XML schema to describe hardware functions as applied to software that uses or interfaces to that hardware. (Courtesy: Multicore Association.)

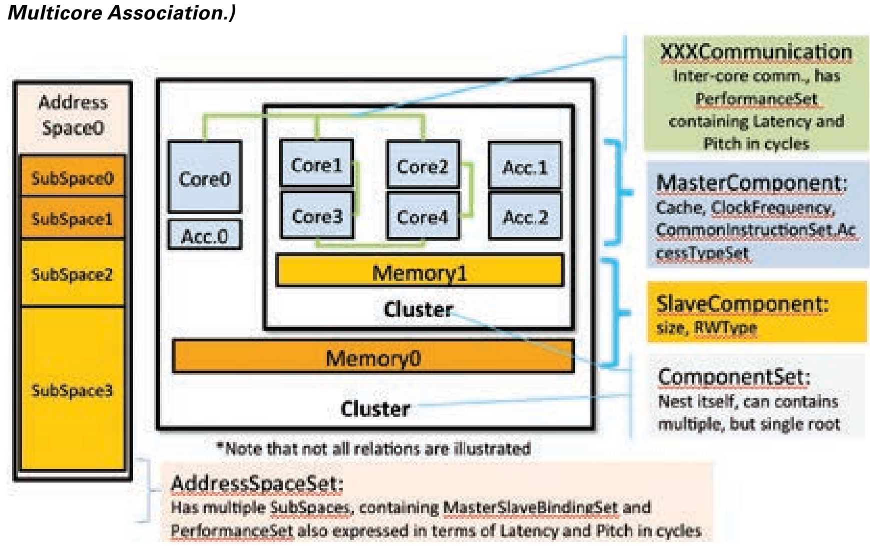

Figure 2: Depiction of how SHIM represents hardware elements in a multicore device. (Courtesy: Multicore Association.)

Use Case Description	Benefit
Performance estimation	Performance information is critical for most design-aid tools
	Examples: auto-parallelizing compilers, other parallelizing tools, performance analysis tools, etc.
System configuration	OS, middleware, run-time libraries need basic info to configure
	Other tools need info to configure
Hardware modeling	May configure H/W model (eg: simulator), device drives, or hardware abstraction layer
	May be useful for architecture exploration
Migration/evolution	May assist in migrating to or porting functions to new hardware or upgrades

Table 1: SHIM use cases. (Courtesy: Masaki Gondo, eSOL and MCA/SHIM Working Group.)

SHIM Defined

Figure 1 shows the relationship between hardware elements—shown here as boards but aimed at the hardware cores on those boards—and types of software tools. The SHIM spec will "define an architecture description standard from the software design perspective." Types of hardware that can be described via SHIM include:

- Processor cores; number of cores
- Inter-core communication channels (for message passing protocols such as MCA's MCAPI)
- Memory sub-system (including hierarchy, topology, coherency, size, latency)
- Network-on-chip and routing protocol
- Hardware virtualization features
- Accelerators
- On-chip peripheral devices
- End-node behaviors (at the hardware level).

It's tempting to think of the SHIM spec as a way to describe the hardware itself; however, that's already defined by the IEEE XACT standard. Rather, SHIM is an interface defined by an XML schema provided by the silicon or IP vendor for use by software developers. On a more basic level, MCA describes SHIM as "an extraction of hardware properties that matter to multicore tools." It's a hardware model described from a software point of view.

Since this is a bit novel, the converse can also help describe SHIM. It's not a functional hardware model because it's descriptive. It's also not intended to describe all of the hardware, either: only those properties that matter to a software developer. Finally, SHIM isn't a tool—its purpose is to help tools vendors create and evolve their tools in phase with the rapid pace of hardware evolution.

Why SHIM?

The fundamental reason hardware and software companies are combining resources to create SHIM is that it's just getting harder and harder for tools vendors to keep up with the hardware. The example above that describes 100-page manuals is real, according to Sven Brehmer, CEO of PolyCore Software (chair of MCAPI and a participant in SHIM). "The complexity [of modern SoCs] basically explodes" as more cores and accelerators are added.

"We need a more sophisticated way to learn how to access and write software [tools] for devices and configure them before run-time," said Brehmer. He described several use cases where SHIM would be an advantage. One is the typical SoC with multiple cores (which may be running multiple operating systems), several accelerators such as a packet processor and DMA engine, and one or more caches (Figure 2).

In this "simple" example—which might be an ARM-based SoC—while the ARM core doesn't change much from vendor to vendor or in subsequent generations of the same SoC family, the accelerators, peripherals and memories can be very different. PolyCore, a provider of tools to simplify multicore programming, would need detailed knowledge of each of these elements in order to port the company's Poly-Platform for each new SoC. The time for a port is at least months for a complicated device. Multiply this by many vendors and devices, and the task for tools providers is unsustainable.

Another use case Brehmer offered is a Texas Instruments Keystone architecture-based C6x device typically used in high performance DSP for wireless and wireline applications. Each device has more than ten custom accelerators. Using TI-provided SHIM-compliant XML schema, PolyCore's tools could more quickly support the TI functions, support them more completely and elegantly, and alleviate the need for the customer programmer to have to figure out each accelerator at the bit/register level. If a SHIM model is available for all hardware functions, that means better tools and more effective SoC use by the end developer. Conversely, hardware not covered by SHIM might force the end programmer to write their own tool—or live with none at all. They'd bit-bang their way through programming and hope for the best.

Figure 3: SHIM XML in detail. (Courtesy: Multicore Association.)

According to Gondo, the SHIM WG Chair, SHIM use cases are shown in Table 1. PolyCore sees three beneficiaries if SHIM is a success. First, the IC vendor benefits by creating broader tools support for their products and hardware features. By making proprietary but XML-compliant files readily available under NDA to chosen tools vendors, competing or complementary tools improve the end programmer's choices and efficiency. This may also aid in IC design wins.

Secondly, the software tools vendors benefit by taking less time to create or port their tools to new hardware or new functions. And it should take much less effort for the tools companies to support hardware, freeing up time to improve tools and/or add creative new developer functions. Finally, the end user programmer is more productive with better (and more accurate) tools that support more complex hardware accelerators, memories or multi/many core devices.

More Details…and the Future

XML was selected because XML data binding allows generating a class library for handling SHIM XML data as data objects, not as XML elements and attributes. Additionally, the MCA states that there are "popular open source implementations of XML data binding tools." An example of hardware elements expressed in SHIM is shown in Figure 2 and a SHIM XML code is shown in Figure 3.

Since a SHIM XML file is merely an abstraction of hardware elements as viewed by software, it's not a complete description of everything the hardware is capable of. Is something lost in the translation? "Nothing is lost," says PolyCore's Brehmer, but he pauses for a moment and continues "but everytime something is an abstraction, there will be something lost." Programmers are now used to this kind of tradeoff as an acceptable way to actually get code written. Early C compilers fought the battle against assembly programmers until their accuracy and ease-of-use was proven. Similarly, argues Brehmer, the goals are getting to a solution on time (a high quality tool released coincident with the hardware launch), and "using as much of the [accelerator] 'candy' in the chip as possible."

The goal of the SHIM working group within the MCA was to have a specification available by the end of 2013, said Markus Levy of the MCA. That didn't exactly happen, but it's not unexpected given the complexity of the task and the way open standards committees work to assure all participants have a say.

As we went to press, the SHIM WG was preparing a formal update for Embedded World 2014 and we were given an exclusive copy of the slides. According to Gondo, the WG will have a draft SHIM release available at the Multicore Developer Conference in May. Additionally, the WG is working on a SHIM editor tool to support Common Configuration File (CCF) to help "ease the authoring of SHIM XMLs". The tool is intended to serve as an open source reference implementation of a SHIM API application.

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.

Telecom Service Provider Innovation Is Not an Oxymoron

The SDN and NFV initiatives are working to provide a lower cost, highly flexible foundation for the future network infrastructure that will enable more rapid innovation while decreasing CAPEX and OPEX associated with network infrastructure.

By Eric Carmès, 6WIND

Around this time every year, industry watchers and analysts write their annual "industry trends" article and conference keynote speakers provide their insights into the challenges and opportunities facing the telecom and data center networking markets. Topics such as technology transitions, costs of network build out, economic factors, government regulation, rapid market growth, emerging markets and new players with different business models are frequently mentioned. Looking back on the prior year's observations, one key challenge appears year-after- year: how to increase service providers' average revenue per user (ARPU). For years, network operators have worked to avoid being pigeonholed as simply providing "pipes" and have tried to introduce added-value services that would appeal to their customers and increase ARPU. However, given that this issue seems to be permanently at the top of the annual list of challenges, we must conclude that increasing ARPU has been a tough nut to crack.

When asked why, network operators will first highlight the fact that markets (particularly in the mobile world) are changing at an incredible pace and demanding more and more services. However, with today's network infrastructure, composed of single-function platforms based on proprietary hardware, the cost to develop, trial and deploy new services is high and the time-to-market for new services is quite long. The launch of a new service typically requires the introduction of 'yet-another-box' into the network with all of the accompanying CAPEX, OPEX and integration implications. In a non-proprietary NFV white paper authored by 13 network operators from around the world, it is stated that the current situation is "inhibiting the roll-out of new revenue-earning network services and constraining innovation in an increasingly network-centric connected world."

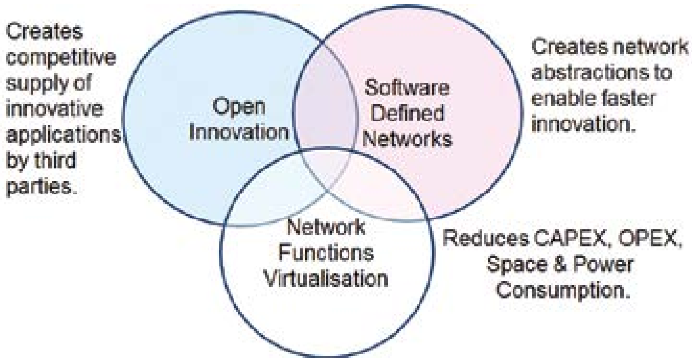

Figure 1 Network Functions Virtualization Relationship with SDN Courtesy of NFV Initiative (http://portal.etsi.org/NFV/NFV_White_Paper.pdf)

The proposed solution: move from today's "hardwired" proprietary appliance approach to a virtualized, software-driven strategy based on commodity hardware. The Software Defined Networking (SDN) and Network Functions Virtualization (NFV) initiatives are working to provide a lower cost, highly flexible foundation for the future network infrastructure that will enable more rapid innovation while at the same time decrease CAPEX and OPEX associated with network infrastructure.

The move away from proprietary, typically ASIC-based, hardware to an industry-standard platform is essential to the success of a software-driven strategy.

Standard Hardware with a Communications Twist

The move away from proprietary, typically ASIC-based, hardware to an industry-standard platform is essential to the success of a software-driven strategy. A common platform eliminates the 'box-per-application' aspect of today's network-appliance approach and allows for the consolidation of workloads onto a single platform. Today's multicore processors such as the Intel® Xeon® processor family are more than capable of handling a mix of control, data plane and application tasks.

CloudNFV™: an early multi-vendor proof-of-concept for NFV
(Participating companies: 6WIND, CIMI Corp, Dell, EnterpriseWeb, Metaswitch, Overture, Qosmos.)

NFV expands the playing field of conventional networking in a number of key ways. First, service components are individually addressed elements, which can raise the complexity of a "service" by creating many more cooperative element relationships and addresses. Second, services are likely to be more extemporaneous, so they have to be easy to build, manage, and tear down as needed.

The CloudNFV initiative is developing a proof-of-concept (PoC), which was recently approved by ETSI and is the first such NFV PoC to achieve this level of approval. It addresses these issues in 3 ways:

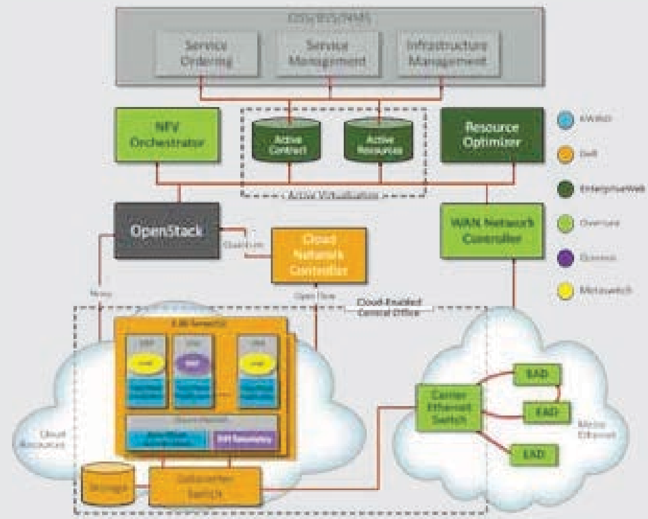

Image courtesy of CloudNFV ©2013

- CloudNFV presumes cloud-based deployment of virtual functions on an OpenStack model (other cloud models can also be supported), so that proven concepts of cloud scalability and tenant separation are incorporated into NFV without duplicating the vast body of work already in place for the cloud. The architecture is so cloud-like that CloudNFV can deploy cloud application components and built "services" that are a combination of network features and application features.

- CloudNFV uses a dynamic management model called "derived operations" that allows fully distributed management of virtual and real network resources in any combination, based on any arbitrary management interface. This means that CloudNFV can be managed by any system currently managing network services, but also that new and more convenient management models can be constructed.

- CloudNFV defines network connections, both within a virtual-function set and between users and service components, using an agile Connection Model Handler concept. With this approach, any connection relationship can be defined in a "recipe" and built using open or proprietary SDN, legacy network, or combined interfaces. For example, any OpenFlow controller can be supported with the proper handler and at the same time supplemented by other network APIs as needed to support the installed equipment.

The CloudNFV architecture is open and extensible, and CloudNFV has an integration program and public tutorials and papers to help educate partners on the process of integration. The work of CloudNFV is being fed back to both the NFV ISG and the TMF to insure conformance with appropriate standards, and also reflects the latest IETF positions on SDN and NFV-related issues.

More importantly, today's processor families frequently include accelerator and offload capabilities specifically targeted at communications applications. For example, the Intel® Xeon® Platform for Large-Scale Communications Infrastructure Systems includes the Intel® QuickAssist Technology that provides improved encryption and compression performance. This technology accelerates security, routing and other workloads such as compression and wireless 3G/4G LTE algorithm offload.

In short, today's "commodity" platform is extremely powerful and includes technologies to enhance communications performance—providing the most cost-effective platform for software-driven networks.

Roles of SDN and NFV

SDN and NFV both contribute to improving network flexibility, CAPEX and OPEX, but address different aspects of the next-generation infrastructure solutions. SDN primarily addresses the issue of network flexibility by moving totally away from traditional hardwired network solutions to a software-driven network architecture. By making each network device programmable and divorcing the control plane from the data plane, SDN allows dynamic orchestration and configuration of the network. Utilizing a centralized view of the network (instead of a node-specific view), SDN enables the network to automatically adjust its configuration to meet changing user demands.

The primary aim of NFV is to leverage standard IT virtualization technology in order to consolidate multiple network functions onto industry-standard, high-volume servers, switches and storage subsystems. By implementing network functions as virtual functions, they can be instantiated wherever and whenever necessary throughout the network and eliminate the "single function per box" aspect of traditional implementations.

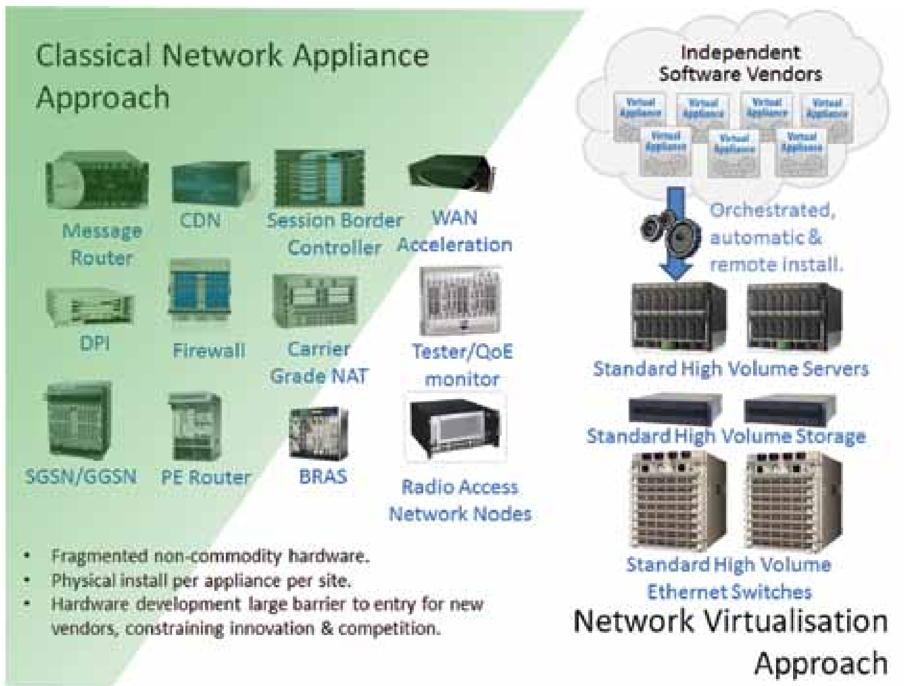

Figure 2 Classical vs Virtual Approaches to Network Appliance
Courtesy of NFV Initiative (http://portal.etsi.org/NFV/NFV_White_Paper.pdf)

NFV Background

Several large carrier organizations initiated NFV with the drafting of the NFV White Paper at the 2012 SDN and OpenFlow World Congress. The fundamental notion was to use standard IT virtualization technology to consolidate many network functions onto standard high-volume servers. NFV gained momentum very quickly, with an ETSI NFV Industry Specification Group (ISG) being formed later in 2012.

The outputs of the NFV ISG are architecture definitions and requirements for both network equipment suppliers and software vendors. Committees and Working Groups within the NFV ISG include Architecture for the Virtualization Infrastructure; Management and Orchestration; Software Architecture, Reliability and Availability; Security and Performance; and Portability.

The key objective of NFV is to achieve simplified, software-defined, virtualized service provider networks. The benefits of such an architecture fall into two categories: reducing operator CAPEX through lower equipment costs (due to the ability to consolidate functions) and reduced OPEX from the resulting savings in power consumption and facilities. The other category is increased innovation.

NFV: The Innovation Enabler

Accelerating innovation requires that the time and cost to develop, test and deploy new services be significantly reduced—and NFV contributes to both of these requirements. New services can be developed using software tools instead of having to develop new hardware. Development and testing can be done using standard platforms, without the need to install new equipment and at significantly lower risk. Similarly, the virtual nature of these services means deployment does not require lots of new boxes to be added to the network. Costs are lowered, development time is reduced and time-to-market is significantly enhanced. The overall effect is to improve ROI on new service investments.

In addition, software-driven solutions tend to enable a variety of ecosystems that can develop and market industry-specific solutions, bringing new services and additional revenue streams online quickly and at much lower cost and risk.

NFV: Improved Innovation through Disruptive Technology

NFV promises major business benefits for service providers, but also presents significant technical challenges. Virtualizing hundreds of networking functions on a single server requires high-performance network bandwidth to be distributed to the virtualized functions. Management and billing systems must be adapted to the NFV architecture. Migration strategies need to be created that support a hybrid environment of traditional network nodes and new NFV-based nodes. And lest we forget, the standards and best practices for NFV are evolving as we speak. It will take several years before the SDN and NFV-based infrastructure is fully realized.

One testament to the potential of SDN and NFV is the fact that the network operators, traditionally a very conservative group, are cooperating to invest in the evaluation and eventual deployment of what are certainly disruptive technologies. The benefits of lower costs will undoubtedly be the goals most desirable in the first years as the trial and deployments of software-driven solutions start. But the crucial goal will be enabling more rapid innovation. Who knows, in a few years maybe "improving ARPU" will fall from its long-standing spot at the top of every year's trend list!

Eric Carmès is a founder and the CEO of 6WIND. Under his leadership, 6WIND was the very first company to understand the impact of multicore architectures on networking software industry and became within three years the undisputed leader for multicore packet processing software. Prior to founding 6WIND, Eric was Head of the Network Department at Thales (previously Thomson-CSF), having also previously served at Dassault Electronique. Eric holds MS degrees from both INSA (French University for Applied Sciences) and ESE (French Electrical Engineering University).

Technology is Reshaping Vehicle Insurance

Usage-based insurance (UBI) is a largely untapped market that is ripe for technology solutions to deliver tangible benefits to consumers, the insurance industry and to society.

By Cyril Zeller, Telit Wireless Solutions

The market for usage-based insurance (UBI) is largely untapped because the way in which automobile insurance risks were assessed stood still for many years. It was based on static, statistical data like age, gender, car model and so on. Applying telematics technology to usage-based insurance is a relatively recent development and it is delivering the means for insurers to make objective assessments of risk profiles based on real-time, dynamic data like mileage, areas travelled, time of day, keeping to speed limits, engine RPM and fuel level as well as driver behavior. This driver-specific information can also be paired with publicly available data to identify road types and weather conditions.

There are two key benefits arising from this shift: for insurers, there's the ability to detect and retain the majority of the lowest risk drivers. For drivers, particularly young drivers, it's a way of getting significant discounts on their premiums. In addition, these assessment mechanisms work like psychological conditioning, where good driving behavior is rewarded and, in a way, bad behavior is punished—with higher premiums or loss of coverage. And since careful drivers have fewer accidents, this becomes a big benefit for society.

> ***Only an embedded device can guarantee accurate X=Y orientation of high g-force data, which is needed to justify insurance claims for whiplash.***

Technology Revved Up and Ready

Automotive SatNav systems employ similar real-time dynamic data and when these applications became available for download and use on smartphones they exploded in popularity. Therefore, creating UBI apps for smartphones was a logical development. Smartphones have the requisite functionality, which includes sensors to detect acceleration, braking and cornering, but insurers have usability and reliability issues similar to those from texting-and-driving.

In the U.S., insurers have concerns about the use of smartphone apps and this is reflected in the fact that solutions based on in-vehicle, on-board diagnostics (OBD) dongles have become the preferred option. These devices plug into the vehicle's OBD-II service port, similar to the one shown in Figure 1a. In the U.S., cars have been equipped with these ports since 1996. (The main reasons for their preference are covered in the following section on "Dedicated in-vehicle devices.") Danlaw's device, which is known as a DataLogger, is shown in Figure 1b. This designation indicates that UBI data is logged (stored) in the device before being transmitted over a cellular network to the insurer.

Nate Bryer, VP of innovation at automotive engineering firm Danlaw says: "Smartphones still lack the ability to provide clean, consistent and accurate driving data that can be used in usage-based insurance programs. As of today, the only way to generate and gather the data needed for a UBI program is an embedded or self-installed device that is tied into a vehicle's electronic system." However, the ability to download a UBI app and run a free trial with an insurance company is filling the awareness gap that's allowing insurers to establish a business relationship with young drivers that can be carried forward when they marry, take out mortgages or need additional insurance policies.

Smartphone Issues

For insurers, smartphone apps are tools to collect data that they can use for risk assessment. Free UBI trials allow smartphone services to be employed as a "teaser" that: (a) introduces the concept; (b) allows drivers to see their driving behavior at the end of the trial; and (c) informs them about the potential reduction in their premium if they drive carefully.

Seen in this context, smartphones are providing a valuable service, but their use in the UBI space is problematic for various reasons. For example, the device may not always be on

Figure 1a (left): Typical OBD-II Port is easily accessible under the car dashboard. Figure 1b (right): Danlaw's DataLogger 7-Series is a small, self-installed, OBD connected telematics solution.

every time the car is driven or the app may not be compatible and certified for use with the phone's operating system or platforms, both of which vary dramatically across the consumer industry. There are vocal advocates for smartphone UBI, but the insurance industry and regulators are key players and particularly concerned about operational reliability and fraud. In addition, offering a UBI policy that requires ownership of a smartphone is unlikely to pass the "fairness" test, as many drivers do not own one of these pricey devices.

As well, regulators are highly critical of the reliability of data delivery. Phones can be removed accidentally or run out of battery power. Users would need to start the UBI app manually and phones could be dropped or become airborne during an impact. Telematics behavioral data such as braking, turning and accelerating is likely to be inaccurate and unreliable because phones are rarely perfectly oriented—data must originate from sensors oriented with the vehicle's travel plane. There are a number of other issues, too, with the most significant being liability exposure.

In almost all international jurisdictions, courts will find companies liable for negligence and damages when there is a "better" solution than the smartphone (e.g., employing a data logger) that the company should have considered. This reliability concern is more critical in life-saving applications such as an emergency response request in case of a crash, but any automotive application is held liable for delivering what it promises, even if the promise is lower insurance rates. Moreover, the smartphone UBI scenario could adversely impact on the validity of fair and honest claims. Only an embedded device can guarantee accurate X=Y orientation of high g-force data, which is needed to justify insurance claims for whiplash.

When designing a smartphone policy, insurers will need to look at three unknowns. The first is the service itself and the reason why the driver will want to use the app. The second unknown concerns the drivers, since it will never be possible to control how and if they use the app. The third is the device itself, since each smartphone is equipped with different sensor grades and quality. There is nothing to legislate the driver's use of the app. The UBI app will gather second-by-second information and build a picture of the driver's behavior, but it will not match what a data logger can do. The insurer has to accept that risk and include it in the model from the start.

Dedicated In-Vehicle Devices

In the U.S., insurers' concerns about smartphones are reflected in the fact that solutions based on in-vehicle OBD-II data loggers have become the preferred option. These robust devices

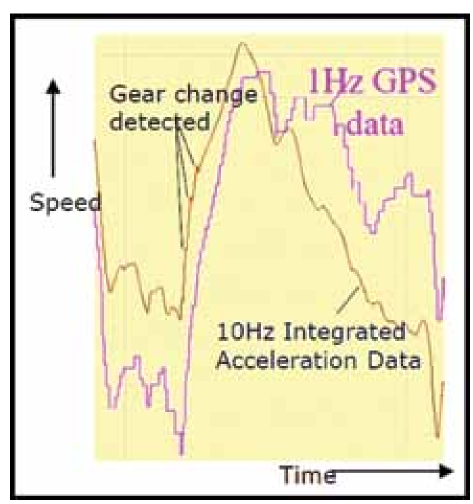

Dedicated in-vehicle devices that employ high-frequency sampling provide accurate observations of a driver's behaviour. They can identify bad behaviour such as tailgating (Figure 2a left) – see how brown line shows excessive speed changes; and detect details like gear changes (Figure 2b right). This type of detailed information is particularly useful when accidents need to be reconstructed. Purple shows data from GPS data at 1 sample per second and Brown is high-frequency sampling of speed sensor data - Graphics courtesy Redtail Telematics.

are unobtrusive and because they have a semi-permanent wired interface to the vehicle's electronic system, they provide accurate driving data. They are relatively cheap, there is no installation cost, nor is there any need to schedule an installation appointment. They can also be used in tandem with a smartphone. This hybrid solution would combine the data quality of the installed device with customer-friendly features like on-screen displays, which could include warning flags about bad driving and the possible financial impact.

UBI solutions based on OBD-II in-vehicle devices address the concerns of the insurance industry and regulators. For example:

Fairness. Regardless of vehicle type, demographics or socio-economic status, all insured drivers are measured the same way.

Reliability. A dedicated hardware solution ensures that the connectivity between the vehicle and provider is controllable and timely. And an embedded in-car system can distinguish a real crash and emergency from the bumps of normal driving (potholes, climbing curbs, speed bumps and road transitions).

If a vehicle were driven 1,000 miles a month, it would typically generate over 190K data points a year.

Security. Dedicated hardware eliminates the potential for fraud, as it is typically a proprietary system that relies on a direct connection to the insured's vehicle.

No-distraction driving. Governments around the world are grappling with the potentially disastrous consequences of cell-phone-induced distraction while driving. In 2010, this type of distraction was directly linked to over 3,000 fatal car accidents in the U.S. Insurers cringe at the thought of being dragged into court if law enforcement finds that drivers are causing accidents while interacting with their UBI smartphone apps.

Storing and Analyzing the Data

UBI can quickly accumulate massive amounts of data at the petabyte level and beyond. If a vehicle were driven 1,000 miles a month, it would typically generate over 190K data points a year. Insuring 1,000 drivers, which is a modest figure, would take this figure to over 190 million data points. You can do the math. However, data loggers can employ processing power-enabling smart filtering of the raw data. Only the relevant data is transmitted over-the-air and sorting can also be employed to ensure that the data plan stays low. Data loggers employing Telit modules already have this capability.

Nevertheless, a lot of UBI "Big Data" is going to be generated in the coming years and insurers need a way of scoring driving behavior and then applying scores to rating algorithms to reward drivers who drive safely. This is not a trivial task but it is one that the larger insurers can accommodate; others will typically partner with companies specialized in data management and predictive modelling.

Storing and analyzing Big Data is a generic issue for other M2M sectors and other industries and we are witnessing a rise in the number of innovative, cloud-based solutions. They include visual analytics, which would allow driving behavior of all insured drivers to be presented in an easy-to-understand, graphical interface.

Conclusions

It is clear that smartphone popularity and penetration will continue to rise in in the coming years with a mind-boggling number of apps that can be downloaded. However, the use of UBI smartphone apps is problematic. The regulatory climate is unfavorable and the legal risks are significant. Dedicated solutions based on proprietary, dedicated hardware that are embedded or connected to the vehicle's OBD-II diagnostic port provide robust results that can accommodate the demanding requirements of the insurance industry and the regulators. As illustrated in figure 2a and 2b, they provide very accurate driving data and the processing power of a dedicated hardware solution enables smart filtering of raw data. Also, they are as easy to install by the insured driver as downloading a UBI app into the smartphone.

Cyril Zeller is senior sales director, Global Telematics Segment at Telit Wireless Solutions. He is responsible for developing and executing Telit's corporate strategy for the worldwide telematics industry, especially in the area of fleet management, stolen vehicle recovery and usage-based insurance. Prior to joining Telit, Zeller served as the vice president of sales and marketing for Mobile Devices Ingénierie, a European-based leader in open-platform telematics technology. Zeller is an expert in the financial and legal issues in telematics.

Smart Antennas Drive Evolution in M2M and Automotive Applications

Smart antennas provide the capability to co-locate the antenna and black-box functionality for better performance and lower cost.

By Jason Furr and Vidhya Dharmarajan, Laird Technologies

The power of machine-to-machine (M2M) technology is driving towards a more connected world. All kinds of machines are getting connected to the Internet of Things to deliver the promise of enhanced productivity and to bolster the performance of businesses.

The evolving realms of technology are revolutionizing many industries. The automotive industry is trending towards intelligent cars that are fully connected in order to achieve vehicle-to-vehicle (V2V) and vehicle-to-infrastructure (V2I) communication for improved safety. Fleet management companies are utilizing M2M systems to monitor operational performance and service efficiency using real-time fleet tracking. Smart technology is enabling improved supply chain systems in industries such as vending and service.

Reliable non-stop connectivity requires that these systems support a wide array of wireless capability. This means machines are getting packed with Bluetooth, Bluetooth Low Energy (BLE) and Wi-Fi for personal or local network connection; Cellular protocols such as 2G, 3G, 4G and LTE for wide-area network connection; and global navigation satellite systems (GNSS) such as GPS, Galileo, Glonass and Beidou to support location-based services. This demand for added capability pushes toward higher system-level sophistication which can create challenges.

Challenges of Traditional M2M Systems

Traditionally, embedded M2M systems come with a black box containing connectivity features, control circuitry, power supply circuitry, wireless radios and other functions. The black box then connects to the bus system and/or a power source within the machine to allow the M2M system to run. Many connected machines in the field are metallic (examples include vehicles, vending machines, generators and containers). On these metallic structures, the antenna for the M2M system must be installed on the outside of the machine to maximize the system performance. As such, the antennas for the M2M system are in a different physical location than the black box and are connected to the black box via RF cables (as seen in Figure 1).

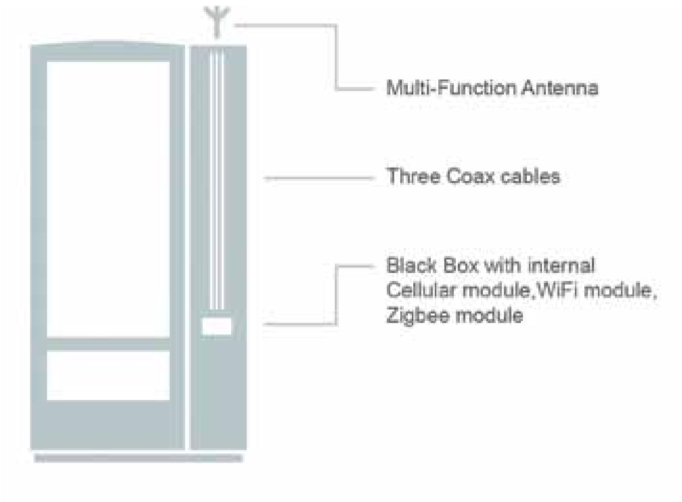

Figure 1: Vending machine in a traditional M2M system

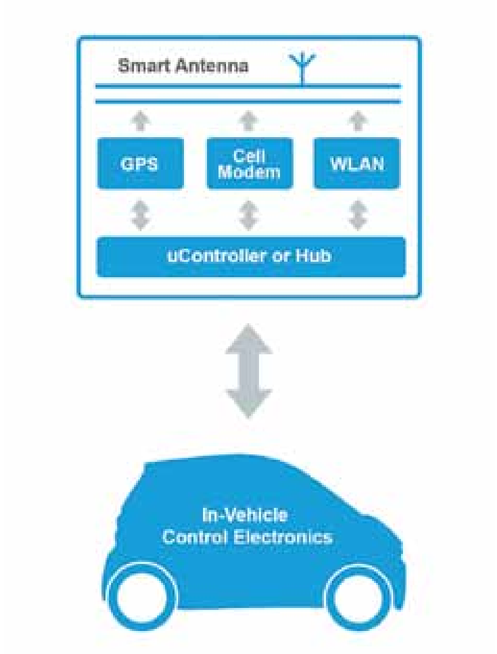

Figure 2: Smart antenna architecture for vehicle

Even though this architecture is used on many M2M systems, it poses several potential system concerns. First, RF cables inherently have loss; the longer the cable, the more antenna gain lost between the black box and the antenna negatively affecting system performance. Second, there are two cost implications. Running multiple RF cables increases the total cost of the system, especially if the distance between the black box and the antenna is large. Additionally, as more technology gets added to the M2M system, more cables need to be run, which also adds cost and complexity. Finally, with more cables and more connectors within the system there are increased installation costs and higher risk of intermittent connection failures. For example, in heavy equipment and automobiles that are prone to vibration, there could be higher instances of connector or cable damage resulting in loss of data.

By moving the wireless functionality into the antenna, the core black box no longer needs to evolve as quickly as the wireless market pushes.

Smart Antennas Maximize Efficiency, Minimize Cost

One way to reduce system cost and improve efficiency in embedded M2M applications is to re-think the architecture of an embedded M2M system. Instead of depending on the black box to perform all of the critical tasks of the system, a system designer could pull the wireless portion of the design in to the antenna. This concept can be referred to as a smart antenna design.

Smart antennas provide an option to co-locate the antenna and some of the black box device functionality (mainly radios) into the antenna package. A single digital cable can then be run from the smart antenna to the depopulated black box. See Figure 2 for an example of how a smart antenna could be architected to work in a vehicle environment.

The embedded M2M system then becomes more efficient as the loss of data between the antenna and the black box is minimized with a digital cable versus multiple coaxial cables. Digital signals can utilize error-correction software and are more resistant to electromagnetic noise and signal attenuation than analog signals. Additionally, cellular systems are evolving quickly and other wireless standards continue to get updated. Customers not only want to upgrade their systems from 2G to 3G and 4G but also want to add the most cutting-edge wireless functionality (i.e., BLE and Wi-Fi Direct). When the technology evolves, it drives hardware and software changes. These hardware and software changes then need to be re-validated and re-certified. By moving the wireless functionality into the antenna, the core black box no longer needs to evolve as quickly as the wireless market pushes. The smart antenna can truly act as the data pipe and the black box can focus on its key functionality.

Finding the Right Fit

While a number of market players provide smart antenna solutions, not all of them are successfully able to address all of the challenges. Getting the design right is the first step.

Because the smart antenna packages additional components into the antenna, the size of the smart antenna may

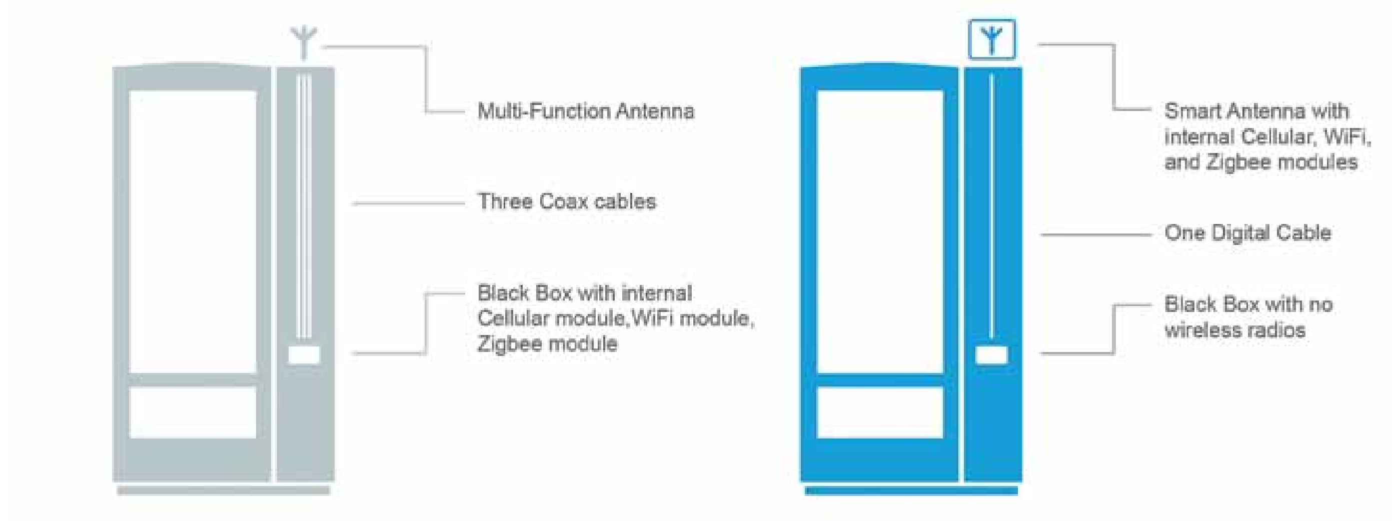

Figure 3: Example of smart antenna for vending machine

be larger than a regular antenna. Space and styling constraints on certain equipment will not allow larger smart antenna packages than typical antenna sizes. For instance, the styling of a vehicle is very important to vehicle OEMs. It is critical that the smart antenna is created with the customer's or system integrator's guidelines in mind.

Depending upon where the smart antenna is located on the subject equipment, there could be more stringent environmental challenges. Electronics that were once hidden inside of the equipment are now located outside of the equipment in the antenna package This may create additional challenges with temperature, humidity, water ingress or a variety of other environmental factors.

Moving the radios from the black box into the antenna could pose firmware challenges. Typically the black box has been seen as the brain of the system with the antenna being the afterthought. When moving some of the electronics away from the black box to the antenna, it is critical to understand how information will be passed from the smart antenna to the black box. No longer is there just an RF signal coming from the antenna. Depending on the complexity of the smart antenna design, there could be RF data, location data and machine bus information passing over a digital line. In this case, there will be some firmware required within the smart antenna. If these issues are not addressed correctly during the system design phase, persuading system integrators to adopt smart antenna architecture will be a challenge.

Although challenges exist, the adoption of smart antenna technology in to the embedded M2M system space could lead to more efficient systems for companies looking to maximize performance while controlling initial capital expense and long term system evolution costs.

Jason Furr is global sales manager for Laird, where he focuses on generating new business in the growing telematics & M2M markets. Jason holds a bachelor's degree in business from The University of Michigan-Dearborn and an MBA from Michigan State University.

Vidhya Dharmarajan is a technical writer for Laird, where she is responsible for creating, refining and maintaining product documents and writing white papers, market research and competitor analysis for the Telematics business group. Vidhya holds a master's degree in VLSI Design from San Jose State University in California.

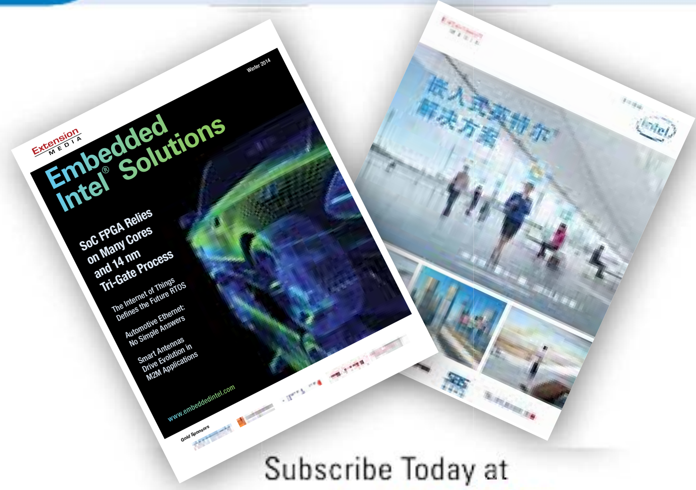

Intel's Open Pluggable Specification (OPS) Takes Hold in Digital Signage

Intel released the Open Pluggable Specification (OPS) to eliminate communication/integration problems between a host computer and its displays, with the goal of easing integration challenges.

By Weihsu Huang, Advantech Co., Ltd.

With its rising importance, digital signage has been praised as the next-generation marketing platform and a competitive tool for companies seeking to extend their global reach and to become more widespread in daily use. In retrospect, however, digital signage development has often been deemed difficult due to a lack of standard specifications. In terms of multimedia specifications, digital server synchronization or integration between multimedia player and display, system integration remains a challenging task.

As application technologies continue gradually to mature, the outlook for digital signage remains promising. However, the integration between the player and the panel has never been an easy task due to a long-standing absence of an industry-wide standard. In response, Intel released the Open Pluggable Specification (OPS) with the aim of eliminating communication/integration problems between a host computer and its display screen(s).

The OPS was developed to simplify the installation, use, maintenance and upgrading of small PCs that are installed with display panels in locations as widely varied as medical offices and fast food order counters.

The idea is simple enough—have manufacturers all develop small PCs with the same dimensions, connectors and mounting characteristics, and have display panel makers build in slots that readily accept and work with those units. It means, for example, an OPS PC that is used for a network that has NEC displays could also work in a sister network that uses Philips display panels, without any modifications.

The first salient point of the OPS is its universal hardware specification, which requires the size of the integrated device and display to be 200 x 119 x 30mm. This enables hardware to be seamlessly integrated. Secondly, the OPS standardizes signal definition. The use of HDMI and the JAE 80-pin connector, which serves as a bridge between the device and the display, facilitates communication and further simplifies device installation, usage, maintenance and upgrades. The OPS enables digital signage manufacturers to deploy interchangeable systems faster and in larger volume, all while lowering costs for development and implementation.

When standards develop and get adopted in an industry, it's a sign the sector is starting to mature, which can reduce the amount of fragmentation.

The concept has a strong list of benefits

Interoperability. Sticking to just one hardware specification for growing, evolving digital signage networks is very challenging. A network's technical team can find itself having to support several different PC set-ups and multiple display panels. It happens through merger and acquisition, consolidation and, often, for budget-trimming reasons.

With OPS, network operators can settle on the standard, and then evaluate PC and display choices based on price and performance, knowing they can switch things around or easily add new units as long as they meet the standard.

Installation and Maintenance. OPS units are slot-loaded devices that slide and lock, or snap, into place in all-in-one panel PCs. They are zero-footprint devices that are integrated within the actual LCD enclosure, meaning no external mounting kit, no cable runs and no special connectors.

All-in-one units install easily, and when on-site servicing is required, one technician just needs to be able to reach up, release the securing mechanisms and slide the PC out of its slot. Upgrades are just as easy, as long as new units meet the same OPS.

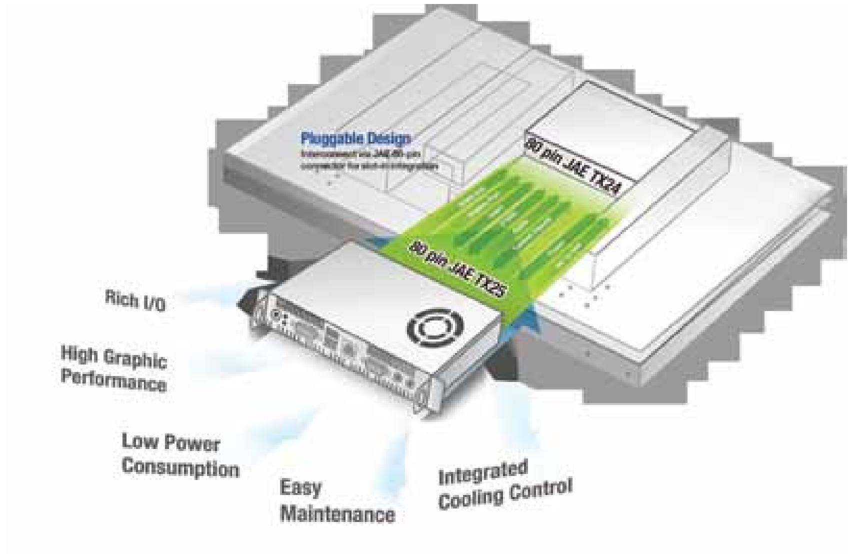

Figure 1. The use of HDMI and the JAE 80-pin connector, which serves as a bridge between the device and the display, facilitates communication and further simplifies device installation, usage, maintenance and upgrades.

There have been all-in-one PC/displays on the market for a few years, but they have custom specifications for swappable PCs or, in some cases, are PCs hard-wired into the units and serviced only by taking the whole unit down.

Costs. Standards introduce predictability, and start to remove some of the uncertainty around developing PC products that need to align with always-shifting market needs. The OPS will start lowering costs for development, and also allow for higher production volumes. Those things tend to reduce end-user costs.

Marketability. As the use of digital signage grows in the vast small-to-medium business marketplace, there is considerable demand and opportunity for products that simplify what can look quite complicated. An all-in-one display and slot-loaded PC, pre-installed with an operating system and digital signage management software, is very attractive to resellers and end users.

Some of the first OPS PCs that came on the market met the technical requirements for dimensions nicely, but didn't necessarily have the CPU and graphics processing power needed to meet the needs of all network operators.

Figure 2: Advantech's ARK-DS262 adopts the 3rd generation Intel® Core™ i7 processor, built-in advanced graphics engine, with support for HDMI output for full HD content display and 3D animation applications. The ARK-DS262 also supports one internal Mini PCIe interface for add-on functions such as wireless cards, so that material can be saved and uploaded to the central server by WLAN, without any USB disk.

The OPS specification released by Intel is an effort to help standardize the design and development of the digital signage industry, featuring a standard slim form factor, standard pluggable interface and standard signal output for easy installation, operation and maintenance.

It was not that long ago when using small form-factor PCs for digital signage reduced the physical footprint of installations but did so by sacrificing some of the performance delivered by larger PCs. Those days are now over. When standards develop and get adopted in an industry, it's a sign the sector is starting to mature, which can reduce the amount of fragmentation. OPS just makes sense, and we see its continued development as very important to the future of the digital signage business.

Weihsu Huang started her career with Advantech in 2008 where she is now the senior product manager for the Embedded Computing Group and focuses on digital signage business development and its global marketing strategy.

ADLQM87PC – 4th Generation Intel® Core™ Processor-based PCIe/104 SBC

ADLQM87PC – 4th Generation Intel® Core™ Processor-based PCIe/104 SBC

The ADLQM87PC PCIe/104 SBC delivers the industry's highest GIPS per in2 and the most densely-packed comprehensive set of features. The ADLQM87PC's 17 in2 footprint features 4th generation Intel® Core™ processors with Intel® 8-series QM87 chipset, onboard Mini PCIe/mSATA socket for on-board WiFi, GPS, or bootable flash storage, an integrated Trusted Platform Module (TPM v1.2), 2xUSB 3.0, up to 8GB 1.35V DDR3L, 4x SATA 6 Gb/s, 2x GLAN, DisplayPort, HDMI/DVI, VGA, 8x USB 2.0, etc...all with a smaller footprint and height than a basic COM-Express module and with greater temperature range and ruggedness.

The ADLQM87PC features a vastly improved HD4600 graphics engine with 2x the performance of the previous 2nd generation Intel Core processors with Intel® HD graphics. It is ideal for space constrained, rugged, or extended temperature applications where multi-core processor performance is critical. It brings unparalleled performance to countless applications such as radar and sonar processing, image signal processing, tactical command and control, surveillance and reconnaissance, transportation and railway.

The ADLQM87PC support includes 3D CAD modeling support and full system design services which include custom chassis design and system integration services. Please contact your sales representative for more information.

Application Areas

- Military and Industrial-Grade, Rugged, Extended Temperature Embedded Systems
- Unmanned Autonomous Systems of All Types Including Ground Vehicles, Aerial, Water Surface, and Underwater for Buoys, Underwater ROVs, UAV, and Others
- High-Performance Computer Vision Applications with Cameralink, Framegrabber, and GigE Interfaces
- High-Performance Mobile Computing
- Mini-Router Applications
- MIL-STD 704/1275/461 Intelligent Systems
- And Many More!

Features and Benefits

- 4th Gen Intel® Core™ i7-2715EQ and i5-4402E processors
- Intel® 8-Series QM87 chipset
- CPU TDP 25W to 47W
- Up to 8 GB DDR3L-1333/1600; 1.35V SoDIMM204 Socket
- Type 1 Bottom-Stacking PCIe/104 V2.01 with Gen2 PCIe x1 Lanes and Gen3 PEG x16
- 4x SATA 6 Gb/s with RAID 0/1/5/10, Backward Compatible
- 2x 10/100/1000 Mbit Ethernet LAN Port
- 2x RS232 COM Ports
- 13x USB 2.0 Total: 8x Onboard, 2x PCIe Connector, 1x Mini PCIe Socket, 2x USB 3.0
- 2x USB 3.0, Backward USB 2.0 Compatible
- Onboard VGA/DisplayPort, LVDS, HDMI / DVI, eDP
- Three Independent Displays
- Input Power = 5VDC, 5VS, 12VDC
- Designed for MIL-STD 810
- Compatible MIL-STD 704, MIL-STD 1275 and MIL-STD 461 Power Options Available
- Ruggedization and -40°C to +85°C Options Available

Availability

January 2014

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

ATCA-7475 40G Packet Processing Blade

The Artesyn Embedded Technologies ATCA-7475 is a 40G AdvancedTCA® (ATCA®) packet processing blade that enables the highest packet processing performance and security features in an ATCA form factor. It allows you to consolidate packet, application and control processing functions in a single blade architecture and benefit from lower development costs and the use of common tool suites. This can get you to market faster and enable you to balance workloads efficiently across the available hardware resources.

The ATCA-7475 is designed to exploit the full capabilities of the two 10-core Intel® Xeon® processors E5-2600 v2 product family and Intel® Communications Chipset 89xx Series, with an optimized balance of processing, memory, I/O, data movement and interfaces.

The fabric interface provides redundant 40G (KR4) capabilities for applications requiring the highest network throughput in the backplane, with the ability to run both fabric interfaces at full speed. The optional crypto accelerator module features two Intel® Communications Chipset 8920 devices to off-load en/decryption, compression tasks from the CPUs. Main memory configuration and mass storage options can be flexibly configured to provide a perfect fit to the needs of your application.

Multiple available rear transition modules provide a flexible combination of storage and I/O, with options for high capacity redundant storage or up to 6x10G Ethernet interfaces. This allows the ATCA-7475 to be easily integrated into different network infrastructures. The blade is optimized to work with Artesyn's range of ATCA system platforms which cover two-slot, six-slot and 14-slot variants that are designed to meet the needs of both telecom central office and network data center environments.

Key Features

- High performance Intel® architecture packet processing blade
- Two 10-core Intel® Xeon® processors, E5-2648L v2, 1.9 GHz or two 10-core Intel Xeon E5-2658 V2, 2.4 GHZ
- Up to 128GB main memory
- Redundant 40G (KR4), 10G (KR) and PICMG(R) 3.1, Option 9, 1 ATCA fabric interface
- Powerful hardware off-loading functions for en/decryption, compression, based on two Intel® Communications Chipset 8920 (optional)
- Multiple 1 and 10Gbps network and storage I/O connectivity options
- Hot-swappable hard disk on optional RTM with flexible choice of storage options and RAID 0/1 support
- Multiple software packages including operating systems
- Designed for NEBS and ETSI compliance

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

LV-67M : a mini-ITX form factor embedded system board with the Intel® QM87 chipset

Taiwan Commate Computer Inc. (brand : COMMELL), launched a mini-ITX form factor, LV-67M, embedded system board with the Intel® QM87 chipset, which adopts the mobile 4th generation Intel® Core™ processor family codename Haswell, manufactured on 22nm processor technology).

COMMELL LV-67M platform provides excellent CPU, graphics, media performance, flexibility, and enhanced security, is ideally suited to applications requiring multi-tasking capabilities, such as gaming, surveillance, medical, defense, transportation and industrial automation application.

The LV-67M Mini-ITX is designed for the 4th generation Intel® Core™ processors in the BGA1364 package supporting two DDR3L SO-DIMM up to 16GB and running 1333/1600 Mhz No ECC function. The platform is based on Intel® HD Graphics 4600 that provides high-end media and graphics capabilities for devices that display videos, 2D/3D graphics and interactive content.

In addition, LV-67M supports VGA, LVDS, DVI-I and one DP port to provide its advanced solutions for imaging, machine vision and digital signage applications. LV-67M offers lots of features including high-speed data transfer interfaces such as 6 x USB3.0 and 4 x SATAIII supported RAID 0/1/5/10, equipped with dual Gigabit Ethernet with Intel® Active Management Technology 9.0 supported, and comes with PS/2 Keyboard and Mouse port, 5 x RS232 and 1 x RS232/422/485, 4 x USB2.0, Intel® High Definition Audio, and one PCI Expressx16 slot, two PCIE Mini card socket.

COMMELL

Taiwan Commate Computer Inc.
+886-2-26963909
+886-2-26963908
http://www.commell.com.tw
e-mail : info@commell.com.tw

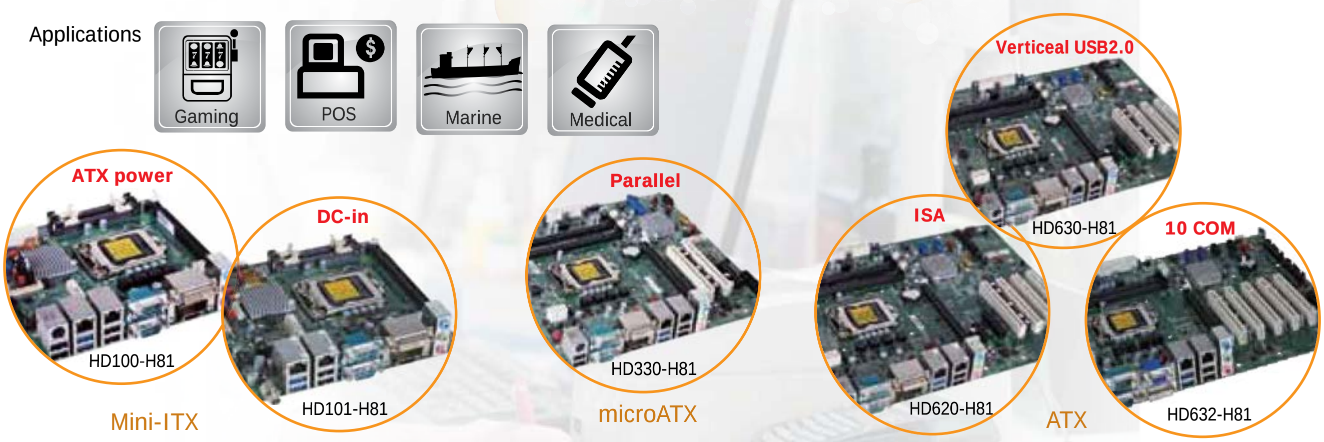

MS-9A61/62/63 15"/17"/19" HMI Panel PC w/ Touch Panel, IP65 Front Bezel

The MS-9A61/9A62/9A63 is a powerful all-in-one HMI panel PC with mainstream mobile Intel® processor. The built-in 15"/17"/19" TFT LCD with 5-wire resistive touch panel make the MS-9A61/9A62/9A63 an optimal HMI platform for automation, kiosk, POI, ATM, digital signage, etc. Designed for diverse HMI applications with high integration flexibility in mind, the MS-9A61/9A62/9A63 allows system integrators to flexibly choose from MSI's Mini-ITX embedded mainboards (under 35W models) according to different application requirements. The rich options of HDD, SSD, mSATA, and/or SATA DOM also satisfy different storage demand of high-capacity or stability. In addition to the built-in display, the 2nd display are also supported for your application demand. Moreover, the IP65 industrial standard front panel and the lockable power jack ensure the reliability of MS-9A61/9A62/9A63's operating in some critical environments.

Designed with high flexibility in mind

This panel PC leads a great concept that allows system integrators to flexibly choose from MSI's Mini-ITX embedded mainboards (under 35W models) according to different application requirements. In order to secure the power connection, the MS-9A61/9A62/9A63 provides a lockable power jack that fastens the power connection. For easy adjustment of the touch screen, the panel PC comes with a removable front frame design. And the friendly mounting design also helps to ease the effort of the engineers' maintenance.

Authentic industrial-grade with fanless design

The MS-9A61/9A62/9A63 offers a powerful computing with dual-core Intel® Atom™ Dual Core D2550 1.86GHz CPU, which wellfulfills all the HMI-related applications in different market segments. To allow the users easily see the monitor, the series uses an LCD with high contrast ratio and wide viewing angle. The 2 independent displays supported also allow users to connect more monitors to extend the application possibilities. For some specific application environments, the IP65 front panel satisfies some dust- and water-proof requirements.

Key Features

- Dual-core Intel® Atom™ Dual Core D2550 1.86GHz CPU
- 15"/17"/19" TFT LCD with 5-wire resistive touch panel built-in
- Removable front frame for quick adjustment/maintenance
- Flexible storage (HDD/SSD/mSATA/SATA DOM) options
- IP65 front panel
- Lockable power jack to secure the power adaptor
- Support DC-in 12V
- Optional stand or open-frame panel PC

MSI (Micro-Star International)
No. 69, Lide St., Zhonghe Dist.,
New Taipei City 235, Taiwan
+886 2 32345599
+886 2 22250272
ipcsales@msi.com
www.msi.com/ipc

Automotive Energy Recuperation

The increasing number of electrical components in a modern car are too much for batteries alone, so future vehicles will feature a form of built-in energy recuperation such as ultracapacitor technology.

By Chad Hall, Ioxus

Today's vehicles are no longer designed just with the intension of getting passengers from point A to B. Engineers and designers understand consumers want practical applications such as power steering, central locking and fuel economy along with being connected to their devices and having uninterrupted access to data and content. This creates a wide range of high-power applications and factors that engineers must account for, as each require low-cost and long operating-life energy-storage systems in order to achieve optimal performance.

The automotive market understands that all the electrical components in a modern car are too much for batteries alone to sustain the long lifecycle needed and they are simply not equipped to respond in time during peak points of demand. This is why all future vehicles will feature a form of energy recuperation built in. This recovery will occur through various energy-storage systems (ESS), which is a component—or network of components—that contain a blend of energy-storage technologies used to collect the energy necessary for various functions.

An ESS is an application of ultracapacitor technology, sometimes referred to as supercapacitors. Either on their own or when combined with other types of energy-storage devices, such as batteries or fuel cells, ultracapacitors meet the requirements of user power demand. Regardless of the combination, ultracapacitors expand the vehicles charge and discharge capabilities into shorter response times, extend the life of other, lower-power elements and open up new options for energy-storage applications.

With government legislation driving higher energy demands, automotive engineers will begin designing vehicles with an electric "Turbo Boost" accessory to allow for smaller engines to be used. Ultracapacitors will provide acceleration and assist the traditional turbocharging process as needed. Automotive manufacturers are concluding that using ultracapacitor technology may result in lighter-weight vehicles with lower costs for consumers due to better fuel efficiency. This combination will go a long way in helping engineers meet the challenges presented by the corporate average fuel economy (CAFE) standards, which call for automakers to lower vehicle emissions and improve fuel efficiency.

The benefits of energy recuperation extend beyond engine productivity. Hybrid capacitors, a mixture of an ultracapacitor and a lithium-ion battery, support the high-cycle, wide temperature range and high-power demands of applications like automotive subsystems. Take memory back-up for example. When a vehicle's battery is disconnected, the computer system loses power. With a hybrid capacitor, the system will retain its power and its information. It's this type of reliability and consistency engineers are looking for. Nowhere is this more crucial than when it comes to driver and passenger safety. Many of the safety features in automobiles today require a high amount of power. As such, engineers must be able to ensure that the power supply feeding these applications can be delivered at a moment's notice.

The market for ultracapacitor is expected to grow rapidly, particularly in Japan and Europe. By 2020, ultracapacitors are expected to be a $7 billion field, which represents a 56 percent compound annual growth rate. The high growth means cost of materials, manufacturing and overall pricing will fall significantly, making ultracapacitors an even more attractive energy-storage solution.

As vehicles and consumers become smarter, automotive engineers need to keep pace. All of the components, devices and connectivity embedded into our cars are for nothing if the energy needed to power them at the time of demand is lacking or non-existent. Energy recuperation via ultracapacitors offers the most reliable and greatest return on investment of any ESS available today. The key issue is to understand how to balance the energy and power demands from the ultracapacitor and the battery to support not only the vehicle itself, but the desirable features we all demand.

Chad Hall is co-founder and VP of marketing and product management for Ioxus, Inc. Find Chad on Twitter @PowerHall_UCAPS and @Ioxus.

Leading provider of Single Board Computers

4th Gen Intel® Core™ Processor-Based Mini-ITX Mainboard

LV-67M17L

- Support the **4th generation** Intel® Core™ i7-4700EQ processor (mobile)
- Intel® QM87 chipset
- DDR3 up to 16GB, LVDS, DVI , VGA, DP ports
- 4 x SATAIII, RAID 0/1/5/10, 6 x USB3.0 & 4 x USB2.0
- 2 x Intel® Gigabit Ethernet controller with Intel® Active Management Technology
- 5 x RS232 & 1 x RS232/422/485, HD Audio
- 2 x PCIE mini card, one PCIE x 16 slot
- 9~24V DC input or ATX P/S

Embedded Boards based on 4th generation Intel® Core™ processors

PICMG1.3 Full-size Computers

FS-A78

- 4th generation Intel® Core™ i7/i5/i3 processors(Desktop)
- Intel® Q87 chipset
- DDR 3 up to 16GB, VGA, DVI, LVDS, Disply port
- 2 x Giga LAN, 4 x SATAIII, HD Audio
- 1 x RS232/422/485 & 5 x RS232
- 4 x USB2.0 & 6 x USB3.0, GPIO, LPT

3.5" Form Factor Boards

LE-37C

- 4th generation Intel® Core™ i7/i5/i3 processors(Mobile)
- Intel® QM87 chipset
- DDR3 up to 8GB, VGA, DVI, LVDS, Display port
- 1 x Giga LAN, 2 x SATAIII, HD Audio
- 1 x S232/422/485 & 1 x RS232, USB 2.0 & USB 3.0
- Mini-PCIe socket, Full range DC 9V~24V input

Mini-ITX mainboard

LV-67N

- 4th generation Intel® Core™ i7/i5/i3 processors(Desktop)
- Intel® Q87 chipset
- DDR 3 up to 16GB, VGA, DVI, LVDS, Display port
- 2 x Giga LAN, 4 x SATAIII, HD Audio
- 1 x RS232/422/485 & 3 x RS232, USB 3.0 & USB 2.0
- PCIE x 16, 2 x PCIE mini card

PCIE mini card & Mini-PCI interface card

We offer a complete family of Mini-PCI and Mini-PCIe interface cards, Includes Giga LAN, RS232/422/485, IEEE 1394A/B, SATAIII, USB3.0, GPIO, Video Capture and more ...

MPX-24794S

- SPI & I2C master
- 16-bit GPIO, Counter

MPX-2515

- CAN 2.0B
- 1Mb/s operation

MPX-25858

- Video capture card
- 8-CH Video & Audio IN

MPX-210D(2)

- Intel® I210-AT
- 1 or 2 Giga Ethernet LAN

MP-6100E

- 4-CH H/W Video capture card
- H.264 H/W compression, SDK

MP-6816D8

- 8-CH S/W Video capture card
- 240fps @CIF, 120fps @D1, SDK

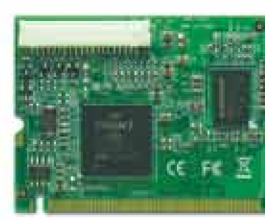

MP-Z9S

- 2D graphic card
- DVI/CRT output

COMMELL
WWW.COMMELL.COM.TW

TEL: +886-2-26963909
Email: info@commell.com.tw , sales@tcommate.com.tw
19F., No. 94, Sec. 1, Xintai 5th Rd., Xizhi Dist., New Taipei City, Taiwan

SUPERMICRO®

Building Block Solutions

7 Year Life Cycle · High-Efficiency · Compact Form Factor · High Performance

Short Depth Server Appliances

SYS-5018A-FTN4 (Front I/O)
+
A1SRi-2758F

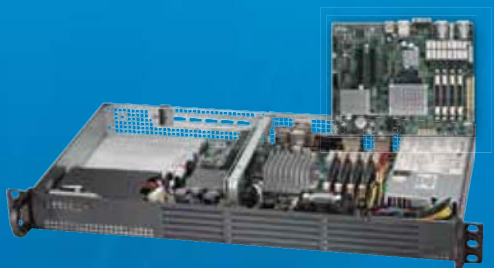

SYS-5017K-N6 (Rear I/O)
+
X9SKV-B915

High Performance Rackmount / IPC Solutions

SYS-5017C-URF (Rear I/O)
+
X9SPU-F

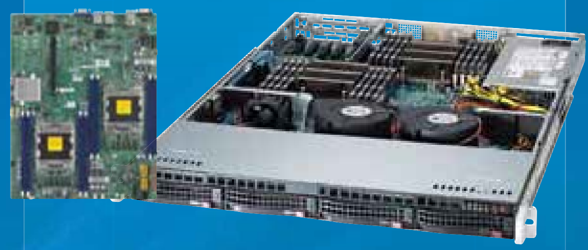

SYS-6017R-TDF+ (Rear I/O)
+
X9DRD-LF

- Standard Form Factor and High Performance Motherboards
- Optimized Short-Depth Industrial Rackmount Platforms
- Energy Efficient Platinum Level and Gold Level Power Supplies
- Fully Optimized, SuperServers Ready to Deploy Solutions
- Remote Management by IPMI
- Worldwide Service with Extended Product Life Cycle Support
- High Performance Intel® Xeon® Processor E3-1200 v2, & E5-2600 v2 product family



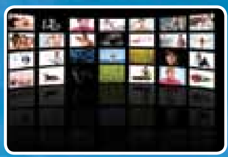


intel® inside™ XEON®

Digital! We Power the Cloud 80 PLUS PLATINUM

www.supermicro.com