Extension MEDIA

Embedded Intel® Solutions

Summer 2015

A Light Bulb Moment for Wearable Technology

IoT Brings Greenfield Opportunity

LAN Bypass for More Robust Networks

www.embeddedintel.com


SUPERMICR●
Embedded/IoT Solutions
Connecting the Intelligent World from Devices to the Cloud
Long Life Cycle · High-Efficiency · Compact Form Factor · High Performance · Global Services · IoT
IoT Gateway Solutions
Compact Embedded Server Appliance
Network, Security Appliances
High Performance / IPC Solution
E100-8Q
SYS-5028A-TN4
SYS-5018A-FTN4 (Front I/O)
SYS-6018R-TD (Rear I/O)
Cold Storage
4U Top-Loading 60-Bay Server and 90-Bay Dual Expander JBODs
SYS-5018A-AR12L
Front and Rear Views
SC946ED (shown)
SC846S
ISC High Performance
2015 International Supercomputing Conference
Frankfurt, Germany
JULY 12-16th
Come visit our booth @ #1130
• Standard Form Factor and High Performance Motherboards
• Optimized Short-Depth Industrial Rackmount Platforms
• Energy Efficient Titanium - Gold Level Power Supplies
• Fully Optimized SuperServers Ready to Deploy Solutions
• Remote Management by IPMI or Intel® AMT
• Worldwide Service with Extended Product Life Cycle Support
• Low Power Intel® Avoton, Rangeley, Quark, Core™ i7/i5/i3 and High Performance Intel® Xeon® Processors
• Optimized for Embedded Applications
intel inside™ XEON®
TITANIUM
We Power the Cloud
80 PLUS TITANIUM
Contact us at www.supermicro.com/embedded
© Super Micro Computer, Inc. Specifications subject to change without notice.
Intel, the Intel logo, Xeon, and Xeon Inside are trademarks or registered trademarks of Intel Corporation in the U.S. and/or other countries.
All other brands and names are the property of their respective owners.

Come On Over to COM!

The computer-on-module (COM) concept isn't new, yet remains relevant as processors upgrade in Tick-Tock fashion.

Chris A. Ciufo, Editor-in-Chief *Embedded Intel Solutions*

The PICMG COM Express standard is unique in that the processor is on the computer-on-module mezzanine board that plugs onto the I/O-laden baseboard. Most of time, it's the other way around.

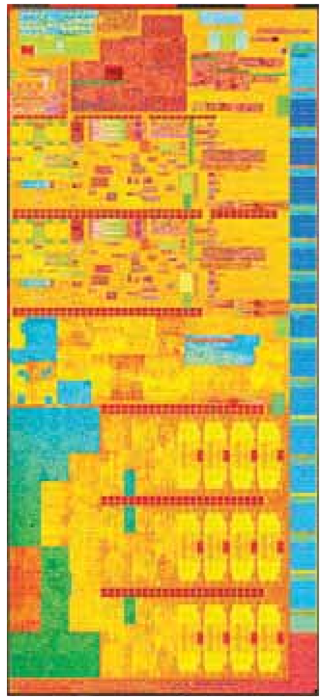

Putting the CPU on a daughter board creates a nice system upgrade path as CPU evolution moves much quicker than the I/O stuffed onto the baseboard. Intel processors, for example, follow the documented Tick-Tock cadence and change at the speed of Moore's Law or faster. The baseboard Ethernet, USB 2.0 or 3.0, A/D, or even the inter-system interconnect like PCI Express 3.0 are all pretty stable; they're not going to change from one system revision to the next.

In the last 1.5 years, Intel has gone from shipping volume orders of the Core i3/i5/i7 Fourth Generation "Haswell" CPUs (Haswell SKUs start with "4xxx") to the recently announced (as we went to press in June 2015) Core i3/i5/i7 Fifth Generation Broadwell CPUs (Broadwell SKUs start with "5xxx"; see figure.). In that timeframe, I think only USB 3.0 Type-C came to market—unless you're counting Intel's Thunderbolt plans (related to USB Type-C, case in point) and lower-power versions of Bluetooth Low Energy (BLE).

The COM Express specification COM.0 dates from 1995 but isn't the only COM concept. A spec called PrPMC (ANSI/VITA 32-2003) from VITA puts a CPU onto a PMC mezzanine module connected via PCI. Curtiss-Wright C4 Solutions, a COTS defense contractor of rugged boards and systems, puts a lower-power crypto CPU on the company's PMC-110 Secure Network mezzanine while freeing up the baseboard for its high-power/high-performance GPU processors. "This 'upside down' configuration," says Greg Sikkens, Product Manager for Graphics and Safety-Certifiable at Curtiss-Wright, "allows us to run full clocks on the [baseboard] GPU."

Another COM concept comes from General Micro Systems (GMS), one of the original suppliers of VME boards. Started in 1979, GMS was putting CPUs on mezzanine cards in the mid-1980s. Using a COM concept, its SPARC-based single-board computer (SBC) was similar to its VME SBC, which later easily evolved into its CompactPCI and VPX SBCs. The GMS COM "engine" is a 4.5 x 4.5 inch conduction-cooled mezzanine Swiss Army knife based on Intel's Core i7 "Haswell" CPU that is so thermally efficient it can operate up to +85 °C without derating the CPU due to temperature. Likely to be upgraded soon to the latest "Broadwell" CPUs, the "engine" is Intel's *ninth-generation* COM.

The point of all this is simple: sometimes designers need to ignore the conventional wisdom of putting the CPU on the baseboard, turn the concept upside down and use the COM mezzanine to power a system. COM on!

Editor's disclosure: The author used to work for companies acquired by Curtiss-Wright.

Embedded Intel® Solutions

www.embeddedintel.com

Vice President & Publisher
Clair Bright

Editorial

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

Managing Editor
Anne Fisher
afisher@extensionmedia.com

Managing Editor - China
Jane Lin-Li

Senior Editors
Caroline Hayes
David Bursky
Gabe Moretti

Creative/Production

Production Manager
Spryte Heithecker
Graphic Designers
Nicky Jacobson
Caldin Seides
Media Coordinator
Stephanie Bradbury
Senior Web Developers
Slava Dotsenko
Mariam Moattari

Advertising / Reprint Sales

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

Marketing/Circulation
Jenna Johnson

To Subscribe
www.embeddedintel.com

Extension MEDIA

Extension Media, LLC
Corporate Office

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

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

Human Resources / Administration
Darla Rovetti

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

SUMMER 2015

DEPARTMENTS

SPECIAL FEATURES

LAST WORD

PRODUCT SHOWCASE

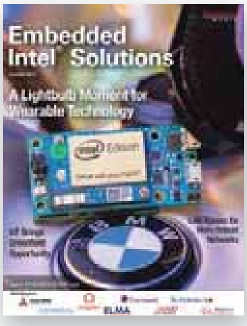

On the Cover:
Attendees at Mobile World Congress 2015 in Barcelona, Spain, experienced a BMW Motorcycle that integrates Intel's Edison microprocessor, enabling a smarter riding experience, using voice commands with the helmet.

(image and caption courtesy of Intel®)

Innovative Solutions

RTD's Embedded Systems and Enclosures

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

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

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

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

Navigating Thermal Challenges for High-End ATCA Server Blades

Reliability concerns don't allow for active CPU cooling on ATCA blades, but even as core counts rise and performance enhancements proliferate, designers are finding ways to tame high temps.

By Chris Engels, Artesyn Embedded Technologies

An ATCA server blade introduced in 2010 typically had a capacity of 64 to 96GB main memory, while today's products can support up to 512GB of memory. The 2010 blade had six cores per CPU, while the latest Intel® Xeon® Processor E5-2600 v3 product family raises the usable core count on an ATCA blade to as high as 14 cores per processor. Future processor generations will likely bring the core count even higher.

Other enhancements in that short time (relatively speaking, in terms of deployed telecom systems) include larger main CPU cache, enhanced CPU instruction set, new application accelerating functions such as encryption for security or 256 b-t integer vector instructions for image processing and signal analysis. Intel has also improved the throughput of data paths inside CPUs, between CPUs and to the main memory, added support for DDR4 memory, upgraded I/O connectivity with PCIe Gen3 and integrated the host controllers directly into the CPUs.

The only performance factor that has remained relatively stable over the various processor generations is the average clock frequency. However, performance can be optimized by clocking individual processor cores at higher frequencies, if the application can benefit from this approach. Available on ATCA server blades, these advances have increased performance per ATCA slot and optimized energy efficiency. As a result CAPEX and OPEX measured per subscriber or per network packet can be greatly improved.

Power Predicament: It Could be Worse

Power consumption of the processor and subsystems has grown in the wake of added performance. In order to compensate for the increased power usage, processor manufacturers introduce new technologies that allow for the shrinking size of silicon dies. Together with lowering of the general operating voltages applied to silicon, these techniques allow for significant reduction of power losses. Lower power and cost are mostly exploited by new classes of processors intended for mobile and appliance applications.

For processors targeted at high-performance servers, adding more capabilities gobbles up the positive effects in the next processor generation. In fact the average power consumption of the Intel Xeon processor family has increased over recent years. Power efficiency could have gotten worse. But with performance improving much faster than the average power consumption in the same time period, power efficiency as a function of performance versus power has risen.

The power increases in the last four years are in the range of 15 to 45 W per Intel Xeon processor. This doesn't sound much, but if a given cooling environment does not improve in the same timeframe, the increases can be significant.

For example, consider a scenario where compute blades are replaced in the field without upgrading the enclosure's cooling subsystem. Insufficient cooling could then result in lower achievable CPU performance. If the desire is to also exploit higher clocked CPUs, the processor wattages can grow even more. This puts a lot of pressure on board designers.

Another concern for designers can be the silicon's thermal specifications. Silicon manufacturers typically define the maximum temperature a silicon device is allowed to operate at under normal conditions. This is either specified for the surface of the device package (TCase) or for the silicon substrate (TJunction). Thermal specification of silicon devices with high power dissipation such as CPUs can be tight. Designers must pay special attention when designing the processor cooling solution. On ATCA, where active CPU cooling is not an option due to reliability concerns, cooling becomes a major design challenge.

Problems Leading to Lower Achievable Performance

Board real estate and cooling capabilities of the target enclosures set the stage for board designers. Typically,designers have to balance board functionality with achievable performance. The ideal processing board has hundreds of processor cores, runs at 5 GHz, features multiple disk drives, provides lots of default I/O and still has sufficient real estate for modular I/O extensions.

In reality, if performance is maxed out, board functionality and configuration flexibility can be limited. By the same token, adding a high degree of configuration flexibility—such as mezzanine cards or storage devices—limits the size of the cooling solution on remaining board real estate. This results in lower achievable performance.

Both the high-performance board and the multifunctional board have their advantages, but they can't co-exist satisfactorily on the same design. Board designs that ignore this fact will create problems when they are integrated into a given shelf solution. The promised performance range may then not be fully achievable.

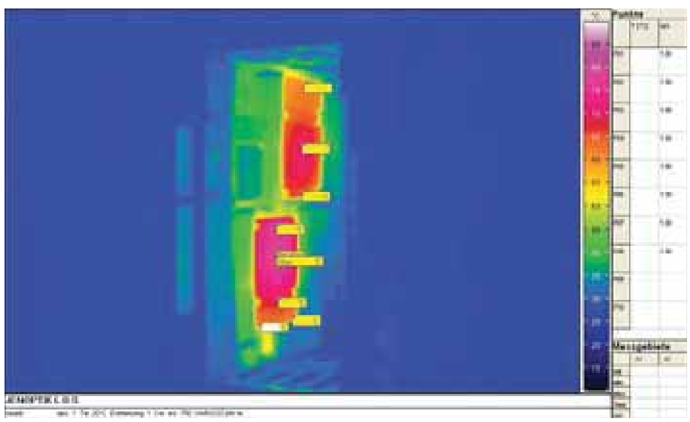

Figure 1: Artesyn ATCA-7480 thermal simulation results, showing blade surface temperature.

Processors running at full pace that are exposed to higher ambient temperature and higher software load will eventually overheat, which is not acceptable. It questions the stability of operation, reduces the component's lifetime, exceeds air exhaust temperature above the allowed limits and in the extreme can cause serious damage.

In order to prevent this, Intel CPUs typically contain a mechanism called clock throttling. The device measures the on-die temperature and, when certain limits are reached, periodically gates the processor clock frequency for some time. This allows the device to dissipate less power and stay within the thermal limits. There is a further protection available that ultimately terminates operation if even the temporary clock gating doesn't bring down the temperature.

While these are useful protection mechanisms, they are counterproductive to the overall system performance. The developer could proactively reduce clock frequency or use fewer of the available cores. The effect remains the same; the performance provided with the processor variant can only be partially exploited by the application.

Avoiding the Costs of Complete Shelf Replacement

Board products that pretend to have both a high degree of flexibility and promise high performance should be treated with caution. Integrators should carefully determine whether the promised performance is achievable in the target shelf environment under all circumstances, such as maximum ambient temperature or high software load. If processing performance is not crucial, it may be the right product. If excellent performance is paramount, it is probably the wrong product decision.

In recent years, new shelf products have become available that provide very strong cooling capabilities in order to support high performance applications. Such shelves are necessary for supporting blades with 400W or even higher power dissipation and enable the ultimate performance experience. Having said this, not every new compute blade that is delivered into communications markets will be integrated into such a high-performance shelf. In fact, there are thousands of installed shelves that have served at key locations in communications networks for many years. From an economic standpoint, it is understandable that service providers replace equipment as-needed and extend their infrastructure capabilities with the growing demand for bandwidth and service capabilities.

Upgrading individual payload blades is a much leaner approach compared to a complete shelf replacement. Shelves are therefore often kept in service for many years; they likely don't have the cooling capabilities the new shelf generation is able to provide. Such new shelves often comply with the CP-TA B.4 cooling class (now managed by PICMG), which defines airflows of up to 40 CFM per ATCA slot. The Artesyn Centellis™ 4440 is an example of this product class. As such shelves are still commonly in use, it is important that the latest server blades provide satisfactory performance in such installations.

A key aspect of designing for communication applications is designing for NEBS Level 3 compliance. The NEBS requirements cover different aspects of the design such as safety, EMC compliance, earthquake or thermal requirements. NEBS requirements define a maximum ambient air temperature of 40°C during normal operation and up to 55°C ambient air temperature during exceptional operation for a limited amount of time per year (such as during a loss of the room air conditioning). As CP-TA B.4 shelves are still commonly in use, it should be a primary goal that ATCA blade products fit in these environments while providing outstanding performance. It is also paramount that there is no degradation of performance across the entire NEBS L3 temperature range.

For example, the Artesyn ATCA-7480 packet processing blade is based on the most recent family of Intel® Xeon® Processors E5-2600 v3. It can host two processors with 75 or 105 W thermal design power. These variants have extended temperature range capability and allow a blade to be fitted into a CP-TA B.4 compliant shelf under NEBS L3 conditions.

With 12 cores per socket and 1.8 or 2.2 GHz clock frequency respectively and up to 512GB of main memory, the blade is optimized for excellent performance.

Furthermore, high-power CPUs (120W) are supported by the blade design. This means the blade can be integrated into high-performance shelves such as the Artesyn Centellis™ 8000 family, and enables data center class performance with ATCA technology. Processor derivatives with up to 14 cores per processor or up to 2.5GHz clock frequency can also be supported.

Artesyn undertook extensive thermal design and simulation efforts to enable the performance envelope for the different cooling environments. The ATCA-7480 supports the operation of the board in a maximum ambient air temperature of 55°C as defined by NEBS, when installed in a CP-TA B.4 compliant shelf. The design goals have been proven by thermal qualification during design verification and testing. The selected thermal solution makes the product sufficiently robust in shelves with CP-TA B.4 air cooling without compromising the available compute performance. It also adds sufficient headroom for squeezing the ultimate performance out of the product when installed in a shelf with enhanced airflow.

Along with the increasing performance of the Intel® Xeon® processor family comes the need for higher power provisioning and tighter thermal specifications. As a result, cooling becomes more challenging for board and system designers, particularly when products need to comply with the more stringent physical requirements of telecommunications equipment. When done right, designs can make use of the full performance capabilities under stringent thermal conditions as defined by NEBS.

Chris Engels is a senior technical marketing manager for the Embedded Computing business of Artesyn Embedded Technologies. In his current role, Engels focuses on AdvancedTCA®, server architectures and product definitions, in addition to inbound and outbound technical marketing. He also represents Artesyn in several standards development organizations.

Engels has 18 years of experience in the embedded computing industry. Prior to assuming his current role, he was responsible for the development of embedded computing and board-level products for industrial and telecommunications applications. He is a graduate of Germany's RWTH Aachen University, where he earned his diploma for electrical engineering.

LAN Bypass for More Robust Networks

Intel® processor microarchitectures are helping drive the creation of ever-more nimble networks—Now LAN bypass for circumventing harm to the vital data on those networks is also becoming more agile.

By Todd Sirois, WIN Enterprises, Inc.

Trends in Networking

Advances in networking such as the integration of workloads of different types with 10-40 Gb/s speeds, virtualization, power management of multicore CPU's, etc. are ushering in a new generation of networks. LAN Bypass isn't a tool that's much thought of when thinking about the new 'agile' networks. It is more a technique to ensure network robustness. Survivability depends on implementing the new and the tried and true.

Circumvention to the Rescue

Few things in our modern world have the unquestioned value of enterprise data. Data can be critical to the operation of the organization or even represent its lifeblood in on-line transaction processing (OLTP) systems. A fault in an in-line networking system can shut the network down and prove costly. Faults can originate from software, hardware, or power supply issues.

If the network is architected with a LAN bypass the problem can literally be circumvented. The datastream intended for the affected in-line networking node is shunted around the device until the device can be brought back up.

Implementation of LAN bypass comes in two basic approaches: integrated into the networking device or as a dedicated LAN bypass switch. Within these two types, the LAN bypass, depending on the network, will be either copper- or fiber-based. The mechanics of both are basically the same in shunting network traffic. LANs are linked into pairs taking where a LAN port takes in the network datastream and its like companion shunts it out again to the next in-line device. This ensures the preservation of important network data.

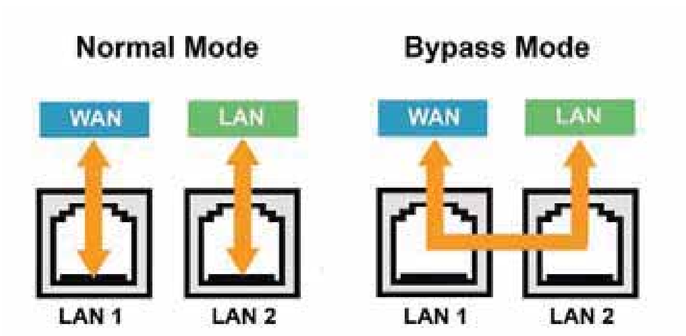

Figure 1. To configure a LAN bypass, two LAN ports are linked together during bypass mode operation in order to shunt network traffic around a faulted node and back to the network datastream and onto the next inline device. Bypass is often a feature of the network node itself or it can be added to the network as a dedicated LAN bypass device.

Another significant categorization is non-latching and latching bypass. A latching bypass is able to maintain its state when the network node's power is off, not so with non-latching bypass. This means the latching LAN bypass is not impacted by power failures, but the non-latching versions will only be helpful where the appliance is able to maintain power through battery power or some other back-up methods like an uninterruptable power supply (UPS). The non-latching method retains its value if the fault was from a software or hardware malfunction not impacting power.

Fiber-based LAN bypass systems require that power be retained to LAN bypass switches. WIN Enterprises utilizes technology from Agiltron to accomplish this (see Figure 2). Agiltron manufactures a 2x2 (i.e., two pairs) single-mode opto-mechanical fiberoptic switch with 2 simultaneously activated bypass switches in a single compact format. The device connects optical channels by redirecting incoming optical signals into selected output fibers. The mechanism uses a proprietary opto-mechanical configuration that is activated via an electrical control signal. Latching operation preserves the selected optical path after the drive signal has been removed. The design significantly reduces the moving parts of traditional switches to provide higher stability. Agiltron

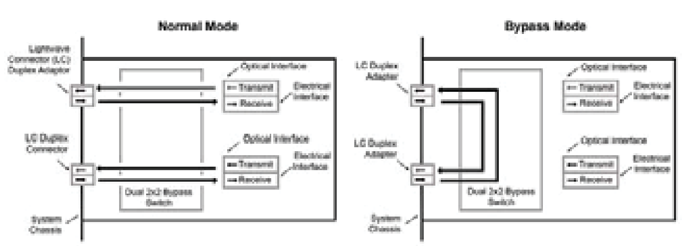

Figure 2. One approach to avoiding the moving parts associated with conventional switches is to opt for switch activation via an electrical control signal.

offers a family of bypass switch products available in multiple configurations for both single-mode and multi-mode optical fiber and latching and non-latching operation.

And soon to be on the market is a 1x or 2x dual-port fiber, single LAN bypass module. Designed expressly for Intrusion Prevention Systems (IPS), it will provide complete visibility to network traffic. It will also act as a fail-safe point should the unit lose power, the cable fail, or the application freeze.

Dedicated LAN Bypass Switches

Dedicated LAN bypass devices often support the most current 40-10 Gb/s LAN technology. These devices use a microsecond heartbeat detection system to detect a fault in an appliance. LAN bypass is implemented immediately upon this detection. When the heartbeat signal is returned again, the node is immediately brought back online.

The Overall Advantages of LAN Bypass

- Network traffic keeps flowing when the in-line appliance fails. Flow of critical data and transactions are maintained.
- If the feature is provided by a dedicated LAN bypass switch an impacted appliance can be removed or serviced without impacting network traffic.

Disadvantages

- In the case of LAN bypasses that are integrated into the appliance, the device will operate in an 'open mode' during bypass. This means any security operation the appliance normally does is suspended. This can be overcome through network design and software redundancy at critical points. The security requirements of a particular application should be considered before including a LAN Bypass feature in order to weigh any potential security versus data value trade-offs.
- Bypass capability adds cost to the solution, but will save cost in the long run by increasing network's uptime.
- A node on an in-line network represents a single point of failure. The bypass switch moves this point of failure from the appliance to the bypass switch. However, as a much simpler device, the LAN bypass is not apt to fail.

LAN Bypass Options

COTS appliance with integrated LAN bypass. Fiber-based solutions always use a LAN Bypass add-on card. Copper-based units have both on-board (i.e., the motherboard) or add-on card. Appliances can then be modified to meet an OEM's more exact specifications.

Mezzanine card with LAN bypass as an addition to an existing solution. These cards can connect to COTS boards or appliances.

Figure 3. R 316A expansion module with 2x SFP+ 10 GbE ports, non-latching bypass, Intel® Ethernet Controller XL710 10/40 GbE (Fortville). Supports the Intel® Xeon® E5-2600 processor v3 (EP) with up to 12 cores in a 1U rackmount appliance. This processor is significant because of its ability to scan smaller packet sizes to more practical security applications. The appliance features redundant power.

Figure 4. PL-10610 with Intel® Xeon® E5-2600 processor v3 (EP) demonstrates how a card can be added to an existing appliance to create a new product capable of LAN bypass. In this case, the R 316A card (Figure 3) was added to a revised PL-10590 chassis to create the PL-1061A. The high-performance device can be used for edge security with by-pass capability.

Fully custom LAN bypass device. WIN Enterprises can also work with the OEM to design and manufacture a fully custom LAN bypass device or appliance with integrated capability. WIN Enterprises welcomes the opportunity to share some of its unique expertise.

Todd Sirois has held the position of Technical Project Manager/FAE at WIN Enterprises since 2006. This position incorporates Project Management, Sales/Sales Support and Technical Concept Development. In addition, Sirois helps manage the end-of-life (EOL) process for WIN embedded products. He has been employed in the PC hardware and software industry since 1998.Previous experience includes QA for AAA software titles at the Vivendi/Universal subsidiary Papyrus Design Group, Irrational Games (now 2K Boston).

IoT Brings Greenfield Opportunity: Q&A with Shahram Mehraban, Intel

How IoT Can Link Brownfield Devices to Greenfield Opportunities

By Anne Fisher, Managing Editor

According to Shahram Mehraban, Marketing Director at Intel's Industrial & Energy Solutions Division, the Industrial IoT's health hinges on the robustness the framework put in place for interpreting data based on the use case. Expressing both enthusiasm and optimism about the changes the IoT is bringing to the Industrial and Energy sectors, Mehraban also did not gloss over the substantial challenge communicating with and securing "brown field" devices poses. Edited excerpts of our interview follow.

Embedded Intel® Solutions: How is edge-based processing in the nodes changing the Industrial IoT?

Shahram Mehraban, Intel: If you are developing a new product you have to think of your product as being part of a bigger ecosystem. It does not come down to product-specific features and functionalities any more, but [rather] what are the types of use cases that that product is going to enable down the value chain.

Security, manageability, connectivity—these are the things Intel talks about —those are big tenets of IoT and the table stakes that you need to have as part of your platform and to be able to take advantage of these use cases and to be able to create them in a scalable fashion.

Especially in the industrial and energy sectors there are billions of devices that are what we call brownfield devices that have been deployed in the field over many decades. Many of them were designed and developed over 40-50 years ago [at a time when] security wasn't even a word. If you talked about security people would probably think you meant physical security.

And now, everybody is talking about the value creation opportunities around connecting these legacy devices to the network and being able to extract data and optimize it, optimize energy consumption, do predictive maintenance—all that cool stuff.

But we have to keep in mind that these devices were designed 50 years ago with no security in mind, and as soon as they get connected to ANY type of network the threat levels are going to go higher. So you have to assure you are putting the right infrastructure around that device to make sure that that device is trusted and that the data that comes out of it is trusted, so that at the end of the day, whatever you are analyzing, you can rely on.

Figure 1. An industrial PC embedded into a wind turbine can aid with determining what repairs are required, especially helpful when visits to diagnose the problem require transporting technicians by boat or helicopter. Courtesy wikimedia.org

Embedded Intel® Solutions: What's involved in securing a brownfield device?

Mehraban, Intel: One approach is the Intel® IoT Gateway. The gateway connects that legacy device to the network, and that gateway needs to have the right types of hardware and software security elements tightly integrated so that only the right type of data is transmitted to the network. It acts as a firewall to the rest of the network, and as well that gateway is managed so that you can download new security policies, or new software, for example, to make sure it stays up to speed.

The second way is around operating that actual legacy device—still running that legacy application, but in a new environment maybe using virtualization. That relates to the introduction earlier this year of Intel® Security Critical Infrastructure Protection (CIP).

CIP uses a combination of virtualization hypervisor as well as different security elements so that you can still run that legacy application that you have in a virtual machine (VM), but all the data and traffic to that legacy application is going to a separate VM that comprises all your firewall and all your security elements.

Embedded Intel® Solutions: Are there marked differences between one company's mix of software and silicon to assure security and anothers?

Mehraban, Intel: When we speak about security from Intel's perspective, there are three elements. One, the device is secure and trustworthy itself, and there are hardware and software elements to enable that, things like trusted platform module (TPM) as well as application whitelisting to make sure that device, once in the field, only executes the application that it is meant to execute so people cannot run harmful applications on it.

Securing data on the network is another element. This involves intrusion detection and prevention and doing analytics on the data that is transmitted to the network for insights into what's happening that potentially could be harmful.

It also means being able to isolate that device if indeed it is attacked and (consequently) begins sending harmful data onto the network. And the third element is the actual overall threat level assessment from a global view.

Embedded Intel® Solutions: For the industrial IoT, do we need a combination of cloud-centric and edge-centric?

Mehraban, Intel: We believe you need a combination of cloud-based, centralized analytics and edge analytics and you need to have a framework that enables both. It really comes down to the use case that you are going to enable. [Intel] isn't saying this because we like more edge analytics. No matter where analytics happen, you need more computing performance, so it is good for all the silicon vendors.

If it is a use case where you are analyzing data, for example, on a wearable device that is connected to a factory worker that, as he is walking through a factory or maybe a refinery, that wearable device detects noxious or toxic gas and that data is sent back to a cloud to be analyzed and provide feedback, the factor of time-delay and [need for] real-time response [means] you need to do that analytics as close to that worker as possible to enable real time alerts concerning the hazardous conditions.

But then if you are a factory like Intel and you are looking at production data that is coming from factory floors across the globe, running cloud-based analytics off that data is going to give you insights that, had you only looked at one factory, you may not have seen the big picture and what the trends are.

In order for IoT to be successful you need to have an analytics framework where you can do analytics on the data based on the use case, where and when it is needed on the network—right at the edge on that gateway or back in the data center.

Embedded Intel® Solutions: What are the elements that should be in that framework?

Mehraban, Intel: For example, with regard to the IoT gateway mentioned above, one of the value propositions of an Intel IoT-based gateway is the scalability that we offer, where you can have a gateway from an Intel® Quark™ to Intel® Atom™ to Intel® Core™ and potentially an Intel® Xeon® processor-based platform, all based on the same sort of architecture running the same software stack, so then depending on what type of analytics you want to do at the edge, you can choose what kind of gateway performance you need.

If it is simple protocol translations from sensors and devices through higher-level IT protocol and maybe some data filtering, maybe an Intel Quark processor-based gateway is sufficient for your needs. But if you are looking at doing some edge analytics right on that gateway, whether it is for energy management, vibration, or other kinds of analytics, and you want to do some machine learning close to the gateway so you do not transmit all the data up to the data center over the network, then you would probably need something more high performance than the Intel Quark processor—maybe you would go to an Intel Atom- or Intel Core processor-based platform.

Embedded Intel® Solutions: How much do you think IoT innovation will depend on no single architecture dominating?

Mehraban, Intel: It comes down to what kind of end user experience are you going to enable with that scalable platform. A lot of companies don' t have the complete solution from one single OEM. And you probably don' t want to buy everything from one supplier. So you might buy your gateways from one supplier, and your servers for the data center from another supplier, but having both on a unified architecture means you know that the security elements embedded in that gateway can work hand in hand with the security intervals that are enabled in your server because they are based on the same architecture.

I think a lot of these industrial use cases are going to be very interesting. We see that there are technologies that we have developed from the consumer side from our PC and Client group over the years that have found their way into the industrial space where, for example, technologies like the Intel® RealSense™ 3D camera that we announced and that has been in a number of different tablets can find really, really unique uses in the industrial space.

Embedded Intel® Solutions: What is Intel doing to nudge that synergy along?

Mehraban, Intel: Take for example, Intel® vPro™ technology, a solution encompassing virtualization, security, as well as remote manageability, and Intel® Active Management Technology (Intel® AMT), which was originally developed in the PC side of the business targeted at large IT departments managing hundreds of thousands of PCs. We have found there is a big applicability of these technologies into the industrial space and the embedded space.

For example, we have customers who run offshore wind turbines, and sometimes these wind turbines malfunction (Figure 1). So you have to rent a helicopter or boat and do a site visit to diagnose what the problem is. It may require another trip if you find you do not have the right tools. And bad weather conditions could cause the individuals sent to assess the problem to become temporarily stranded offshore.

Intel vPro technology can fit nicely into this scenario where you can log into an industrial PC that is embedded into those wind turbines and try to diagnose what the problem is, try to solve it remotely, and if you can't, you can [nevertheless] send the right technician with the right tools to solve it and that can eliminate a lot of the operating costs of running a wind turbine.

We are working very closely with our PC groups in terms of providing Industrial and Embedded IoT-G related use cases to our product planning team so that as they design these technologies they have our customers in mind and they design technologies that are applicable to them.

Embedded Intel® Solutions: What role are standards playing in the emergence of the Industrial IoT?

Mehraban, Intel: Part of the issue with existing implementations in the industrial domain is that in the 1990's there was the battle of the fieldbuses, and everybody wanted to use Ethernet but everybody recognized that for small factory automation purposes you need real-time Ethernet and every single manufacturing facility came with their own flavor of Ethernet, be it EtherCAT or ProfiNET or ModBus, you name it.

Manufacturers who wanted to automate their factories had to pick their vendor and basically get locked into that fieldbus technology. They are learning now with IoT that it makes sense to implement a solution based on standards and open platforms so you are not tied into one single OEM or one walled garden approach and one ecosystem. The broader it is, the total cost of ownership is going to come down, and more things are going to be user friendly in terms of plugging in and integrating new solutions and upgrading and things like that.

People learn from their mistakes. I went to both the Control System Integrators Association (CSIA) event in Washington, D.C. about a month ago and also to the Automation Conference in Chicago a couple of weeks ago and among the questions fielded at both events was "how do we merge all these fieldbuses into just one common standard?" and "how does the industry work together to enable the user experience that the automation engineer is demanding?"

Embedded Intel® Solutions: Are the fieldbus wars still alive and well?

Mehraban, Intel: I think it is still a war basically between Siemens and Rockwell and other automation vendors. You still have to pick which fieldbus you want to integrate your automation technology on. There are definitely standards in the works, like time synchronized networking (TSN) (IEEE). That hopefully is going to address some of these going forward in the future, but TSN is not going to be ratified for another year or so.

Once that is ratified that is going to help. Also, some vendors are talking about software-defined automation and doing things at a more abstract level and farther away from the device itself. I think that is also going to be beneficial, but it's not going to be overnight.

Even if there were one standard today we've got billions of devices already deployed that talk "legacy fieldbus" and you still have to communicate with those and nobody is going to come and toss out the existing tools and automation solutions they have because now, all of a sudden, there is a standard that everyone agrees to.

And is this is where the opportunity is around this whole IoT gateway concept, where you can bridge those legacy systems, do the protocol translations, do the protocol conversions which are necessary and provide an abstracted, unified protocol to the upper level where you can do the data analysis, and that is what is getting everybody excited—the opportunities around that.

A Light Bulb Moment for Wearable Technology

The hardware/software Intel® Edison Board platform and the dime-sized Intel® Curie™ Module take different approaches to wearables' size, weight and power challenges.

By Caroline Hayes, Senior Editor

What do the American entrepreneur and inventor of the light bulb and the phonograph, Thomas Alva Edison and Marie Sklodowska-Curie, the Polish physicist and chemist, the first woman to win a Nobel prize have in common? Born 20 years apart, the two scientists are honored by Intel in eponymous technologies, designed to encourage creative development of wearable technology. It was back in 1991 that consumers began to take an interest in what was inside the PCs they were buying, when Intel drew their attention to which processor was used with the logo "Intel Inside." Now, 20 years on, the company is interested on being on the outside—as wearable products using the Intel® microcomputer, based on its processors, begin to grab attention.

At the 2014 Consumer Electronics Show (CES), the company announced Edison technology, a hardware/software platform, which is designed to be combined with sensors for Internet-enabled products.

Quashing the Queue

The mini-computer targets the wearable market, offering a general compute platform in a small form factor, and with built-in wireless capability.

The initial system was the same size and shape as a secure digital (SD) card, i.e. 1.260 in. x 0.945 in. x 0.083 in., or 32.0 mm x 24.0 mmx 2.1 mm). Later that year, the company switched the dual core, Intel® Quark™ processor x86, running at 400 MHz, for a 22 mm Intel® Atom™ Z34XX processor (Tangier SoC), made up of two Intel Atom processors based on Silvermont microarchitecture, running at 500 MHz, and an Intel Quark processor, running at 100 MHz.

Figure 1: The Intel Edison mini-computer is now slightly larger than an SD card, but with increased I/O, more software support, and higher computing performance as compared to the original, launched at CES 2014.

Low-power Silvermont microarchitecture features dynamic execution (also known as out-of-order execution). This allows it to process instructions as and when it is able to, rather than queuing tasks until one has been executed before beginning the next one.

This second version has components on both sides of the board, and as a result is thicker than the original. It is slightly larger all around, with dimensions of 1.398 in. x 0.984 in. x 0.154 in., or 35.5 mm x 25.0 mm x 3.9 mm.

It also increased storage, with 1GB of (Low Power Double Data Rate version 3 (LPDDR3) RAM and 4GB of embedded MultiMedia Card (eMMC) Flash memory.

Designed to Expand

Crucially for developers, Intel Edison board technology is compatible with the Arduino Integrated Development Environment (IDE) and Wolfram language, also known as Mathematica. Wolfram is a multi-paradigm programming language to build functions, graphics and structures as well as solve equations; it is installed in every Raspberry Pi, the credit-card sized Single Board Computer (SBC) developed by the UK's Raspberry Pi Foundation to promote computer science education. It also supports existing Intel® architecture (IA)-based programming tools and has a Linux OS based on Yocto, the open source for templates, tools and methodologies to create custom Linux-based systems for embedded products. The OS allows users to compile C/C++ files, or run codes such as Python or Node.js.

Retaining the Bluetooth Low Energy (LE) and dual-band 802.11a/b/g/n WiFi support, the second version of the Intel Edison mini-computer has increased I/O capability, boasting 30 interfaces via a compact, 70-pin connector, maintaining the small

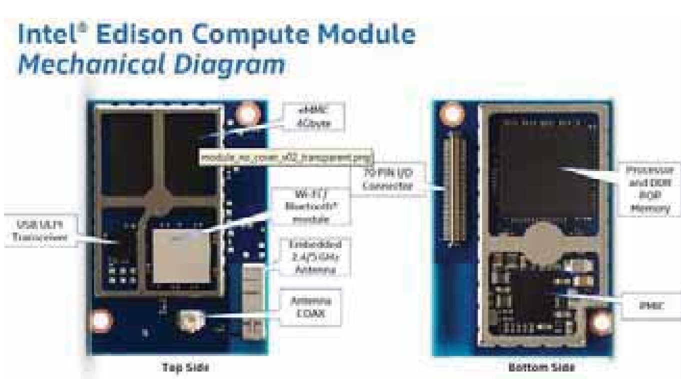

Figure 2: Using both sides of the board, the second version of Intel® Edison board technology increases compute performance and functionality options.

SPECIAL FEATURE SMALL FORM FACTORS

form factor, but increasing design flexibility. The connector contains SD, Universal Asynchronous Receiver/Transmitter (UART) Serial Peripheral Interface (SPI), General Purpose Input/Output (GPIO) and USB 2.0 On The Go (OTG) signals as well as support for Arduino add-ons.

The Intel Edison board software/hardware platform is clearly designed for development teams, using an SoC-based platform with a choice of expansion boards to develop specific, targeted applications. Intel has also released the Intel® IoT Developer Kit, which contains a development board, Yocto Linux system, Eclipse and Intel® XDK IDE, IoT Cloud Analytics and libraries to get projects started.

Created for the experimentally minded developer, the Intel Edison board offers a direct path to productization for consumer as well as light industrial IoT projects. Its connectivity options elevates it above, for example, the Intel® Galileo development board, which is described as a maker board for IoT product development, via prototype and design kits.

Wearables are Buttoned Up

What a difference a year makes. At this year's CES, the company unveiled the prototype of the Intel® Curie™ module, a button-sized microcomputer. This too, is targeted at wearable product development, but this time for design teams that may have never used silicon before, said Mike Bell, vice president, general manager, new devices group, Intel.

The circular, low power module is tiny—approximately 18mm in diameter (or 0.71 inch); about the size of a dime. Yet, its size belies its potential for extended battery operation in small products that will sense the wearer's movements or body conditions.

It is based on a 32-bit Intel® Quark™ SE processor. This is the company's first purpose-built SoC for wearable devices. Within the small form factor, it packs 80kB SRAM, 384kB Flash memory for storage, WiFi/Bluetooth LE support and a DSP sensor hub with a pattern-matching, six-axis sensor. This motion sensor combines an accelerometer and a gyroscope. The pattern matching identifies different activities, which can be used to assess fitness or health levels, in small, unobtrusive medical or fitness devices, or in an app to send data to a smartphone.

Unlike the Intel Edison board, the Intel Curie module is designed to be a standalone technology for wearable devices. It has an integrated Power Management Integrated Circuit (PMIC) that enables it to operate from a coin cell battery for long periods of time—days, even months, said CEO Brian Krzanich, during his keynote at CES 2015.

New Applications

The size of the Intel Curie module opens up possibilities for it to be integrated into small pieces of jewelry, from watches, to pendants and rings as well as smartwatches and even as buttons on clothing.

Figure 3: At CES 2015, Intel took wearable technology forward again, with a button-sized micro-computer. The Intel® Curie™ Compute Module is supported with a raft of smartphone-compatible and Cloud-capable software.

The tiny Intel Edison board is targeted at the development of wearable technology with expansion boards to add functionality for specification applications, rather than with an eye to small, near-invisible wearable technology, offered by the Intel Curie module.

There is one word of caution. What was unveiled at CES 2015 was a prototype, and the Intel Curie module has not been authorized by the rules of the Federal Communications Commission at the time of going to press. The company appears confident that these conditions will be met, as it is expected to be available in the second half of this year.

When it does ship, the module is expected to be bundled with Intel® IQ Software kits. These are readymade programs, with features for wearable applications. They include embedded software that runs on the module, algorithms, and apps for smartphone, cloud management and cloud analytics. The Body IQ Kit addresses biometric design, with functions for step-counting and recognizing the wearer's activity in real-time. The Social IQ Kit includes connectivity features.

From a focus on its processor technology inside computing devices, Intel recognizes the opportunities in wearable technology, where it's what's on the outside (of the body) that counts. In two successive years, it has introduced two form factors, for two different genres of developer of wearable technology products.

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

4G LTE Spurs Mobile Connectivity Growth with Multimedia-Rich Apps and Services

Users fell in love with using their smartphones for Internet access. Now new revenue is springing from reliable, high-quality mobile broadband, for human-to-human, human-to-machine or machine-to-machine applications.

By Charles Sturman, u-blox

The mobile communication industry originated in the late 1980s as an offshoot of the century-old telephone industry. The Global System for Mobile Communications (GSM) standard was first defined with voice as the only service and, later on, a low priority Short Message Service (SMS) was "added on," becoming the first global wireless data service. It will perhaps come as a surprise that this first mobile data service generated around 100 billion dollars in global revenues last year.

However, since then, much has changed. Voice telephony has become a commodity and is, in fact, being supplanted in society by various Internet data services, such as instant messaging and social networking. Today, the big growth is in multimedia-rich applications and services. Although, initially, operators struggled to monetize data during the days of 3G, today the value of reliable, high quality mobile broadband is seen to drive new revenue and it is why operators are rapidly rolling out 4G LTE, the fastest growing telecommunications standard ever, as shown in Figure 1.

LTE delivers a quantum leap in bandwidth and quality of service for both mobile devices, whether they're designed for human-to-human, human-to-machine or machine-to-machine applications. And affordable, high-bandwidth, high-quality and low-latency wireless connectivity is creating many new business opportunities that were unimaginable only a few years ago.

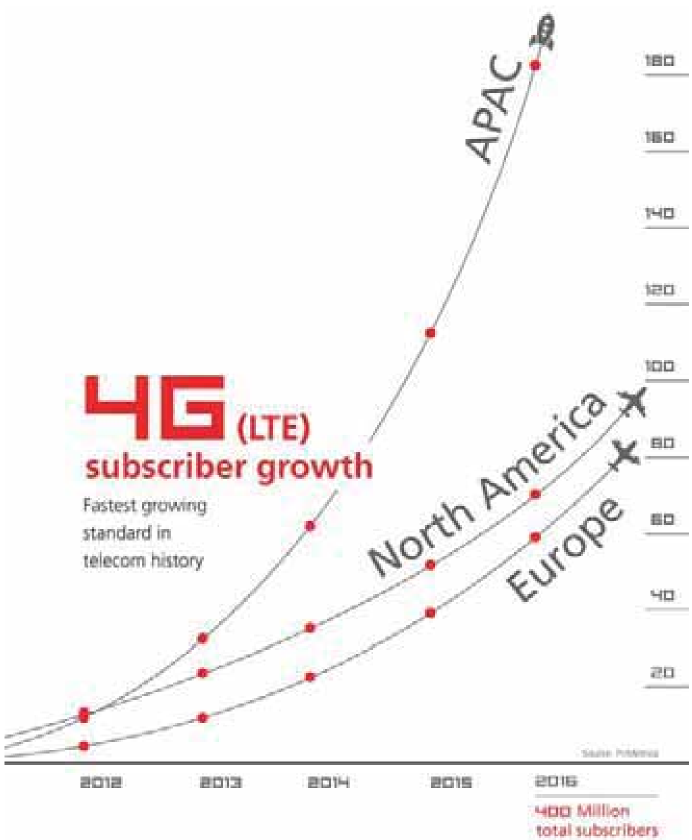

Figure 1: Data is now the primary driver of mobile connectivity, and why LTE demand has grown so fast.

Revenue Opportunities for Operators

Traditionally, operators fought for time-based voice revenues and dabbled with pay-per-use models for data traffic. However, the success of the mobile phone revolution drove average customer revenues lower in a frenzy of competition for market share.

Data services were an opportunity to reverse falling voice revenues, but operators started charging for data on a "per use" basis and produced their own branded content and services, most of which spectacularly failed to impress consumers. Then, as high-speed DSL and cable networks started to proliferate, there was an astonishing increase in Internet traffic and applications. This multimedia-rich broadband experience quickly came to play a central role in the everyday life of hundreds of millions of people in the developed world.

When Apple launched the iPhone in 2007, it was the first phone to allow full access to the Web and it kick-started the mobile Internet revolution. The first generation iPhone only supported EDGE (at 236 kb/s), but even though the bandwidth was relatively low, because it offered a full web experience —and bypassed mobile operators' walled-garden services—it was a huge success.

Suddenly the operators' sales pitch changed to "all you can eat" data subscriptions and this, together with a bewildering array of new apps becoming available, launched the smartphone revolution.

Several iPhone generations later, mobile Internet users expect an experience comparable to that of being connected to wired networks, either directly or over high-speed Wi-Fi.

Video from YouTube, IPTV, Netflix, and Hulu etc., is now the largest consumer of Internet bandwidth, followed by file sharing from the likes of Google Drive, MediaFire, FilesTube and RapidShare. Other traffic, including web surfing, VoIP, and gaming, together made up only 20 percent of total bandwidth in 2014, as demonstrated in Figure 2.

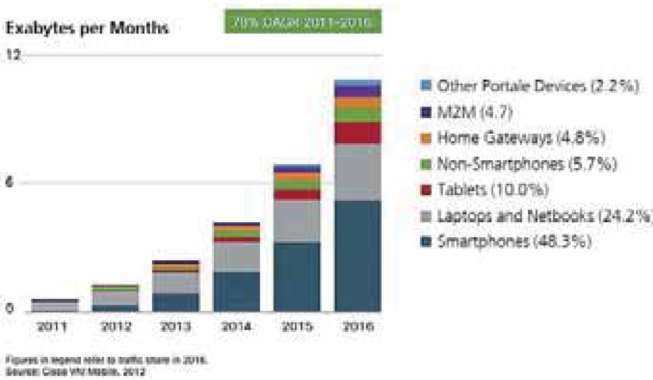

Figure 2: Internet traffic growth (2011 – 2016), Source: CISCO

2G networks are now inadequate and many are being re-farmed in favor of 4G. 3G networks struggle because of inherent protocol limitations and so operators are rapidly rolling out 4G LTE infrastructure to bandwidth-hungry subscribers.

High Bandwidth—Just the Start

LTE has four key advantages over earlier technologies:

1. Designed from scratch for high bandwidth, with up to 150 Mb/s download and 50 Mb/s upload in single carrier, or 300Mb/s + 100Mb/s in 2x20 MHz Carrier Aggregation mode. It is sufficient to stream 8 simultaneous HDTV channels—a level of demand that is rarely required in practice.

2. Bandwidth can be shared efficiently between large numbers of subscribers. This allows cost-effective data allocation to different devices and simultaneous handling of high-bandwidth human-to-human and low-bandwidth machine-to-machine traffic.

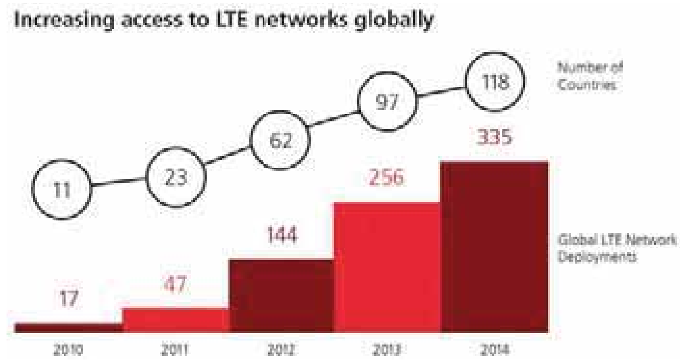

Figure 3: Global LTE network rollout, Source: GSMA

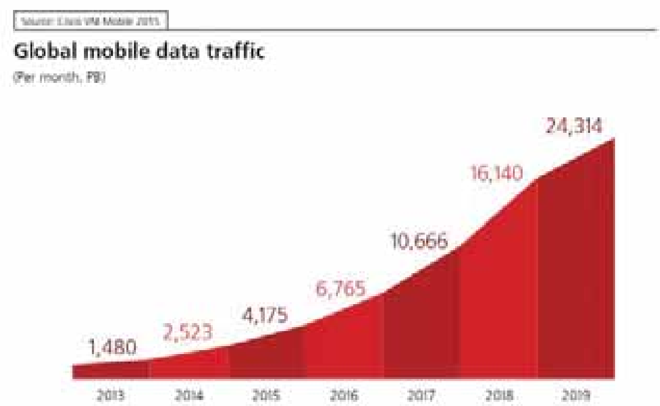

Figure 4: Global mobile data traffic, Source: CISCO VNI Mobile 2015

3. Operating Expenditure (OPEX) is reduced thanks to a simpler network architecture and improved spectral efficiency (double that of HSPA and 30X that of UMTS).

4. Real-time latency enables teleconferencing, video streaming and gaming applications. The maximum latency guaranteed over LTE is around 10 ms—a level below the threshold of human perception.

These attributes combine to bring users affordable mobile data services with high quality of service and satisfy responsiveness supporting content-rich, interactive applications such as on-demand video, social media, voice, web, gaming and cloud-based applications.

LTE Creates the Wireless "Last Mile"

LTE delivers the cable/DSL broadband experience to mobile users. As costs fall, LTE "last mile" connections will replace cable and DSL connections over ageing copper telephone wires. This will create an almost exclusively wireless environment on the periphery of the Internet.

Already, LTE is found in mobile routers that create Wi-Fi hotspots in public transport, at events, in city centers and apartment buildings. What's more, LTE is in remote locations where no other communications network infrastructure may exist.

Capitalizing on GPS and More

Although users typically access similar content when mobile or stationary, the combination of location data enables new usage models, such as location aware search and augmented reality. By exploiting an accurate GPS/GNSS (or other) positioning system, LTE is able to bring a variety of new, video-rich, location-aware services to mobile consumers. These will include information services for shopping and tourism, and multimedia navigation that allow users to see and interact with destinations before they arrive. New social media applications will appear too, and those that are based on text and images today will increasingly use video as the medium becomes more easily accessible.

Cloud Computing Meets the Mobile Terminal

Professional and consumer data storage and applications are migrating to the cloud, as demonstrated by everything from Google Drive to SalesForce and GoToMeeting.

Many consumer applications allow users to interact with each other and offload storage and data processing to the cloud. Servers can be thousands of miles (or kilometers) away, so fast access and minimal delay (latency) during both upload and download is critical to creating a response that is perceived as instantaneous. LTE is the first cellular network to achieve the required performance in this respect. As a result, cloud computing and LTE are complementary technologies, the availability of each driving the adoption of the other.

LTE Specifications for Machine-to-Machine Evolve

In contrast to 'human to human' interactions, machine communication typically involves relatively low data rates, with a strong emphasis on low power and low cost of deployment. For this reason, new LTE specifications are being created to address LTE for Machine Type Communications aka LTE-MTC. In this configuration, the new LTE Category 0 specification provides a stripped down system with data rates closer to the traditional 2G GSM technology that has been used for M2M to date. However, LTE Cat0 and the next generation Cat 00 specification will provide better performance, power consumption and cost over time whilst fitting into operators' existing LTE network infrastructure and spectrum.

Through a combination of broadband LTE and narrowband LTE-MTC, some new M2M use-cases are already emerging:

- Vending machine monitoring and delivery of video commercial displays
- Retail terminals displaying product demonstrations or live helpdesks
- Remote wireless surveillance cameras
- Telehealth terminals providing remote diagnosis and healthcare
- Multimedia signage and advertisement displays
- Facial recognition systems for homeland security
- Wireless routers for in-car infotainment systems
- Unmanned video drones used for security, traffic, accident, crowd control and fire monitoring

The guaranteed low-latency of LTE is essential for time-critical applications in industrial control, vehicle safety, traffic control and financial systems. Here, split-second reaction times are as crucial for industrial robots as for automated financial transactions. Soon, devices connected over cellular radio networks will exceed human users.

With no cabling needed, high-bandwidth enterprise-grade LTE networks can be set up quickly and easily. They can then be managed remotely with minimal hardware configuration or associated IT costs.

Despite all this new capacity and the exciting opportunities it presents, as with every previous mobile standard, traffic will grow to occupy all available LTE bandwidth, but with many new techniques under study, such as carrier aggregation, multi-antenna MIMO and beamforming this is happily some way off; as far as we can tell.

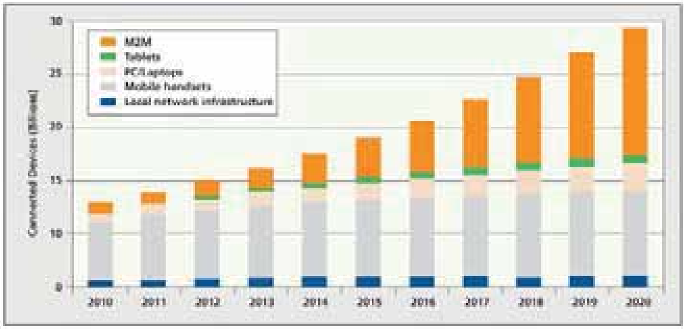

Figure 5: The number of wirelessly connected M2M devices is forecast to exceed all other types of connected devices by 2020 (Source: GSM Association)

Future-proof Connectivity

LTE is destined to be the long-term future of all cellular networks as 2G and 3G networks are gradually phased out. Many M2M operators already consider LTE as the only choice for long-term service applications and Fierce Wireless reports that there were already 497 million LTE subscribers in December 2014.

Where cellular wireless-enabled devices are used in remote locations, retrofitting modems to hundreds or thousands of units is expensive so, at least as far as possible, it's wise to design with future technologies in mind. This means either designing with LTE modem technology now, whether you need it or not, or at least future-proofing hardware design so that modem upgrades are as economical as possible.

As a major manufacturer of cellular modems, u-blox addresses the future-proofing challenge by providing its customers with PCB mounted modems based on a nested design (Figure 6). This means that 2G, 3G and 4G cellular modem modules all drop onto the same printed circuit board footprint. Customers do not need to keep changing their PCB designs whenever u-blox introduces an improved version of its modules. In addition, u-blox offers a range of pluggable modules in both mini PCIe and the newer PCIe M.2 form factors. Furthermore, portfolio-wide standard software APIs ensure upgrades and evolution with minimal effort.

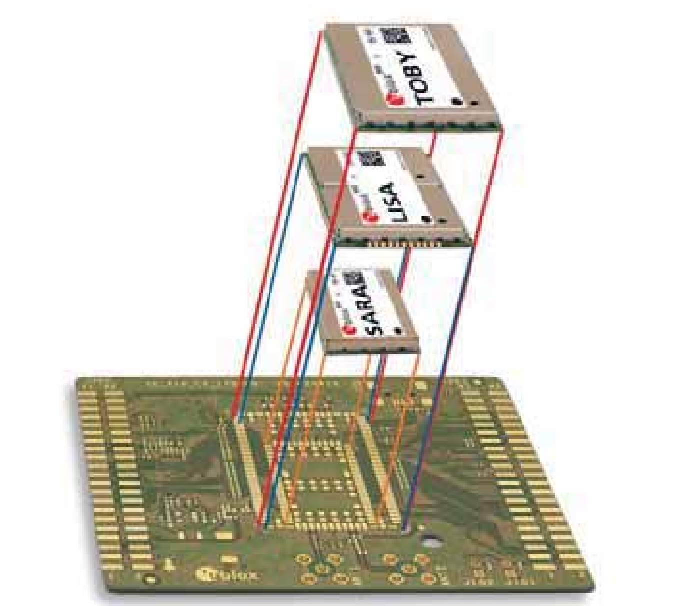

Figure 6: u-blox nested design philosophy: SARA 2G, LISA 3G and TOBY 4G module series are footprint compatible

What to Expect from the Latest 4G LTE Modem Modules

Now available are LTE modem modules designed for Cat.4 operation. (Figure 7) The LGA-format surface mount modules, of which there are several variants, measure 24.8 x 35.6 x 2.8 mm, making them very easy to install, even where space is limited. Depending on the version selected, they can be used on most major North American, European and Asian LTE networks and on some in South America.

Figure 7: TOBY-L2 series modules enable product manufacturers to quickly add reliable LTE connectivity.

What's Next

The rapid deployment of LTE infrastructure has laid the foundation for the next revolution in mobile connectivity. The ready availability of small, economical, network-certified modem modules has made it easier than ever for multiple devices to be connected to the Internet and take full advantage of the low latency and wide bandwidth now available, even when moving around the globe. New products and services are emerging that deliver tangible benefits to consumers, enterprises and public bodies as 4G becomes the norm and 5G dangles the promise of an even more exciting connected future.

Charles Sturman is senior principal product strategy, u-blox. He is leading u-blox' product marketing activities in the consumer space, responsible for product strategy and market development. Prior to this he was a founder and executive VP Sales & Marketing at Cognovo, responsible for business, product and market strategy together with sales planning and leadership.

Test Access versus Crowded PCBAs

The drive to smaller form factors, rising bus speeds and other pressures mean cramped quarters for PCBAs. And while denser and denser PCBAs can put the squeeze on test access, today there are test methodologies that break from the conventional to help maximize test coverage.

By Mark Lau, Keysight Technologies

Electrical test, unlike imaging test, requires access and physical connection to the electronic circuit's nets to provide the stimulus and measure the expected results. A test probe contacts a test pad of the net on a printed circuit board assembly (PCBA). Electrical test makes it possible to validate the electrical path of a net whether the trace and solder joint on the PCBA is visible or hidden—as for example with the ball grid array (BGA) or microBGA IC packages popular for SoCs.

Test access can be reduced as PCBA density increases to accommodate:

- Smaller (and smaller) form factors
- Less sockets as more devices are soldered down (CPU, memory)
- IC package on IC package assembly
- Increasing bus speeds
- High density interconnect (HDI) for soldered down LPDDR3 memory
- Cost pressures

To recover test access, it's now possible to use a methodology that specifies how test targets, or bead probes, can be placed directly onto copper signal traces. This approach yields superior test coverage by providing test access points virtually anywhere on a board layout. Selecting a test access technology that does not use traditional test pads offers the advantage of demanding negligible additional real estate. And the bead probes can be placed on high-speed traces without affecting signal fidelity[1]

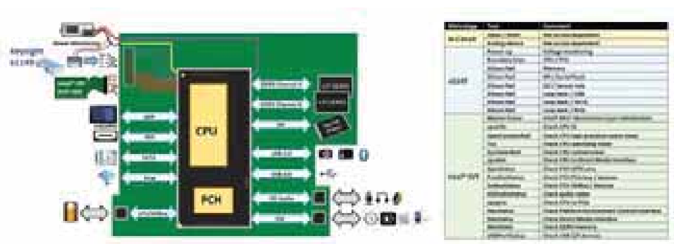

Figure 1: Block diagram of a Broadwell CPU design , based on Intel's Broadwell microarchitecture, with the deployed test strategy using Intel® Silicon View Technology (Intel® SVT). Courtesy Keysight Technologies.

Ultra Large Scale Integration Test Efficiency

IEEE 1149.1, commonly referred to as Boundary Scan, is a reliable test methodology to identify the IC package and validate the connectivity of the IC package to the PCBA. Boundary scan only requires test access to four pins of the IEEE 1149.1 test access port (TAP) regardless of the number of pins of the IC package. It is a very efficient and productive test methodology for ultra large scale integration (ULSI) devices like SoCs.

Boundary scan analyzers are available which have small footprints and maximize test coverage by using boundary scan test extensions. It's also possible to test non-boundary scan devices by using devices which are boundary scan-enabled to act as drivers/receivers, doing away with the need for physical test access or nails.

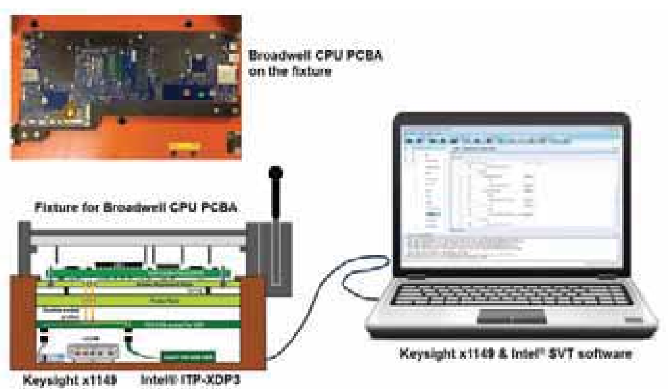

Figure 2: Small footprint solution comprised of a Broadwell CPU PCBA fixture, Intel® ITP-XDP3, x1149 Boundary Scan Analyzer and software on a Windows 7 PC.

Testing Intel® Microarchitecture-Based Processor Designs

Intel® Silicon View Technology (Intel® SVT), a proprietary test technology from Intel, enables the test for Broadwell CPU designs if test access is constrained by PCBA real estate or

[1]*Intel course ID # 9492 Manufacturing with Intel® components - Keysight® Bead Probe Technology MA*

high-speed signal fidelity. It requires test access to the debug port2 and uses Intel® DFx abstraction layer to securely access the Broadwell CPU silicon to verify its function and that of surrounding devices.

Intel SVT requires the BIOS to be set up according to Intel's BIOS Writer Guide3, reserving a designated register to host the results of the Intel SVT test. During manufacturing test, the PCBA under test has to be powered up safely to run the BIOS and Intel SVT will post the results of its test into the designated register. The content of the register is then compared to known good board values to assess whether the PCBA passes or fails the set of tests.

Intel SVT will test via the Broadwell CPU (and BIOS) the:

- Platform hub controller (PCH)
- Memory
- Graphics, e.g., VGA, HDMI, eDP
- High-speed I/O (HSIO) e.g. PCIe, SATA, USB3
- Communication interfaces, e.g., LAN, USB2
- I/O peripherals, e.g., keyboard, audio

In-Circuit Test

In-circuit test should be implemented for nets with available test access to ensure these nets are devoid of open and short and the components loaded are correct. Preferably these nets should include power and ground nets to increase the confidence of successfully sequencing power to the PCBA prior to Intel SVT test. Test-access-challenged designs should prioritize test access to the debug port, power and ground nets and critical components.

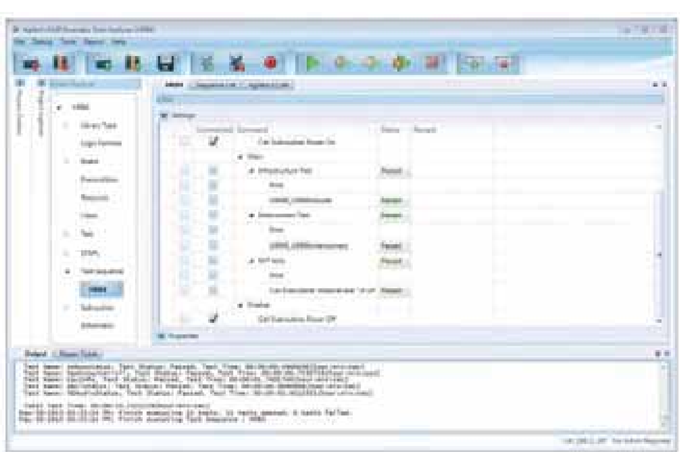

Figure 3: Screenshot of Keysight x1149 and Intel® SVT software test execution.

After completing the unpowered in-circuit test, advanced In-circuit testers (ICTs) or boundary scan analyzers can sequence the power to the Broadwell CPU PCBA and measure the on-board voltages. The ICT will shut off the power to protect the PCBA if the on-board voltages are not within the safe limits. After successfully powering up the PCBA, boundary scan tests will verify the solder joints of the CPU (and PCH in 2 chip designs). If an open or short is detected, the test will abort; otherwise it will continue to tests to validate the memory, sensor hub and high-speed bus using loop-back connectors.

At the completion of these tests, the ICT or x1149 will launch the Intel SVT tests using Intel DFx abstraction layer to verify the PCH, memory, graphics and HSIO (USB 3.0, SATA, PCIe,…) and peripherals.

At the conclusion of a successful combined ICT or x1149 with Intel SVT test, confidence will be high the CPU and peripherals on the motherboard are in good working condition, ready to be assembled into the final product. Figure 1 illustrates a Broadwell CPU design and the recommended test strategy with Intel SVT.

Mark Lau began his career as a field applications engineer and has taken on several roles in sales and business development within the electronic manufacturing test equipment business. Prior to joining Hewlett-Packard in 1993, he worked at Schlumberger Systems and GenRad.

He currently has a Product Marketing role focusing on new test technologies with Keysight's Measurement Systems Division. MSD products include electronic manufacturing test equipment like In-Circuit Tester and Functional Testers.

IoT: Definition, Standards, and Security

By Gabe Moretti, Senior Editor

Almost every day you can read something about Internet of Things (IoT). This market segment is defined as the next big opportunity for the electronics industry and thus for EDA. Yet many questions remain and some of them, if not answered correctly, will become stumbling blocks. Not surprisingly I found that there are professionals that share my concerns.

There are a number of issues to be solved in IoT, and unfortunately they will be solved piecemeal, as problems arise since humans are better at solving problems than at avoiding them. To address all of the possible issues and potential solutions the article would have turned into a book, so I decided to limit the topics discussed to just three for now.

IoT Still Needs a Full Definition

I asked Drew Wingard, CTO of Sonics if IoT was sufficiently well defined. His answer "Yes and no. If you think of IoT as an umbrella, it covers an incredibly wide variety of disparate applications. What is needed is some kind of characterization or taxonomy below the umbrella. For example, the I IoT is the Industrial Internet of things. IoT also includes wearables–medical wearables, other wearables…it's multi-dimensional." This is a great answer. it is great because it is honest, no marketing spin here. What it says to me is that at the local node level IoT is well defined. We can design and build devices that collect information, make decisions, and provide feedback in the form of controls at the local level. What is missing is the experience and thus the strongly defined architecture that defines how and where the heterogeneous data is processed into actable information. Also we need to define if and how diverse segments of IoT like wearables and automotive for example, should and how communicate and interact.

Developers of IP have of course turned their attention to the needs of IoT systems. The Tensilica group of Cadence has just announced a new product, Fusion DSP, to deal with communication and computational requirements in local IoT nodes. During the discussion they showed me how diverse IoT opportunities for IP are projected to be as shown in Figure 1.

Lauro Rizzatti, a well known verification consultant told me that " If you ask 10 people for a definition of the Internet of Things, you'll get 10 different descriptions. One thing is certain: IoT chip designs will need sophisticated verification tools to fully test the functionality." This is, unfortunately, no surprise. We are still at the point where verification is a key component of design. It is human nature to fix problems, not avoid them.

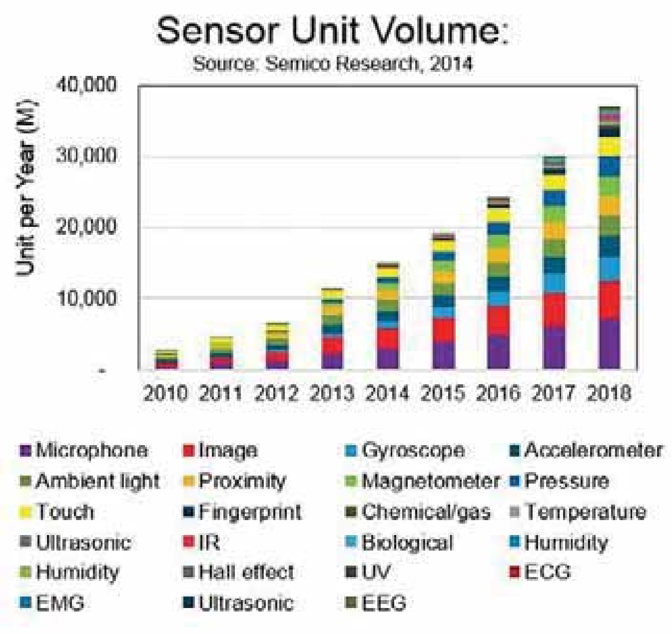

Figure 1. The various applications of IP devices in IoT systems

Omri Lachman, co-founder & CEO of Israeli wireless charging startup Humavox, points out that "IoT is a term that is defined yet still far from resonating with end users. People may be using IoT associated devices or products but for most, the term IoT means nothing. History shows that in order to bring a revolution you need to have all relevant stakeholders in line with one overall objective in mind. IoT is probably one of the biggest life changers we're going to see in the coming time. It is all about personalization and optimization of technologies/products/services for us as people. Connecting humans, homes, transportation with medical, industrial and enterprise environments is a huge objective to take on. In order for this revolution to succeed, all stakeholders should be involved in the education of the people. Visual aids should be created to help individuals of all life categories to easily connect the dots. Consumers need to be better educated about the endless opportunities that can fill their life by adopting IoT. Ultimately, the key here is better consumer education, better selling of the vision that IoT is expected to deliver and the creation of visual aids so anyone can easily grasp the concept."

Vic Kulkarni, SVP & GM, RTL Power Business, ANSYS-Apache Business Unit described the IoT architecture by breaking the system in three functional parts: Sensing and Processing, Connectivity, and Storage and Analytics. The first part must deal with MEMS and RFID issues, the second with Network, Gateway, and Supervisory Logic design and verification, and the third deals with processing at the cloud and data center level. Vic thinks that revenue from IoT will divide almost equally between consumer and industrial segments, with a slight advantage of the industrial sector (52% to 48%). In the consumer segment Vic places wearables, connected cars, and connected homes. While connected cities, healthcare, oil and gas, transportation, and the industrial internet make up the bulk of the industrial segment. ANSYS is addressing the market in its electronics and semiconductor business units by providing design and analysis tools for IC, PCB, MEMS/ Antenna, Thermal, and Physical Impact.

From what Dr. Kulkarni is saying it is clear that IoT is not just an electronics system, but an heterogeneous collection of diverse parts that must be assembled into a system in order to design, verify, and build the product. EDA already provides tools for power and signal integrity, but has either not yet addressed or not completely addressed Structural Reliability, Thermal, and Regulatory Compliance.

IoT Needs Standards

One of the most creative portion of my engineering career was spent creating standards within consortia and the IEEE. As the discussion about IoT heated up I became concerned with the absence of standards to interconnect the "things" to the conglomerating nodes and these to the cloud. And then I heard about the Open Interconnect Consortium (OIC).

International Data Corporation expects that the installed base of IoT will be approximately 212 billion "things" globally by the end of 2020. This is expected to include 30.1 billion installed "connected (autonomous)" things. Today, these devices are connecting to each other using multiple, and often incompatible approaches. The members of the Open Interconnect Consortium believe that in order to achieve this scale, the industry will need both the collaboration of the open source community and industry standards to drive interoperability of these devices.

Guy Martin of Samsung describes the purpose of the consortium this way: "There's a lot of great work going on in different areas of the IoT – you've got digital health, obviously smart home is huge, you've got in-vehicle – but there's nothing that does a really good job of connecting all of those things together. We believe that while you may have a lot of good things going on in those individual communities, the next big thing in IoT is going to be the applications that span multiple verticals. What we're really trying to develop is the framework for that."

OIC is the sponsor of the IoTivity Project, an open source software framework enabling seamless device-to-device connectivity to address the emerging needs of the IoT. The Consortium is recruiting other industry leaders to collaborate and join the efforts. The goal is to define a comprehensive communications framework to enable emerging applications in all key vertical markets. You can read more about the consortium at http://openinterconnect.org.

Ron Lowman, strategic marketing manager for IoT at Synopsys believes that standardization of communication protocols especially at the thing to local conglomerator nodes is either already here or will happen in a short time. He thinks that: "Everyone has their own definition of the concept of IoT, and the market has a lot of great semiconductor products for IoT including many microcontrollers with mixed-signal IP, such as 12-bit 5Msps ADCs, Bosch Sensortec & PNI's sensor hubs, and the Intel® Curie™ module, all of which will be used in everything from wearables, smart homes and cities, and building and factory automation. Kickstarter is a great example of where to find a sample of the limitless opportunity that IoT creates. What will actually define IoT, and what is currently missing, is the massive adoption of connected products and we're just on the brink of this larger adoption in 2015."

It is curious that the architecture uses the term "Internet" since it does not look like the Internet protocol will be used locally, like in the intelligent home and certainly not in wearables. The natural question for Ron was: "Are local protocols already standardized? If so what are they?"

Ron responded: "Wearables obviously have seen the adoption of Bluetooth Smart as a de facto standard for a couple reasons. Companies such as EMMicro have benefited from that with their low power Bluetooth capabilities. The cost of implementation including die size, stack size and power budget, is significantly better in Bluetooth Smart than WiFi and it's available on our most personal devices (mobile phones and tablets). Ethernet and WiFi protocols weren't initially designed for "things" and the protocols defining "the field bus wars," such as Modbus, were not designed to be streamed to websites, however there are a myriad of standards organizations that are tackling this problem very proactively. The important thing to note is that these standards organization's efforts will provide an open source platform and open source abstraction layer that will enable developers and designers to focus on their key value generation to the market. Interoperability will be a reality. It will not be a single solution but a small array of solutions to fit the different needs for each IoT subsegment."

IoT Needs Security

No one disputes that security is of paramount importance in IoT applications. When everything is connected the opportunities for mischievous and illegal activities are just too great. During my discussion with Vic Kulkarni he recalled how in 2008 it was shown that pacemaker devices could be hacked at a range of a few centimeters, that is less than one foot, but recently MIT graduate students hacked a pacemaker device at a range of 1,524 centimeters, or approximately 50 feet. Such capability enables electronic murder perpetrated by a totally anonymous killer.

Two of the most obvious reasons for hacking are: collection of information and illegal control of functionality. Vic also provided information on automobile vulnerability both to the control of an individual vehicle and to car-to-car communication for collision avoidance function.

Jason Oberg, CEO at Tortuga Logic, observes that: "With the advent of IoT, we are going to see a drastic shift in the security landscape. Attacks have already been demonstrated on embedded devices such as pace makers, automobiles, baby monitors, and even refrigerators. Most companies are trying to solve this problem purely with software security, but this is a constant cat-and-mouse game we cannot win. As IoT grows, we are seeing more software being pushed down into hardware and our modern chipsets are growing in complexity. This is driving attackers to begin focusing on hardware and, without ensuring our chipsets are built in a secure manner, these attackers will continue to succeed."

When thinking about security I generally think about software based hacking, but breaches that use physical techniques are just as damaging. The Athena Group, Inc., a provider of security, cryptography, anti-tamper, and signal processing IP cores, has introduced a comprehensive portfolio of IP cores with side-channel attack (SCA) countermeasures, based on advanced differential power analysis (DPA) countermeasure approaches pioneered by the Cryptography Research Division of Rambus.

DPA is a type of SCA that involves monitoring variations in the electrical power consumption or electromagnetic emissions from a target device. DPA attacks are non-invasive, easily automated, and can be mounted without knowing the design of the target device. Unlike invasive tampering, electromagnetic attacks can even be performed at a distance. As an example, attacks on cell phones have been demonstrated at a range of 30 feet. DPA countermeasures are essential to protect devices that use cryptographic keys, especially sensitive defense applications that require strong anti-tamper protection of advanced electronics and commercial devices that perform high-value processing, including mobile devices and IoT endpoints.

Although I am not privy to any official information from government agencies I can develop an example of security treats from published articles, both in print and on the net. The network of cell phones is a good example of a candidate IoT. If one wants to gather information on the location and use of individual cell phones and the relationship between and among two or more such devices it can use the cell phone networks. My cell phone, for example, gathers environmental information, location, and behavioral profile as I go about my daily activities. It also records and submits to my service provider who I call, how long I talk, what data I download, what pictures I upload, and so on. Without security such information is available to any one capable and willing to build and use a tracking system to collect and analyze all that data. Can my cell phone be disabled remotely? Can an app be installed on it without my knowledge? The answer is yes for both questions.

Conclusion

In spite of what some editors and analysts have written, there is not a clear, generally shared definition of IoT that can be used as a base for architectural design at all hierarchical levels of IoT. So in this article I chose to write about IoT definition as well as two issues that are not much talked about: standards, and security. Obviously there is much more to say about IoT, and I am grateful to all those who have sent a large volume of input for this article. What I have learned I will not keep for myself and I will share more information about IoT in the near future.

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

No Resting On Laurels for ATCA: Q&A with PICMG's Joe Pavlat

Sometimes more is more. The efforts PICMG is making on behalf of ATCA and other embedded standards include more board real estate (PICMG spec 3.7), more speed (100 Gbps simulations proceed apace) and even more confidence in the part standards organizations play in creating meaningful specifications.

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

Editor's note: It was after an informative interview with my friend Joe Pavlat, president and chairman of one of the primary open standards organizations, PICMG, that I wrote PICMG Picks Up Steam—All the Way to the Red Planet *to catch readers up on the evolution of the COM Express, CompactPCI and ATCA specs. The backdrop to my most recent talk with Joe includes—among much else—the Internet's march toward heterogeneous local intelligence, faster Ethernet, and attempts to topple Intel server space hegemony. Edited excerpts follow:*

Chris (C2) Ciufo: Joe, what's the latest on faster Ethernet (100 Gbps), ATCA and COM Express?

Joe Pavlat: Work on speeding up AdvancedTCA to provide 100G Ethernet signaling is moving forwards and the spec is due to be released by the end of 2015. The simulations are going well. PICMG has long believed multi-vendor interoperability is important, whereby a user can buy different bits and pieces of a system from different vendors and when put together, everything works. The IEEE, which controls the Ethernet standard, didn't provide any guidance on allocating budgets for things like cross talk and insertion losses to the various boards, backplanes and the connectors the signals must travel through. We are paying careful attention to this and think we have it right.

PICMG also recently released PICMG 3.7, or AdvancedTCA Extensions. This spec is a major enhancement to ATCA. It fully defines double-wide boards capable of dissipating up to 800W each. It also defines double-deep systems where two systems are interconnected in a front-to-back fashion. This dramatically increases board real estate and minimizes the amount of valuable vertical real estate in a rack. PICMG 3.7 also provides a great deal of thermal design information and is, I believe, the best thermal design guidebook for rack-based systems there is.

Not a lot is going on development-wise in the very small form factors like COM Express, although at last year's Embedded World trade show in Nuremburg the general consensus of the manufacturers I talked to is that about 2 million COM Express units are shipping per year. That number is expected to increase. The increasingly popular CompactPCI Serial standard is just about to be released in a new version that provides more flexibility in implementing board-to-board Ethernet capability without the need for a switch. It also better defines rear panel I/O, which is important for critical conduction-cooled applications.

C2: We can't go too far these days without reading about Network Functions Virtualization (NFV) and Software Defined Networking (SDN), whether from an Open Plane (S/W) standpoint or the evolution of the Internet, or ways in which customers are trying to break the lock of the equipment providers. Give us your opinion on these.

Pavlat: Software Defined Networking and Network Functions Virtualization are important tools that can allow customers to rapidly add features to their networks or to reconfigure them without upgrading a lot of hardware. These tools are useful in some places and not in others. There are really two networking worlds out there, and they are very different. The Communication Service Providers (CSPs) like Verizon, AT&T, Comcast, etc., have a very different set of requirements than the Enterprise Service Providers (ESPs) such as Amazon, Facebook and Rackspace. The CSPs must provide data across their large networks with little latency—think mobile video—and they have reliability requirements that are much higher than the ESPs. When you post a message to Facebook, it doesn't really matter if it takes 1 millisecond or a hundred to transfer a packet of information. In the CSP world, where voice and video must be delivered seamlessly, latency and reliability are critically important. For some functions like 911 services, reliability and "up time" are mandated by law. It is a common misconception that high-reliability telecom networks can be built with cheap commodity servers like the type the ESPs use.

In fact, Verizon and AT&T claim that only 10% of their future equipment requirements can be met with commodity "pizza boxes."

C2: A presentation from the Open Server Summit openly predicts that standards such as ATCA are on their way out: too expensive, too much/too late, and just not fast enough technological evolution to keep pace with the market. Your thoughts?

Pavlat: That presentation's conclusion notes the need for an ecosystem of COTS hardware tuned to meet the needs of the communications service providers (rather than enterprise) [See above]. PICMG and other standards organizations play a key role in making sure that multivendor interoperability will be maintained in this emerging ecosystem.

ATCA is a long way from being on the way out. It remains the most popular open hardware platform used in telecom today. What is changing is the software that runs on it. SDN and NFV allow more flexibility without replacing or adding hardware. PICMG began work about two years ago to define a next-generation platform that was bigger, faster, more extensible and generally meaner than ATCA. We talked to a number of CSPs, and they indicated their highest priority now isn't new hardware per se, but rather increasing their networks throughput. Video, which is expected to be over 80% of telecom network traffic in a few years, is a data hog when compared to simple voice communications. Verizon told us that it needs to increase its network throughput 60% per year for the foreseeable future. That is a staggering 10x increase in just five years and 100x in 10 years. So PICMG is focusing on improving the existing platforms like ATCA through technologies like video acceleration and 100G data transfer. At 40G, the current top speed specified for ATCA, a full mesh, 16-slot system is capable of handling just over 10 terabits of information per second. As I mentioned earlier, high availability and high system reliability are proven features of ATCA, and there aren't any other open standards out there that can provide that.

C2: I've recently noted how the Internet's evolving from big core/central servers to disaggregated and heterogeneous local intelligence and processing—as with Intelligent Gateways (Intel' s term). The Internet's data-packet centric architecture's changing. What are the ramifications to standards groups like PICMG?

Pavlat: The Internet is a big place with many technologies and players. The trends that you mention are only a few of the explorations that are going on as operators strive to reduce expenses and provide new revenue streams and services.

The need for standards organizations does not diminish with the acceleration of innovation. To the contrary, industry specifications promote multivendor interoperability and help prevent ecosystem fragmentation. PICMG's role as a standards organization is to bring industry players together in a safe harbor where meaningful specifications are developed.

C2: Is PICMG's day in the sun over?

Pavlat: Quite the opposite, in fact. PICMG has been active and producing specifications wanted by our members for over 20 years. Unlike many of the "one trick pony" standards organizations that pop up, release one or two things, and then disappear, PICMG is stable. Our processes are well proven and our intellectual property policies are strong. It is true we are putting most of our efforts on new revisions of existing standards right now and not working much on very new and exciting stuff. That's a cycle that changes every few years.

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.

Big Data Delivers on its Promises when Embedded Designers are IoT Enablers

If the IoT is to be more than just a concept and yield benefits to industrial, medical and consumer and other applications with actionable information, embedded designers will need to wield devices and systems that simplify the acquisition and management of remote data.

By Elizabeth Campbell, ADLINK Technology

The Internet of Things (IoT) isn't in itself a tangible thing. It is a concept—a concept that helps to define a complicated system of networked, distributed, communicating embedded "devices" that provide important data sets (big data) to the cloud for access from anywhere (with an Internet connection) and for integration into back-end systems in order to be tracked and analyzed and turned into actionable information that leads to better business processes (lower costs) and improved customer journeys (increased revenue).

Big data can offer detailed demographics and interaction patterns for better marketing, individual buying behaviors for improved sales, network usage patterns of bandwidth and storage for optimized information technology (IT) planning, inventory and maintenance requirements for more efficient logistics, and patterns of revenue fluctuations for more accurate financial projections; and big data can also reveal how all of these issues affect the overall company for better operations.

Analysis to Action

The problem is that getting from Point A to Point Z in an IoT deployment takes a village. Embedded designers can build it, but the big data manipulations may not come. And without a commitment to long-term data analysis and turning that analysis into action, then the IoT is really just a concept.

The IoT ecosystem consists of myriad role players and is growing every day. To make the IoT a reality, all of these role players within the ecosystem must work together to create standards and build solutions that enable businesses to more easily get from Point A to Point Z. As industrial technology vendors, our focus is to provide the enablers for the delivery and access of big data, while other role players focus on how an individual company can use that big data to make business improvements.

IoT Technology Enablers

When the IoT first became a buzzword, there was much talk of it being machine-to-machine (M2M) 2.0. But there is an enormous expansion of application scope and anticipated benefits when moving from M2M to the IoT.

M2M communication refers to the remote device access inherent in IoT deployments. Traditional M2M solutions typically relied on point-to-point communication using embedded hardware modules and either cellular or wired networks, but that has evolved over the years to include the IP-based network communication more commonly used to interface device data to a cloud or middleware platform. However, M2M is really just the communication enabler in the IoT concept.

Intelligent Endpoint Devices

So what are the other IoT technology enablers? The ground floor enablers—the ones most representative of the physical manifestation of the IoT—are the intelligent endpoint "devices" and sensors that allow the extraction of data that is being communicated. The reason for enquoting the term "devices" is to signify the variance in its definition. An IoT device is essentially any computing engine in an endpoint solution that enables that solution to be intelligent and connected. It doesn't have to be handheld or mobile, but it can be. Devices can be anything from an embedded board installed in factory floor machinery to an actual handheld field service device.

Connecting to remote devices to extract collected data can be done in different ways, but all require hardware, firmware, and software components. A good place to start is with a dedicated board management controller (BMC). Initially designed for power sequencing tasks, the BMC has evolved to include many new and useful features for board management and control. Measuring the supply current to get a snapshot of the system's power consumption is only one example of these new capabilities. And compatibility with the latest Embedded Application Programming Interface specification (EAPI) reduces design efforts to port existing calls to the BMC.

Providing the interface from the hardware to the operating system is one of the most important functions of the device's remote management system. The BMC first collects all relevant information from the chipset and other sources. Utilizing the System Management Bus, the application layer fetches the data and presents it to the user, displayed either in the BIOS menu or a user-friendly dashboard suitable for supervision and troubleshooting.

Cloud Computing

Cloud connectivity takes today's intelligent middleware a step further than previous generations of remote management technology. By employing a cloud server architecture and an M2M stack on top of the intelligent middleware, embedded devices can connect to the cloud without additional design requirements (Figure 1). For example, the M2M stack pushes system data to the user's cloud server via any kind of TCP/IP connection. System managers have easy access to data and analytics through any commercial cloud portal, using any device (e.g., PC, tablet, smart phone).

When systems are available, operators can observe their performance. Cloud-based remote management furthers that process by enabling observation anytime, anywhere.

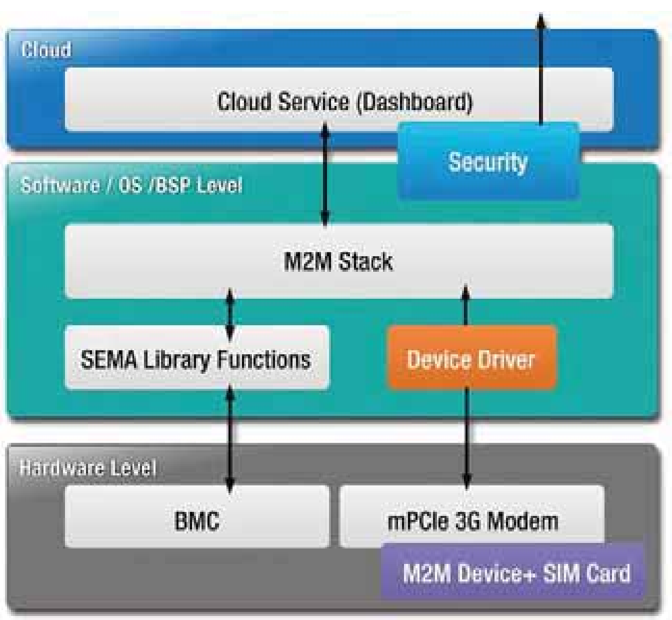

Figure 1. By employing full connectivity, from edge to cloud to end application, SEMA-enabled embedded devices can connect to the cloud without additional design requirements. Being a holistic solution, ADLINK's Smart Embedded Management Agent (SEMA) Cloud offers users the entire infrastructure required. Customers do not need to develop their own cloud solution, avoiding laborious checking of hardware compatibility, finding a suitable cloud server, implementing data encryption or developing proprietary communication protocols.

Embedded management agents may be used to continuously upload data through an encrypted Transport Layer Security (TLS, the successor protocol of Secure Sockets Layer or SSL) connection.

Data to the cloud enables operators to verify, monitor and manage system performance from a single, central location—providing immediate improvements to reliability and reduction in management costs. But the larger promise is realized when this big data is integrated into back-end business systems, such as enterprise resource management (ERP) and customer relationship management (CRM).

Machine-to-Cloud Communication

And this brings us back to M2M communication. With the IoT, M2M communication has become more about machine-to-cloud, or machine-to-machine-to-cloud. A critical challenge to deploying industrial IoT solutions is the lack of a single standard of network connection. For example, brownfield devices use a scattered variety of proprietary protocols; using an IoT gateway that supports cross-communication protocols and can connect with IP-based networks can greatly simplify an IoT deployment.

The ideal gateway platform integrates routing and data collection functions, conversant with existing protocols. Strong support for widely used fieldbus protocols provides bi-directional communication and acquisition (Figure 2). Simple IoT gateways are capable of accepting a wide range of connectivity protocols for analog and digital data and can provide communication for a wide array of industries, taking advantage of the valuable data existing in their hardware assets.

Figure 2. ADLINK's embedded IoT gateway platforms feature fieldbus control interface, as well as superior WiFi/LTE/3G enabled edge device/ cloud connectivity. Secure, scalable computing gateways with fieldbus interfaces enable seamless connection, aggregation, filtering, and data transmission to the cloud with confidence. Looking beyond the gateway, ADLINK's SEMA Cloud includes a secure cloud server architecture and machine-to-machine (M2M) stack on top of its intelligent middleware. This built-in service means the ADLINK IoT gateways can connect to a commercial cloud without additional design requirements.

Intel has invested much time and effort in an Intel® IoT Gateway platform that addresses the need for incorporating both legacy devices and newer, more open systems into IoT deployments. The Intel® IoT Gateway platform integrates an Intel® processor-based third-party hardware with an Intel® software stack. This software platform bundles together Wind River Intelligent Device Platform* (IDP) XT and McAfee* Embedded Control to provide a complete, pre-validated communications and security solution.

Wind River IDP XT provides a software stack for communicating with local equipment and the cloud. Its extensive connectivity choices include Wi-Fi*, Bluetooth*, ZigBee*, and short-range wireless protocols widely used in smart buildings. Wind River IDP XT supports the MQTT protocol for data transportation, and remote management protocols such as Technical Report 069 (TR-069), CPE WAN Management Protocol (CWMP) and Open Mobile Alliance Device Management (OMA DM). For developers, the Wind River IDP XT software stack provides Lua, Java, and OSGi application environments to enable rapid, reusable application development.

Wind River IDP XT delivers built-in security features as well—such as a hardware root of trust—to secure the communication channel, the data, and the end device. In addition, McAfee Embedded Control adds dynamic whitelisting. This technology locks the system down to a known good baseline so no program outside the authorized set can launch. McAfee Embedded Control also contributes a policy-based change control feature that monitors files and prevents unexpected changes.

Embrace the Role of Enabler

Embedded, connected systems can generate and collect a vast amount of system performance data – and embedded designers can sleep well knowing that the role they play in the creation of a successful IoT application is critical. By developing devices and systems that simplify the acquisition and management of remote data and by focusing on the movement to develop industry standards to simplify IoT deployments for end customers, embedded designers are enabling big data. They can then confidently pass the baton to the next level ecosystem partner to figure out what to do with it all.

Elizabeth Campbell is general manager, Americas, ADLINK Technology. Campbell has more than 30 years of sales and general management experience, the last 12 with ADLINK Technology via its acquisition of Ampro Computers in 2008. She began her career in management consulting, advising high-tech clients worldwide on organizational effectiveness techniques and sales methodologies.

ATCA-7480 Packet Processing Blade

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

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

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

With enabling software for Artesyn's SDN/NFV solutions, the ATCA-7480 blade includes support for the Intel® Data Plane Development Kit with accelerated OpenVSwitch, OpenFlow and OpenStack plug-ins for managing virtualization services on the computing platform. The blade will also support Wind River's Titanium Cloud.

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

SharpStreamer Video Accelerator Card

The Artesyn SharpStreamer PCIE-7207 high-density video accelerator enables service provider networks to offer video transcoding services quickly and dynamically. As an add-on card, the SharpStreamer PCIE-7207 offers quick and scalable integration with existing and standard server architectures to meet the demands of ISPs and MSOs who want to use existing servers and cloud infrastructure to support new video transcoding services.

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

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

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

Artesyn Embedded Technologies
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

ADLMES-8200 Rugged Modular Enclosure Systems

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

Key Features

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

Potential Applications Include:

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

Embedded Solutions™
Embedded. Works Anywhere.™

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

Power-Saving & Rugged Embedded Solutions

15" Light Industrial Touch Panel PC
Intel® Atom™ Processor E3800 Product Family

KS150-BT

- 2 DDR3L SODIMM up to 8GB/1 DDR3L SODIMM up to 4GB
- Rich I/O connectivity: 2 GbE, 4 COM, 3 USB
- Available in both resistive and capacitive touch screen
- Supports a wide temperature range for harsh working environments
- True-flat touch screen with IP65 front protection

3.5" Embedded SBC
4th Gen Intel® Core™ Processor U Series

HU551/HU553

- 1 DDR3L SODIMM up to 8GB
- Three independent displays: 1 LVDS, 2 DP/HDMI
- 2 Mini PCIe (1 supports mSATA), 1 SIM
- 2 GbE, 3 COM, 4 USB 3.0, 2 USB 2.0
- Wide voltage: 15~36V DC-in

DFI®

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

intel IoT Solutions Alliance

conga-QKIT/IOT

Qseven IOT Gateway Development Kit

congatec

- ~ conga-QA3 Intel® Atom™ processor-based Qseven module
- ~ Full featured Qseven IOT Mini Carrier Board
- ~ 7″ Single Touch Display with cable set
- ~ Intel® Dual Band Wireless-AC 7260 card & Antenna
- ~ Power supply 5V/4A, 20W, 4pin jack
- ~ Intel® IOT Gateway Solution OS (Windriver IDP trial)

intel IoT Solutions Alliance

Find more details at: www.congatec.us

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