

Embedded COMPUTING DESIGN

Connecting Silicon, Software, and Strategies for Intelligent Systems

Available for download on
iTunes App Store and Kindle.

FOLLOW US:



DECEMBER 2013
VOLUME 11 • NUMBER 8

www.embedded-computing.com



BETWEEN MACHINES:

INSIGHTS INTO TOOLS AND TECHNOLOGIES FOR M2M COMMUNICATIONS >> PG. 12

>> **INDUSTRY-UNIVERSITY
RESEARCH CENTERS DRIVE
EMBEDDED INNOVATION**

PG. 8

PLUS
+

GET VERIFIED

VERIFYING EMBEDDED DESIGNS
WITH CLOUD COMPUTING

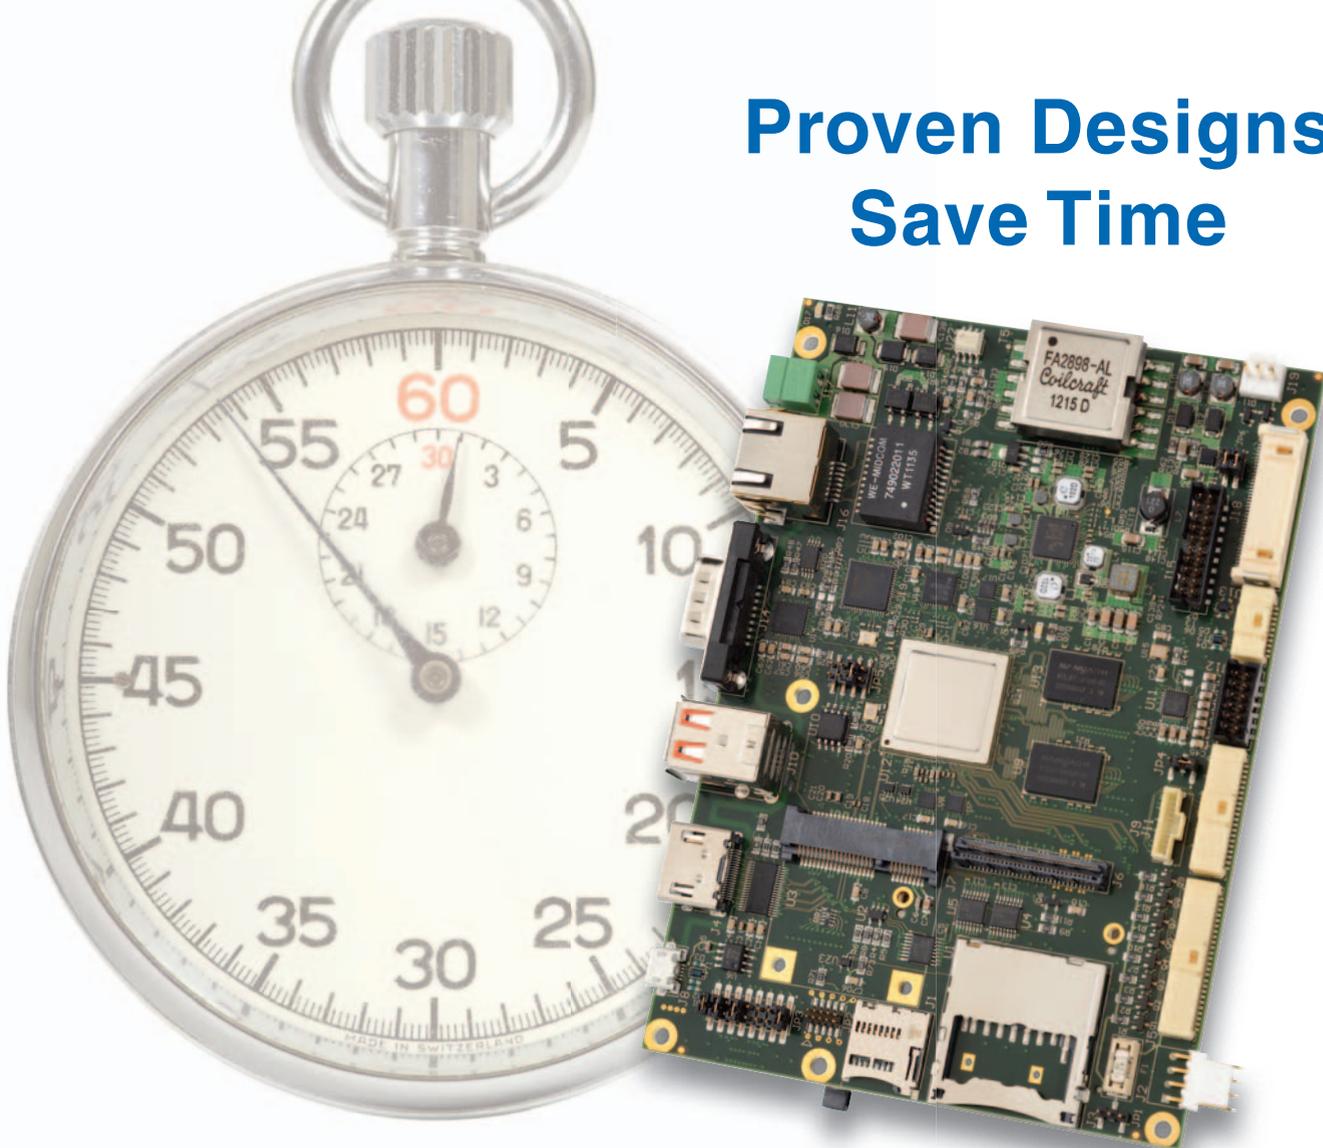
PG. 30

WHAT'S HOT

2014 EMBEDDED COMPUTING TRENDS

PG. 33

Proven Designs Save Time



Save Selection Time

- x86 and ARM® SBCs
- I/O expansion modules
- Power supplies
- Engineered for durability

Save Development Time

- Off-the-shelf availability
- Easy startup with OS images
- Knowledgeable technical support

Save Manufacturing Time

- Production-ready designs
- Long-term availability
- Custom configurations available

With more than 31 years in the industry, WinSystems focuses on embedded computing solutions so you can focus on your unique applications.

*Speak to an Applications Engineer
at 817-274-7553 or visit our website
www.WinSystems.com/E2*



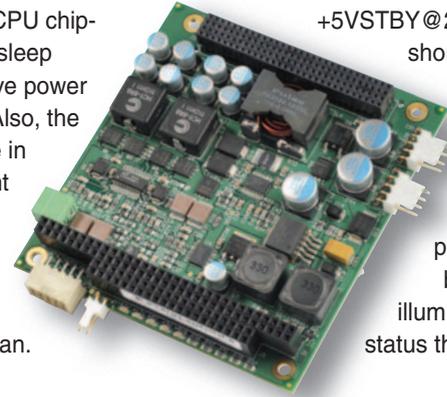
715 Stadium Drive • Arlington, Texas 76011
Phone 817-274-7553 • FAX 817-548-1358
E-mail Info@WinSystems.com

Single Board Computers • I/O Modules • Panel PCs • Power Supplies

ATX-compatible DC/DC Power Supply offers Wide Input Range and -40° to +85°C Operation

WinSystems' PPM-DC-ATX is a PC/104-Plus DC/DC power supply for PC/104, EPIC, and EBX single board computers (SBCs) that support ATX power controls. It features a wide voltage input range from 10 to 50 volts, which allows the unit to operate with 12, 24, or 48 volt battery-operated or distributed DC power systems. It generates five regulated DC output voltages from one common DC input, plus supports the software controlled shutdown and power monitoring for SBCs

with advanced CPU chipsets employing sleep modes and active power management. Also, the unit can operate in a +85°C ambient temperature environment using normal convection cooling and no fan. The outputs are +5V@10A, +3.3V@10A, +12V@3A, -12V@500mA, and



+5VSTBY@2A. Each output is short circuit protected and current limited. A minimum load is not needed to bring the the supply into regulation. When power is applied to the board, five LEDs will illuminate providing a visual status that power is available.

WinSystems, Inc. (817) 274-7553
WinSystems.com/PPM-DC-ATX-E2

Rugged Single Board Computers with Industrial I/O and Expansion

Designed for demanding applications and long-term availability, WinSystems single board computers feature low power x86 and ARM® processors with options for expansion and customization.

Long-life CPUs

- Freescale™
- AMD
- Intel®
- DM&P

Industrial I/O

- Ethernet
- Serial
- USB
- GPIO

Expansion

- PC/104
- Mini-PCIe
- IO60

Rugged Designs

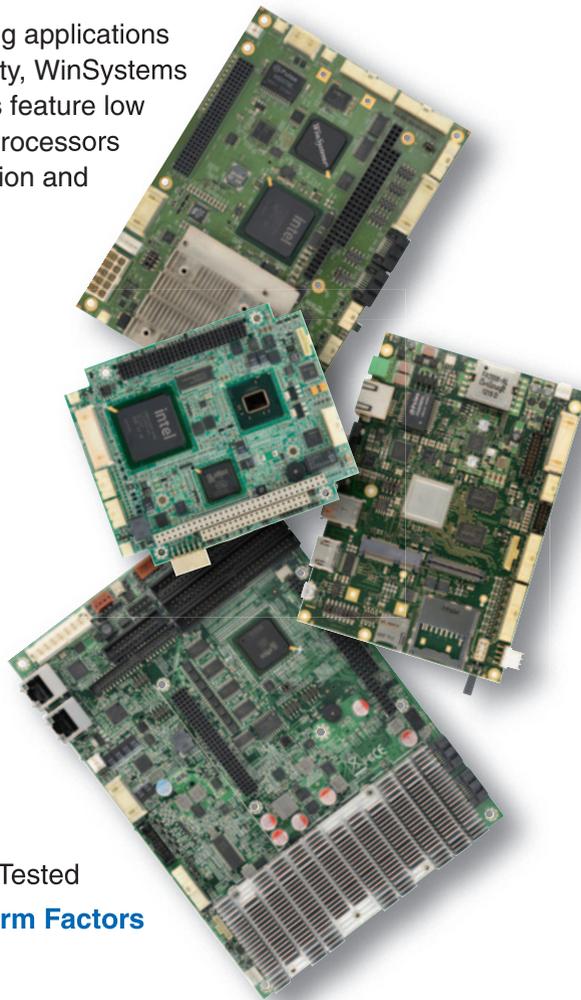
- -40° to +85°C
- Shock & Vibration Tested

Industry Standard Form Factors

- PC/104
- EBX
- EPIC
- 3.5-Inch SBC

Quick Start Support

- Windows® Embedded
- Linux
- Android™



Call 817-274-7553 or

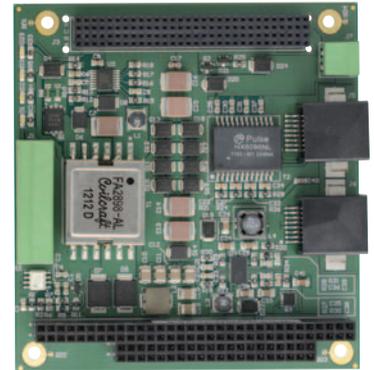
Visit www.WinSystems.com/SBCs-E2

Ask about our evaluation program.



PD Power Supply PC/104 Module for PoE Applications

WinSystems' PPM-PS397-POE-1 is an isolated 25W, 802.3af-compliant, Power over Ethernet (PoE) module. It powers a PC/104-Plus single board computer stack from DC power extracted from the CAT5 cable. It is designed for use in remote areas where neither AC nor DC power is close and/or available to be used.



The PPM-PS397-POE-1 accepts 42 - 57 VDC and converts it to three isolated outputs: +5VDC@5.0A, +12VDC@1.0A, and -12VDC@1.0A. Each output is short circuit protected and current limited. A minimum load is not needed to bring the supply into regulation.

WinSystems also offers this module configured for PC/104 and standalone systems. No fans or heat sinks are required to meet its extended operating temperature range of -40° to +85°C.

WinSystems, Inc. (817) 274-7553
WinSystems.com/PPM-PS397-POE-E2

ON THE COVER

With the rapid growth of a wide variety of connected devices and the Internet of Things, it's important to develop connectivity, interoperability, security, and other strategies to make effective and secure M2M systems. In this issue we also talk trends in our special advertising section which seeks to answer "What's Hot in 2014 in Embedded?"



7 Tracking Trends in Embedded Technology

(Optical) bonding with your panel PC

By Rory Dear, Technical Contributor

8 Research Review

Industry-university research centers drive embedded innovation

By Monique DeVoe, Assistant Managing Editor

Silicon

M2M connectivity



22

Between machines: Insights into tools and technologies for M2M communications

By Curt Schwaderer, Technology Editor

12

Balancing interoperability and security in an M2M deployment

By Carl Cohen, Avnet Embedded

14

M2M integration platforms enable complex IoT systems

By Robert Andres, Eurotech

22

Software

Integrated development environments



26

Code reviews take a page from social networks

By Curt Schwaderer, Technology Editor



Strategies

Cloud computing



30

Verifying embedded designs with cloud computing

By Dr. Raik Brinkmann, OneSpin Solutions



OpenSystems media.

© 2013 OpenSystems Media ®

© 2013 Embedded Computing Design

All registered brands and trademarks within *Embedded Computing Design* magazine are the property of their respective owners.

iPad is a trademark of Apple Inc., registered in the U.S. and other countries.

App Store is a service mark of Apple Inc.

ISSN: Print 1542-6408, Online: 1542-6459

ENVIROINK

The inks used to print the body of this publication contain a minimum of 20% by weight, renewable resources.

Embedded Computing Design



E-CASTS

<http://ecast.opensystemsmedia.com>

Automotive SPICE compliance made easy – Case studies and lessons learned

Presented by: IBM

Go to E-cast: <http://ecast.opensystemsmedia.com/439>

Catch the security breach before it's out of reach

Presented by: Klocwork

Go to E-cast: <http://ecast.opensystemsmedia.com/440>



**BY 2017, MOBILE DEVICES
WILL STORE OVER
64 BILLION GAMING APPS.**

Be ready with the storage
that's ready for the future.

The universe of apps is exploding. Mobile downloads are skyrocketing. And the need for storage has never been greater. That's why for 25 years, SanDisk has been expanding the possibilities of storage and enabling companies to create game-changing products. sandisk.com

Advertiser Information

Page Advertiser/Ad title

- 36 **AdaCore Technologies** – What will be hot in 2014 for embedded and mobile software development?
- 32 **ADLINK Technology, Inc.** – nanoX-BT: Mini COM Express type 10 module
- 32 **ADLINK Technology, Inc.** – cExpress-BT: Compact COM Express type 6 module
- 38 **AMD** – x86 and ARM enable next-gen intelligent embedded devices
- 40 **AMD** – World's first x86 quad-core SoC!
- 39 **Annapolis Micro Systems, Inc.** – WILDSTAR OpenVPX ecosystem
- 10 **Cogent Computer Systems, Inc.** – Low power, high performance AMD x86 SoC
- 28 **COMMELL Systems Corporation** – We offer advanced and reliable IPC products
- 15 **Creative Electronic Systems** – High-performance solutions
- 25 **Dolphin Interconnect Solutions** – Need a fast network
- 9 **Elma Electronic – Systems** – Reliable, watertight fanless PCs
- 29 **Embedded World** – Summit for innovations
- 13 **Innovative Integration** – The real source of your super powers
- 19 **Kontron** – When your network cannot fail
- 37 **Kontron** – Scalable HEVC delivery and video transcoding in the cloud
- 17 **Micro Digital, Inc.** – Your solution is here
- 5 **SanDisk** – By 2017, mobile devices will store over 64 billion gaming apps
- 16 **Sealevel Systems, Inc.** – Serious advantages for serious challenges
- 23 **Sensoray Co., Inc.** – Embedded electronics experts
- 17 **Technologic Systems** – TS-4710 high end CPU module
- 35 **Vitesse Semiconductor Corp.** – An eye on end-to-end wireless security
- 20-21 **Wind River Systems, Inc.** – It's time to make some future
- 34 **Wind River Systems, Inc.** – NFV and IoT are forging new paths for intelligent connected systems
- 2-3 **WinSystems, Inc.** – Proven designs save time
- 11 **X-ES** – 4th generation Intel Core i7 from X-ES

Get your free digital edition at embedded-computing.com/subscribe

available on
kindle fire

Available on the
App Store



Embedded COMPUTING DESIGN

OpenSystems media.

ECD Editorial/Production Staff

Monique DeVoe, Assistant Managing Editor
mdevoe@opensystemsmedia.com

Curt Schwaderer, Technology Editor
cschwaderer@opensystemsmedia.com

Rory Dear, Technical Contributor
rdear@opensystemsmedia.com

Steph Sweet, Creative Director
ssweet@opensystemsmedia.com

Sales Group

Tom Varcie, Senior Account Manager
tvarcie@opensystemsmedia.com

Rebecca Barker, Strategic Account Manager
rbarker@opensystemsmedia.com

Eric Henry, Strategic Account Manager
ehenry@opensystemsmedia.com

Ann Jesse, Strategic Account Manager
ajesse@opensystemsmedia.com

Susan MacLeod, Strategic Account Manager
smacleod@opensystemsmedia.com

Kathleen Wackowski, Strategic Account Manager
kwackowski@opensystemsmedia.com

Christine Long, Vice President, Online Business
clong@opensystemsmedia.com

Asia-Pacific Sales

Elvi Lee, Account Manager
elvi@aceforum.com.tw

Regional Sales Managers

Barbara Quinlan, Southwest
bquinlan@opensystemsmedia.com

Denis Seger, Southern California
dseger@opensystemsmedia.com

Sydele Starr, Northern California
sstarr@opensystemsmedia.com

Reprints and PDFs

republish@opensystemsmedia.com

EMEA

Rory Dear
Technical Contributor
rdear@opensystemsmedia.com

Gerry Rhoades-Brown
Account Manager – Europe
gerry.rhoadesbrown@husonmedia.com

Christian Hoelscher
Account Manager – Europe
christian.hoelscher@husonmedia.com

Lauren Palmer
Account Manager – Europe
lauren.palmer@husonmedia.com

OpenSystems Media Editorial/Production Staff



Curt Schwaderer, Technology Editor
Embedded Computing Design
xTCA and CompactPCI Systems
cschwaderer@opensystemsmedia.com

Monique DeVoe, Assistant Managing Editor
Embedded Computing Design
DSP-FPGA.com
mdevoe@opensystemsmedia.com

John McHale, Editorial Director
Military Embedded Systems
jmchale@opensystemsmedia.com

Joe Pavlat, Editorial Director
xTCA and CompactPCI Systems
jpavlat@opensystemsmedia.com

Jerry Gipper, Editorial Director
VITA Technologies
jgipper@opensystemsmedia.com

Amanda Harvey
Assistant Editor
VITA Technologies
Military Embedded Systems
aharvey@opensystemsmedia.com

Brandon Lewis
Associate Editor
xTCA and CompactPCI Systems
PC/104 and Small Form Factors
Industrial Embedded Systems
blewis@opensystemsmedia.com

Konrad Witte
Senior Web Developer

Steph Sweet, Creative Director

David Diomedea, Art Director
Joann Toth, Senior Designer

Corporate

Patrick Hopper, Publisher
phopper@opensystemsmedia.com

Karen Layman, Business Manager

30233 Jefferson
St. Clair Shores, MI 48082
Tel: 586-415-6500

www.opensystemsmedia.com

Rosemary Kristoff, President
rkristoff@opensystemsmedia.com

Wayne Kristoff, CTO

Emily Verhoeks, Financial Assistant
16626 E. Avenue of the Fountains, Ste. 201
Fountain Hills, AZ 85268
Tel: 480-967-5581

Subscriptions

Tel: 586-415-6500

www.opensystemsmedia.com/subscriptions

subscriptions@opensystemsmedia.com



Tracking Trends in Embedded Technology

By Rory Dear, Technical Contributor



rdear@opensystemsmedia.com

(Optical) bonding with your panel PC

Panel PCs and industrial displays marketed as conforming to IP65 and/or being waterproof naturally lead one to surmise a sealed, airtight unit. Whether this is true could be irrelevant, but in certain operational environments it can be very relevant indeed. Here I discuss how system integrators often overlook environmental suitability criteria and assume IP65 is sufficient for usage in outside environments with heavily varying temperature and humidity conditions.

IP rating translating

First, we should understand what is meant by IP65. Whilst “IP,” denoting “Ingress Protection,” is fairly self-explanatory, the numbers that follow are not. Our interest lies with the second number, which defines the protection against liquid ingress, where “5” is protection against low-pressure jets of water. We should note the rating does not express that no ingress is possible, moreover that any ingress cannot be harmful to the functionality of the panel PC. To achieve a product that is immersible in water, thus considered “watertight” – though not “airtight,” – you will need a minimum of an IP67/68 rated unit.

Now for the interesting part: In any interior environment, temperatures and humidity tend to stay fairly constant. Whilst of course there are fluctuations, they are nothing compared to what the unit will experience *al fresco*. It is probably valid to argue that unless you are required to wash down your panel PC with a pressure washer or immerse it in water (indoors), standard IP65 will meet your needs entirely. Unfortunately, issues can arise when one believes such a unit is equally suitable for being left to the wrath of the great outdoors.

Here I concentrate on IP65, as this is the most common rating, but what follows is

equally true of anything less than a special type of IP68.

Exterior inferior?

Two issues exist with such units outside the “phenomenon” of condensation.

Most users’ first reaction to seeing condensation on the inside of the display is “How did the water get in there?” which initially seems a valid question. Actually, when one considers the units are far from airtight, this becomes far easier to understand.

Water molecules are far larger than air molecules, so, as mentioned, a watertight system isn’t necessarily airtight – though vice versa, of course, this does hold true. For a system to be airtight, it must be hermetically sealed, which even for IP68 certified panel PCs isn’t common, and the additional costs involved aren’t really necessary.

Now we understand that almost no unit is airtight nor contains a vacuum; it’s easy to see how water can exist in the air inside. Some will be trapped at build time, though these units tend to be built in dehumidified rooms to minimise this. Also, because of the change in external air pressure when units are flown across the world at high altitude, air – and the moisture contained within – is pulled into the unit.

In a fairly constant temperature environment, this humidity stays “hanging” in the air inside of the unit; this isn’t evident on systems located indoors, or a significant quantity outside, actually. If they are exposed to the full elements, they experience often freezing nights and relatively hot days, possibly in direct sunlight. In these conditions, the moisture separates from the air and finds the coolest surface to condense upon, which

is the inside of the glass of the display; there it rests to form the condensation.

It’s worth mentioning that this is usually temporary. Once the ambient temperature rises sufficiently, the condensation will again evaporate. Interestingly, the compass orientation in which the unit faces varies the severity of the effect, because of the angle of the sun.

The other issue when sitting these units outside is reflection, given the layer of air between display glass and touch screen. A resultant zigzag effect of reflecting light can render a display illegible.

Solution evolution

Thankfully, a solution exists for both the condensation and reflection issues – a solution that does not require the level of seal a nuclear sub would be proud of.

Optical bonding is the process of gluing the touch-screen membrane directly to the display, using optical-grade adhesive to avoid any image distortion. This solves the condensation issue, as no layer of air exists that the moisture can use to access the glass front and thus condense. The reflection issue is also solved, as no zigzagging of light can occur between the two layers. Also, optical bonding strengthens the display from breakage as it has universal support across its face from the layer behind.

“Why isn’t this standard?” Of course there is a cost attached to this that isn’t insignificant, and which would be an extraneous cost to the majority using such systems in a more temperature-controlled environment. It’s also down to user experience: The condensation and reflection don’t cause any functional detriment to the unit; it will work perfectly well without optical bonding, so it really becomes a decision based on the application.



RESEARCH *REVIEW*

Monique DeVoe, Assistant Managing Editor

Industry-university research centers drive embedded innovation

Embedded systems are facing growing demands on performance, power, size, cost, time to market, and increased capabilities in general. In order to keep up with these demands, industry research

needs to be done in order to make the necessary technological advances – a lot of often time-consuming, inefficient, and expensive research. That is, if companies go it alone.

The Center for Embedded Systems (CES) – an Industry-University Cooperative Research Center (I/UCRC) through Arizona State University's Ira A. Fulton Schools of Engineering and a sister site at

<p>Concurrency and Scheduling Analysis of Real-time Embedded Software on Multi-core Processors <i>Arizona State University</i></p>	<p>Towards Predictable Execution of Safety-Critical Tasks on Mixed-Criticality Multi-Core Platforms <i>Southern Illinois University Carbondale</i></p>	<p>Ground Work for Embedding a Field Oriented Motor Controller into a Single System on a Chip <i>Southern Illinois University Carbondale</i></p>	<p>Parallelization of Embedded Control Applications on Multi-core Architectures: A Case Study <i>Arizona State University</i></p>
<p>Visual Interface for Metric Temporal Logic Specifications <i>Arizona State University</i></p>			<p>Improving Usability of Multi-core DSPs <i>Arizona State University</i></p>
<p>Performance Optimal Control of a System of Interconnected Components Under Thermal and Energy Constraints <i>Arizona State University</i></p>			<p>Achieving Energy-efficient Mobile Computing Through Explicit Data Communication and Global Power Management <i>Arizona State University</i></p>
<p>Spintronic Threshold Logic Array <i>Arizona State University</i></p>			<p>Fundamental Research Project: Collaborative Research: Synthesis and Design of Robust Threshold Logic Circuits <i>Arizona State University and Southern Illinois University Carbondale</i></p>
<p>I2AV: Integrate, Index, Analyze, and Visualize Energy Data for Data-driven Simulations and Optimizations <i>Arizona State University</i></p>			<p>Registration and Fusion of EVS and SVS Runway Images for Embedded Systems <i>Southern Illinois University Carbondale</i></p>
<p>Readout Integrated Circuit for Fast Imager <i>Arizona State University</i></p>			<p>CORBI: Collaborative Research: Consortium for Embedded Systems <i>Southern Illinois University Carbondale</i></p>
<p>Adaptive Compressive Sensing Techniques for Low Power Sensors <i>Southern Illinois University Carbondale</i></p>	<p>An Effective Test Strategy Based on Coverage Driven ATPG <i>Southern Illinois University Carbondale</i></p>	<p>Reliable Wireless Communications in Aircrafts and Other Challenging Environments <i>Southern Illinois University Carbondale</i></p>	<p>Multicore Simulator Critical Path Analysis <i>Southern Illinois University Carbondale</i></p>

Figure 1 | The Center for Embedded Systems, currently in its fifth year, is working on a total of 18 projects between its two university partners that focus on a variety of embedded software and hardware technology developments.

Southern Illinois University Carbondale – is partnered with 11 leading electronics companies worldwide including Intel, Toyota, Qualcomm, and Marvell to conduct embedded computing research for the benefit of everyone involved.

Embedded research focus areas

Now in its fifth year, CES's research projects are focused on power, energy, and thermal aware design; Electronic System-Level (ESL) design and technologies; multicore architectures and programming; software systems; cyber-physical systems; and integrated circuit technologies, design, and test. The current year's projects – nine at ASU, eight at SIUC, and one joint project – focus on various aspects of multicore hardware and software, low power and thermal management, and hardware performance, among other areas of embedded technology (Figure 1).

Partnership structure

As an I/UCRC, CES is a program organized by the National Science Foundation (NSF), who sees importance and benefits in university-industry-government partnerships and provides a framework to formalize the relationship. A majority of funding comes from the sponsoring companies through an annual \$50,000 membership fee. Of the membership fee, 90 percent of the funds goes directly toward research and 10 percent goes to managing the CES as overhead. In contrast, if a company wanted to fund university research on its own, the facilities and administrative cost is more than 50 percent of the total research award, says ASU CES Center Director Sarma Vrudhula. On some qualifying projects, the companies' research funding can also be matched with federal funds through the I/UCRC arrangement.

In addition to funding research more directly, member companies each have one representative to serve on the Industry Advisory Board (IAB). Through the IAB they can screen project proposals from expert faculty at an annual meeting and vote on which projects to fund. The results of the research projects and all resulting Intellectual Property (IP) are shared equally between members. Vrudhula says this can be a turnoff for some companies, but it's

an unwarranted fear. "Competitors like Intel and Qualcomm participate together, and other centers have lots of competitors who participate. Everyone gets something out of it."

Benefitting from the results

Vrudhula explains that there are three methods of beneficial "tech transfer" that take place through CES.

"The most effective way [of tech transfer] is through people," Vrudhula says. Companies can participate in internship

programs and hire the students who graduate out of the CES research programs. The number of participating faculty and students varies per year, but typically there are about 12 faculty members and 26 graduate students participating through both academic sites.

Companies can also use the software and tools produced by CES and have the option to license them royalty-free. Members may also request exclusive licensing in inventions resulting from CES work.

ELMA
Your Solution Partner

Reliable, watertight fanless PCs

Elma's SFF-IP68 is a compact, fanless, rugged computer for demanding environmental conditions. Designed to meet IP68 protection from continuous water immersion and dust penetration, it also offers high shock and vibration resistance. The box ships with a conformal coated Intel® based single board computer loaded with SDRAM and NAND Flash, I/O ports for PCI Express, Ethernet, serial and COM ports. The system is tested for temperature operation range from -40°C to +85°C. The conduction cooled enclosure ships off the shelf with the above features. Tailored configurations can be easily accommodated.

Find out why Elma is the authority in embedded computing platforms, systems & components.
www.elma.com | sales@elma.com | 510.656.3400

One of the most valuable and cost-effective ways that companies use the CES research compared to internal research, according to Vrudhula, is to take what's been done at the university and expand upon it in their own organizations, without disclosing this to competitors in the CES. Many members will use tools and software from at least one project in their own projects because they have worked closely with the research teams and, in some cases, provided data to help create the tools. Companies can also change directions because of the results of the university's projects.

In this way, companies can use presentations and technologies from the CES research as a method of "precompetitive research," Vrudhula says. "They don't have to chase after several possible solutions because they know what works. They can take [the solution] and modify it to their liking, or not try it because they know it won't work. They're ahead of where they would be if they didn't participate."

Companies join the CES by signing a standard membership agreement used by all members, paying the membership fee, and meeting with the center director to initiate a research project. Member companies participate in two annual meetings of the IAB each year to review and vote on research projects and conduct center business.

Future plans

The CES is looking to add new universities to the group, including Northeastern University in Boston in 2014, Vrudhula says, as well as potential international partner universities in Europe and Asia sometime in the future.

In 2014 *Embedded Computing Design* will be covering CES's various research projects and findings as well as other areas of research into techniques and designs for the embedded systems of the future.

For more information about the CES visit ces.asu.edu.

CENTER FOR EMBEDDED SYSTEMS INDUSTRY PARTNERS:

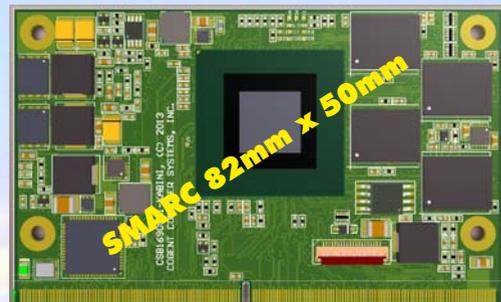
Alphacore, INC
General Dynamics
Intel (Arizona)
Intel (California)
Intel (Oregon)
Johnson Controls
Marvell
Qualcomm
Raytheon
Rockwell Collins
Toyota
UTC Aerospace Systems

Low Power, High Performance AMD x86 SoC



CSB1890T10-Kabini

1.5Ghz Quad Core AMD G-Series SoC
4GByte 64-Bit DDR3-1600 w/ECC
PCIe x4 and Dual x1 GEN 2, GIGe
Dual SATA Gen 2, Dual USB3, USB2 x 8
Dual Display (DP, eDP or HDMI)



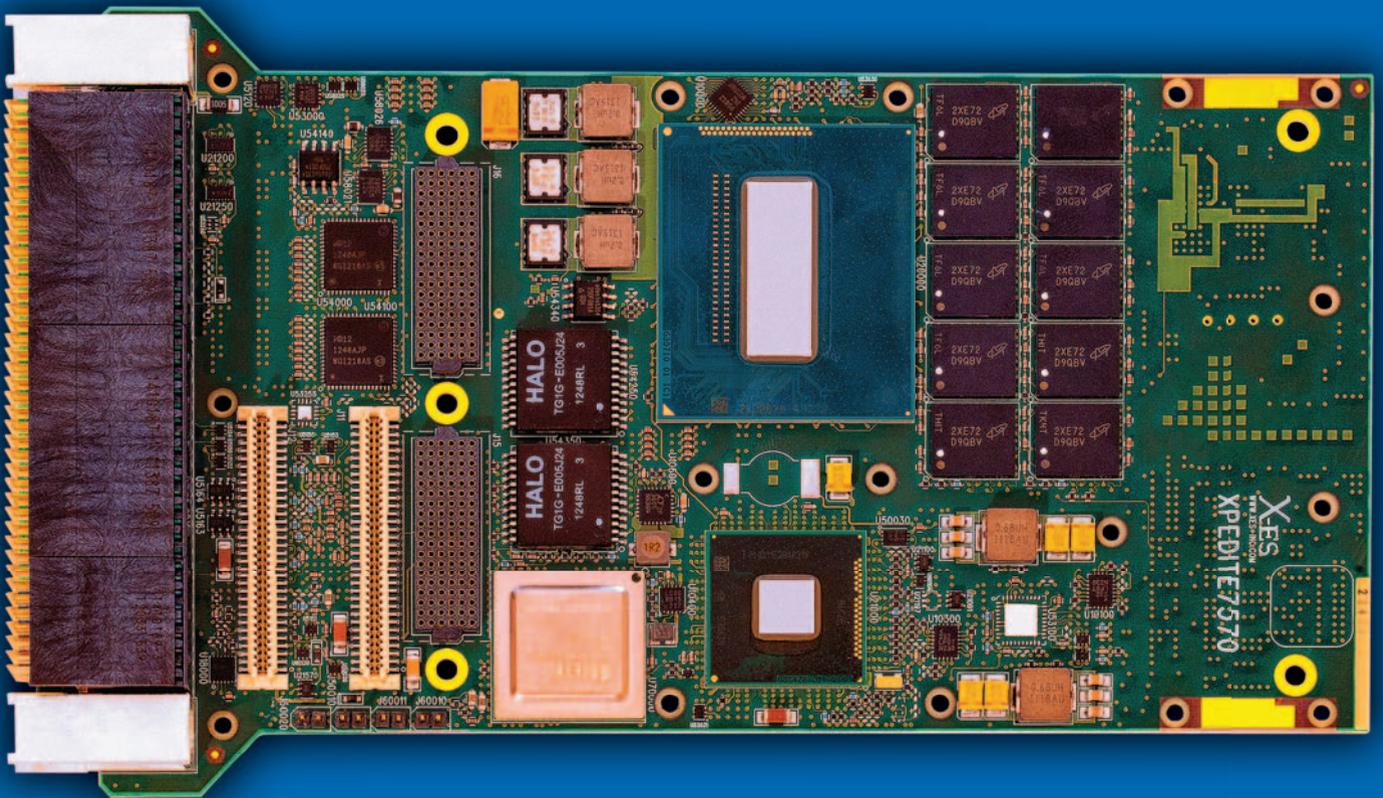
COMING SOON: CSB1690SM1-Kabini

1.0Ghz <6W Dual Core AMD G-Series SoC
4GByte 64-Bit DDR3-1066
Three PCIe x1 GEN 2, GIGe
SATA Gen 2, Three USB2
Dual Display LVDS and HDMI
Worlds First x86 SMARC Solution



For more information
Tel: 401-349-3999
Email: sales@cogcomp.com
Web: www.cogcomp.com

4th Generation Intel® Core™ i7 from X-ES



VPX · cPCI · VME · XMC · COM Express · Custom

High-performance, rugged, and versatile 4th Generation Intel® Core™ i7 solutions from X-ES

X-ES

Extreme Engineering Solutions
608.833.1155 www.xes-inc.com



100% Designed, manufactured, and supported in the USA



Between machines: Insights into tools and technologies for M2M communications

By Curt Schwaderer, Technology Editor

As M2M systems become more varied and widespread, new networking and platform-building methods can help solve communication, implementation, and scalability challenges.

Independent wireless analyst Berg Insight asserts that the number of cellular network connections worldwide used for Machine to Machine (M2M) communications was 47.7 million in 2008. They forecast that this number will increase to an astounding 187 million by 2014, according to *The Global Wiress M2M Market* report by Berg Insight.

Why M2M and wireless? Because it simplifies things. Often, M2M involves communications with remote sensors or devices in locations where there is no infrastructure. Wireless M2M provides the ability to monitor remote applications without dealing with or deploying local infrastructure.

I spoke with Dave Richkas, MCU product marketing manager at Microchip, about M2M history, uses, and trends.

M2M beginnings

Initially, Microchip developed an M2M platform that was GPS/GSM/GPRS based. It was pretty simple to “load and

go” – buy a SIM card and deploy it. Richkas mentioned that shortly after the release, AT&T announced the shutdown of the 2G network by the end of 2016 and Verizon guaranteed a 10-year life for CDMA applications. New requests started coming in for CDMA networks since they couldn’t move forward with the GSM-based networks.

“ Wireless M2M provides the ability to monitor remote applications without dealing with or deploying local infrastructure. ”

Microchip partnered with Twistthink to create the CDMA platform and to achieve cellular network certification. This system used a PIC32 32-bit microcontroller with 512K bytes of flash and

128K bytes of RAM running at 80 MHz, proving that a higher priced processor is not necessary for M2M applications. In addition, they worked with Exosite to provision the platform on the cellular network and enable intuitive user-developed web portals to monitor their M2M applications.

M2M and cellular application examples

Richkas cited an early M2M example involving vending machines:

“An application and early driver of the design involved a bank of vending machines at an airport. One hub was used to collect stocking and health information from the vending machines using ZigBee. The hub then communicates over the cellular network to the centralized system where health and stocking status can be monitored and acted upon.”

I also brought up M2M in a corporate network – for example, perhaps a fitness

machine is communicating from a fitness center and logging information for use by doctors and other medical professionals. This intersection of multiple corporate networks and the IT hoops involved can be very challenging. However, offering an M2M solution integrated with the fitness equipment over a wireless network can avoid issues with data traveling on the fitness center corporate network. The data can be uploaded to a portal using the cellular network where the individuals can monitor themselves and track activity levels. If the health club wanted to monitor machine usage levels, hubs could be installed at the fitness centers which would gather the information from the center and relay that information to the cloud where the health club could see consolidated information on equipment usage, diagnostics, and availability for all sites.

More than a platform

Of course, implementing M2M requires more than a platform and there are a number of partners involved beyond just the remote platform itself.

The Microchip MPLAB IDE development suite enables development of the sensor and health application for the embedded device remote platform. Portals are often used as the sites where the remote information is collected, stored, analyzed, and presented to the administrator. In this case, Microchip worked with Exosite on a portal to monitor a spreadsheet snapshot of multiple devices. The portal offers a "drill down" style of information where the top layer provides the overall view, but individual information rows provide additional information about specific remote nodes.

These machines are also sending data across FCC-regulated wireless connections. Getting these systems certified by the FCC and Verizon can be no small task. This is where companies like Twistthink can provide pre-validated systems and manage the certification process to make things simpler and faster for the M2M customer.

M2M scalability

Scalability can also be a consideration in M2M. However, scalability has many

facets – number of remote nodes, the volume of information transferred within a time period, and perhaps even the variety of I/O and connectivity options.

Richkas mentions that there is plenty of bandwidth available to process many things, from diagnostics to remote control and beyond. The entry point platform prices at around \$380, but there is I/O expansion available on the platform itself. Developers can add their own board, access the analog-to-digital converters, and get access to the 5 V onboard power. This platform

expansion coupled with the additional portal and interface options provides a complete M2M solution that has room for growth.

The future of M2M

M2M is here and being deployed in a wide variety of systems for an even wider number of uses. Platforms, portals, systems, and analytics must all work together in order to provide robust M2M solutions now and into the future. **ECD**

For more information, contact Curt at cschwaderer@opensystemsmedia.com.

The real source of your Super Powers

mini-K7

Features

- Industry-standard COM Express CPU module with an FMC I/O Module in an extremely compact, stand-alone design
- Programmable Kintex-7 325T/410T and Spartan-6 FPGAs
- Small form factor: 4" H x 7" W x 10" D
- Conduction cooled design: Fins or Cold-plate
- Stand-alone operation: Able to operate headless, booting from SSD
- Windows & Linux support / RTOS available
- Dual VITA 57 FMC I/O Module Sites - add anything from RF Receivers to Industrial Control Modules
- FMC I/O Site delivers >3000MB/s to host memory**
- Integrated Timing and Triggering support for I/O includes GPS, IEEE1588 or IRIG -Disciplined Clock
- Supports Innovative Integration & 3rd-party FMC Modules for Private Data Channels, Triggering & Timing features
- USB 3.0 x2, USB 2.0 x2, 1Gb Ethernet x2, SATA x4, DisplayPort, Touch Screen
- Up to 2 SSD (2.5 in.)
- Flexible AC or DC operation

Customization Available

DOWNLOAD DATASHEETS NOW!
 #FMC

mini-K7 is the choice for
Wireless, Industrial Control
Medical Imaging
Telecom/Intelligence and
Test & Measurement
applications.

Innovative Integration
a subsidiary of Interconnect Systems, Inc.
... real time solutions!

805.578.4260 phone
www.innovative-dsp.com

Balancing interoperability and security in an M2M deployment

By Carl Cohen

M2M systems incorporate an increasing number of disparate devices and connectivity options, but they must still maintain interoperability and mitigate security risks from potentially compromised points in the system, all while meeting fast time to market goals.

It is nearly impossible to go a day without reading about the explosive growth of devices being connected to the Internet. Industry forecasters project astronomical numbers in terms of connected machines and the volume of data they will transfer. These figures pose both a significant opportunity, and considerable challenges to our industry.

In particular, this paradox exists in balancing the interoperability of disparate devices and protocols, from edge to enterprise, with the need to secure the data being passed from point-to-point. Each node along the continuum must be optimized for performance and power with minimal overhead for the “what if” scenarios created by not knowing exactly what else the device may need to communicate with. This is further complicated by different communication standards, bus protocols, and regulatory requirements. In light of all of this, designers are

challenged to meet aggressive time to market requirements in an environment where future-proofing the application is at best uncertain due to lack of industry standards.

Several functions must be considered for developing M2M systems for applications that depend on interoperability. In solving interoperability challenges, security risks must also be addressed.

Interoperability: Discrete devices, one system

It is nearly impossible to find a single, purpose-built system employing M2M from one vendor that covers the gamut from edge to enterprise. Even a modular approach where multiple functions are integrated, such as a sensor and control unit, still requires more elements to be added to the overall solution. Figure 1 shows the basic functional elements in a typical M2M deployment.



Each use case requires a different set of hardware, software, and connectivity configurations depending on whether the asset at the edge is mobile or fixed, the distance between devices, and whether there are regulatory standards that must be adhered to. For example, a way-finding kiosk located in a shopping mall would be considered an edge device. The kiosk might contain proximity and motion sensors to trigger power up when an individual approaches the unit, along with CMOS image sensors for detecting the individual's demographics to be used for targeted content delivery and a touch screen interface for display navigation. The kiosk may also contain additional internal environmental sensors to monitor the health of the system such as a temperature sensor for heat buildup within the enclosure. This combination of sensors is designed to work in conjunction with one another to create the

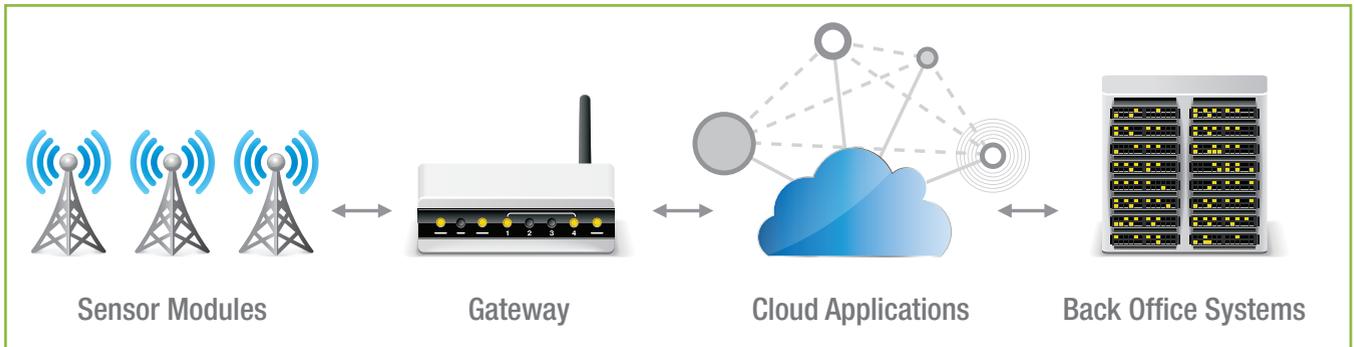


Figure 1 | Four basic functional elements in a typical M2M deployment are sensor modules, the gateway, cloud applications, and back office systems.

user experience and efficiently connect to the balance of the overall solution.

A variety of connectivity options are available, depending on the floor plan layout at the mall and how the mall is managed and operated. The kiosk may be a standalone unit that interacts, in concert with other standalone units, with a central gateway. Alternatively, the kiosk may contain the gateway within its physical structure. The con-

nectivity also varies depending on whether the kiosk is being hosted by the mall's property management company or by an independent third party, such as a marketing agency, that is prohibited from connecting to the mall's network. In addition to these considerations, the designer would have to take into account the type of content being transmitted, which in this example would be both data and video, and the modality of the transmission, which

could be Wi-Fi, cellular, or Ethernet. With all of these variables to consider from the onset it is no wonder that interoperability and standards become a challenge.

Securing M2M systems

Security is also a critical consideration when reviewing the various component options in the solution. The risk is both within the device as well as between devices. If the sensors are

High-Performance Solutions FROM CES

RIOV-2440

The latest SBC from CES provides 192 GFLOPS calculation power with the 24 processing threads of the brand new Freescale® QorIQ® T4240 processor. With high-speed interconnects matching the 3U OpenVPX™ form-factor, it provides a rugged high-performance computation engine for data-intensive applications in the embedded avionic and defense markets.



Headquartered in Geneva, Switzerland, CES - Creative Electronic Systems SA has been designing and manufacturing complex high-performance avionic, defense and communication boards, subsystems and complete systems for thirty years (such as ground and flight test computers, ground station subsystems, radar subsystems, mission computers, DAL A certified computers, video platforms, as well as test and support equipment). CES is involved in the most advanced aerospace and defense programs throughout Europe and the US, and delivers innovative solutions worldwide.

For more information: www.ces.ch



compromised, the system is likely to fail or perform in an unexpected manner. If the sensor is unexpectedly turned off or the connection being used to transfer the data from the sensor is hacked, the data can be missing or altered, which in turn impacts how the data is analyzed and acted upon. Imagine if an image is deleted from the video being fed from a security camera to a processor designed to monitor changes in the viewable area from frame to frame. This disconnect between what the sensor (security camera) captures and the processor could allow for theft, vandalism, or undetected hazards. Designers can prevent this by using secure connections that can be established in a variety of ways. Options include creating a Virtual Private Network (VPN) with communication carriers or with a secure cloud-based application. Costs and deployment complexity are considerations when choosing from these options.

Efficiently creating an ecosystem

Taking all these criteria into account and the balance of the functional elements such as connections to cloud-based applications like analytics or storage, a developer could lose valuable time reaching the market in the evaluation process for each device option. Equally inefficient would be trying to design everything themselves from the board on up. While this would ensure each functional block was designed for a specific purpose and intentionally tied to the corresponding functional blocks, the process would be time consuming and expensive. Similarly, taking the time to test individual Commercial Off-The-Shelf (COTS) hardware items would cost precious time and require a significant resource allocation. Working with a solution provider can save development time.

A system architect should first assess their company's core competency and

ways to leverage this in order to create differentiation. They would then select a solution partner to deliver an integrated solution in terms of hardware, software, and related services for the balance of the application. Given the wide range of use cases and associated technologies able to support them, it is important to explore all architectures and levels of integration. Quite often there is more than one way to deliver a solution to market, and factors such as cost, engineering competencies, and time to market will be key to the deployment strategy.

For example, two companies seeking to deploy a digital signage solution may go about it in completely different manners. One, whose experience is creating general purpose desktop computers, would look to differentiate through hardware features such as number of screens supported or means of connectivity and use a commercially available,



SERIOUS ADVANTAGES FOR SERIOUS CHALLENGES

Built for harsh, mission critical environments, Sealevel's COM Express solutions offer you the advantages of custom solutions with the convenience of COTS.

- Speed to Market
- Scalable Processing
- Application Specific I/O
- Flexible Mechanical Configuration
- Vibration Resistance
- Long-term Availability



Learn more about our complete COM Express solutions by visiting www.sealevel.com/ecd123/com or by scanning the QR code to the left.

SEALEVEL[®]

sealevel.com • 864.843.4343 • sales@sealevel.com

standard content management system software package. The other company, an Independent Software Vendor (ISV), would potentially look to deliver customized software that optimizes content management with viewer analytic software. Here a preconfigured media player would satisfy the hardware requirements and align with the company's core competency.

A partner who has the ability to support both models gives designers flexibility and allows for internal and external resource optimization. A good partner would provide a robust set of technology options and be unbiased as to the best combination for meeting the requirements to maximize system capabilities. During the partner selection process, it is also critical to ensure the solutions integrator has domain-specific experience, which will go a long way toward future-proofing for interoperability, and has earned the necessary industry and

agency certifications to address security and compliance. An example here would be in healthcare. Remote patient monitoring is a booming application within the telehealth segment of the medical market. Consider an OEM or ISV whose core competency lies in the analytics associated with patient monitoring. They will need a purpose-built appliance and wireless connections for collecting and managing the data flow. In this example they would benefit from a partner that is both ISO-13485:2003 certified for medical devices and has experience delivering Health Insurance Portability and Accountability Act (HIPAA) privacy and security-compliant solutions. Beyond these capabilities, a partner must also be in a position to explain things like the tradeoffs between securely transmitting over the public Internet as compared to creating a virtual private network for the overall application. The tradeoffs can be cost, total cost of ownership, support models, and level of security provided

Your solution is here.

Save time – and money – with embedded software solutions built to run right out of the box. Get development started quickly, with no integration required and full support for popular tools. With Micro Digital you have low-cost, no-royalty licensing, full source code, and direct programmer support. So get your project off to a great start. Visit us at www.smxrtos.com today.



800.366.2491 sales@smxrtos.com

www.smxrtos.com

ARM • Cortex • ColdFire • PowerPC • CodeWarrior • CrossWorks • GCC • IAR EWARM



Processors Supported: www.smxrtos.com/processors

Free Evaluation Kits: www.smxrtos.com/eval

Free Demos: www.smxrtos.com/demo

TS-4710 High End CPU Module

pricing starts at **\$138** qty 100

\$155 qty 10



TS-4710 shown mounted on TS-8160 baseboard with PC/104 bus

TS-4710 Features

- Up to 1066MHz CPU w/ 512MB RAM
- User-Programmable 8K LUT FPGA
- Boots Linux in under a half second
- Robust DoubleStore Flash storage
- LCD video output up to WUXGA
- USB2, Ethernet, PCIE, SPI, 6 UARTS
- Touch Panels available

Other TS-SOCKET CPUs

- TS-4200: Atmel ARM9, super low power
- TS-4600: 450MHz at very low cost
- TS-4712: like TS-4710 + 2 ethernet
- TS-4800: 800MHz iMX515 with video

TS-SOCKET Benefits

- Simplifies custom embedded systems
- Rapid design process with CPU Cores
- COTS development boards available
- Design your own baseboard or use our design services
- Interchangeable for future upgrades

➔ Design your solution with one of our engineers

- Over 25 years in business
- Never discontinued a product
- Engineers on Tech Support
- Open Source Vision
- Custom designs with excellent pricing and turn-around time
- Most products ship next day



We use our stuff.

visit our TS-7800 powered website at

www.embeddedARM.com

(480) 837-5200

as compared to perceived value in the market, time to market, and opportunity for differentiation. All of these considerations come into play before determining the bill of material or component selection.

System deployment: The last piece of the puzzle

Once the design and production of the application is complete, deploying the solution in the field is where interoperability and security are most critical. Connecting the devices to the network, provisioning devices (often across multiple carriers), and effectively tunneling through firewalls are just a few of the tasks that may be needed for a successful deployment. These tactical details, optimized for the user experience, create a solution that makes the technology “invisible” to the end user and allows them to focus on the benefits the product provides. For example, people expect their wearable health monitoring device to automatically connect to Wi-Fi in the home and cellular outside of the home without disruption of service. They would not want to have to stop and change settings because of a location change. The ability to make this transparent is a key element to a successful deployment. Not only are these critical steps in the process, but they also provide opportunities for new revenue streams through services provided to the end user.

Additional features can be designed into the device itself to allow for monitoring the health or security of the unit and notification if something goes awry. Features such as temperature sensors, shock and vibration sensors, or the ability to record the amount of time the system has been running can be added and analyzed against calibration points (such as Mean Time Between Failure or MTBF) to proactively schedule routine maintenance and parts replacement. Solutions integrators can provide these services along with others such as technical call centers, warranty repair, and installation. Figure 2 offers a blueprint for the ranges of services required to support a deployment through its life cycle.



Figure 2 | A wide range of services are required to support an M2M deployment through its life cycle.

Deploying today in the absence of tomorrow’s standards

There are as many potential solutions for deploying an M2M system as there are use cases to address. For companies seeking to capitalize on the growth potential offered by the Internet of Things (IoT), a degree of scale is needed to enable profit maximization. As the deployment of M2M models expands beyond its legacy roots in telematics, standards will start to emerge to make interoperability among the enabling elements of an M2M signal chain common across a given market segment at a minimum. Organizations like oneM2M are striving to develop these standards but, in the interim, working with a partner is the most practical, effective, and efficient way to move forward. Choosing a partner, such as the Rorke Global Solutions (RGS) unit of Avnet Embedded, who brings a wide range of technology options and different levels of integration, along with the complimentary life cycle of services to implement the solution is critical. Ensuring that partner is not predisposed to one solution versus considering all and that the partner has domain-specific experience exponentially increases the probability of success. Taking the time to assess and select the right partner will lead to faster deployment, better end customer satisfaction, and optimized revenue in the long run. **ECD**



Carl Cohen is Director of M2M and Emerging Technologies Solution Architect at Avnet Embedded.

Avnet Embedded
www.avnet.com

Follow:

» When your network cannot fail «

Trust Kontron carrier grade rackmount servers for HA network communication applications.

Kontron's proven 'long-life' carrier grade communication servers are never your weakest link when deploying telecom network applications. Each server is expertly engineered to supply you with exemplary Intel® performance, massive I/O bandwidth, lights-out management, innovative "floating" vibration suppression, NEBS-3 and ETSI-compliance, plus many more carrier-critical features. Talk to a Kontron carrier grade server specialist to find out how to get started.



CG1200 – 1U Server

- » Intel® Xeon® Single or Dual 8-Core Processor E5-2448L
- » Optimized for Gen3 PCI-E IO card
- » Up to four hot-swappable 2.5" SAS HDDs/SATA SSDs



CG2200 – 2U Server

- » Dual Intel® Xeon® 10-Core Processor E5-2600 v2 Series
- » Optimized for PCI-E IO cards; quad rear GbE NIC ports
- » Hot-Swap 2.5" SAS/SATA HDDs

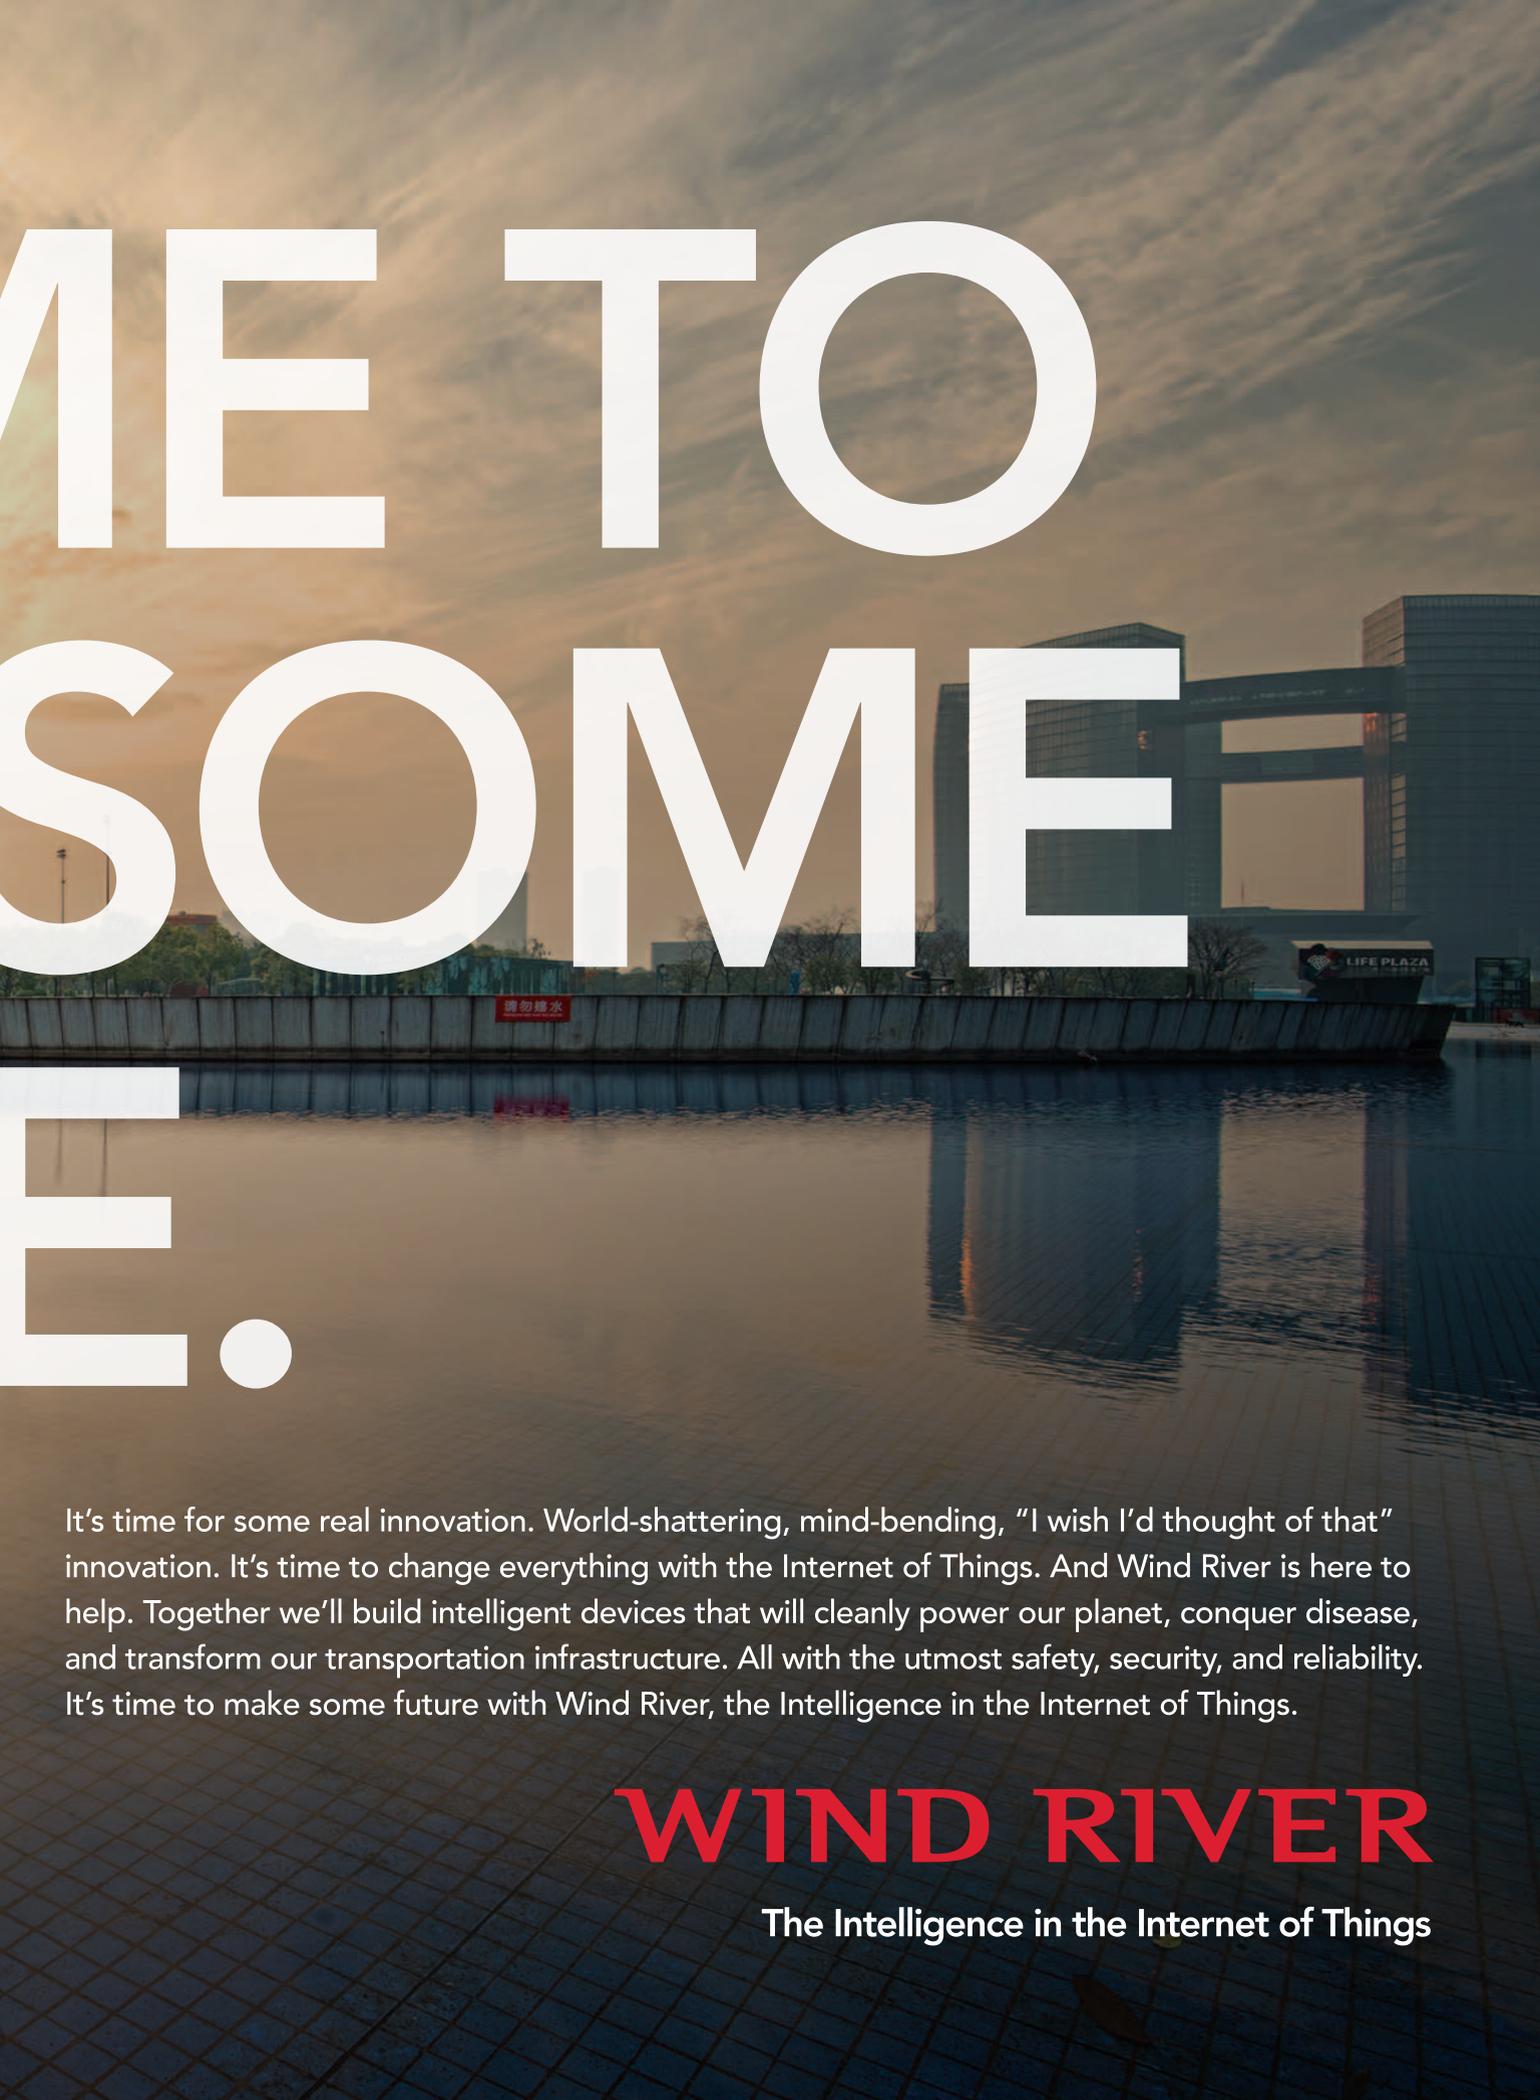
Integrated platforms that bring network infrastructure to life kontron.com/communications

The pulse of innovation



IT'S TIME TO MAKE OUR FUTURE

To find out more visit us at iot.windriver.com



IT'S TIME TO MAKE SOME THINGS.

It's time for some real innovation. World-shattering, mind-bending, "I wish I'd thought of that" innovation. It's time to change everything with the Internet of Things. And Wind River is here to help. Together we'll build intelligent devices that will cleanly power our planet, conquer disease, and transform our transportation infrastructure. All with the utmost safety, security, and reliability. It's time to make some future with Wind River, the Intelligence in the Internet of Things.

WIND RIVER

The Intelligence in the Internet of Things



M2M integration platforms enable complex IoT systems

By Robert Andres

While some M2M solutions can be relatively simple, with a single application operating upon data from a single type of device over one type of communications system, many solutions are becoming more complex as devices proliferate through the world in the age of the Internet of Things (IoT). Advanced M2M solutions may have multiple services on the edge node, several different connectivity requirements, various data consumers, and geographically diverse services and devices. These distributed systems require an intermediary system, also called an M2M integration platform, to connect them seamlessly to business applications.

Some Machine to Machine (M2M) solutions look fairly simple, with a single business application, one connectivity option, one type of service gateway, and one device or sensor on the other end (Figure 1). However, these simple M2M systems are being replaced by more complex systems as devices are quickly added to the Internet of Things (IoT).

Today M2M applications and cloud computing are combining to create significant new capabilities in which input – from machines, people, sensors, video streams, maps, newsfeeds, and more – is digitized and placed onto networks. These inputs are integrated into systems that connect devices, people, processes, and knowledge to enable collective awareness, efficiencies, and better decision-making in the Enterprise. As a

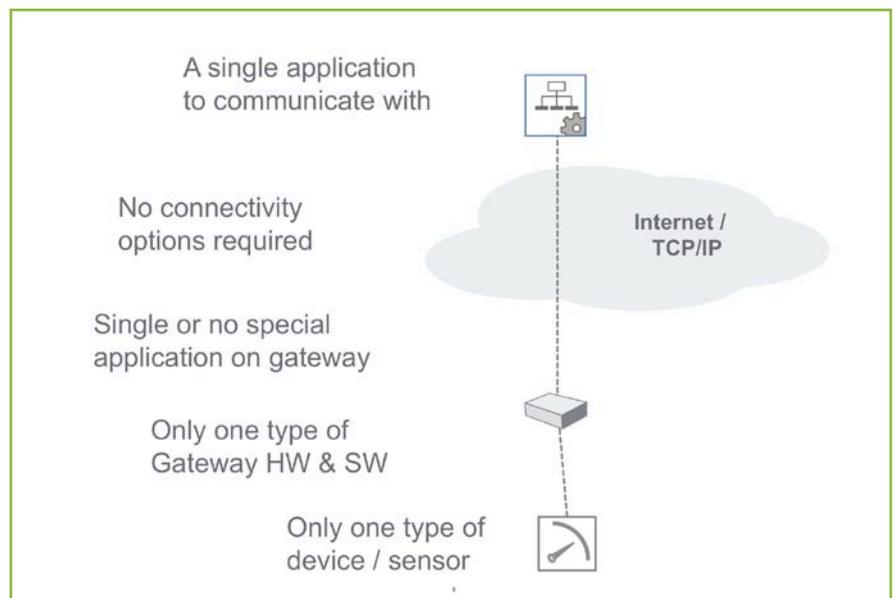


Figure 1 | Simple M2M projects connecting a single application with one type of device are becoming rarer.

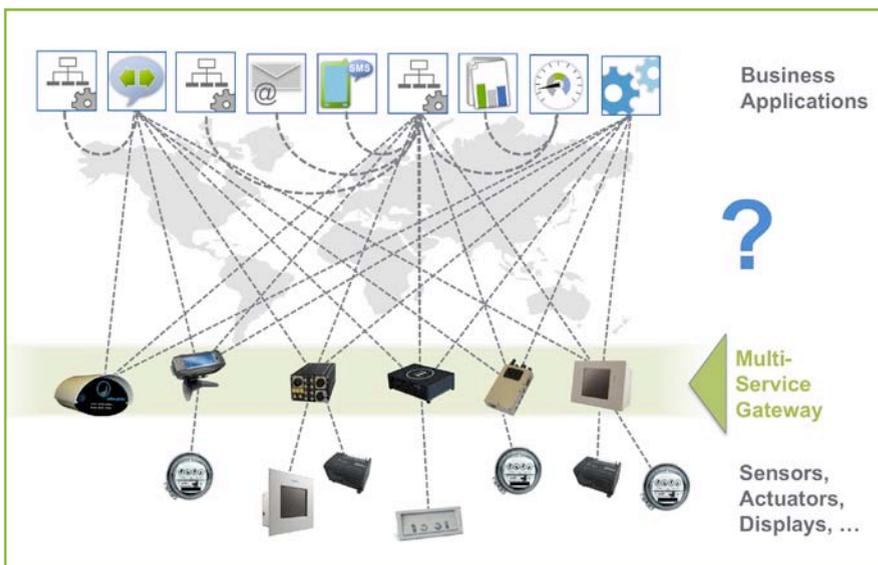


Figure 2 | Complex M2M projects combine several business applications, connectivity options, gateways, and back-end displays.

result, M2M applications are much more complex, with multiple services on the edge node, various connectivity options and customer specific business logic not only in the data center but also embedded into edge devices. Plus, the services may be geographically dispersed and have several device data consumers (Figure 2).

Thus far, the M2M marketplace has been full of hundreds of “piecemeal” technologies that can be cobbled together from disparate vendors that have, for the most part, never really spoken, communicated or learned from each other. Industry tools with open protocols that support interoperability over long application development lifecycles are necessary to unite these diverse technologies.

An M2M integration platform designed to act as an intermediary system between the distributed devices and the applications making use of the data can reconcile the varied technologies found in complex M2M projects. An effective M2M integration platform must do several things:

- Act as an operating system for the IoT, enabling the transfer of device data independent of any programming language, platform, or operating system
- Offer the means to perform effective lifecycle management of the devices in the field
- Integrate seamlessly with the Enterprise IT world, using IT best practice and architecture approaches while implementing optimal M2M technologies

Integration platform as a service

In the enterprise IT world, IT research and advisory firm Gartner has introduced the concept of integration Platform as a Service (iPaaS, Figure 3 on following page) as a specific category within the PaaS offerings, using the following definition:

An iPaaS offering provides users with a combination of cloud services – collectively called integration platform services to develop, execute, and manage integration flows. Integration flows running on iPaaS can connect, in a many-to-many fashion, any combination of on-premises and off-premises applications, services, processes, and data. (Source: Gartner, Inc. Enterprise Application & Architecture Summit, “Cloud Services Integration: How to Get Greater Business Value From Your Cloud Investments,” Massimo Pezzini, March 2013)

iPaaS offerings usually combine cloud services for protocol bridging, messaging transports, transformation, routing, service virtualization, adapters, orchestration, partner community management, managed file transfer, registry/repository, development tools, and others.

PRODUCT SPOTLIGHT

SENSORAY

embedded electronics experts

Made & supported in the USA

Specializing in the development of off-the-shelf and **custom OEM embedded electronics** for medical, industrial and security applications.



Model 953-ET Rugged, Industrial Strength PCIe/104 A/V Codec

- 4 NTSC/PAL video input/outputs
- 4 stereo audio inputs/outputs
- H.264 HP@L3, MPEG-4 ASP, MJPEG MJPEG video; AAC, G.711, PCM audio
- Ultra-low latency video preview concurrent w/compressed capture
- Full duplex hardware encode/decode
- Text overlay, GPIO

www.SENSORAY.com/EM12/953



Model 2224 | Small Form Factor High Quality, USB H.264 HD-SDI A/V Encoder

- High Profile H.264 HD/SD encoding
- Transport Stream Format
- HD (1080p/1080i, 720p) and SD (NTSC, PAL)
- Audio from SDI or external
- Blu-ray compatible
- HD/SD uncompressed frame capture
- Full screen 16-bit color graphics/text overlays with transparency
- Low latency uncompressed preview

www.SENSORAY.com/EM12/2224



Model 826 Measurement & Control

- 6 advanced 32-bit counters
- 16 analog inputs, 16-bit, 300 ks/s
- 8 analog outputs, 16-bit, 900 ks/s
- 48 digital I/Os with edge capture
- Supports incremental encoders, pwm/pulse generation
- Watchdog timer

www.SENSORAY.com/EM12/826

SENSORAY.com | 503.684.8005

An iPaaS connects M2M solutions to the distributed systems in the field and the enterprise application and IT management world. The distributed device “network” is one end of an integration flow, the enterprise application is the other.

iPaaS offerings are superior to previous methods such as when data from embedded devices is downloaded manually or through some proprietary means within a dedicated infrastructure. Without an iPaaS it can take weeks, months, and even years to plan, procure, and deploy IT infrastructure to connect embedded devices to the network and capture valuable data. An integration platform, on the other hand, simplifies the project and shortens the development lifecycle by quickly connecting devices to the cloud.

An operating system for the Internet of Things

The ideal M2M integration platform is middleware that functions like an operating system for the Internet of Things – an intermediate system between the distributed devices and the applications making use of the data coming from these devices. Any communication must be two-way in nature, allowing those applications to control and manage the devices where required. The system must enable the transfer of device data independent of any other language, platform, or operating system to accommodate the complex nature of M2M projects.

Developers can think of this middleware partly as an Enterprise Service Bus (ESB)

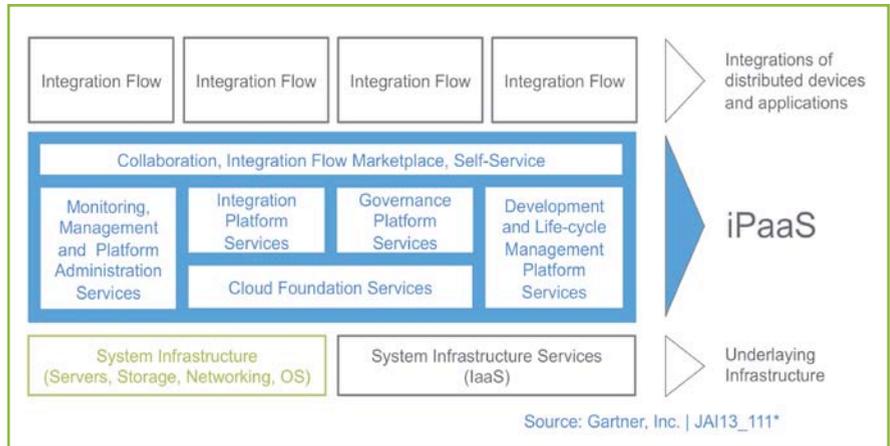


Figure 3 | As defined by Gartner, integration Platform as a Service (iPaaS) offerings provide users with services to manage integration flows.

for machines. An ESB is a software architecture model used for designing and implementing the interaction and communication between mutually interacting software applications in a Service-Oriented Architecture (SOA) manner using a lightweight, ubiquitous integration backbone. ESB solutions hide complexity, simplify access, and allow developers to access and interact with other software components in a generic way, while handling complex details in the background. An “ESB for Machines” can be implemented to connect distributed systems to business applications while effectively separating the consumers and producers of data.

With this type of unifying platform, M2M solutions can provide easy integration of different device data systems and applications to enable the delivery of data to the enterprise.

Devices and applications publish data onto the bus depending on local conditions and business logic. Data consumers receive it in real time, based on prior selections of topics of interest, while concurrently the data is stored in a self-configuring database. This basic functionality is exposed via a set of standard web service APIs providing a set of connectors with the ability to quickly add more as new services and technologies emerge.

On the device side, the platform provides tools for device management including software, firmware, and configuration. Managing the devices is just as important as managing the data, and the ability to achieve this across a geographically dispersed device population without the need for site visits by skilled engineers produces potentially thousands of dollars per month in the

People-counting demo shows power of integrated IoT applications

Throughout JavaOne and Oracle OpenWorld 2013, the movements of attendees were tracked in an impressive people-counting demonstration called IoT in Motion to show that powerful applications can be rapidly put together and gather a host of data with relatively little expense or effort. IoT in Motion was a collaboration based on Eurotech hardware, Hitachi SuperJ Applications Ecosystem, Oracle Java SE Embedded, Oracle Fusion Middleware, and Oracle Business Intelligence products.

Throughout the conference the people-counting application efficiently counted and tracked conference attendees in various locations to reveal the power and utility of end-to-end data collection and management technologies. Trend data and analytics of those traffic patterns were then revealed during the JavaOne Strategy Keynote and Oracle OpenWorld keynote presentations. With the system, Oracle was able to see when people arrived for events, when they left, and how that correlated with the event schedule in real-time.

One session included a coffee break, and data gleaned from the IoT in Motion demo showed that most attendees left for the coffee break and then returned for the session without a significant drop-off. This type of data can help event organizers plan for future events and even make last minute changes to ensure maximum attendance and plan for traffic flow.

Development of IoT in Motion took less than six weeks, highlighting the ease of implementation, time to market, and return on investment possibilities of a solid IoT integration platform.

ongoing cost of running and maintaining the system.

Eurotech's Everyware Cloud is an iPaaS specifically designed to provide "the glue" for IoT or M2M solutions and remote systems.

Increasing M2M integration efficiency
Without a unifying platform in an M2M solution, systems have multiple integration methods that lead to inconsistency and higher costs of management and change. No matter which solution is adopted, an integration platform is essential to solving the complex M2M application issues customers face today. **ECD**



Robert Andres
is Corporate
Marketing Director
at Eurotech.

Eurotech
www.eurotech.com

Follow:     

Embedded
COMPUTING
DESIGN
Editor's Choice

Editor's Choice



Big MCU performance in a small package

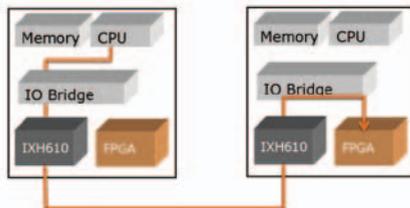
Embedded devices increasingly require more with less – more features and performance with lower cost and fewer MCUs. In response, Microchip Technology Inc. introduced the high-performance PIC32MZ family of 32-bit MCUs that boasts performance benchmarks of 330 DMIPS and 3.28 CoreMarks per MHz in a package as small as 9 mm x 9 mm.

The MCUs have a spacious 2 MB of live-update flash and 512 KB RAM for simultaneous operation of multiple protocol stacks. For mixed-signal applications, they include a 12-bit 28 MSps ADC. Peripheral pin select allows for mapping I/O, and a hardware crypto engine runs in parallel with the core to free it for other tasks. The PIC32MZ uses the Imagination Technologies MIPS microAptiv core that provides 159 new DSP instructions that execute DSP algorithms at up to 75 percent fewer cycles than Microchip's previous MCU families. See more features on the block diagram: <http://opsy.st/27Ua8Cn>

Microchip Technology Inc.
www.microchip.com
www.embedded-computing.com/p9916299

Need a Fast Network

PCI Express® Networks



Fast Data Transfers

Ultra-low latency PCI Express® networks deliver a fast private network for transferring data between systems. FPGA, high availability, and transaction processing based systems benefit from the real-time performance of PCI Express®. Dolphin delivers a complete PCI Express® network solution.



Learn how PCI Express® improves your application's performance

www.dolphinics.com





Code reviews take a page from social networks

By Curt Schwaderer, Technology Editor

Code reviews are an important part of software development, but the usual methods are often inefficient and disconnected from how developers typically work. Integrating social media-like collaborative features in the IDE can improve reviewing efficiency and effectiveness.

It's quiz time – stop and think for a moment. What constitutes the biggest waste of time in your software development cycle? When I first pondered this question, my mind immediately went to one topic: code reviews. Or perhaps I should rephrase the question – what's the most often ignored component of your software development cycle? It's not uncommon for development teams to simply ignore code reviews altogether!

With a skeptical chuckle I began my search for development tools relating to code reviews. Among all the information relating to code repositories and agile methodology information, I discovered a fresh, new approach to code reviews.

The trouble with code reviews

I open my mail – “Subject: code review tomorrow.” Ugggh – a shiver rolls down my spine. My mind immediately shifts to how many of my tasks aren't going

to get done this week now. And sitting in a room for one to two hours when I only really have feedback on about 20 percent of the code feels like a waste of time.

I don't think I'm alone in this sentiment. But that's not to say code reviews are useless – quite the contrary if done effectively and efficiently. Some considerations to optimize the effectiveness of code reviews include:

- Pick the right amount of code being reviewed. Too many lines of code and it's difficult to expect the reviewers to spend the time required to have good feedback. Too few lines of code and it takes too many meetings to get through all the code.
- Providing ample time for the reviewers. Reviewers usually write code too. The value of the feedback is limited if the reviewer simply

doesn't have time to do a decent job reviewing the code within the time allotted.

- Don't waste time during the meeting. Running effective code reviews can be difficult. Often reviews devolve into the code owner going through dozens of files line by line explaining what the code does. Sometimes issues get brought up for the first time in the code review by the owner, which doesn't give the reviewers adequate time to think through options and determine the best solution for the application.

One overall problem is that code reviews require reviewers to multi-task. Multi-tasking can introduce inefficiencies, especially when the multi-tasking requires the developer to stop their development, review code, then be somewhere at a fixed time to listen and comment on the code. However, it is possible to add some efficiency back in.

A new approach

The many difficulties and inefficiencies with code reviews requires a significant re-think of the process. Klocwork has applied a collaborative social media paradigm to the code review process, then integrated this new paradigm within popular integrated development environments. The product is called “Cahoots” and the paradigm promises to streamline code reviews, maximize efficiency, and eliminate wasted time.

How it works

What does it mean to “socialize” the code review process? If you’re familiar with how social media works (who isn’t these days?), it’s a community of individuals that interact by posting opinions and information while also commenting on others’ spaces.

Developers are constantly writing code under time pressures and typically do their best to implement clean and efficient algorithms within the time constraints. Multiple times during code development, a developer thinks, “I know George has experience with this kind of algorithm” or “this isn’t the best algorithm, but I’ll come back to it later.” Unfortunately, the best of intentions often ends in no action as the developers move on to the next important implementation task.

Let’s apply a social media paradigm to code development. The implementer’s source files represent their “message board.” While writing the code, any questions or issues that arise can be annotated without leaving the environment. The implementer may already do this by adding comments in the code that point out possible problem areas or areas that need refactoring or optimization later. But the problem with comments is they have no proactive properties – they sit in the code and the implementer must keep notes and/or remember the areas they annotated.

Bringing these annotations to life at the same time the code is being written has some powerful benefits. If these annotations behave more like social media broadcasts by the implementer, followers can become immediately aware of these thoughts and respond more quickly and on their own

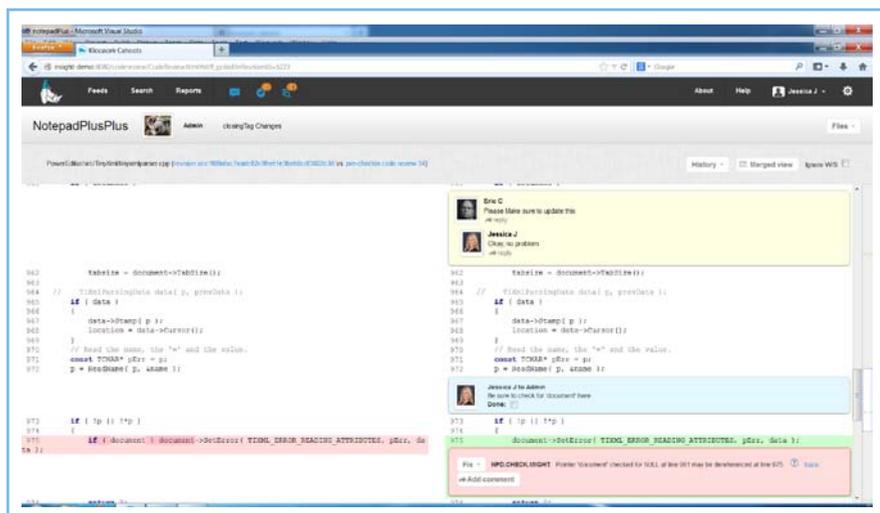


Figure 1 | Notice the social media feel inside the IDE – users comment and respond to highlighted code blocks until the review area is cleared.

timeline. This maximizes the opportunity for a reviewer to study the code, ask specific questions, or provide additional options. This kind of environment reduces a large group code review task into many small request/response tasks that can be responded to as time allows. The system can also keep track of these comments, tasks, and responses and not pass the implementation phase of the project until all of the annotations are cleared.

Cahoots nuts and bolts

The Cahoots environment can be used from within a variety of popular Integrated Development Environments (IDEs) or by using a command line tool. The tool can be used with any text-based file and has syntax highlighting and special options for C, C++, C#, Python, XML, Java, and other popular programming languages. The tool doesn’t require the implementer to switch back and forth between their code and some logging/notes tool, which saves time and makes adoption much easier.

The user interface is intuitive and allows implementers to launch a review and get everyone participating. Adding re-viewers and tags directs the right people to specific areas of the code to lower the learning curve for everyone and get re-viewers with the right expertise spending quality review time on the right code (Figure 1).

The tool also integrates with bug tracking systems, email servers, and websites.

Reporting, status, and metrics

It can be difficult to measure the effectiveness or coverage of a code review. Sure, you know what files are covered – but where was the time spent? What files were glossed over? The Cahoots system enables anyone to filter review information based on code revision, actions, and comment/response annotations by creating full text and keyword searches. This way each reviewer can efficiently identify specific areas where their expertise is required. Changes in the code can be looked at in-line or side-by-side. The drag-and-drop reports show the effectiveness and frequency of code reviews for each developer. Project leaders can quickly identify if there are people who need help managing their reviews and people who are making frequent use of the code review tool.

The reports also enable searching of reviews, how long they are taking, and if any reviews are stuck. Reporting tools help developers keep track of what they are finished with and what they need to do next (Figure 2 on following page).

Roles

The Cahoots environment defines multiple roles within the environment – the originator/developer, the reviewer, or a moderator. The developer is the person implementing the initial code. At any point (perhaps they’ve created the overall stubbed out solution, or maybe they are implementing specific classes or algorithms) the developer can select the “Create Cahoots Code Review” and begin the process. Reviewers can be



COMMELL
We offer advanced and reliable IPC products

Intel® 4th generation Core™ CPU
LV-67N Mini-ITX

- Support Intel® Haswell Desktop Core™ i3/i5/i7
- Intel® Q87 Express chipset, DDR3 up to 16GB
- VGA/DVI/DP/LVDS, 2 x Giga LAN, HD Audio, SATAIII
- USB3.0, RS232/422/485, PCIE x 16, Mini-PCIe

LV-67M Mini-ITX

- Support Intel® Haswell Mobile Core™ i3/i5/i7
- Intel® QM87 Express chipset, DDR3 up to 16GB
- VGA/DVI/DP/LVDS, 2 x Giga LAN, HD Audio, SATAIII
- USB3.0, RS232/422/485, PCIE x 16, Mini-PCIe

Mini-PCIe & Mini-PCI Card

MPX-25858	MPX-2515	MP-6100E
Video Capture	CAN 2.0B card	4-ch H/W H.264 video capture
8-CH Video in	Microchip 2515	MP-954
8-CH Audio in	ISO-11898	
240fps @D1	1 Mb/s	
Windows SDK	API & SDK	RS232/422/485

Embedded Boards with Intel® 3rd generation Core™ Processors

HS-B70 PICMG 1.3 Half-size Computer

- Intel® Ivy Bridge mobile Core™ i3/i5/i7
- Intel® QM77 Express chipset
- 2 x DDR3 SO-DIMM up to 16 GB
- 2 x Giga LAN(iAMT), SATAIII
- RS232/422/485, USB3.0, Mini-PCIe
- VGA, LVDS, DVI, HD Audio

LS-37B 3.5" Miniboard

- Intel® Ivy Bridge mobile Core™ i3/i5/i7
- Intel® QM77 Express chipset
- DDR3 up to 8 GB, VGA, LVDS, DVI
- Giga LAN(iAMT), HD Audio
- RS232/422/485, USB3.0, SATAIII
- SIM socket, GPIO, Mini-PCIe, 9~24V DC

LV-67K Industrial Mini-ITX mainboard

- Intel® Ivy Bridge mobile Core™ i3/i5/i7
- Intel® QM77 Express chipset, LVDS
- DDR3 up to 16 GB, VGA, DVI
- 2 x Intel® Giga LAN, SATAIII
- RS232/422/485, USB 3.0
- HD Audio, Mini-PCIe socket

LS-576 5.25" Network Computers

- Intel® Ivy Bridge mobile Core™ i3/i5/i7
- Intel® QM77 chipset, VGA, DVI, LVDS
- DDR3 up to 16 GB, SATAIII
- 6 x Intel® Giga LAN(iAMT)
- HD Audio, USB 3.0
- 1 RS232 & 5 RS232/422/485
- Mini-PCIe, ATX or 9~24V DC IN

www.commell.com.tw
Taiwan Commate Computer Inc.
General Information: info@commell.com.tw
sales@tcommate.com.tw

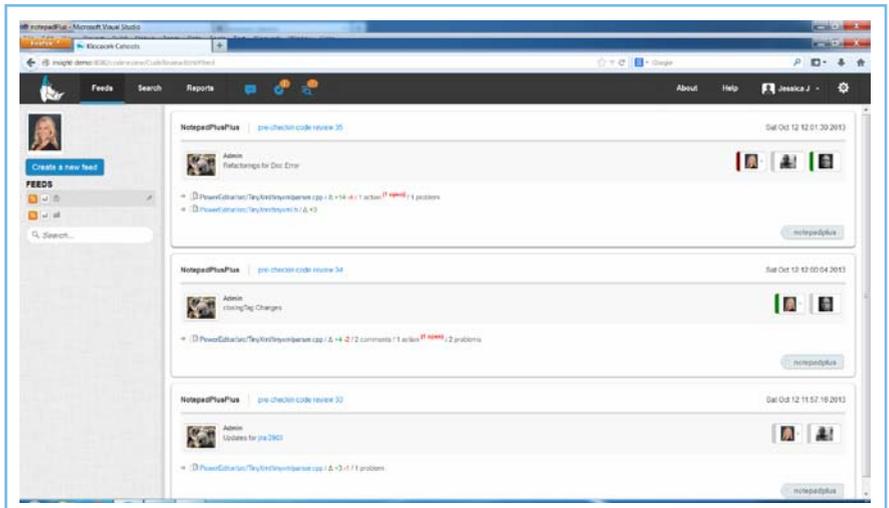


Figure 2 | Code review reports can identify what was reviewed when and by whom. Color coded reviewer information on the right indicates who has reviewed what and if they are ahead or behind schedule.

added within the code review creation window and code can be tagged based on project and feature. Each role can see all file changes and comments can be added in the code from the Cahoots code reviewer window. The reviewer role can "create action" that recommends something be changed, refer the developer to external resources, or actually change the implementation if desired.

Agile development synergy

For those that advocate things like Agile development and pair programming, perhaps the initial thought might be "this takes away from the collaborative nature of face-to-face interaction." And it does remove an element of this. Perhaps the optimal strategy is to do a mix of both strategies.

For example, let's say I'm implementing user story XYZ within our agile development process that breaks down into five tasks. In general, it would be natural to assume code reviews on each task at a minimum. There might be specific code block development where expertise is needed outside a formal code review. These areas can be quickly annotated and feedback can be given without disturbing the process and wasting the overall team's time. Doing small partial commits and performing team member pre- and post-commit reviews might be another time to use the Cahoots

environment. Cahoots is also valuable to identify areas of the code that generate "why is the code like this?" kinds of comments. Those kinds of comments naturally identify areas of the code where a team code review meeting could have significant value. Using this kind of paradigm, then knowing when to launch a code review meeting can be invaluable to maximizing the benefit and efficiency of code reviews.

Improving code reviews

Code reviews can be a sore spot for many development organizations for a variety of reasons. Applying a social media interaction paradigm within the IDE can be a solution that maximizes team member efficiency and effectiveness. And just maybe it alleviates some of the dread associated with being a participant in the code review process. **ECD**

Informational links

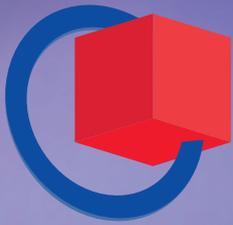
- <http://download.klocwork.com/resources/cahoots/videos/ide-code-reviews>
- <http://download.klocwork.com/resources/cahoots/videos/command-line-code-reviews>
- <http://download.klocwork.com/resources/cahoots/videos/participate-in-a-review>

For more information, contact Curt at cschwaderer@opensystemsmedia.com.

Register now and
secure your ticket!
embedded-world.de

Nürnberg, Germany

February 25 – 27, 2014



embedded world 2014

Exhibition & Conference

... it's a smarter world

Summit for Innovations

The world's largest event for Embedded-Technologies is inviting you. The international Embedded-Community trends of tomorrow will be set right here! A must attend event, to establish contacts and expand your network.

Media Partners

elektroniknet.de

computer-automation.de

energie-und-technik.de

karriere-ing.de
for Engineers and Technicians

Markt & Technik
die unverzichtbare Wochenzeitung für Elektronik

elektronik
Elektronik
Fachmedium für industrielle Anwender und Entwickler

Elektronik
automotive
Fachmedium für professionelle Automobilelektronik

ENERGIE & TECHNIK
Fachmedium für Energieeffizienz

Computer
AUTOMATION
Fachmedium der Automatisierungstechnik

MEDIZIN & elektronik
Fachmedium für Elektronik in der Medizintechnik

DESIGN & ELEKTRONIK
KNOW-HOW FÜR ENTWICKLER

Organizers

Trade Show

NürnbergMesse GmbH
Tel +49 (0) 9 11.86 06-49 12
visitorservice@nuernbergmesse.de

Conference

WEKA FACHMEDIEN GmbH
Tel +49 (0) 89.2 55 56-13 49
info@embedded-world.eu

NÜRNBERG MESSE



Verifying embedded designs with cloud computing

By Dr. Raik Brinkmann

Many industries have recognized the value of cloud computing both in terms of cost reduction through shared resources, and new capability offered by the medium. However, adoption of the cloud has been slow in the electronics segments, partly due to security concerns. This article discusses an approach to cloud usage in electronic design that breaks the security adoption barrier and highlights capabilities previously unavailable through traditional solutions.

The electronic design community has been slow to embrace a cloud computing strategy for its design projects, mostly due to concerns about IP security. An emerging approach could change that reluctance and allow more designers to leverage the cloud as a compute platform for verification. This approach has particular benefits for embedded designs making use of embedded processor subsystems and other IP from various sources.

Formal verification technologies, which are becoming the basis for more and more verification solutions, also hold the key to mitigate cloud security concerns. Design security can be assured by the technology's ability to break a verification problem down into multiple, abstract mathematical problems, and transfer them to the cloud to be solved (Figure 1), with all design detail removed. The formal verification process accomplishes this by extracting specific states to be analyzed and comparing with required properties. The design detail is retained and reapplied locally to the results of the cloud operation. This process eliminates the need to transfer IP out of the office.

The resulting solution harnesses a collection of cloud-based servers to provide a

broad range of verification capabilities, offering easy adoption and usage, on-demand performance trade-offs, and a pay-per-use business model. This is particularly relevant for IP-based embedded designs where the IP is utilized by a third party who may not have ownership of the formal technology.

Acquiring verification software through the cloud

Cloud computing verification software makes use of a "Client Tool" – that is, a component that operates on a local host computer and acts as a front end for the cloud engines, creating the mathematical proof problem sets from the design code for transmission to the cloud. The client tool includes a "linting" capability – a way to find design errors in Hardware Description Language (HDL) code – and manages communication with the cloud solution. It also allows any required debug of the results of the cloud processing to be performed back on the local host computer.

Traditionally, verification software is licensed to run on a local host computer and the entire operation from design input to results debug is accomplished using the single product. For the cloud

scenario the client tool makes the user experience identical, as if the complete product was running locally, providing a transparent feel to the cloud use model.

Practical applications

The combination of a transformational, pay-per-use business model, access to a limitless number of compute engines, and on-demand verification applications provides some interesting new benefits.

For example, this software is ideal for new users who want to add powerful formal design checks to their existing simulation-based verification flow with a minimal learning curve and setup process.

In the most basic cloud verification process, the design Register Transfer Level (RTL) code is checked locally by the included lint capability. Then, assertion synthesis is used to automatically create powerful tests for many facets of the design. The automated tests include a range of fatal design error checks, potential mismatch issues between simulation and synthesis, register and signaling initialization and toggling problems, code and Finite State Machine (FSM) coverage, and many other fault scenarios. By

providing this mechanism in the cloud, new users can take a “try and see” attitude without a time-consuming evaluation process, quickly testing their design with no painful learning curve, and getting a taste of formal technology, unlike the more traditional approach where a full tool licensing and installation process must be accomplished, usually by a vendor’s engineer visiting onsite, often coupled with a sales effort by the vendor.

Advanced users can also receive considerable benefit. Formal verification is inherently a concurrent process with the mathematical proof problems being run in parallel. The cloud solution provides a limitless number of compute servers and, as such, a verification run may be executed across the number of machines to provide optimal parallel execution. Even the cost remains static – 10 hours on a single server would cost the same as one hour on 10 servers. In the software licensing approach, enough licenses would have had to be purchased up-front to cover the full parallel usage, placing a usually small cap on the number of operations that may be run together. Parallel operation is often severely limited in most verification installations due to this.

The solution also enables on-demand usage for specific purposes, such as IP integration or a means for verification service providers who need to have access to the tools while working in a foreign environment. In each case, the software can be leveraged without the need to purchase it up front by the end-customer, ideal when a third party is involved who does not leverage formal solutions.

Cloud advantages for embedded designs

Embedded designs have some specific issues that can be mitigated through cloud application-based verification. The nature of embedded designs is that IP will be leveraged from a variety of sources. This IP will have a range of verification metrics applied to it, and may use complex interfaces for interconnection purposes (Figure 2).

Formal verification can be used to great effect in this scenario, as discussed by a leading semiconductor company at the

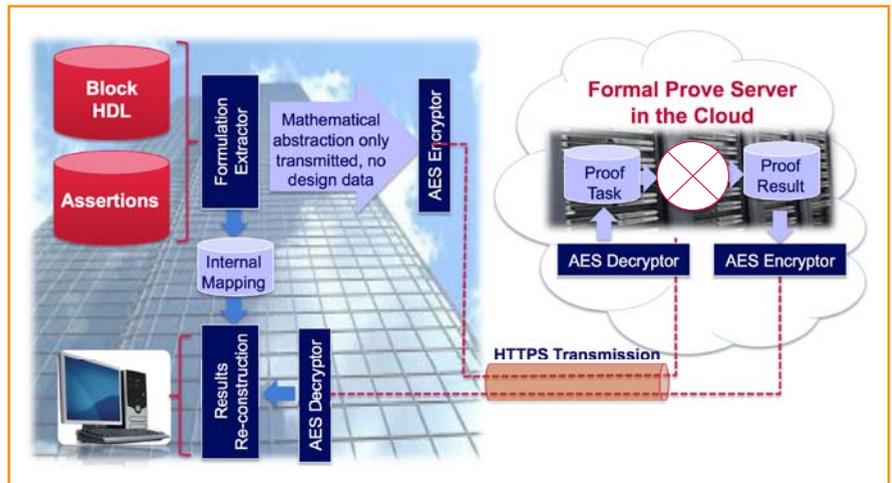


Figure 1 | Formal verification divides abstract mathematical problems, removes design detail, and transfers them to the cloud to be solved.

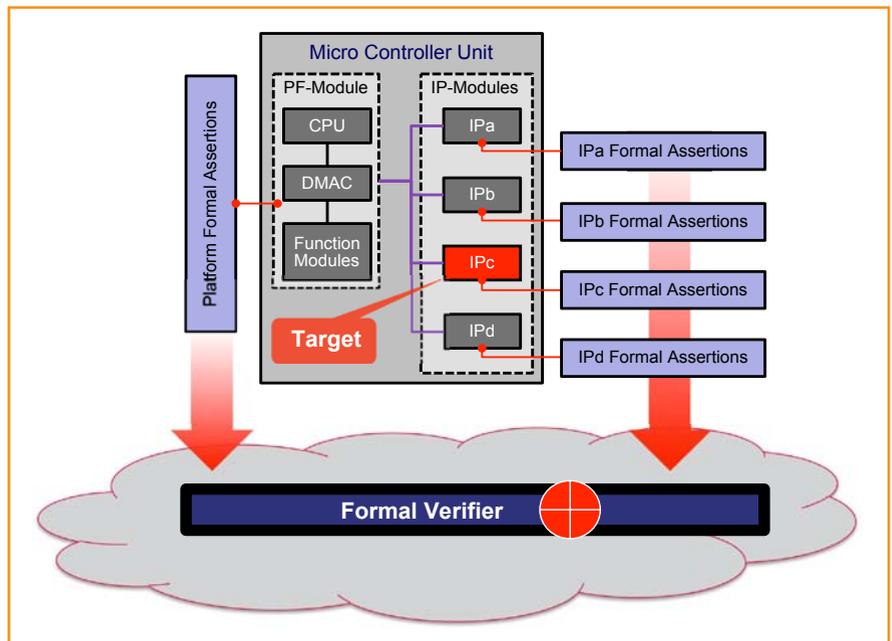


Figure 2 | Formal verification can provide a rigorous integration test environment for embedded designs leveraging IP from various sources.

recent Design Automation Conference, to provide a rigorous integration test environment. Assertions are used to specify interfaces and the expected information to be passed between the IP and platform. This investment in assertions is considered worthwhile given the reusable nature of the IP and the importance of ensuring that it is interconnected correctly in a foreign environment. Of course, if this interconnect is through a standard bus protocol – the AHB standard from ARM, for example – then a standard set of protocol assertions may be used to ensure its validity. The use of assertions in this manner has been proven to increase quality and reduce integration time, as, given the reusable nature of the IP, more

investment may be applied to ensuring it is fully tested, and the need to rewrite the assertions will have been eliminated.

What has this got to do with the cloud? The IP producer might utilize assertions in a formal environment to test IP interfaces, for example, to ensure that the communication protocol applied by the IP consumer is as specified. However, this does not mean that the IP consumer has access to formal technology. In this case, it is unlikely that the consumer will want to purchase tools and endure a lengthy evaluation process. The cloud allows IP integration analysis to be leveraged without the expense and overhead of unnecessary tool ownership by the IP consumer, saving as much as 95 percent



of the cost, simply by accessing the cloud solution for those IP checks on a pay-per-use basis.

Of course, the same is true if other formal static checks suitable for embedded design are used. For example, protocol analysis, register checking, and other System-on-Chip (SoC) style analysis are available in the cloud and might require a small number of applications in the design process. This is ideal for the embedded design team who might want to quickly create a hardware platform to utilize either an emulator or virtual models, and check that it is correctly implemented without employing a complex verification process, so they can get on with software design.

Verification within budget

One of the most significant benefits of the cloud solution is the business model and its implications on tool budgeting and embedded designer control. Traditional Electronic Design Automation (EDA) tool licensing, for example, normally requires an upfront investment, either a time-based or perpetual license. This requires the purchaser to have a good understanding of tool resource requirements, particularly

difficult to estimate with any verification solution as the use model will depend on coding quality, complexity, and other hard-to-predict factors. This is generally true regardless of application area, and becomes more complex with potential project size and team structure.

In most verification scenarios, there is a bulge of tool usage that gets larger as more code is completed and checked into the design database. As RTL coding nears completion, the demand on formal verification software will be extensive, sometimes by as much as four to five times average for complex designs such as communications and multimedia platforms, and drop off as the design is synthesized and silicon laid out. The extent of this verification bulge is partially dependent upon code quality, coverage achieved, and other metrics. As such, a certain predictable proportion of the licenses will be used throughout the entire design flow. A somewhat unpredictable number of licenses will be applied during the bulge period.

The cloud computing solution enables a combined business model where some licenses may be purchased up front and others applied on-demand using a

pay-per-use scheme. The design team is in control of its resource requirements, leveraging exactly what it needs for the variable components of the verification process and eliminating redundant licenses. This model also allows for greater financial control, shifting some of the tool expense burden from capital budgets to more appropriate operational or project money sources. **ECD**



Dr. Raik Brinkmann is President, Chief Executive Officer, and Co-founder of OneSpin Solutions. He brings to the CEO role more than

15 years of broad expertise in management, verification and industrial R&D, having served at companies in both Europe and the U.S. Dr. Brinkmann holds a Diplom Informatiker from the Clausthal Technical University, Germany and a Dr.-Ing. from the Department of Electrical Engineering at the University of Kaiserslautern, Germany.

OneSpin Solutions
www.onespin-solutions.com
Follow:

SPECIAL ADVERTISING FEATURE

Bay Trail Product Spotlights

**nanoX-BT:
Mini COM Express®
Type 10 Module**



- Latest single-, dual-, quad-core Intel® Atom™ or Intel® Celeron® processor
- Up to 4GB soldered dual channel DDR3L at 1333MHz
- Two DDI channels (optional LVDS)
- Three PCIe x1, GbE
- Two SATA 3 Gb/s, one USB 3.0, seven USB 2.0
- Extreme Rugged™ version with operating temperature range of -40°C to +85°C



ADLINK Technology Inc.
www.adlinktech.com
info@adlinktech.com
408-360-0200

**cExpress-BT:
Compact COM Express®
Type 6 Module**



- Latest single-, dual-, quad-core Intel® Atom™ or Intel® Celeron® processor
- Up to 8GB dual channel DDR3L at 1333MHz
- VGA & two DDI channels (optional LVDS)
- Three PCIe x1, GbE
- Two SATA 3 Gb/s, one USB 3.0, seven USB 2.0
- Extreme Rugged™ version with operating temperature range of -40°C to +85°C



ADLINK Technology Inc.
www.adlinktech.com
info@adlinktech.com
408-360-0200

What's **Hot** in 2014 in Embedded?

2014 is looking to be an exciting year for embedded computing. Many markets are demanding more embedded technology in their applications to deliver more performance at lower power and in smaller packages, and the industry is responding. The embedded industry has made possible the intelligence and powerful capabilities of the current state of electronics, and we're looking forward to seeing what trends and developments are next.

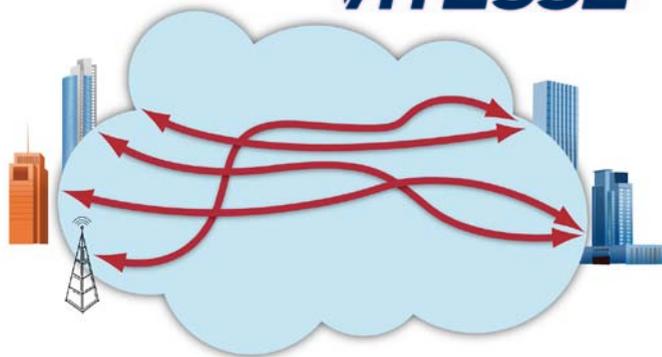
In this special advertising section, leading technology companies wrote in about what they see as the hot technologies to watch for in 2014. Last year's hot topics included automotive infotainment, Machine-to-Machine (M2M) connectivity, FPGA and System-on-Chip (SoC) integration, and software monetization. This year, smart, connected systems continue to be the hot trends, with vendors writing in about Network Functions Virtualization's role in the connected world of the Internet of Things, using ARM and x86 for next-gen intelligent devices, and securing wireless devices, along with improving mobile software and strengthening video performance in the cloud.

These trends present opportunities for embedded designers to achieve faster performance, better connectivity, tighter security, and a multitude of other features that address the current and future demands for intelligent embedded devices and systems.

Embedded COMPUTING DESIGN
What's Hot
 in 2014 in Embedded?



VITESSE

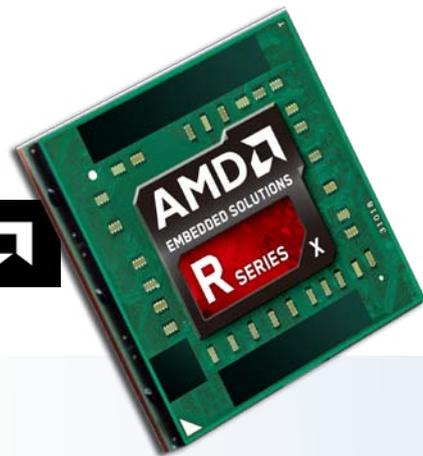


AdaCore
 The GNAT Pro Company

WIND RIVER



AMD



NFV and IoT are Forging New Paths for Intelligent Connected Systems

WIND RIVER

Wind River

Network functions virtualization (NFV) and the Internet of Things (IoT) are two major concepts that will continue to be hot areas moving forward.

The telecommunications industry is facing serious challenges from the exponential growth in data traffic due to the flood of devices connecting to an increasingly strained network. At the same time, operators are eager to roll out advanced services to generate new revenue streams and drive higher average revenue per user and device. They are also looking to reduce the cost of scaling and operating the network by increasing efficiency in hardware utilization and energy consumption.

Given this environment, virtualization is increasingly being viewed as an important approach. Operators are looking toward NFV to support the transition to scalable platforms that enable flexible deployment of network services. Ultimately, through virtualization, operators can flexibly run intelligent services anywhere on the network, increasing network efficiency and significantly lowering operational costs.

Currently, software platforms for NFV have to be built by integrating multiple products from different vendors. The work to integrate and test various products creates new obstacles, delays and increased development costs. In addition, existing virtualization options today were not originally intended for the telco environment and yield performance that is lacking for the rigorous demands of this industry. Instead of a force fit, what's needed is virtualization that is designed to meet carrier grade needs.

There is a new category of NFV vendors emerging to fill these gaps. For example, Wind River® offers open source-based, real-time kernel virtualization technologies, such as Wind River Open Virtualization, that offer high performance, with near-native hardware performance speeds. With this type of technology designed specifically for the telco world, companies can achieve real-time virtualization that supports the rigorous SLAs of a carrier network and enables them to gain the flexibility, scalability, and cost and energy benefits cloud data centers already enjoy.

The other key concept that cannot be ignored is IoT. This new paradigm is being driven by the convergence of increasingly connected devices, cloud economics for compute and data, and the acceleration of big data analytics to extract value from data.

The IoT transformation is happening now, but the term actually describes a network approach that Wind River has been practicing and refining for decades: powering interconnected, automated systems. With more than 30 years of embedded leadership and innovation, Wind River software is at the heart of more than 1 billion embedded computing devices around the world. Wind River has translated its extensive embedded experience into deep cross-sector expertise, making it a trusted partner for customers in a wide range of industries. Few companies are better equipped than Wind River to help organizations determine how to migrate to IoT right now.

Companies are increasingly looking for innovation, embracing end-to-end solutions that move beyond the device to address the entire data cycle from device to cloud and back. For example, companies can turn to Wind River technology and expertise to build machine-to-machine (M2M) applications and devices that communicate with the cloud. The company's software platform for IoT, Wind River Intelligent Device Platform, is a complete software development environment that provides ready-to-use components to secure, manage, and connect IoT gateways.



IoT is generating new opportunities that can be seized right now by companies to develop differentiating services, enhance productivity and efficiency, improve real-time decision making, solve critical problems, and develop new and innovative experiences. We are witnessing the beginning of an exciting transformation.

Wind River
www.windriver.com

An Eye on End-to-End Wireless Security

Vitesse Semiconductor

VITESSE®

2014 will be a defining year for wireless security in the telecommunications industry. According to Cisco's most recent Visual Networking Index, there will be more than 10 billion mobile devices/connections by 2017, including over 1.7 billion M2M connections. Increasing the number of network connections by definition increases the vulnerabilities. And by "security," we're not just talking about data security and encryption. Authorization, authentication, and auditing of access and connections are equally critical, especially at the network edge. A significant operator concern in 2014 will be network-wide security, especially in mobile networks.

Security has long been an important issue. The continued strong growth in the number of mobile Internet connections will bring more challenges in 2014 and force operators to utilize the most advanced technologies available. New mobile devices with bandwidth-hungry applications, and the small cell networks needed to accommodate them, exponentially multiply the number of network elements required to support users. The days of network equipment residing exclusively in physically secure locations like a central office or a macro base station are gone. New types of access points directly exposed to users pose the obvious security concern. The BYOD trend introduces a layer of vulnerable access points for enterprises to protect. Small cells are also particularly susceptible to hackers, as they are often installed outdoors at street level or indoors in easy-to-reach locations.

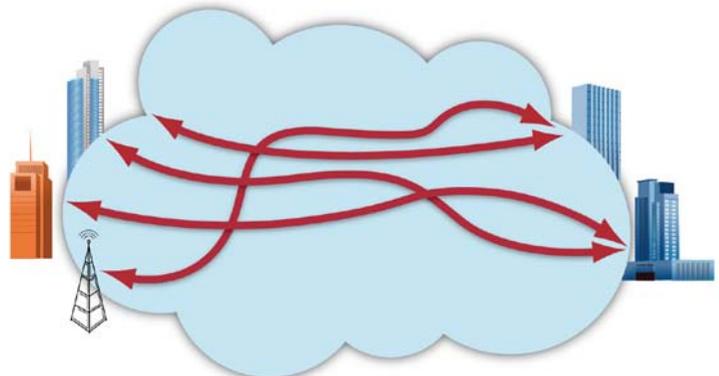
Securing the backhaul network is equally important and will require advanced encryption technologies. While these technologies must prevent security breaches, they also must ensure network timing. Small cells are typically connected via microwave and millimeter wave, which require highly accurate IEEE 1588 precision timing protocol to maintain user QoE. Securing these wireless links is pointless if the chosen encryption method is incompatible with the nanosecond-accurate timing that next-generation networks demand.

Currently, the most common network-based security protocol is IPsec, which operates at Layer-3 (L3) of the Internet protocol hierarchy. Frequently used for routed networks, IPsec usually requires dedicated standalone or embedded encryption engines. IPsec is also prohibitively expensive to deploy network-wide, and since it operates at L3, it is unable to encrypt common Layer-2 (L2) based control protocols.

In 2014, expect operators to increasingly demand more L2 encryption, also called IEEE 802.1AE MACsec, at the network edge. With service providers now offering more Carrier Ethernet L2 VPN services (such as MEF CE 2.0 and MPLS-based L2VPNs), encryption directly at the L2 Ethernet layer makes the most sense. Fortunately, new MACsec technologies make it a viable option for wired and wireless WAN security.

Vendors now have options to implement MACsec security encryption with the 1588 protocol in mind – a particularly important development for TD-LTE and LTE-A networks, which require extremely accurate timing. With the right technology implementation, OEM equipment can now accurately predict delays and correct the time stamps on a packet-by-packet basis to allow operators to secure wireless backhaul networks without sacrificing timing.

With advanced equipment, MACsec will offer peace of mind for mobile networks in 2014. As operators continue their network upgrades to accommodate the rising demand for mobile Internet, they will need strong MACsec encryption end-to-end to optimize their investments. Expect operators to require more L2 security in 2014 – MACsec over the WAN is here to stay.



Vitesse Semiconductor Corporation
www.vitesse.com



What Will be Hot in 2014 for Embedded and Mobile Software Development?

Adacore

By S. Tucker Taft, VP and Director of Language Research

The first trend we see is a convergence between *dynamic* and *static* approaches to software verification. Traditionally, computer scientists favored using static analysis and formal proof to ensure that software was safe, secure, and correct. Quality engineers and many certification authorities preferred testing to verify that the software met its requirements. But now, these approaches are being combined. So far this has mostly been seen in the academic sector, where conferences like ISSTA (<http://issta2013.inf.usi.ch/program>) are including more papers where static methods such as model checking and symbolic execution are augmenting dynamic methods such as run-time monitoring and targeted data-race detection.

In the commercial sector the integration has thus far been largely cosmetic. However, we expect more meaningful integration to occur over the next year, making this important capability more available to the software developer. One current example is the Code Contracts extension to Visual Studio from Microsoft Research (<http://research.microsoft.com/en-us/projects/contracts/>) where abstract interpretation is used to infer code contracts (pre/postconditions) from the existing code, which can then be inserted as run-time checks to support testing. AdaCore has similar contract inference capability in its CodePeer technology (see <http://adacore.com/codepeer>), and in 2014 the company will be releasing a new toolset for the verifiable SPARK 2014 language which brings a unique combination of test and formal proof to the embedded software world, based on the executable contract-based programming approach now provided by Ada 2012 (<http://spark-2014.org> and <http://ada2012.org>).

The second trend we see is toward lighter-weight *work-stealing*-based parallelism with *lock-free* synchronization mechanisms, being built on top of the underlying operating systems. This trend is driven by the steady move toward multicore, manycore, and GPU-based systems. The old, O/S-based explicit thread and lock models are simply too static and heavy-weight for an algorithm expected to scale automatically from a single-core micro-controller, to a multicore vectorizing CPU, and on to a 1,000-core GPU. We see this in new capabilities added to older languages, like the *fork/join* concurrency model in Java and the *atomic* synchronization primitives of C/C++11, as well as in newer languages like *Cilk+*, *Go*, *Rust*, and AdaCore's *ParaSail* (<http://parasail-lang.org>) and *Sparkel* (<http://sparkel.org>) experimental languages. In these languages, thousands of very light-weight potentially parallel computations (variously termed *picothreads*, *tasks*, or *goroutines*) can be created for a given algorithm, and then a work-stealing scheduler with a separate double-ended queue per processor core balances the load while preserving better locality of reference than typical symmetrical multi-processing schedulers. See <http://supertech.csail.mit.edu/papers/steal.pdf> for the classic paper on work stealing.

The final trend we see heating up in 2014 is *automated* storage management without the *pauses* associated with classic garbage collection. The ease of use of automatic garbage collection can speed up initial development and debugging of complex applications. But in the resource-constrained embedded or mobile world, these early simplifications can cause a high price to be paid later, as the application faces stringent response time demands. Alternatives are emerging in several new languages, and newly extended languages like C++11 in the form of specialized pointers that carry storage *ownership*. In C++11 these are called *unique* pointers, in the Rust language these are called *owning* pointers, and in ParaSail and Sparkel ownership is intrinsic to their normal large-object models. Another alternative is *automatic* reference counting (ARC), based on compiler support (*manual* reference counting is a long-known technique). ARC has become the standard storage management methodology for Objective-C development on iOS and Mac OSX, and is part of C++11 as *shared* pointers. Unfortunately, reference counting has measurable overhead in sequential programs, and can become a real bottleneck (or source of race conditions) in concurrent programs. Hence, for newer languages like Rust, ParaSail, and Sparkel where parallelism is a fundamental part of the language, using unique ownership seems like a better alternative, and we foresee that becoming an important new part of the embedded software developer's arsenal in 2014.

Scalable HEVC delivery and video transcoding in the cloud



Kontron

Broadcasters, cloud service providers, cable companies, carriers, mobile operators and content delivery network providers are all part of a video services eco-system. Together, they strive to ensure the creation, production, delivery and consumption of video is both cost-effective and available, anywhere.

Providers today are less concerned about 'broadcasting' and more about content, namely video. Market analysts Ooyala forecast that taken together, the share of tablet and mobile video grew by 19% in the first quarter of this year and a quarter of total tablet viewing time was spent with content more than 60 minutes long.

As both traditional mobile service providers and cable operators continue to position themselves for growth with HD multiscreen content delivery, the arrival of 4k television to the home and Ultra HD / HEVC to the mobile device is making the challenge increasingly complex.

They have to cope with a wider array of outlets – whether online, social, mobile, TVs, PCs or tablets – and a variance of video quality, file sizes, and compression codecs to stream content to. Therefore, providers need more video processing power than ever before.

The most convincing response has been to deploy most supporting functions for media delivery in the cloud. Producers of broadcast and TV programming, b2b, or live events streamers can spin up hundreds or thousands of servers in minutes and scale them down after use. They know that with the cloud, they no longer have to plan and procure IT infrastructure weeks or even months in advance.

Cloud providers, in turn, are looking to video processing hardware specialists such as Kontron to provide the powerful transcoding units to handle H.264 HD streams, and even the HEVC codec for UltraHD and 4k television, in the most efficient manner. While HEVC is in its early evolutionary stage, it will quickly become the de facto codec as the Ultra HD era catches on. HEVC processing for high-definition mobile phone content requires up to 20 conventional video processors to deliver a single UHD stream, and 80-times more video power to deliver UHD-quality video.

Kontron has harnessed the integrated graphics processing power of 4th generation Intel Core technology in its scalable SYMKLOUD Media platform. Leveraging the Intel Media SDK, video transcoding software vendors can achieve a tremendous cost versus performance potential to make both HD and Ultra-HD/HEVC in cloud infrastructure exceptionally viable.

Compared to existing server architecture with massive core count CPUs, the modular and highly scalable SYMKLOUD framework is considerably more dense, providing upwards of 50% cost savings and 20-times more performance.

This is only one half of the SYMKLOUD story. While it can cost-effectively supply a new order of magnitude of video channel density that broadcasters, CDN and cable providers require to meet new video workload demands, SYMKLOUD offers many other key advantages.

High on the list for cloud providers are system, power, and cloud management. Kontron helps with the various tools to better manage nodes that are under-utilized; or, conversely, re-utilize nodes for other applications, such as Big Data video analytics, during low traffic periods. Lastly, SYMKLOUD has been tested to support open source cloud computing orchestration software, such as OpenStack, for massively scaled provisioning of video workload instances.

Moreover, SYMKLOUD Media maximizes 42U rack spaces with its integrated load balancer and switching in its 2U high and 21-inches deep footprint. This reduces or eliminates the need for extra top-of-rack equipment and 2x to 4x Fewer Fiber/Copper Cables. As an "infrastructure-in-a-box" platform, it can seamlessly cluster 8x 2U platforms as single platform. SYMKLOUD Media also helps save IT and infrastructure energy costs with its multifaceted approach to control power.



Kontron
www.kontron.com

x86 & ARM enable next-gen intelligent embedded devices

By Kamal Khouri, Director of Embedded Products



In 2014, embedded customers will be given more choice than ever before when it comes to high performance, power efficient processors with the introduction of ARM's 64-bit ARMv8 architecture. The ARMv8 architecture is a ground-up redesign that combines the trademark power efficiency with high performance features such as 64-bit memory addressing that will enable it to tackle use cases such as compute and storage servers and next generation networking equipment.

The benefits of ARM's ARMv8 architecture has already been recognized by a number of silicon vendors with Altera, Broadcom and Samsung stating they will have ARMv8 processors, while AMD recently announced its "Hierofalcon" 64-bit ARM SoC. Only AMD, however, will offer customers a choice between 64-bit x86 and ARM processors from an established semiconductor vendor.

ARM's current generation ARMv7 architecture has powered the smartphone and tablet revolution, enabling device makers to not only create revolutionary devices but new multi-billion dollar markets. The ARMv7 architecture is an evolution of an architecture that was designed from the ground up with one purpose, to be as energy efficient as possible. While the need for power efficient processing has not diminished, customers are demanding more computation from their devices, whether it be a smartphone, tablet or servers, a need that the ARMv8 architecture will address.

The ARMv8 architecture's headline feature is its 64-bit memory addressing, which means processors based on the architecture will be able to address more than 4GB of system memory. Even in 2012, consumer devices such as Google's Nexus 10 tablet were shipping with 2GB system memory, highlighting the inevitable and fast approaching need for more memory to facilitate greater multitasking and computing larger datasets.

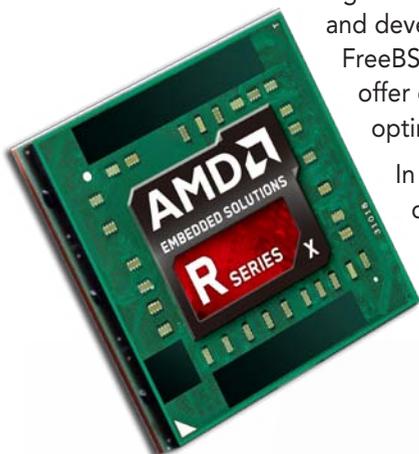
Being able to address more than 4GB memory is the most visible feature of ARMv8 but the architecture packs numerous under-the-hood improvements while retaining compatibility for the existing 32-bit ecosystem. The architecture brings performance improvements courtesy of an increase in general purpose registers and cryptography instructions have led to the first consumer-grade implementation of the ARMv8 architecture in Apple's iPhone 5s consistently performing higher than those based on previous generation architectures in independent tests.

As with any new processor architecture the software ecosystem is vital to its adoption. The software ecosystem around the ARMv8 architecture is rapidly being developed with software engineers, including those from silicon vendors such as AMD, actively participating in open source projects. The Linux kernel already supports the ARMv8 architecture and the Red Hat sponsored Fedora Linux distribution announced it will support the architecture while Canonical's Ubuntu already has an image ready to download. Developer tools are also coming together with GCC and the GNU C library already available for the ARMv8 architecture.

The ARMv8 software ecosystem will continue to mature at great pace as the architecture is designed to meet the needs of numerous embedded markets. For all the disruption ARMv8-based processors are set to cause, the embedded industry is cautious to accept a new architecture and its software ecosystem, and rightfully so. In 2014, customers will be looking to experiment with ARMv8 processors, while maintaining the need for processors that provide the highest performance and value.

While the ARMv8 architecture will move the ARM architecture into new markets, processors based on the x86 architecture will continue to give customers access to an extremely mature software stack, which includes applications and development toolchains that can run on a multitude of operating systems such as Linux, FreeBSD and Microsoft Windows. Processors based on the x86 architecture will continue to offer excellent performance thanks to the architecture's maturity enabling developers to optimize their code.

In 2014, customers looking to build intelligent embedded devices will be given a choice to pick the right product for their needs, allowing them to fully realize their vision. In 2014, only AMD will offer customers the choice between 64-bit x86 and 64-bit ARM, paving the way for customers to push to the boundaries of what is possible in an exciting time for the industry.



Advanced Micro Devices (AMD)
www.amd.com

Annapolis Micro Systems

The FPGA Systems Performance Leader

WILDSTAR OpenVPX Ecosystem

FPGA Processing Boards

1 to 3

Altera Stratix V or
Xilinx Virtex 6 or 7
FPGAs per Slot

Open VPX Storage

Up to 8 TBytes Per Slot

4 - 8 GBytes
Per Second

Input/Output Modules

Include:

Quad 130
MSps
thru

Quad 550
MSps A/D

1.5 GSps thru
5.0 GSps A/D

Quad 600
MSps D/A

Dual 1.5
GSps
thru
4.0 GSps D/A

1 to 40 Gbit
Ethernet
SDR to FDR
Infiniband

GEOINT,
Ground Stations,
SDR, Radar,
Sigint, COMINT,
ELINT, DSP,
Network
Analysis,
Encryption,
Image
Processing,
Pattern Matching,
Oil & Gas
Exploration,
Financial and
Genomic
Algorithms,

Open VPX Switch

1 to 40 Gbit
Ethernet
SDR to FDR
Infiniband



Chassis
4, 6 or 12 Slot
Up to 14G



**High Performance Signal and Data Processing
in Scalable COTS FPGA Computing Fabric**

190 Admiral Cochrane Drive, Suite 130, Annapolis, Maryland USA 21401
wfinfo@annapmicro.com USA (410) 841-2514 www.annapmicro.com



WIN BIO

World's First x86 Quad-Core SOC!



Inspiring Embedded Innovation.

Introducing the AMD Embedded G-Series SOC.

- 113% increase in CPU performance
- 20% increase in graphics performance
- 33% footprint reduction
- Stunning HD multimedia
- Low power design
- Scalable, flexible platform

*Compared to the two-chip solution of the AMD G-Series APU.