

# Embedded COMPUTING DESIGN

Connecting Silicon, Software, and Strategies for Intelligent Systems

Available for download on  
iTunes App Store and Kindle.

FOLLOW US:



OCTOBER 2013

VOLUME 11 • NUMBER 6

[www.embedded-computing.com](http://www.embedded-computing.com)

## EASY DOES IT:

64 BIT ARMV8 ARCHITECTURE PROVIDES SMOOTH  
MIGRATION FOR 32-BIT OSS AND APPS

Q&A WITH KEITH CLARKE,  
VP OF EMBEDDED PROCESSORS, ARM  
P. 10



## ARM OR X86?

OR BOTH? AMD ADOPTS  
"AMBIDEXTROUS" STRATEGY

P. 18

## STRATEGIES:

IMPROVING PERFORMANCE,  
ACCURACY, AND RELIABILITY IN  
TOUCH-SCREEN BASED  
APPLICATIONS

P. 32

World's First x86 Quad-Core SoC!



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.



**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](http://sandisk.com)

## ON THE COVER

As ARM adoption picks up speed, *Embedded Computing Design* talked with Keith Clarke, VP of Embedded Processors at ARM, about the upcoming 64-bit architecture, the state of 32-bit ARM, and what's next for ARM. This issue also features ARM products and partners in addition to AMD's ARM/x86 strategy.



## 8 Tracking Trends in Embedded Technology

ISA isn't dead

By Rory Dear, Technical Contributor

## 36 E-community Post

Joining the embedded conversation

By Monique DeVoe

## Silicon

ASIC shortcuts, FPGA SoCs



24

Design management system eliminates ASIC shortcut risk

By Bob Smith, Uniquify

20

Hitting the wall in FPGA SoC verification

By Thomas L. Anderson, Breker Verification Systems

24

## Software

Embedded virtualization



28

Embedded virtualization protects legacy investment

By Robert Day, LinuxWorks



## Strategies

Touch-screen control



32

Improving performance, accuracy, and reliability in touch-screen based applications

By Jeff Erickson, Cypress Semiconductor



## ARM and AMD: A view from the top

Easy does it: 64-bit ARMv8 architecture provides smooth migration for 32-bit OSs and apps

Q&A with Keith Clarke, VP of Embedded Processors, ARM

10

ARM Featured Products

14

ARM TechCon Company Guide

17

ARM or x86? Or both? AMD adopts "ambidextrous" strategy

18

By Brandon Lewis, Associate Editor

## E-CASTS

<http://ecast.opensystemsmedia.com>

A Better Way to Cloud: Why DSPs are optimal for high-performance computing

Presented by: Texas Instruments

E-cast: <http://ecast.opensystemsmedia.com/424>

On the Road to Space-Age Vehicles: Developing for the Convergence of consumer Electronics and In-Vehicle Entertainment

Presented by: Polaron, IBM, Vector Software, Mentor Graphics

E-cast: <http://ecast.opensystemsmedia.com/420>



© 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





# IT'S TIME TO MAKE SOME FUTURE.

It's time for some real innovation. World-shattering, face-melting "Holy Moly, 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

To find out more visit us at [iot.windriver.com](http://iot.windriver.com)

## Advertiser Information

Page	Advertiser/Ad title
2	<b>AMD</b> – World's first x86 quad-core SoC!
9	<b>American Portwell Technology</b> – Embedded with unlimited applications
12	<b>AMP Inc. Accelerated Memory Production</b> – The SATA 3 Rg SSD, ruggedized with AES Crypto Engine
39	<b>Annapolis Micro Systems, Inc.</b> – WILDSTAR OpenVPX ecosystem
33	<b>ARM, Inc.</b> – Keil MDK-ARM, ARM DS-5
22	<b>Cogent Computer Systems, Inc.</b> – Low power, high performance COMe solutions
33	<b>COMMELL Systems Corporation</b> – We offer advanced and reliable IPC products
34	<b>Dolphin Interconnect Solutions</b> – Make the right connection
13	<b>Elma Electronic – Systems</b> – Reliable, watertight fanless PCs
19	<b>Innovative Integration</b> – The biggest challenges deserve the smallest solutions!
31	<b>Kontron</b> – Innovative computing platforms based on 4th generation Intel Core processors
35	<b>Kontron</b> – Massively scalable content delivery with SYMKLOUD media platform
29	<b>Micro Digital, Inc.</b> – Your solution is here
3	<b>SanDisk</b> – By 2017, mobile devices will store over 64 billion gaming apps
21	<b>Sealevel Systems, Inc.</b> – Embedded RISC system delivers performance and I/O
11	<b>Sensaray Co., Inc.</b> – Embedded electronics experts
25	<b>Technologic Systems</b> – TS-4710 high end CPU module
5	<b>Wind River Systems, Inc.</b> – It's time to make some future
27	<b>Wind River Systems, Inc.</b> – Let's rewrite the way we think about energy
40	<b>WinSystems, Inc.</b> – Industrial ARM single board computers
7	<b>X-ES</b> – 4th Generation Intel Core i7 from X-ES

## ARM Featured Products

Page	Advertiser
14	<b>Artila</b> – Matrix-516
14	<b>MSC Embedded Inc.</b> – MSC Q7-TI8168
15	<b>PLDA</b> – XpressRICH3-AXI PCIe Gen3 IP
15	<b>Toradex Inc.</b> – Colibri Vybrid VF50 & VF61

## ECD Editorial/Production Staff

Curt Schwaderer, Technology Editor  
 cschwaderer@opensystemsmedia.com  
 Sharon Hess, Managing Editor  
 sharon\_hess@opensystemsmedia.com

Monique DeVoe, Assistant Managing Editor  
 mdevoe@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  
 ehenny@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  
 Sharon Hess, Managing Editor  
*Embedded Computing Design*  
 sharon\_hess@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  
 Will Gaines, Graphic 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

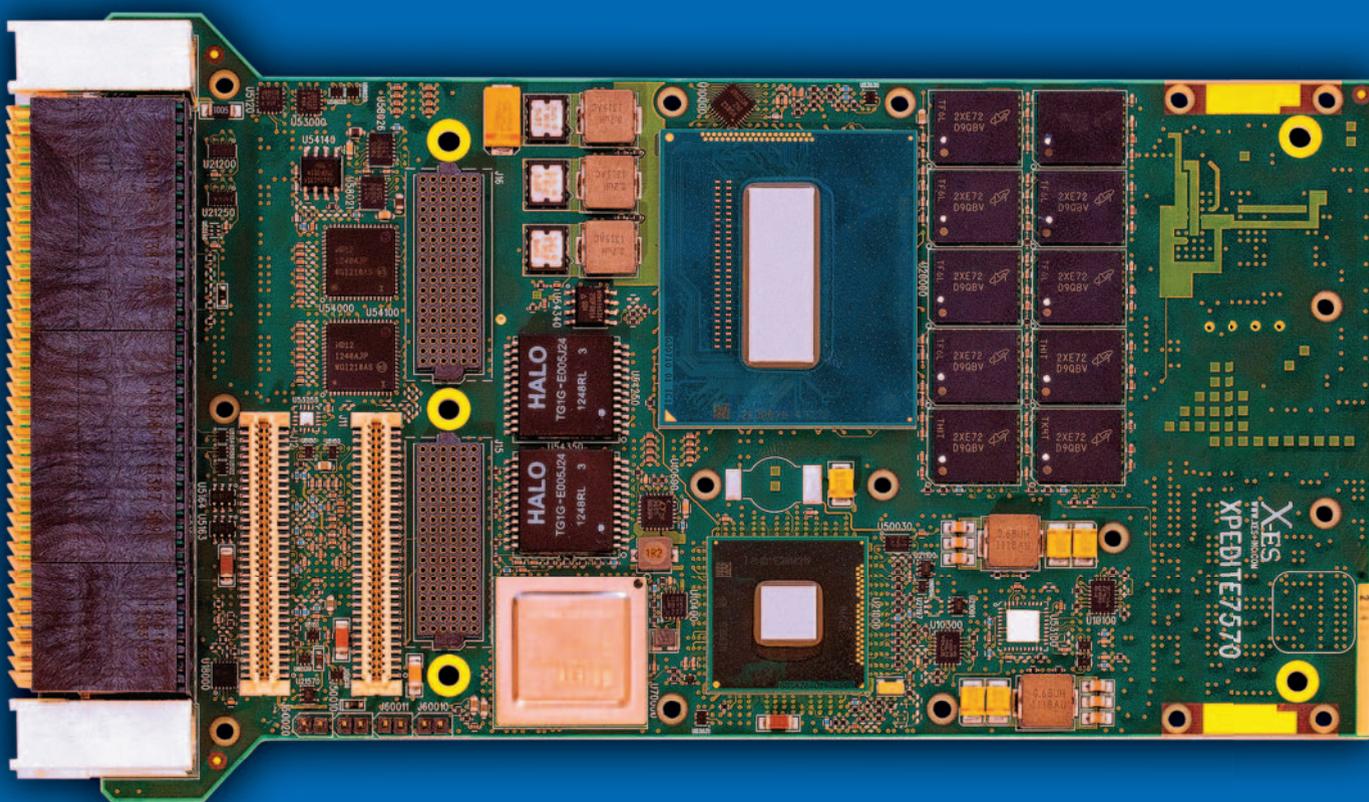
Rosemary Kristoff, President  
 rkristoff@opensystemsmedia.com  
 Wayne Kristoff, CTO  
 Emily Verhoeks, Financial Assistant  
 16626 E. Avenue of the Fountains, Ste. 201  
 Fountain Hills, AZ 85268  
 Tel: 480-967-5581

## Subscriptions

www.opensystemsmedia.com/subscriptions

Tel: 847-763-4946

# 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](http://www.xes-inc.com)



100% Designed, manufactured, and supported in the USA



# Tracking Trends in Embedded Technology

By Rory Dear, Technical Contributor



[rdear@opensystemsmedia.com](mailto:rdear@opensystemsmedia.com)



## ISA isn't dead

We all know the embedded computing industry is some years behind mainstream desktop computing. Longevity being such a critical factor, we take pride in demonstrating products where we've maintained availability, unchanged, for more than a decade.

Whilst this is impressive enough, even more amazing is the support maintained for technologies that most outside of our industry would consider archaic, possibly even extinct. One of these is the Industry Standard Architecture (ISA) bus – currently 32 years mature (and counting).

Originally developed by IBM in 1981 as an 8-bit system, by 1984 the 16-bit ISA bus variant was released, opening doors around the globe with the first truly universal expansion bus.

ISA peripherals are historically typically packaged as slot cards, mating with a motherboard or backplane perpendicularly – typically functioning as display adapters, sound cards, modems, or even early Ethernet. When the desktop world moved onto PCI, PCI-capable Integrated Circuits (ICs) satisfying these popular functions became widely available quickly and ISA faded away into the sunset.

This process was far from smooth in the embedded marketplace, and still continues today. With our embedded marketplace naturally many times smaller than enterprise computing, the commercial viability of PCI replacement ICs for often highly specialized functions simply wasn't there. Even where it was commercially viable for the IC manufacturer to migrate, I've experienced

countless scenarios where, for the customer, it wasn't. Often, the approvals required for military equipment are so constrained that such a migration could cost millions. Additionally, through the passage of time of stricter regulation in all industries, this aspect increasingly applies to all vertical markets.

I've seen custom ISA slot cards costing upwards of \$30k each; this unit cost should sound alarm bells of the indicative background costs involved. Recognizing this, industrial computer manufacturers strived to offer longevity whilst industrializing the vibration-unfriendly slot card format – eventually standardized in 1992 as "PC/104." PC/104 permitted a compact method of integrating one's ISA-based, application-specific, intellectual property. By enabling stacking rather than mounting perpendicularly, a vibration- and shock-proof platform was born, and again, standard PC-based functionality was soon spawned as PC/104 peripheral modules in many guises.

The parallels of ISA's slot and PC/104 formats are obvious, whilst commercial functionality was easy to replicate by other means: Expensively developed and approved customized expansion boards had to stay, and stay unchanged.

### ISA's commercial retirement

Manufacturers have to varying degrees supported ISA as it matured; variations are easy to see in where a manufacturer draws the line of plugging new technology versus supporting legacy technology, and I would place myself firmly in the latter camp. Today, most have dropped support for the ISA bus architecture entirely, concededly often pushed by a lack of native ISA support

found within new CPU chipsets – but there are ways around this.

The Low Pin Count (LPC) bus was introduced as a replacement for the ISA bus in 1998, predominantly to continue support of the internal legacy peripherals device, though with a significant limitation with it inherently only supporting 8-bit data transfer, whilst ISA peripherals are often based on a 16-bit bus.

The requirement from hardware developers for full ISA bandwidth legacy support was so strong that as part of my role as Technical Sales Manager, I had to react to this. We soon developed effectively an LPC>16-bit ISA converter that has extended the lifetime of hundreds of legacy products. The converter utilizes parallel processing to emulate exactly what once was natively commonplace and incredibly easy to design with and debug.

Additionally, whilst the major commercial CPU manufacturers have long forgotten the ISA bus, embedded CPU manufacturers – specializing in mid-range performance and predominantly low power for a wide range of control and monitoring applications – continue to embrace the full ISA interface natively, which will be music to the ears of many an engineering, or indeed production, manager.

### Waking the dead?

The time always comes when one must concede a past technology should now be buried, but we are years away from this eventuality with the ISA architecture. I'm personally proud to be a part of increasing product maturity and longevity far into the future with ISA.

# Embedded with Unlimited Applications



Portwell's flexible business model and optimized SFF boards & systems have helped numerous Medical, Networking, and Industrial OEMs reduce costs, save space, accelerate development schedules and extend system longevity. Contact us to learn how we can tailor our embedded solution offering to your exact needs!



**Small Form Factor Boards**

- Small footprints to save space
- 120x120mm or 102x146mm
- Low power consumption
- Single-/ dual-core Intel® Atom™ and Intel® Core™ processors
- Extended temp: -40 to 80°C with Intel® Atom™ processor Z510PT Z520PT
- On board dual display, GbE LAN, PCIe, IDE or SATA and more



**Small Form Factor Systems**

- Compact turnkey solutions
- Fanless system with Intel® Atom™ or ULV Intel® Core™ processors
- All-in-one with LCD and touchscreen
- Ideal for digital signage, telematics, industrial automation, control and communications
- Expansion I/O options



**Computer-on-Modules**

- Wide selection of COM Express and Qseven modules
- The latest Intel® Core™ i5/i7 and single-/ dual-core Intel® Atom™ processors
- Custom carrier board design and manufacturing
- Quick time-to-market solutions
- Extended temp and ECC memory support

Portwell's extensive product portfolio includes single-board computers, embedded computers, specialty computer platforms, rackmount computers, communication appliances, and human-machine interfaces.

We provides both off-the-shelf and versatile custom solutions for applications in the medical equipment, factory automation, retail automation, semiconductor equipment, financial automation, mission critical and network security markets.

American Portwell is both an ISO 9001:2008 and ISO 13485:2003 certified company.



Q &amp; A



Interview with Keith Clarke, VP of Embedded Processors, ARM

## Easy does it: 64-bit ARMv8 architecture provides smooth migration for 32-bit OSs and apps

ARMv8 introduces 64-bit processors to make the architecture more competitive in data center and server environments, as well as high-end tablet and smartphone applications. But what about everyone who does not need or is not ready to migrate their 16- and 32-bit Operating Systems (OSs) and apps? In this Embedded Computing Design interview, Keith Clarke, ARM's Vice President of Embedded Processors, explains how ARMv8 lays out a smooth migration path to 64-bit for designers working on lower level architectures, as well as how his company is still advancing 32-bit processor technology into segments like automotive and industrial control. Edited excerpts follow.



**ECD:** When setting out to develop a new ARM processor architecture, what are the top design considerations?

**CLARKE:** The processors we design have to be fit for a purpose. ARM has a very large range of different processors, from the Cortex-M0, which can be as small as 10,000 gates, to the Cortex-A57, based on the ARMv8 architecture, which is designed to deliver very high performance. Implementation size, performance, and low power consumption are key attributes of the ARM architecture.

**ECD:** The 64-bit processors are a pretty big deal. When will companies be shipping these processors?

**CLARKE:** We are pretty confident 64-bit processors will begin shipping from ARM partners in 2014. The technology has been delivered to our partners; initially they have to go through the cycle of developing silicon and getting that ready for production. [Our partners] and their customers all may be on slightly different time frames, so we expect a number of introductions starting in 2014, and that will accelerate greatly into the future. Eventually, the industry will come to a point where the majority of new designs will be 64-bit capable.

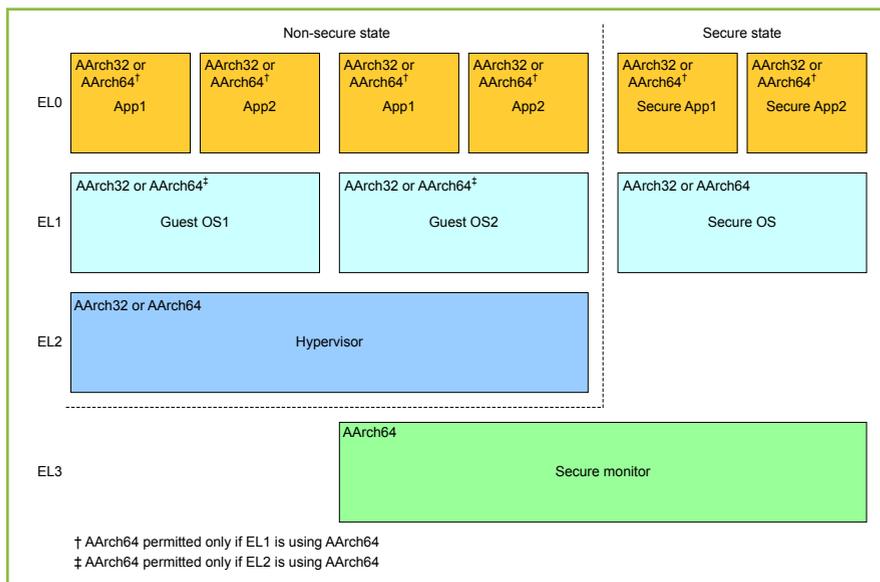
**ECD:** Do you think we will see the 64-bit designs anywhere besides a data center and server markets?

**CLARKE:** Yes, we have seen the trend to have OSs running 64-bit, and it is not just an issue of "big iron" servers where the data sets are very large. We've seen demand for OSs with 64-bit support even if they are running 32-bit applications, which is possible with the ARMv8 architecture. It allows them to upgrade the OS so it is more capable and still run all of the legacy applications and software they currently have. And then on a subsequent generation, as there are more 64-bit applications, they can successfully support those as well. So we definitely see this as applicable to high-end tablets and smartphones in the initial phase.

**ECD:** Will we see 64-bit designs with GPUs?

**CLARKE:** Due to the computing and graphics needs of the high-end tablet and smartphone markets, GPUs will eventually be in the mix when it comes to 64-bit architecture.

In general, we see a smooth transition to 64-bit. We designed the ARMv8 architecture such that even 32-bit applications will just run and, in fact, the processors we have released will run existing 32-bit OSs and software as-is with no alteration (Figure 1).



**Figure 1** | The ARMv8 architecture provides a 64-bit silicon design that is also capable of running 32-bit operating systems, and 32-bit and 64-bit applications, enabling a smooth migration path for embedded systems.

## 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/EM10/953](http://www.SENSORAY.com/EM10/953)

### Model 2253S

CE



Also available as OEM board  
**Model 2253**

### Small Form Factor for Streaming Video

- Simultaneous encode/decode
- Low preview latency; Text overlay
- H.264 HP@L3, MPEG-4 ASP, MJPEG video compression
- CC support for H.264 streams
- MP4 files can be edited in Adobe Premiere with A/V sync preserved
- Extended temp. version available

[www.SENSORAY.com/EM10/2253](http://www.SENSORAY.com/EM10/2253)



### 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/EM10/826](http://www.SENSORAY.com/EM10/826)

[SENSORAY.com](http://SENSORAY.com) | 503.684.8005

Companies may have silicon that is capable of running 64-bit software, but they do not need to. Some device makers will release products that have 64-bit capable silicon and runs today's 32-bit operating systems and then on a subsequent release choose to switch to a 64-bit-capable OS and eventually 64-bit apps. It will be a nice, smooth transition with exactly the same silicon. We did not want to force a difficult transition for people in the marketplace;

they are getting a higher performance processor from ARM that is also capable of running a 64-bit OS and applications if they choose to do so at some stage.

**ECD:** *What are the trade-offs when designing different processors?*

**CLARKE:** When developing processors, there's always a trade-off between how much logic you put into the design versus the area you want to achieve. There is always a balance: to get the very

highest possible performance you need to put more silicon area down, and you have got to balance that with the target market you are aiming for. So the considerations are how many pipeline stages, how much cache is there, and how complex the logic is you use to predict branches, as well as the complexity of the pipeline in terms of in-order execution or out-of-order execution. So there are many, many trade-offs you can make in processor design, which ultimately lead to what performance and what power consumption you can achieve; we do all of this with a target process in mind.

For example, in the past we may have designed a processor to target 28 nm manufacturing process technology. Now, for newer processors, we may be targeting 10 nm FinFET. Since we're designing something that won't immediately hit the market, there are a lot of considerations to be made. We have to make some judgment on what that would look like in a number of years' time when our product is ready to be delivered to the market. In the meantime, we're doing things like R&D, verification, and implementation.

**ECD:** *What is on the horizon for 32-bit processors? Are there any changes there?*

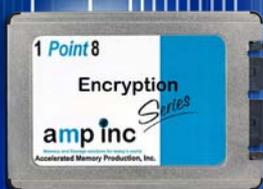
**CLARKE:** Generally, the distinction between 32-bit and 64-bit is something that we do not think about throughout the whole range of processors. In what we call the "application processors," yes, we are making this transition to be able to support 64-bit, but we also have microcontroller cores and real-time cores – our Cortex-M and Cortex-R class of processors – which we intend to continue developing as 32-bit processors. The kind of applications those processors are being designed into today – it could be a sensor, a Bluetooth controller, a microcontroller, or a baseband controller for cellular, for example – those applications just do not require anything beyond the 32-bit in the near future.

We do see many more developments in our range of processors where 64-bit

## The SATA 3 Rg SSD, Ruggedized with AES Crypto Engine

The SATA 3 Rg SSD, ruggedized with AES crypto engine of 256 bits, is the key to encrypt the entire SSD. This is a flash based solid state disk drive with SATA 3.0 compliant interface. It provides fast read and write speeds; high reliability and its data protection make it an ideal storage solution for the server and mobile environment. Built-in ECC and EDC ensure error-free transactions for the most demanding applications.

- > Density up to 480 GB (20 grams)
- > User defined form factor
- > Host Interface 2.0 & 3.0
- > Programmable Hardware Encryption
- > Military Erase (Multiple Protocols)
- > SLC/MLC
- > Extended Temperature
- > Made In The USA
- > Extended Life Cycle



**For more information, call 800-778-7928**

Accelerated Memory Production, Inc.  
1317 E. Edinger Ave., Santa Ana, CA 92705  
714-460-9800 | 800-778-7928

[www.ampinc.biz](http://www.ampinc.biz)





Of the overall microcontroller market, ARM processors now have an 18 percent market share that grew from 15 percent from last year.



isn't needed. Our real-time products are being used in automotive more, so we are definitely looking at developments that will enhance our processors' capability for automotive and industrial control. That is definitely something that we are working on for the future, and we have some good technologies there that we think will significantly enhance the ARM processors' capability in those applications.

One of the misconceptions of ARM is we do mobile and nothing else, but last year alone, 4.5 billion ARM chips were sold in what we call the "embedded segment." ARM is very, very strong in 32-bit microcontrollers for applications like smart cards and automotive.

Of the overall microcontroller market, ARM processors now have an 18 percent market share that grew from 15 percent from last year. We are going to continue developing that microcontroller range to help our partners win even more volume and transition from 8-bit and 16-bit processors. There are definitely new technology introductions to come from ARM and from ARM's partners in those markets.

**ECD:** *What will the next generation of ARM processors look like?*

**CLARKE:** In our Cortex-R real-time range, we are going to be introducing technology that targets specific challenges that automotive and industrial

control face – where they need to have very real-time capability, as well as dealing with different functionality merged onto the same chip. For example, in automotive, a high-end car today may have as many as 200 Electronic Control Units (ECUs). Car manufacturers would like to consolidate and reduce the number of separate ECUs, but in order to do so, they need to run multiple functions on each ECU. Our technology will target helping car manufacturers deal with that challenge. That's the kind of technology we will be talking about in more detail at ARM TechCon. **ECD**

*Keith Clarke is VP of Embedded Processors at ARM.*



**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. It 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

# Artila

www.artila.com

## Artila's Matrix-516

**Artila's Matrix-516** industrial box computer is a small form-factor, low power consumption, and Linux-ready computing platform. With a 400MHz ARM9 CPU, 64MB SDRAM, and 128MB NAND Flash inside, Matrix-516 ensures system high performance. Matrix-516 is equipped with multiple I/O interfaces, including 2 LANs, 8 x isolated RS-485 serial ports, 2 x USB hosts, 1x micro-SD and 21 x GPIOs. It is easy to develop pure C/C++ programs or Web-based applications (ex: SQLite+ PHP) to run on top of the Matrix-516.



## FEATURES

- > ATMEL 400MHz ARM9 AT91SAM9G20 processor
- > Linux kernel 2.6.29 with file system
- > 64MB SDRAM, 128MB NAND Flash
- > One micro SD socket inside, upto 32GB capacity
- > Two 10/100M Ethernet ports
- > Eight Isolated RS-485 serial ports
- > Two USB hosts, 21x GPIOs
- > Less than 3W power consumption

**Artila Electronics Co., Ltd. | Tel: +886.2.86.67.23.40**

**Contact:** sales@artila.com  
www.artila.com

embedded-computing.com/p9915999



www.msembedded.com

## MSC Q7-TI8168

The innovation of the MSC Q7-TI8168 is the combination of an ARM® Cortex™ A8 RISC MPU with TI's C674x VLIW DSP core offering up to 8000 MMACS. This powerful product is especially suited for Medical/Industrial Vision Systems, High-end Test & Measurement, Tracking and Control as well as Medical/Biological Imaging.

The MSC Q7-TI8168 MPU Module incorporates a high performance C6A8168 MPU @ 1.2 GHz with DDR3 memory, GbE and industrial interfaces including Flash memory. The module fully supports the Qseven Standard Rev. 1.2 interface, allowing simple integration with our range of Qseven baseboards or custom developments.

For evaluation and design-in of the MSC Q7-TI8168 module, MSC provides a number of different development platforms.



## FEATURES

- > Texas Instruments TMS320DM8168:
  - ARM Cortex™ A8 CPU up to 1.2 GHz
  - DSP Subsystem C674x up to 1.0 GHz
  - 1GB DDR3 SDRAM
  - 2GB Flash SSD soldered on board
  - Gigabit Ethernet
  - 1x PCI Express x1 port
  - HDMI/DVI up to 1920x1080 resolution
  - Single Channel LVDS 24 bit up to 1280x720 resolution
  - Dual Independent Display support
  - Two SATA-II interfaces
  - Six USB 2.0, HD Audio

**MSC Embedded Inc. | 650-616-4068**

**Contact:** info@msembedded.com

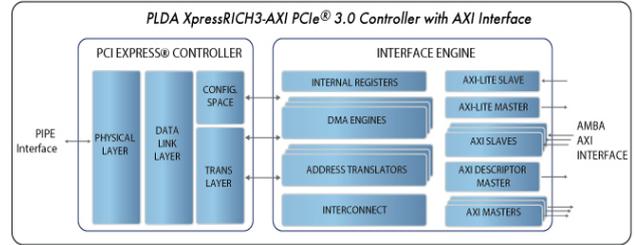


**PLDA XpressRICH3-AXI PCIe Gen3 IP**

PLDA's **XpressRICH3-AXI** is a configurable PCIe® 3.0, 2.0, 1.1 Controller IP for ASIC/SoC with AMBA® AXI user interface. The XpressRICH3-AXI IP seamlessly interfaces to ARM's high-performance CoreLink CCN-504 Cache Coherent Network in order to address demand for energy-efficient solutions for the enterprise market. The IP supports x16, x8, x4, x2, x1 at Gen3, Gen2, Gen1 speeds and the is tightly integrated, optimized and proven in the ARM® architecture.

PLDA's XpressRICH3-AXI IP enables:

- PC-compatible direct access to PCIe configuration space (ECAM)
- PCIe-AXI deadlock prevention through optimized ordering rules
- Customized error management (INT/MSI notifications)



**FEATURES**

- › AMBA® AXI user interface
- › Up to 8 independent DMA engines/channel
- › Supports Rootport and Endpoint configuration
- › Proven PHY IP integration with multiple vendors and support for most foundries
- › Packet-based or block-based data transfers and configurable data path width
- › Flexible and configurable scatter-gather/DMA chaining, automatic and customizable address translation and automatic clock domain management

**PLDA | 408-273-4528**

2570 No. First Street, Ste 218 • San Jose, CA • 95131-1036

Contact: sales@plda.com



www.Toradex.com

**Colibri Vybrid VF50 & VF61**

**Colibri VF50 and VF61** are the latest additions to the standardized and industry proven Colibri product family.

Powered by Freescale Vybrid embedded SoC, the VF61, for example, features a Cortex-A5 processor running at 500MHz, alongside with a Cortex-M4 processor clocked at 167MHz. This unique heterogeneous dual-core system enables you to implement a stunning user experience while running hard real-time and safety critical software on the M4 core.

With its small form factor, the Colibri Vybrid VF50 and VF61 are world's most cost effective computer modules and hence, the ideal platform to upgrade any typical microcontroller system to the requirements of the emerging Internet of Things.



**FEATURES**

- › World's most Cost-Effective Computer Modules
- › Asymmetric Multicore Processing System
- › 2x CAN, Analog In/Out, Camera Input, SD Card
- › USB 2.0, Ethernet, UART, LCD, SPI, I2C, PWM, etc.
- › WinCE and Linux BSP provided
- › Industrial Temperature option available (-40 + 85 °C )
- › Pin compatible with Colibri Product Family
- › Long Term availability until 2028

**Toradex Inc. | 800-871-6550**

Seattle@Toradex.com  
 Twitter.com/Toradex  
 www.Facebook.com/Toradex

## HIGH-PERFORMANCE IMPLEMENTATIONS FOR ARM 64-BIT CORTEX-A57/A53 PROCESSORS



PRESENTED BY: ARM

GO TO E-CAST:  
ECAST.OPENSYSTEMSMEDIA.COM/419

Mobile and server markets present unique challenges for the implementation of ARM 64-bit Cortex-A57 and Cortex-A53 processors. This E-cast will discuss techniques to achieve high-performance and power-optimized implementations, and how to leverage ARM POP IP core-hardening acceleration technology to achieve faster time to market with less risk.

## POWER INTEGRITY ANALYSIS OF AN ARM CORTEX-A15 QUAD CORE HARD MACRO



PRESENTED BY: ARM

GO TO E-CAST:  
ECAST.OPENSYSTEMSMEDIA.COM/424

The Seahawk hard macro is a quad-core Cortex-A15 implemented in the TSMC 28 hpm process designed to operate at frequencies close to 2 GHz and has DVFS and retention capabilities built in. This presentation describes the power integrity checks that were performed on the design to ensure proper functioning and reliability.

## A BETTER WAY TO CLOUD: WHY DSPS ARE OPTIMAL FOR HIGH-PERFORMANCE COMPUTING



PRESENTED BY:

TEXAS INSTRUMENTS  
GO TO E-CAST:  
ECAST.OPENSYSTEMSMEDIA.COM/425

Cloud app designers have specialized processing needs best addressed with targeted processors and server elements. TI's multicore SoCs based on its KeyStone architecture featuring high-performance ARM Cortex-A15 MPCore processors, TI's latest DSP cores, and built-in Ethernet switching can enable more efficient servers for a wide variety of industries.

## READ THE DSP-FPGA.COM 2013-14 ANNUAL RESOURCE GUIDE!



### IN THIS ISSUE:

- :: DSP-FPGA.com talked with Texas Instruments and CEVA about benefits and challenges of using ARM + DSP cores to optimize SoCs for high-performance apps
- :: Silicon Labs explains how ARM Cortex-M4-based microcontrollers bring energy efficiency and high performance to intelligent applications
- :: CommAgility discusses floating-point multiprocessing with a TI architecture that uses quad ARM Cortex-A15 MPCore processors along with DSP cores

DSP-FPGA.com



2013-2014  
RESOURCE GUIDE

Page 18

**Advanced Micro Devices, Inc.**

Booth #606

[embedded-computing.com/news/amd-embedded-product-roadmap/](http://embedded-computing.com/news/amd-embedded-product-roadmap/)**Aldec, Inc.**

Booth #702

[embedded-computing.com/p9910471](http://embedded-computing.com/p9910471)**Algotochip Corporation**

Booth #821

[embedded-computing.com/p9915900](http://embedded-computing.com/p9915900)**Apache Design Solutions**

Booth #716

[embedded-computing.com/p9910477](http://embedded-computing.com/p9910477)**Arasan Chip Systems, Inc.**

Booth #609

[embedded-computing.com/p9915873](http://embedded-computing.com/p9915873)**ARIUM**

Booth #603

[embedded-computing.com/p9915901](http://embedded-computing.com/p9915901)**Atmel**

Booth #207

[embedded-computing.com/p9915902](http://embedded-computing.com/p9915902)**Breker Verification Systems**

Booth #714

[embedded-computing.com/p9911990](http://embedded-computing.com/p9911990)**Cadence Design Systems**

Booth #600

[embedded-computing.com/p9915565](http://embedded-computing.com/p9915565)**Carbon Design Systems**

Booth #615

[embedded-computing.com/p9915903](http://embedded-computing.com/p9915903)**Ceva Inc.**

Booth #313

[embedded-computing.com/p9915875](http://embedded-computing.com/p9915875)**Chips & Media, Inc.**

Booth #605

[embedded-computing.com/p9915897](http://embedded-computing.com/p9915897)**The DINI Group**

Booth #1110

[embedded-computing.com/p9915896](http://embedded-computing.com/p9915896)**Dorado Design Automation, Inc.**

Booth #416

[embedded-computing.com/p9915904](http://embedded-computing.com/p9915904)**Express Logic**

Booth #412

[embedded-computing.com/p9911859](http://embedded-computing.com/p9911859)**Freescale**

Booth #500

[embedded-computing.com/p9911822](http://embedded-computing.com/p9911822)**GLOBALFOUNDRIES US, Inc.**

Booth #704

[embedded-computing.com/p9915906](http://embedded-computing.com/p9915906)**Green Hills Software, Inc.**

Booth #512

[embedded-computing.com/p9915907](http://embedded-computing.com/p9915907)**HardKernel Co Ltd**

Booth #213

[embedded-computing.com/p9915908](http://embedded-computing.com/p9915908)**IAR Systems**

Booth #620

[embedded-computing.com/p9915876](http://embedded-computing.com/p9915876)**Imperas**

Booth #520

[embedded-computing.com/p9915909](http://embedded-computing.com/p9915909)**INSIDE Secure**

Booth #621

[embedded-computing.com/p9915910](http://embedded-computing.com/p9915910)**Infineon Technologies**

Booth #519

[embedded-computing.com/p9911848](http://embedded-computing.com/p9911848)**Jasper Design Automation**

Booth #712

[embedded-computing.com/p9910605](http://embedded-computing.com/p9910605)**Lauterbach, Inc.**

Booth #1105

[embedded-computing.com/p9915911](http://embedded-computing.com/p9915911)**LSI Corporation**

Booth #315

[embedded-computing.com/p9915914](http://embedded-computing.com/p9915914)**MAGILLEM**

Booth #1104

[embedded-computing.com/p9915877](http://embedded-computing.com/p9915877)**MatrikonOPC**

Booth #616

[embedded-computing.com/p9915915](http://embedded-computing.com/p9915915)**Mentor Graphics**

Booth #201

[embedded-computing.com/p9915497](http://embedded-computing.com/p9915497)**MontaVista Software, LLC**

Booth #316

[embedded-computing.com/p9915878](http://embedded-computing.com/p9915878)**OneSpin Solutions**

Booth #809

[embedded-computing.com/p9915916](http://embedded-computing.com/p9915916)**Open Silicon**

Booth #801

[embedded-computing.com/p9915917](http://embedded-computing.com/p9915917)**Oracle**[embedded-computing.com/p9915918](http://embedded-computing.com/p9915918)**PLDA**

Booth #320

[embedded-computing.com/p9911922](http://embedded-computing.com/p9911922)**Qt by Digia**

Booth #607

[embedded-computing.com/p9915919](http://embedded-computing.com/p9915919)**Rambus Inc.**

Booth #200

[embedded-computing.com/p9915920](http://embedded-computing.com/p9915920)**Real Intent**

Booth #812

[embedded-computing.com/p9910474](http://embedded-computing.com/p9910474)**S3 Group**

Booth #321

[embedded-computing.com/p9915921](http://embedded-computing.com/p9915921)**SEGGER Microcontroller**

Booth #521

[embedded-computing.com/p9915922](http://embedded-computing.com/p9915922)**Silicon Image**

Booth #221

[embedded-computing.com/p9915923](http://embedded-computing.com/p9915923)**Sonics Inc.**

Booth #717

[embedded-computing.com/p9915924](http://embedded-computing.com/p9915924)**Space Codesign**

Booth #421

[embedded-computing.com/p9915925](http://embedded-computing.com/p9915925)**STmicroelectronics**

Booth #413

[embedded-computing.com/p9915879](http://embedded-computing.com/p9915879)**SYSGO AG**

Booth #420

[embedded-computing.com/p9914403](http://embedded-computing.com/p9914403)**Target Complier Technologies**

Booth #317

[embedded-computing.com/p9915926](http://embedded-computing.com/p9915926)**TechNexion**

Booth #713

[embedded-computing.com/p9915927](http://embedded-computing.com/p9915927)**Toradex**

Booth #613

[embedded-computing.com/p9911828](http://embedded-computing.com/p9911828)**True Circuits, Inc.**

Booth #708

[embedded-computing.com/p9915928](http://embedded-computing.com/p9915928)**TSMC**

Booth #808

[embedded-computing.com/p9915929](http://embedded-computing.com/p9915929)**Uniquify, Inc.**

Booth #700

[embedded-computing.com/p9915880](http://embedded-computing.com/p9915880)**WolfSSL**

Booth #516

[embedded-computing.com/p9915881](http://embedded-computing.com/p9915881)**Xilinx, Inc.**

Booth #515

[embedded-computing.com/p9915882](http://embedded-computing.com/p9915882)**Xively**

Booth #212

[embedded-computing.com/p9915930](http://embedded-computing.com/p9915930)

# x86 or ARM? Or both? AMD pursues “ambidextrous” product strategy

By Brandon Lewis, Associate Editor

*In September, AMD revealed the “ambidextrous” product strategy that it will begin implementing in 2014, producing both x86 and ARM-based processors targeted at the embedded market. As the traditional lines separating ARM and x86 continue to blur, AMD hopes to bring added flexibility to their customer base while hedging their bet – on both sides of the fence.*

There was a time when choosing an embedded processor was more simple: If you wanted horsepower for high-end embedded applications, you typically perused the available x86 offerings; if you were looking for flexibility and low power, a variety of ARM designs were at your fingertips. If you haven't been paying attention, all of this has started to change.

Since the start of 2012, x86 processors have made a concerted effort to push into the mobile space, with Intel targeting its Medfield architecture at smartphone applications and AMD releasing tablet reference designs for the Jaguar core. From the other side, the 64-bit ARMv8 architecture squarely positioned the RISC processors for use in networking and datacenter applications. So, if your money is on embedded, which one is a safe bet? If you ask AMD, the answer is “both.”

At a September press event in San Francisco, AMD unveiled the first chips on the company's “ambidextrous” embedded product roadmap, which includes both x86 and ARM offerings. Along with two x86 Accelerated

Processing Units (APUs) – one targeted for high performance and the other for low power – and a discrete Graphics Processing Unit (GPU), AMD announced that it will begin shipping the “Hierofalcon” System-on-Chip (SoC) based on the 64-bit Cortex-A57 core in the latter half of 2014 (Figure 1). Aside from shedding light on the company's 2012 partnership with ARM and hiring of former P.A. Semi and Apple

chip designer Jim Keller, becoming the first company with both x86 and ARM solutions gives perspective on the embedded processor market.

“The embedded market is extremely diverse, and we have to address multiple application areas. You go on a visit to an embedded customer and you have to be conscious of that,” says Kamal Khouri, Director of Embedded

## A FIRST LOOK: “HIEROFALCON”

- ▲ High-performance 64-bit ARM CPU SOC
- ▲ Up to 8 ARM Cortex™-A57 CPUs
- ▲ 10Gb KR Ethernet
- ▲ PCI-Express Gen 3
- ▲ Enhanced security with ARM TrustZone™

## FIRST ARM SOLUTION FROM AMD FOR EMBEDDED MARKETS

**Figure 1** | The Hierofalcon SoC is AMD's first ARM-based processor, and will be manufactured in 4- or 8-core Cortex-A57 core variants. The 64-bit ARMv8 design supports 10 Gigabit Ethernet (GbE), PCI Express Gen 3, and Error Correct Code (ECC) with a Thermal Design Power (TDP) of 15 W to 35 W, making it well suited for networking and datacenter applications.

Product Marketing and Management, AMD. "We cannot just print out chips anymore. We have customers that have a lot of money invested in a software architecture, and we can now say that we have a solution."

### Ambidextrous mirrors market trends

With the chip wars ramping up, AMD's decision to enter an ambidextrous strategy is one that reflects the state of the industry. Despite the disparities in volume and traditional application area, a 2012 IDC Research report found that x86 processors are in a near dead heat with ARM chips for embedded processor market share (Figure 2). By executing on existing IP, AMD hopes to carve out portions in both segments, Khouri explains.

"When you talk about leveraging IP and leveraging investment of the Jaguar core and how it can function, the basic IP building blocks are the basic IP building blocks for any architecture. And, as you look at the gates and transistors, things start to look the same," Khouri says. "With the choice of the 64-bit Cortex-A57 architecture, we are trying to leverage the investments we have made in IP across multiple markets, and we are going to use APUs and ARM to get there."

Since establishing its Embedded Solutions Group and co-founding the Heterogeneous Systems Architecture (HSA) Foundation last summer, mixing ARM into their portfolio gives AMD a unique position in the embedded space, Khouri says. For the embedded market, this provides flexibility in a wide range of applications that enables end-users to distinguish their products from the competition, he continues.

"When you look at the fact that we were first to market with ARM and x86 solutions, the openness of the HSA ecosystem, and then you tie in our involvement with OpenCL, I believe that gives us a lot of leeway against protectionism," Khouri says. "Our goal is to provide a lot of options that still allow our customers to differentiate their products."

Oh the times, they have a-changed. Again. **ECD**

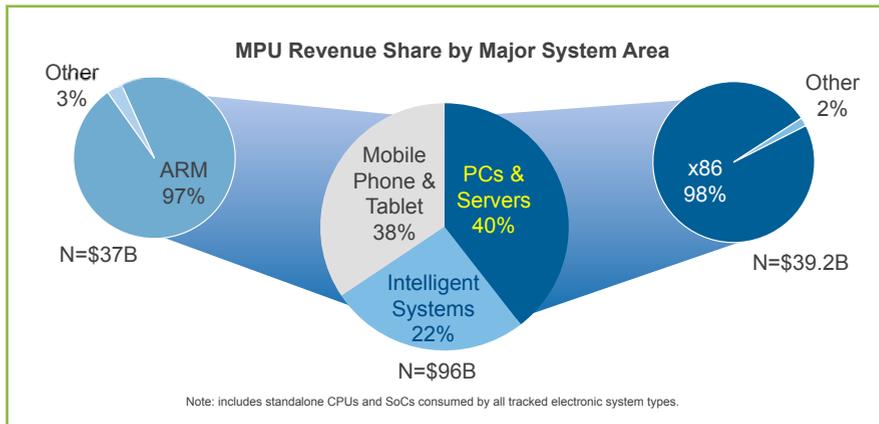


Figure 2 | First displayed in May at SMART Technology World 2013, this 2012 market report from IDC Research depicts the market share for x86 and ARM processors in the embedded space. Graph courtesy IDC Research.

## The Biggest Challenges Deserve the smallest Solutions!



Introducing ePC-K7, the best solution for many big challenges in a variety of markets including Wireless, Industrial Control, Military Hardware, Medical Imaging, Telecom/Intelligence, and Test & Measurement

### Features

- Industry-standard COM Express CPU Type 6 module with dual FMC I/O Modules in a compact, stand-alone design
- Programmable Kintex-7 325T/410T and Spartan-6 FPGAs
- Small form factor: 5" H x 8" W x 11" D
- Conduction cooled design: Fins or Cold-plate
- Stand-alone operation: Able to operate headless booting from SSD
- Windows, Linux and RTOS support
- Dual VITA 57 FMC I/O Module Sites - add anything from RF Receivers to Industrial Control Modules
- FMC I/O Sites deliver 6400MB/s sustained to FPGA memory, 3200 MB/s continuous streaming to host memory
- Integrated Timing and Triggering support for I/O includes GPS, IEEE1588 or IRIG -Disciplined Clock
- Supports Innovative Integration and 3rd-party FMC Modules for Private Data Channels, Triggering and Timing features
- USB 3.0 x2, 1Gb Ethernet x2, eSATA x2, DisplayPort, Touch Screen
- Up to 4 SSD or HDD (2.5 in.)
- Flexible 8 to 36V DC operation

Marsh environment? Check out ePC-K7 Rugged

[DOWNLOAD DATASHEETS NOW!](#)

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

805.578.4260 phone • www.innovative-dsp.com

ASIC design management systems rely on standards and consistency to be a useful reference, just as libraries need the same qualities to be reliable.

José Vasconcelos Library photo source: <http://www.flickr.com/photos/lwy/5470629800/>

## Design management system eliminates ASIC shortcut risk

By Bob Smith

*Taking shortcuts can be a risky endeavor when designing an ASIC. However, there are avenues to streamlining a design project and reducing the risk when using shortcuts. A design management system can help achieve consistent results in the shortest amount of time, critical to a project team's confidence that a new project will proceed to completion on time and on budget.*

Seasoned ASIC engineers will likely tell a junior member of the project team that there are no shortcuts, especially if challenging factors come into play, for example: if the design is to be implemented on a modern, aggressive process node; or if the project team is relying on semiconductor Intellectual Property (IP) that can account for as much as 70 percent of the chip; or if the ASIC has 500 million or more gates; or if the project team delegates some or all of the design to a contract design house, relieving themselves of much of the detail but placing their trust in a separate organization. Plenty can go wrong when taking a shortcut in any of these scenarios (Figure 1).

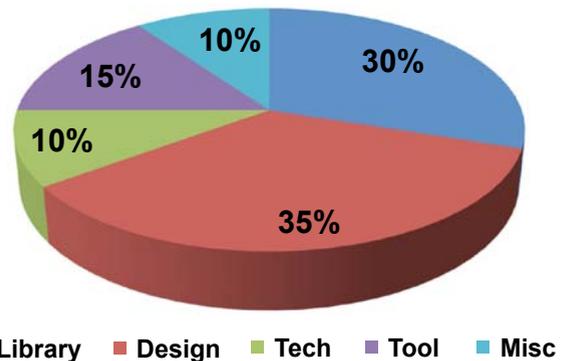
That does not mean, however, that there are no avenues to streamlining the design process. Ensuring project success each and every time requires the use of a design management system based on years of hands-on project management experience to ensure every project is on track and achieves its goals. Accordingly, the following focuses on the decision to devise an ASIC design management system and implement what has become an important resource for engineers and managers.

### Design management system streamlines process

A design house has numerous goals when designing a chip, from achieving the highest-performance design to completing it in the shortest possible time while maintaining high quality.

ASIC designers often have more to consider than FPGA designers. For example, a large percentage of the design can be made up of third-party IP or previously implemented pieces

### ASIC/SoC Implementation Issues



**Figure 1** | Most implementation issues occur during the design phase, or are caused by issues with library files.

of code. Further, process technologies are getting smaller, while gate counts continue to rise, making verification (or verifying that the ASIC works as intended) a necessity.

The goal of an ASIC design management system, then, is to provide project consistency with minimum overhead so that all designs progress in a predictable manner.

As new semiconductor process technologies evolve, each design is subjected to an ever-more-complex tool chain. One difficulty of project management, and a threat to consistency, is that each engineer will have a different idea of how the design

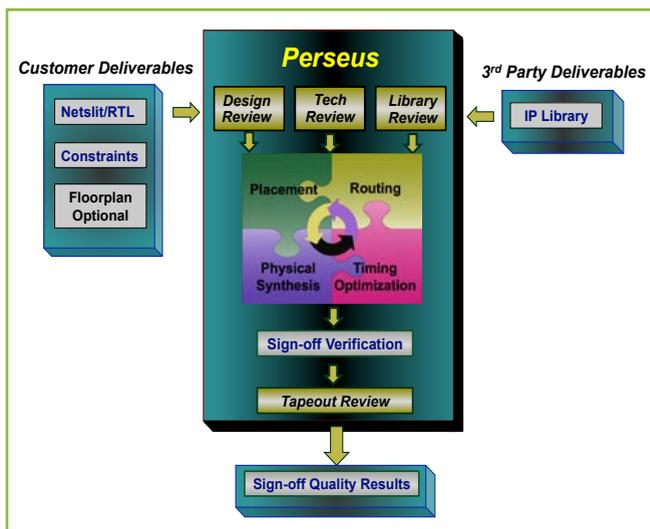
should flow through the tools. For example, individual engineers like to automate the flow of their portions of an ASIC design through scripts that invoke tools and manage results. He or she may have a scripted routine that helps simplify output results of a detailed timing analysis. If not managed properly, this can lead to "résumé dependency" where each project is managed according to the history, skill set, and whims of the engineers. If this happens, every project is at risk if engineers join or leave the company and it becomes hard to move engineers from one project to another.

An ASIC design management system is a software-based platform that offers a standard, consistent way of doing all designs (Figure 2), balancing the needs of all projects through the use of different modules, including a data manager, a build manager, an analyzer and a monitor.

### Files, files, and more files

During the life of a project, many files will be consumed and created, and many tools invoked with results analyzed and used to determine what to do next. A design management system automates and manages this process by providing a consistent methodology and flow throughout the chip implementation process.

It's easy for things to go awry in ASIC design and tempting to place faith in prior work that has been done by someone else. Proceeding without checking everything, from scripts to code to library files can produce problems that can result in rework.



**Figure 2 |** An ASIC design management system offers a consistent format for each design. One of these systems, for example, includes a data manager, a build manager, and an analyzer and monitor.

The wide variety of library files is a good example, especially if third-party IP is used with numerous files reflecting different process variants and corners. Even with established processes, it is not unusual to find a software bug in those files, and there's a real risk that such a bug might not be found until late in the project. With so many variants of library files and data, it is easy to mistakenly use the wrong one and, thus, the need for a way to manage revision control.



## EMBEDDED RISC SYSTEM DELIVERS PERFORMANCE AND I/O

Sealevel's new Relio R9-4100 embedded computer combines powerful RISC processing power with versatile serial and digital I/O in a rugged, compact form-factor. Perfect for control and monitoring applications, the R9-4100 provides eight optically isolated RS-485 ports to guarantee reliable serial communications even in electrically noisy environments and over long cable runs.

#### R9-4100 Features:

- 400MHz Atmel ARM9 Processor
- Dual 10/100 Base-T Ethernet
- 1 USB 2.0 Device and 1 USB 2.0 Host Ports
- 8 RS-485 Serial Ports with 1.5KV Port-to-Port Isolation
- 12 Optically Isolated Digital Inputs/Outputs
- 12VDC Power Input
- -10°C - +60°C Operating Temperature Range
- Windows CE and Linux OS Support

The R9-4100 is the ideal platform for embedded applications requiring small size, wide operating temperature range and flexible I/O connectivity. Contact us today to discuss your application or visit [www.sealevel.com](http://www.sealevel.com) to learn more.



Learn more about Relio R9-4100 by visiting [www.sealevel.com/ecd103/r9-4100](http://www.sealevel.com/ecd103/r9-4100) or by scanning the QR code to the left.

**SEALEVEL®**

[sealevel.com](http://sealevel.com) • 864.843.4343 • [sales@sealevel.com](mailto:sales@sealevel.com)

Typically, these library files are managed manually by copying from an original location into the data structure for the project, an error-prone effort. File names or directory names might indicate which files, yet those names may have no bearing on what's inside the file, and may actually be wrong. This is where a design management system can help by automatically locating, reviewing, and storing all of the project's input files. It parses (analyzes, organizes, and distributes) the files' contents to ensure that mistakes can't happen. It enforces a standard file structure so potential problems can be discovered early in the design phase, not later when they might impact schedule and costs.

Automated analysis plays a critical role in early reviews of the design. All of the design and IP files are scanned to extract the hierarchy, trace the clocks, and inspect module connectivity. Information from these early analyses help design engineers focus their efforts on the more problematic areas of the design.

**Automatic optimization**

Once the pre-design analysis is complete, work on the design can proceed. The build manager and analyzer modules are important components of this phase. The build manager encapsulates the entire ASIC flow, ensuring that every project has the same structure. Individual engineers don't maintain

their own build scripts. Instead, the design management system automatically generates build scripts (sets of programmatic instructions) and creates working directories for all of the various files needed by the EDA tools used for designing the ASIC.

ASIC design is not a linear process, and it might not be obvious which of the many design strategies might be the most effective. The build manager creates and runs a series of different design options so the project team can review them and find the best solution.

The analyzer, meanwhile, performs results aggregation and optimization (Figure 3). The various electronic design tools that range from synthesis and simulation to verification and test will generate a large number of different results files that, for a

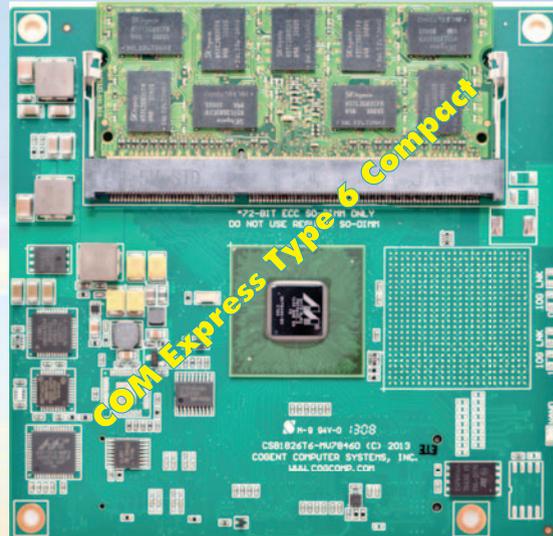
```
cell Rise(internal_nand2_x2a_a9t1) {
  index_1 {"0.01800, 0.03200, 0.06100, 0.11700, 0.23100, 0.45800, 0.91200 "};
  index_2 {"0.00085, 0.00201, 0.00474, 0.01117, 0.02633, 0.06205, 0.14622 "};
  values {" 0.01800, 0.03200, 0.06100, 0.11700, 0.23100, 0.45800, 0.91200",
" 0.00085, 0.00201, 0.00474, 0.01117, 0.02633, 0.06205, 0.14622",
" 0.03429, 0.04008, 0.05098, 0.07256, 0.11953, 0.22846, 0.48624",
" 0.04960, 0.05747, 0.07247, 0.09878, 0.14857, 0.25936, 0.51540",
" 0.07186, 0.08455, 0.10588, 0.14251, 0.20222, 0.31688, 0.53333",
" 0.10320, 0.12227, 0.15559, 0.20800, 0.29011, 0.42424, 0.64545",
" 0.14387, 0.17361, 0.22529, 0.30713, 0.42913, 0.61495, 0.92321"};
}
```

Figure 3 | The analyzer module extracts results from files and summarizes them in one place.

# Low Power, High Performance COMe Solutions



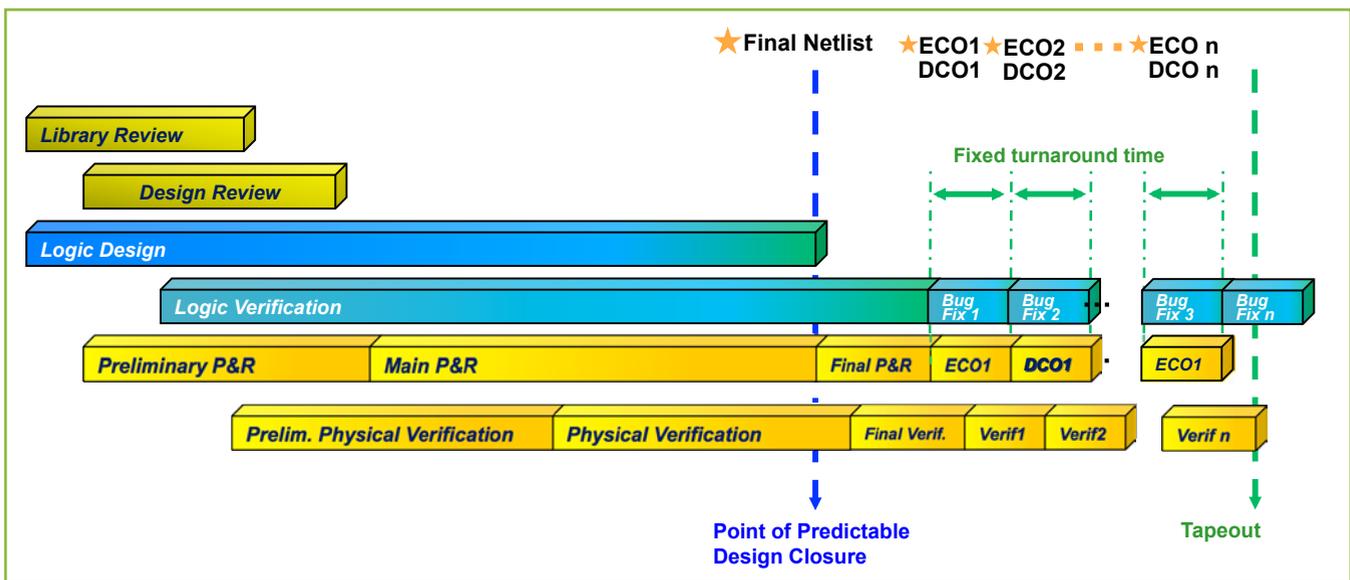
**CSB1890T10-Kabini**  
**1.5Ghz Quad Core AMD G-Series SoC**  
**4GByte 64-Bit DDR3-1600 w/ECC**  
**PCIe x4 and Dual x1 GEN 2, GIGe x 1**  
**Dual SATA Gen 2, Dual USB3, USB2 x 8**  
**Dual Display (DP, eDP or HDMI)**



**CSB1826T6-MV78460**  
**1.6Ghz Quad Core Marvell Sheeva (ARMv7)**  
**4GByte 64-Bit DDR3-1333 w/ECC**  
**PCIe x4 and Dual x1 GEN 2**  
**Dual SATA Gen 2 and USB2 x 6**  
**Quad Gigabit Ethernet, 1 x Copper,**  
**3 x SGMII (two are 2.5G Capable)**



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



**Figure 4** | A design management system can accommodate last-minute changes and incorporate them into the design.

human, are not easy to read, making it hard to identify the most pressing problems. The analyzer extracts results from the files and summarizes them in one place, making it easier to review the state of the current build.

### Implementing the ASIC top level

With a design management system, individual blocks in an ASIC design can be straightforward to implement. Each engineer is assigned a piece of real estate from the floorplan, and place and route within that area can proceed uneventfully, assuming there is enough silicon area to work in. It takes additional logic and interconnect to tie the blocks into a single design at the top level, which can be challenging.

When designing the floorplan, it is critical to leave area for the top-level ASIC integration to ensure the top-level logic (components) can fit into the channels. These added components must be stuffed into the channels provided around and through the finished blocks.

Depending on the size and complexity, a chip design project may involve a handful to dozens of engineers, all working on different blocks at the same time. The project manager will want visibility into how different aspects of the project are coming along.

As its name suggests, the monitor module has constant visibility into the current state of each design activity. Since it works in the background, it gives management visibility into progress without disturbing and slowing the design work. Any identifiable issues will be visible. Because they're caught early, they can be fixed early, making schedule slips less likely as the project progresses.

### Getting to the finish line

Tapeout is the most critical time in the project and where any final changes to the design must be addressed. These include ECOs (changes to the netlist) and DCOs (changes to the RTL code) (see Figure 4). In the ideal world, all changes to the design would be in place before tapeout. In reality, it is an expected

part of the design process that final design changes will be identified and must be addressed before the final tapeout.

A design management system should be able to accommodate last-minute changes and incorporate them into the design without having to go back and start over. Since the design management system has direct access to all project data, it can quickly accept design changes and automatically rerun the design with the new data.

### Winning the high-stakes ASIC design game

The stakes in ASIC design are enormous and range from lost market opportunities and revenue to a company going out of business as a result of a design failure. Too many things can go wrong when taking the risk of shortcuts. By establishing a standard project process and structure, a project team can engage each new project with confidence that it will proceed in a predictable, orderly fashion – a much better option than taking shortcuts. That process should include the implementation of a design management system. Uniquify, for example, has developed and implemented Perseus, a design management system that forms a project's backbone. Such design management systems are a critical resource for designers and managers. **ECD**



**Bob Smith** is Senior Vice President of Marketing and Business Development at Uniquify. He began his career in high tech as an analog design engineer working at Hewlett Packard. He has spent more than 30 years in various roles in marketing, business development, and executive management primarily working with startup and early-stage companies. He received a Bachelor of Science degree in Electrical Engineering from U.C. Davis and a Master of Science degree in Electrical Engineering from Stanford University.

**Uniquify**  
www.uniquify.com

Follow:

Taking a deep, detailed look is often the key to diagnosing diseases, and the same is the case with finding, diagnosing, and fixing bugs in FPGA SoCs through verification.

## Hitting the wall in FPGA SoC verification

By Thomas L. Anderson

*The ability to fix bugs by reprogramming has been a huge benefit for FPGA designers. Many have made do with little or no verification, preferring instead to debug the design in the bring-up lab. This method is breaking down in the era of FPGA SoCs with embedded processors. FPGA teams can greatly increase the quality of their verification by adopting a method used by many leading-edge ASIC and custom SoC developers: automatic generation of multithreaded, multiprocessor, self-verifying C test cases.*

Nobody ever said that FPGA design was easy. However, designers of programmable devices have long had a major advantage in verification over their counterparts developing ASICs and custom chips. That advantage, of course, is that design bugs missed during the FPGA verification process can be fixed without having to refabricate the chip. Not long ago, many FPGA designers didn't verify at all; they could just "blow and go" directly into the bring-up lab.

In the lab, the FPGA was inserted directly into a prototype of the final system. The team focused on product validation, bringing up the hardware and software (if needed) in the context of the product's end application. Any design bugs that slipped through could be tedious to find in the lab. Once one was found, it was a simple matter to reprogram the FPGA and continue bring-up. This

serial process worked well as long as the number of bugs remained fairly small.

As programmable chips grew larger and more complex, the number of bugs being found in the lab and the time it took to find them both increased significantly. To keep bring-up schedules reasonable, FPGA development teams realized that they had to do a better job of verifying their designs before going into the lab. In response, FPGA verification teams arose and started looking much like their ASIC cousins.

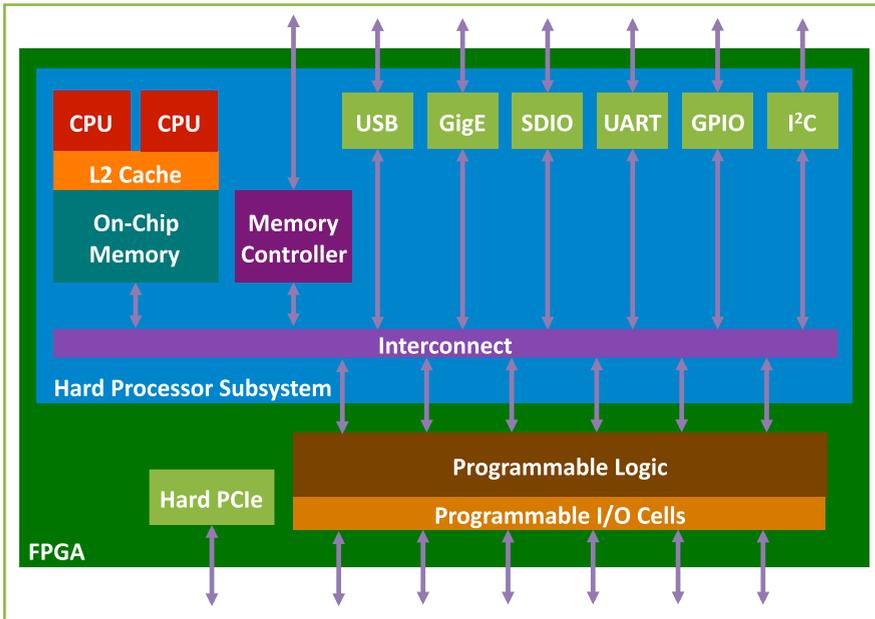
Register Transfer Level (RTL) simulation remains at the heart of all chip verification, and FPGA teams moved from simple handwritten vector files to more automated testbenches in simulation. Some adopted the constrained-random capabilities of the Universal Verification Methodology (UVM) standard. Static

analysis tools for checking clock domains and low-power structures started appearing, and some advanced FPGA teams even began using formal analysis.

More sophisticated verification approaches became more common to reduce the time spent in the bring-up lab and to accelerate time-to-market for the end product. However, FPGA designers still had the fallback position of being able to reprogram a device to fix a bug that slipped through verification and was found in the lab. With the advent of FPGA System-on-Chip (SoC) designs, even this escape mechanism increasingly is unavailable.

### The challenges of FPGA SoC verification

Figure 1 shows a representative FPGA SoC, based on publicly published block diagrams from multiple vendors. A large



**Figure 1** | Today's FPGA SoC contains multiple processors and standard interfaces in addition to traditional user-programmable logic.

portion of the chip remains traditional programmable logic available for the end product and its applications. However, a hard processor subsystem is included to provide SoC-class power. This subsystem typically includes at least two embedded processors, on-chip memory, and a wide variety of internal and external interfaces.

There are three reasons why FPGA teams are hitting the verification wall as they move into the SoC era with this type of complex chip. The first is the time to recompile and reprogram a huge FPGA. Once a bug is found and the source RTL code is changed, the process of creating a new image can take an entire day on a high-end personal computer. Then the image must be downloaded to the FPGA, which can take as long as several hours.

Secondly, the process of debug in the lab is also a big factor in the time it takes to make an FPGA design functionally correct. Once the chip is mounted into a real target system, it becomes hard to manipulate inputs or read outputs. Protocol analyzers are available for standard buses, but there are almost always FPGA ports with custom or ad hoc interfaces. It is not at all unusual for a team to spend days or even weeks in

the lab trying to track down the source of a single elusive bug.

Depending upon the FPGA architecture, the team may have to make multiple compile/program passes to bring out internal signals for debug purposes. Once the bug is found, it might take even more compile/program passes to try out possible fixes before verifying that the bug has been addressed. Generally, the team will verify the bug and test the fix in simulation before reprogramming. This is a smart move, but adds more time to the debug cycle.

The third aspect of the problem lies in the architecture of the FPGA SoC itself. By definition, an SoC has at least one embedded processor. It may have several, or many, homogenous or heterogeneous processors. The key to an SoC is that the processors are in charge, controlling the data flow between the many functional blocks, memories, and I/O ports. An SoC can do very little without software running on its embedded processors.

The main consequence of this is that some form of software must be available to run on the FPGA SoC's processors in the bring-up lab. End-product software usually is not ready when the

## 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 ethernets
- 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](http://www.embeddedARM.com)

(480) 837-5200

FPGA has been designed, so often the development team has to create special diagnostic software to test the device. This adds a resource burden to the project because this software must be developed in parallel with the hardware design.

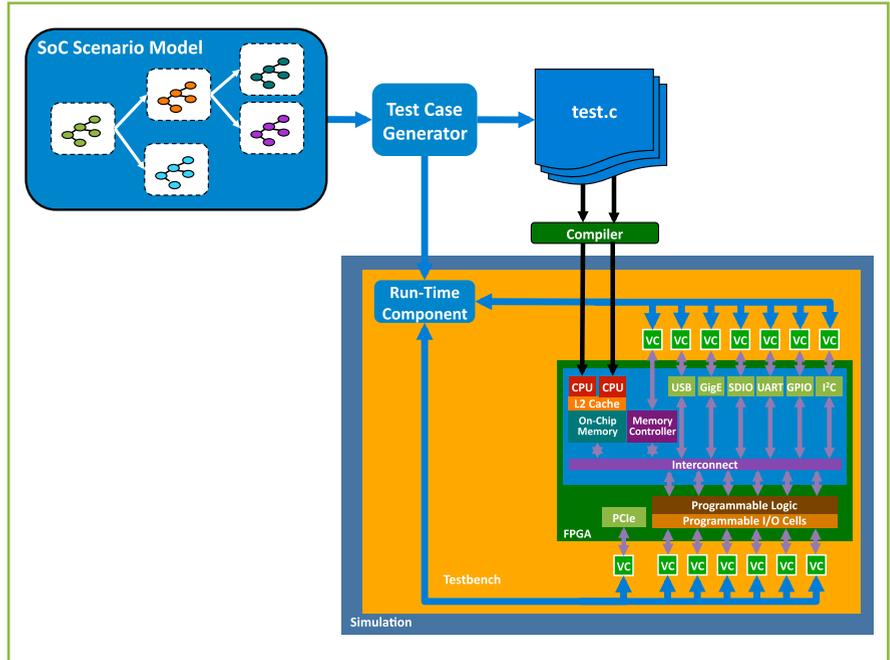
Handwritten diagnostic code is time consuming and expensive to develop, hard to maintain, and limited in functionality. Humans are not good at thinking in parallel, so it is rare for diagnostics to stress concurrency in the design, coordinate across multiple threads or multiple processors, or string together blocks into realistic end-user applications. The result is that design bugs may lurk in the FPGA until found at final system integration, or even by a customer.

**A solution from the non-FPGA SoC world**

To address the diagnostic software code dilemma, FPGA SoC developers must take another page from the book of ASIC and custom chip verification. They can benefit from a method that automatically generates multithreaded, multiprocessor, self-verifying C tests that stress system-level behavior in the SoC. These tests can be loaded into the embedded processor and run in simulation or in hardware acceleration. Figure 2 shows how this method works.

The source for the test case generator is a graph-based scenario model that captures both the intended chip behavior and the verification plan. The generator analyzes the graph to determine the capabilities of the design and then generates a set of test cases that verify these capabilities using the embedded processors. The C code is compiled and downloaded into the processors and run in simulation or simulation acceleration just as any other software would be.

These test cases are designed to stress the FPGA design, running multiple threads on multiple processors in parallel to test concurrent functionality. Since some of the test cases will pull data from the FPGA inputs or send data to its outputs, this approach includes a



**Figure 2 |** Multithreaded, multiprocessor, self-verifying C test cases can be generated automatically from a graph-based SoC scenario model.

runtime component in the testbench that coordinates the processors and the I/O activity. It is easy for the verification team to connect to standard UVM Verification Components (VCs).

Creating scenario models is straightforward since they reflect the data flow in the design and resemble SoC block diagrams. This initial investment enables the generation of virtually unlimited test cases to run in simulation. If a suitable connection to the I/O pins is available, it is even possible to run these test cases on the programmed FPGA.

This generation approach provides FPGA developers with a vast improvement over traditional “burn and churn” reprogramming cycles as bugs are found one by one in the lab. Automated test cases save development time, provide more thorough verification, and save resources because embedded programmers don’t have to develop throwaway diagnostics. The result is a faster and more predictable FPGA development schedule for even the most complex SoC designs. Such a test case generator is available today: the TrekSoC solution from Breker Verification Systems. **ECD**



**Thomas L. Anderson** is VP of Marketing for Breker Verification Systems. His previous positions include Product Management Group Director of Advanced Verification Systems at Cadence, Director of Technical Marketing in the Verification Group at Synopsys, VP of Applications Engineering at 0-In Design Automation, and VP of Engineering at Virtual Chips. He holds a Master of Science degree in Electrical Engineering and Computer Science from MIT and a Bachelor of Science degree in Computer Systems Engineering from the University of Massachusetts at Amherst.

**Breker Verification Systems**  
www.brekersystems.com

Follow:



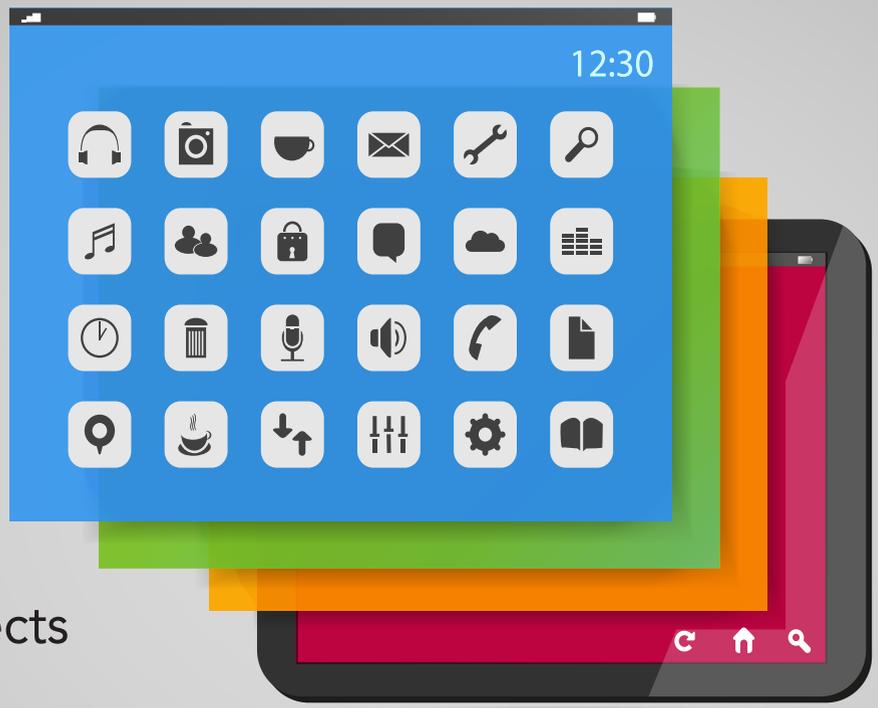
# LET'S REWIRE THE WAY WE THINK ABOUT ENERGY.

Meeting the world's growing demand for energy means everyone and everything must rethink how energy is generated and delivered. Enter Wind River. As the Intelligence in the Internet of Things, we've recently been selected by one of the world's leading energy management providers to be the backbone of their next-generation smart substations. So now simple routers are active decision-makers capable of calibrating demand, adjusting production, and communicating with other points in the grid to intelligently ensure a consistent supply of greener, cleaner energy. It's a huge stride that's leaving a small footprint.

**WIND RIVER**

To find out more visit us at [iot.windriver.com](http://iot.windriver.com)

The Intelligence in the Internet of Things



## Embedded virtualization protects legacy investment

By Robert Day

*Embedded developers can now take advantage of a real-time embedded virtualization technology that can help protect their legacy software investment as they move to new hardware platforms. This gives embedded developers the same benefits as IT organizations that have been running legacy Operating Systems (OSs) and applications on new server platforms for years. This investment is particularly key as hardware platforms shift from single-core to multicore 32- to 64-bit architectures and add new connectivity such as Wi-Fi, Bluetooth, or cellular communications.*

As embedded systems take on a new lease on life with the advent of the Internet of Things (IoT) and as they take part in the very relevant subset of Machine to Machine (M2M) communications, they are evolving faster now than ever over their relatively short history. These traditionally dedicated systems now have to evolve onto new architecture platforms that have more than one core, and often have or need 64-bit architectures rather than only 32 bits of address space. Adding new communication mediums such as Wi-Fi, Bluetooth, or cellular and new GUI requirements such as touch, swipe, and drag brings a whole new set of complications for embedded designers.

However, the virtualization paradigm can now be used by embedded developers as they try to evolve their legacy systems to meet these new requirements without a complete rewrite or redesign of their

existing systems; hence, virtualization can help developers manage the cost and time of bringing their systems into today's connected world.

### **An elegant solution to a tricky problem**

There is one technology that allows embedded developers to preserve their legacy software system even though the world around it – including the hardware that it runs on – is in a world of change. The technology is virtualization. This technology has been used in the enterprise software world for many years, allowing the IT departments to run multiple versions of IT applications and operating systems long after they are effectively obsolete. It is not widely used in the embedded world yet, mainly because the embedded system refresh rate is not as fast as in the enterprise world, but also because the enterprise hypervisors do not meet the embedded

size and performance requirements in much the same way a desktop OS is less frequently used than a Real-Time Operating System (RTOS). What is needed for the embedded world is an embedded hypervisor that gives the benefits of virtualization but satisfies the key embedded requirements of small memory and code size and real-time deterministic performance.

Before examining how an embedded hypervisor can benefit the embedded developer, let's examine virtualization in a little more detail. Virtualization is a technology that offers a "virtual" representation of hardware to the software running on top of it, including the processor, memory, and devices. A good example of this in our everyday world is the hypervisor that runs on an Apple Mac and allows Microsoft Windows to run next to Mac OS on the same hardware platform, because it is supplying

Windows a “virtual” PC. This is a good example of a Type 2 hypervisor, where the hypervisor actually runs on top of the native operating system as an application, and then the “guest” operating system runs on top of it. This gives the advantage that the native OS is still in control and the guest OS is running at the same time, as it is essentially running as an application. However, there is a huge performance issue with the multiple layers of software in this solution, and this technology is not very portable for use across all embedded RTOSs.

There is also another virtualization solution that is commonly known as “Type 1” or “bare metal” that doesn’t rely on the native OS, and interacts directly with the hardware to provide the “virtual” hardware to the guest. This is getting closer to meeting the needs of the embedded developer, as it is certainly more efficient than the Type 2. However, as can be seen in Figure 1, it still relies on a “helper” operating system, which is close in size to the native OS in the Type 2 case. And even though it is known as “bare metal,” that does not mean it has the real-time or deterministic properties required by embedded developers.

Thus, a different breed of hypervisors has been developed that could be used by embedded developers to utilize a virtualization solution. Because of some inherent differences in design required to build an efficient embedded hypervisor, a new name was required to differentiate from the Type 1, and so, the embedded hypervisor has been called a Type 0 hypervisor.

Two key differentiators when comparing Type 0 to Type 1 are 1) size and 2) real-time performance. By removing the “helper” OS used in a Type 1 and creating a true bare-metal hypervisor, the runtime memory requirements drop from GB to MB and the static code size drops from MB to KB (see again Figure 1). Also, by removing the heavyweight helper OS and replacing it with a small embedded real-time kernel or “separation kernel” to separate resources, real-time determinism can now be realized. Additionally, by taking advantage of hardware virtualization features found in modern processors, the performance

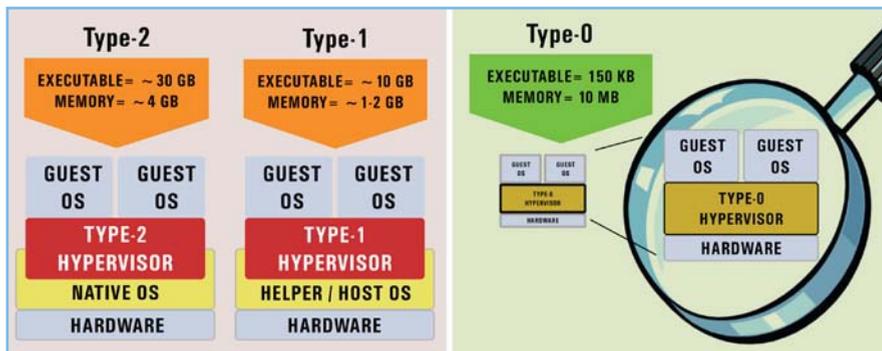


Figure 1 | Embedded Type 0 hypervisor comparison.

overhead of using a hypervisor get to over 95 percent of native performance.

### Real world embedded virtualization use

Having established that an embedded hypervisor can provide a reasonable embedded footprint and real-time performance, we can now focus on the benefits to embedded developers as they look at their migrating their existing system to the next generation of connected embedded devices. The key objectives for these new systems can be the following:

- Move to modern hardware platforms that can be both multicore and 64-bit
- Compatibility with a modern standard user interface
- Increased connectivity, and with that the requirement for increased security

How can an embedded hypervisor help with all of these, and reduce the amount of rewritten code from the legacy system? The hypervisor does a couple of things really well that would be difficult to achieve with just a standard RTOS.

# 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](http://www.smxrtos.com) today.



800.366.2491 sales@smxrtos.com

[www.smxrtos.com](http://www.smxrtos.com)



Processors Supported: [www.smxrtos.com/processors](http://www.smxrtos.com/processors)

Free Evaluation Kits: [www.smxrtos.com/eval](http://www.smxrtos.com/eval)

Free Demos: [www.smxrtos.com/demo](http://www.smxrtos.com/demo)

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



### Abstracting the underlying hardware

The first of these hypervisor capabilities is that of abstracting the underlying hardware from the RTOS and applications that run on it. This really helps when migrating an existing RTOS and applications to a new hardware platform, as the hypervisor can make the virtual hardware look just like the original hardware, and can apply to memory, processor, and devices.

So as the new physical hardware could be a multicore, 64-bit processor with a set of new modern devices, the hypervisor could present a 32-bit, single-core virtual processor with the new physical devices being mapped to virtual versions of the legacy devices. This prevents the immediate move to a new SMP, 64-bit RTOS, and also reduces the need to create new BSPs for the new devices on the board. It is also a relatively elegant way of introducing new connectivity options such as Wi-Fi or cellular networks, as they could be made to look just like a regular Ethernet device by means of device virtualization.

### Multi-OSs facilitate new interfaces, security

The second key capability of virtualization is to allow multiple operating systems to run on a single hardware platform. This allows the embedded developers to keep their legacy system intact, and add new functionality by introducing another operating system running in parallel. This might seem very inefficient, but with the leap in hardware technology, and the cost of memory, a modern multicore system with hardware virtualization support could actually run multiple operating systems without any performance degradation; this is also much cheaper than having to source obsolete hardware components. This multi-OS scenario is another interesting design consideration when bringing new connectivity mediums into the system, as rather than virtualizing the new devices, the new operating system could be used to communicate externally and then pass information via the hypervisor's internal communication scheme to the legacy system.

This multi-OS scenario also has some interesting design benefits for connected embedded systems outside of just bringing in legacy systems, and can help bring in new requirements such as modern user interfaces and additional security. An issue that has been constant for embedded developers has been standard user interfaces for embedded systems, as standard GUI-based OSs have traditionally been too big or too slow for real-time systems. And with RTOSs, the GUI has to be built from scratch and hence does not have all the touch/swipe functionality that we are used to on our phones, tablets, and computers.

By using an embedded hypervisor, developers can have their cake and eat it too. By having an RTOS in one Virtual Machine (VM) and a more traditional GUI OS such as Android in another VM, the real-time part is taken care of and a user-friendly standard GUI is built in; by using a multicore processor, each can be given their own dedicated processor, memory, and resources by the hypervisor. (See Figure 2, which shows the separation of real-time from desktop functions.)

The final piece of the puzzle is security – possibly one of the hottest topics for connected embedded devices, as these once dedicated and proprietary systems are now being connected and controlled across the open Internet; therefore, they are now possible targets for cyber crime and cyber terrorism. By using a separation kernel and embedded hypervisor, different parts of the system can be easily segregated and protected.

For example, the virtual machine that is connected to the Internet is often different from the virtual machine that is controlling something or storing information, which is typically the primary target for malicious attacks. So even if one VM gets infected, that infection cannot spread to another virtual machine as the separation kernel is keeping them apart, in much the same way that they would be if they were running on physically separate pieces of hardware.

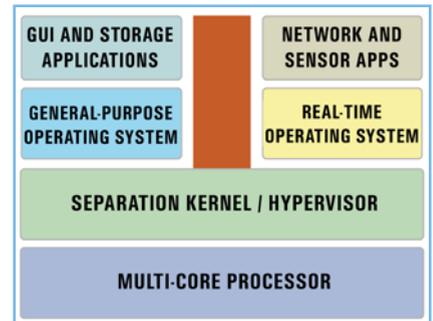


Figure 2 | Multi-OS RTOS and general OS embedded system.

### Is embedded virtualization real?

In conclusion, the virtualization technology known as the embedded hypervisor, based on a real-time separation kernel, can help embedded developers bring their legacy embedded systems into the next generation of connected, multicore systems with user-friendly GUIs and added security to protect against malicious threats. It all seems a little too good to be true. At LynxWorks, we have developed a separation kernel and embedded hypervisor called LynxSecure, and now in its fifth generation, it is helping embedded developers meet their new design goals – on time and on budget. **ECD**



**Robert Day** is Vice President of Marketing for LynxWorks, where he is responsible for all global external and internal

marketing functions. With more than 20 years in the embedded industry, his most recent position prior to joining LynxWorks was heading marketing for Mentor Graphics' embedded software division. Prior to the marketing role, he held a variety of management, sales, and engineering positions for Mentor Graphics and Microtec Research, spanning more than 18 years in total. He is a graduate of the University of Brighton, England, where he earned a Bachelor of Science degree in Computer Science.

**LinuxWorks**  
www.linuxworks.com

Follow:



## » Innovative computing platforms

based on 4th generation Intel® Core™ Processors «

- » Enhanced performance
- » More power efficiency
- » High graphics and media power
- » Improved security and manageability

### Profit from our competencies

- » System- and OS-integration services
- » Customization and ODM services
- » Extended lifecycle management
- » Application and migration support
- » Excellent global technical support



COMe



mITX



cPCI



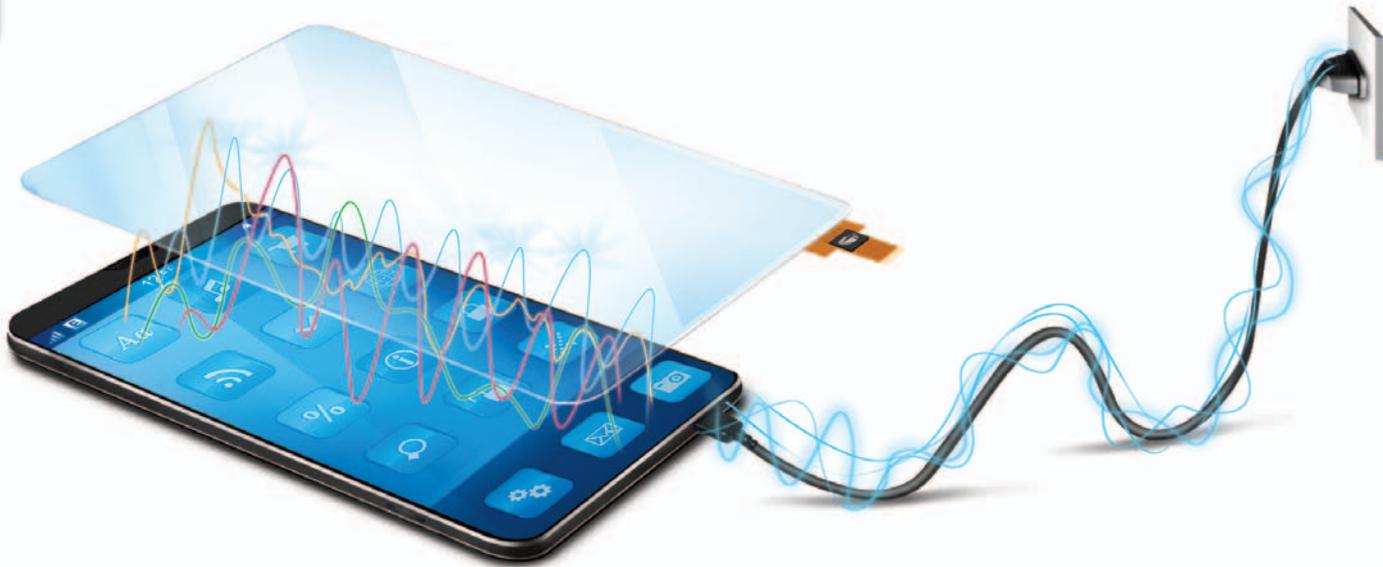
Symcloud

and more to come



Get more information at:  
[kontron.com/next-gen](http://kontron.com/next-gen)





## Improving performance, accuracy, and reliability in touch-screen based applications

By Jeff Erickson

*Challenges arise in providing reliable touch performance in devices when the touch sensor is subjected to large amounts of noise from sources such as displays and chargers. This noise can corrupt reported touch data, degrading the performance and reliability of the device. Today's touch-screen controllers use a variety of features to boost SNR and provide a seamless user experience amidst the noisy environment.*

As touch-screen based devices continue to get thinner, different sources of noise become more problematic to address. Noise can substantially impact the reliability of touch-screen performance and lead to undesirable product returns because of the following issues:

- Higher jitter leads to variability in the location reported for a stationary finger
- False touches are reported when there is no finger on the screen
- There is no reported touch when a finger is touching the screen
- The device completely locks up

Thus, reliability, performance, and the quality of user experience are significantly affected by how the system addresses noise. While some noise is environmental – including radio signals, AC mains, and even fluorescent lights – there are several internal sources that influence performance, accuracy, and reliability on a larger scale. These include common-mode noise from chargers and

noise generated from a display such as an LCD.

### The downside of thinner

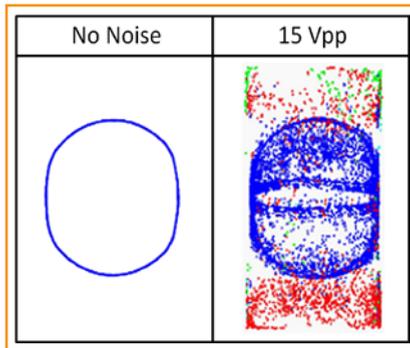
Displays are a major source of noise, and the thinner profile of today's devices places displays closer to touch sensors. Traditionally, air gap and shield layers have protected sensors from coupling in display noise. To reduce cost (for example, the shield layer for a 4" display can cost as much as \$1) and reduce thickness, these layers have been removed, and the touch sensor is laminated directly to the display. As a result, more noise is coupled into the sensor from the display with a substantial impact on performance and reliability. This issue is only expected to become more prominent as new display-integrated stackups become more popular. With display integration, the touch sensor is integrated into the display stackup, placing the sensor's receive electrodes even closer to the display and thus coupling in even more noise.

With thinner touch-screen architectures, the reduced distance between transmit and receive electrodes also increases the parasitic capacitance ( $C_p$ ) of the touch sensor. Higher  $C_p$  increases the time it takes to charge and discharge when scanning, limiting the maximum frequency at which the sensor can be scanned. Since there is typically less noise in the higher-frequency bands, having to scan at a lower frequency means the system will be exposed to more noise. Scanning at a lower frequency also lowers the refresh rate and requires more power. A lower refresh rate means less frequent touch reporting, which can result in reduced performance and reliability.

### Common-mode noise

Another key source of noise for capacitive touch screens is aftermarket device chargers. While the typical charger supplied with a device has been designed to minimize noise – quiet chargers produce noise on the order of 1 to 5 Vpp – many aftermarket chargers are designed to be low cost. As a consequence, these

chargers often use low-cost components and even completely omit key components for reducing common-mode noise, resulting in noise that can fluctuate in the range of 20 to 40 Vpp, causing charge transfer on the order of 10 to 100 pC. With this much noise it becomes extremely difficult to sense if and where a finger, that couples in a relatively miniscule 5 to 500 fC of charge, is touching the sensor. Figure 1 displays the difference in reported touch data when tracking a single finger drawing a circle while operating with no charger noise versus 15 Vpp of charger noise. (Many chargers are much worse than this.) With charger noise injected, the reported location of the finger on the panel (shown in blue) is not only corrupted, but additional false touches are reported on the panel (shown in various other colors).



**Figure 1** | Effects of charger noise on finger tracking

to earth ground, touches the screen and causes the noise to be injected into the touch screen. The amount of charge injected is primarily determined by the peak-to-peak voltage of the noise and affected by the thickness of the touch screen's cover lens and the area of contact between the finger and the touch screen, as follows:

**Thickness:** A touch screen cover lens is the primary factor determining the distance between a finger and the receive electrode. An OEM moving to a thinner lens needs to be aware that a 0.5 mm

Common-mode noise arises when both the power and ground supplies fluctuate in voltage relative to earth ground. Because they maintain the same differential voltage, performance is only affected when a finger, that is coupled

**COMMELL**  
We offer advanced and reliable IPC products

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



- Support Intel® Haswell Desktop Core™ 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

**Mini-PCIe & Mini-PCI Card**

<b>MPX-2585B</b>	<b>MPX-2515</b>	<b>MP-6100E</b>
		
<ul style="list-style-type: none"> <li>Video Capture</li> <li>8-CH Video in</li> <li>8-CH Audio in</li> <li>240fps @D1</li> <li>Windows SDK</li> </ul>	<ul style="list-style-type: none"> <li>CAN 2.0B card</li> <li>Microchip 2515</li> <li>ISO-11898</li> <li>1 Mb/s</li> <li>API &amp; SDK</li> </ul>	<ul style="list-style-type: none"> <li>4-ch H/W H.264 video capture</li> <li><b>MP-954</b></li> <li>RS232/422/485</li> </ul>

**ARM-based 3.5" miniboard**

**LE-340**

- Freescale i.MX535
- DDR3 up to 1GB
- 2 x 10/100 LAN, USB
- VGA/LVDS, Audio
- RS232/485, eMMC
- SATA2, SD, CAN

**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

**HS-774 PCI Single Board Computer**

- Intel® Ivy Bridge mobile Core™ i3/i5/i7
- Intel® QM77 Express chipset
- DDR3 up to 16 GB, VGA, LVDS, DVI
- 2 x Giga LAN, SATAIII, HD Audio
- RS232/422/485, USB3.0, Mini-PCIe
- SIM socket, GPIO, 185mm x 122mm

**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

**AS-C76 Industrial ATX mainboard**

- Intel® Ivy Bridge desktop Core™ i3/i5/i7
- Intel® Q77 chipset, VGA, DVI
- DDR3 up to 16 GB, SATAIII
- 2 x Intel® Giga LAN
- RS232/422/485, USB 3.0
- HD Audio, Mini-PCIe socket
- 5 x PCI, Mini-PCI, PCIe x16, PCIe x 4

**www.commell.com.tw**

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

**Keil MDK-ARM**

MDK-ARM™ is the complete software development environment for ARM® Cortex™-M series microcontrollers.

[www.keil.com/mdk](http://www.keil.com/mdk)



**ARM DS-5**

DS-5™ Professional is a full featured software development solution for all ARM Powered® platforms.

[www.arm.com/ds5](http://www.arm.com/ds5)

**ARM**

+1 800 348 8051

© ARM Ltd. AD370 | 02.13



lens cover will have approximately twice the capacitance – and thus inject twice the noise – when compared to a 1.0 mm cover lens.

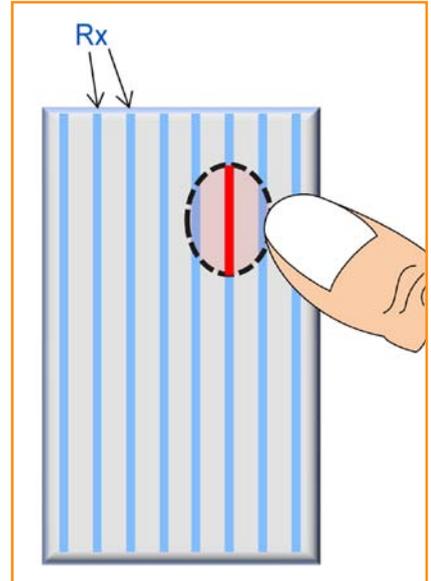
**Finger diameter:** Capacitance increases proportionally as the area of contact between the finger and touch screen increases, and more capacitance means more noise. Since the touch screen's receive electrodes are laid out as narrow rows or columns, finger diameter is the primary consideration when testing a system's immunity to charger noise (see Figure 2). Some OEMs test for immunity using a small finger, on the order of 7 mm in diameter. However, this can give an inaccurate representation of a system's reliability, given that an average finger has a diameter of 9 mm and a typical thumb ranges from 18 to 22 mm. As a 22 mm thumb injects 3X the charge of a 7 mm finger, certain gestures, such as unlocking a phone or scrolling through a list, might perform

unreliably if a system has not been sufficiently tested across the full range of finger sizes.

To be reliable, devices must be able to detect a finger in the presence of massive charger noise, even when the finger signal is several orders of magnitude smaller. Note that chargers are not the only source of common-mode noise. The new Mobile High-Definition Link (MHL) standard designed to connect mobile phones and HDTVs, for example, sources common-mode noise from the TV's power supply and drives it through the HDMI and USB cables connected to the phone.

### Improving SNR

Given the variability of environmental noise that can be present, systems need to be able to adapt to the level and type of noise present at any given time. How effective a system is at this is measured using the Signal-to-Noise Ratio (SNR).

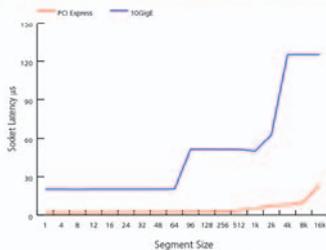


**Figure 2 |** Intersecting area of finger and Rx electrode

Several ways exist to increase a touch screen's SNR and eliminate bad data (that is, detecting false touches or not reporting real touches):

# Make the Right Connection

## PCI Express® Networks



**Higher Transmit (Tx) Voltage:** As signal, and thus SNR, is directly proportional to the Tx voltage, increasing the Tx voltage improves noise immunity. Traditionally, this has been achieved through the use of either an external high-voltage analog supply, which also has the effect of increasing power consumption, or external components such as switching regulators that increase system size and cost. Next-generation touch-screen controllers support a higher Tx voltage on-chip through an internal charge pump.

**Advanced filtering:** Filtering is also an effective way to remove noise. Loading the touch-screen controller CPU with noise-filtering algorithms, however, reduces the refresh rate while increasing power consumption. Alternatively, hardware accelerated filtering boosts the SNR while maintaining the refresh rate. What may take tens or hundreds of milliseconds to compute with the CPU can be completed in less than a millisecond in hardware.

**Adaptive frequency hopping:** When the noise frequencies generated by the LCD are close to the sensor-scanning frequency, touch data can be corrupted.

Using adaptive frequency hopping, the scanning frequency can be changed to one where the noise is low enough to prevent corruption. To be most effective, frequency hopping must support a wide range of frequencies to provide immunity across all types and frequencies of noise. In addition, given that the fundamental frequencies of noise from chargers primarily reside from 1 KHz to 300 KHz, high-frequency scanning in the range of 300 KHz to 500 KHz can completely avoid the highest amplitudes of charger noise as well as the first few harmonics.

**Increased receive range:** Noise levels that are too high can completely saturate a touch-screen controller's receive channel and render the touch screen useless. This can be avoided by dividing down the raw signal passed to the receive channel to reduce the noise. However, this also divides down the finger signal, so care must be taken with this approach. Increasing the range of the receive channel so it can handle larger amounts of charge is also an effective approach, but it comes at the cost of larger capacitors that require greater silicon area.

**Delivering a reliable user experience**  
Next-generation touch-screen based devices offer slimmer profiles and are expected to support any compatible charger. While important to users, these features also increase the amount of noise in which devices must be able to operate in the presence of, and perform reliably. By employing a combination of techniques to increase SNR, touch-screen controller vendors can ensure a consistent, accurate, and reliable experience for users. **ECD**



**Jeff Erickson** is a Staff Product Marketing Engineer and Product Manager at Cypress Semiconductor in the TrueTouch touch-screen controller

group. He holds a Bachelor of Science in Electrical Engineering and a Master of Business Administration, both from the University of Washington. He can be reached at [jeff.erickson@cypress.com](mailto:jeff.erickson@cypress.com).

Cypress Semiconductor  
[www.cypress.com](http://www.cypress.com)

Follow:      

SPECIAL ADVERTISING FEATURE

## IDF Product Spotlights:

### Massively Scalable Content Delivery with SYMKLOUD Media Platform

- Deploy real-time 1080p H.264 and H.265/HEVC encoding/transcoding in the cloud
- Highest video channel density per watt for multiscreen services
- Distributed processing, switch and load balancer in one 2RU, 21-inch deep platform
- Up to 18x 3rd gen and/or 4th gen Intel® Core™ i7 processors
- Capable to support elastic cloud provisioning and management



**Kontron**  
800-354-4223  
[www.kontron.com/symkcloud](http://www.kontron.com/symkcloud)

### Take ECD With You Wherever You Go



<http://appstore.com/embeddedcomputingdesign>



<http://www.amazon.com/gp/product/B009SQJCPQ>



# E-community Post

Joining the embedded conversation

By Monique DeVoe

www.embedded-computing.com

BLOGS | MARKET STATS | INNOVATION | VIDEOS | SOCIAL MEDIA

## BLOG

### Chip Wars 2.0 – Intel’s Quark takes aim at ARM, IoT

By Brandon Lewis

To kick off IDF 2013 in San Francisco, Intel CEO Brian Krzanich unveiled some of his company’s “Quarks,” the first working SoCs based on 14 nm process technology. Although Krzanich and Intel President Renee James offered little by way of schematic details, they did allow that the Quark processors measure in at 1/5 the size and consume only 10 percent of the power of Intel’s Atom line.

Read more: <http://opsy.st/18OFjts>

Read blogs at: <http://embedded-computing.com/blog/>

[To become an Embedded Computing Design ([www.embedded-computing.com](http://www.embedded-computing.com)) guest blogger, send me a one-paragraph abstract for consideration at [mdevoe@opensystemsmia.com](mailto:mdevoe@opensystemsmia.com).]

## WHITE PAPER

### Using the High Performance Memory Subsystems in an FPGA for Innovative Bridging Applications

By Staff, Microsemi

When the advanced features of the High-Performance Memory Subsystem (HPMS) – DDR controllers, embedded non-volatile and SRAM memory, DMA, and an efficient AHB Matrix – are combined with a wealth of FPGA fabric (including advanced DSP and embedded memory blocks), and up to 16 lanes of multi-protocol SERDES, impressive external bandwidth can be matched with extensive internal data-processing and data-management capabilities. This allows even the most bandwidth hungry bridging applications to run continuously at full speed while keeping power, cost, and board space within even the stingiest of budgets.

Read more: <http://opsy.st/11mnD7W>

## E-cast

### On the Road to Space-Age Vehicles: Developing for the Convergence of consumer Electronics and In-Vehicle Entertainment

Presented by Polarion, IBM, Vector Software, Mentor Graphics

The advances in In-Vehicle-Infotainment (IVI) and advanced driver assistance technology encompass a wide variety of connectivity, graphics, and sensor technologies. Time-to-market pressures place significant demands on developers to provide more entertainment options and communications capabilities while continuing to enhance driver and passenger safety and awareness. It’s critical to start with a strong software foundation and utilize a development environment in order to meld these technologies into a safer, more integrated automotive environment. This E-cast will explore software, tools, and techniques that further integrate the driver with the automobile to increase safety and provide a more enjoyable driver and passenger experience.

Go to E-cast: <http://ecast.opensystemsmia.com/420>

## VIDEO



### Raspberry Pi – ARM Processor Modes

Though the Raspberry Pi board is simple enough to use by students and beginners, getting into the guts of the hardware can take some embedded know-how. The Raspberry Pi’s ARM processor has several different modes

that handle things like interrupts and allow the user to do more complex operations. This video explains what the ARM processor modes are and how to use them.

Watch the video: <http://opsy.st/15DBMjf>

See more videos at: [video.opensystemsmia.com](http://video.opensystemsmia.com)

### Roving Reporter blog: Hey, core eyes! Machine vision with Haswell

By Brandon Lewis

As robotics continue to replace humans on the factory floor, Machine Vision (MV) technology has become industrial automation’s new lens. Applications like motion control and quality assurance require high-resolution image analysis from these new “eyes” on the assembly line, and must execute with extreme precision to ensure optimized manufacturing processes.

Read more: <http://opsy.st/1f2yhGk>

## E-cast Intel platform codename Bay Trail: An SoC for Intelligent Systems

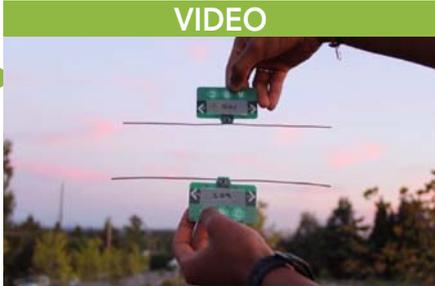
Presented by Intel, ADLINK, congatec

The next-generation Intel Atom processor codenamed Bay Trail is a true System-on-Chip (SoC) family for Intelligent Systems. This product family features the new Intel Silvermont microarchitecture with up to four re-architected CPU cores, an advanced new Intel GMA graphics engine, integrated I/O for various embedded uses, integrated media and security accelerators – all on a low-power 22 nm device that enables smaller, lighter, power efficient designs. This E-cast will examine the features and benefits of the next generation Intel Atom processors.

Go to E-cast: <http://ecast.opensystemsmedia.com/421>

More E-casts: [ecast.opensystemsmedia.com](http://ecast.opensystemsmedia.com)

### VIDEO



#### University of Washington Ambient Backscatter

Power is almost always a challenge for embedded computing devices. Mobile device batteries need

to last long enough to be useful but cannot be too heavy or expensive, and in other applications a reliable power source is not always available. University of Washington students came up with Ambient Backscatter communication that doesn't require any outlets or batteries. You read that correctly. It uses existing TV and cellular transmissions as its power source and communication method between two battery-free devices. This may be something to keep an eye on for future embedded developments.

Watch the video: <http://opsy.st/19eOTWY>

### WHITE PAPER

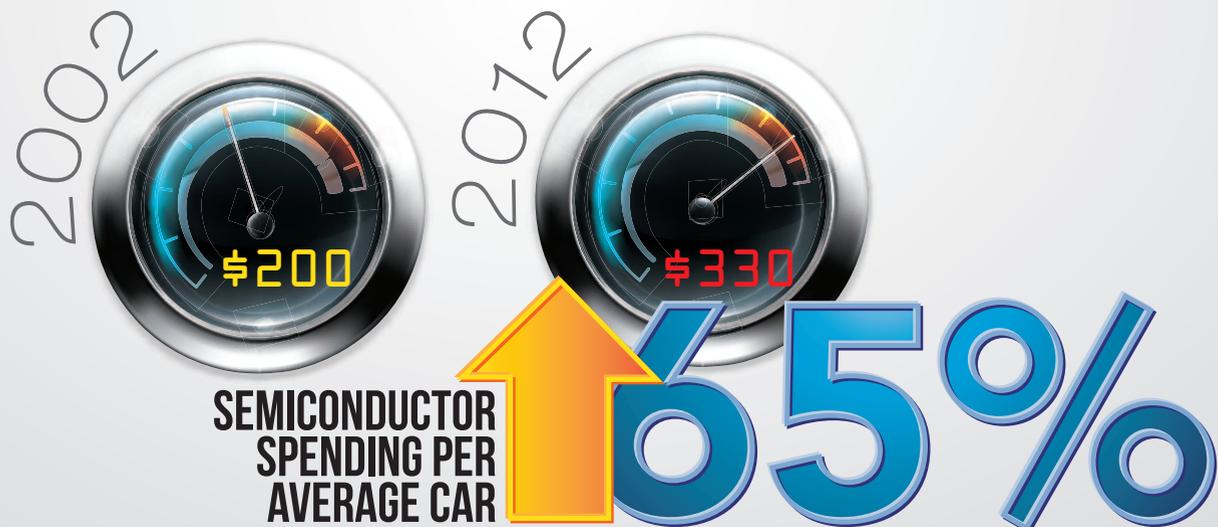
#### A scalable software build accelerator

By John Ousterhout and John Graham-Cumming, Electric Cloud

For organizations that depend on software innovation, a slow software build process can be a bottleneck for the entire company. Slow build times not only impact engineering efficiency, but they can affect product quality and company agility. Furthermore, diagnosing build problems is difficult to impossible as the cryptic output in Make log files can be hard to decipher and it is difficult to relate to the individual build steps in a large build. This paper looks at the impact of slow builds, traditional approaches to improving build performance, and how to combat these issues.

Read more: <http://opsy.st/1ctYJJb>

Read other white papers at: [whitepapers.opensystemsmedia.com](http://whitepapers.opensystemsmedia.com)



Source: IHS Inc. <http://opsy.st/18viGrO>

# OpenSystems Media

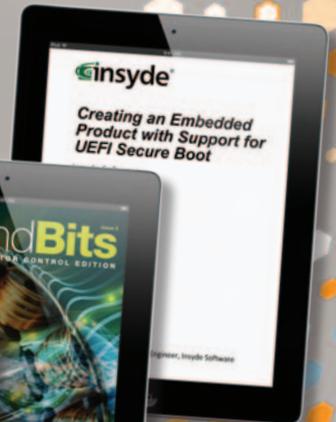
works with industry leaders to develop and publish content that educates our readers.

Check out our white papers.

<http://whitepapers.opensystemsmedia.com/>

## Most popular topics:

- |                              |                    |
|------------------------------|--------------------|
| AdvancedTCA                  | Multicore          |
| Android                      | PCI Express        |
| Avionics Certification       | Radar              |
| Automotive                   | SDR                |
| Deep Packet Inspection       | Static Analysis    |
| GUI Linux in Medical Devices | Switched Fabrics   |
| Internet of Things           | Test & Measurement |
| M2M                          | UAVs               |



# 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

# Industrial ARM® Single Board Computers

## High-Performance Graphics with Industrial I/O and Expansion

### -40° to +85°C Operating Temperature

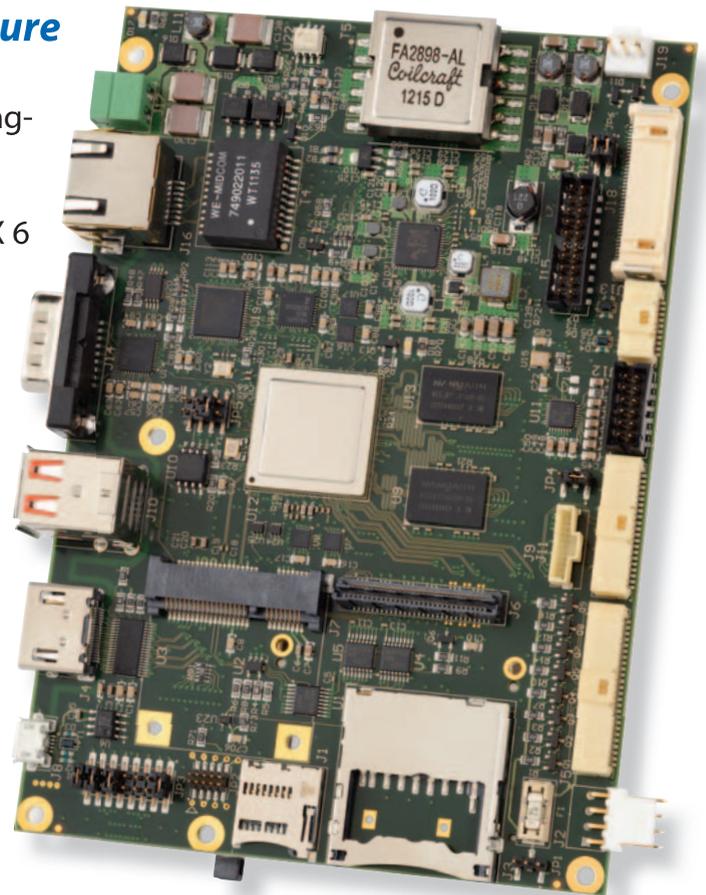
Designed for demanding applications and long-term availability, WinSystems' SBC35-C398 single board computers feature Freescale i.MX 6 industrial application processors with options for expansion and customization.

### Features

- ARM Cortex™-A9 Processors;  
Quad, Dual, or Single Core
- Multiple Graphics Interfaces
- Wide Range DC or PoE Power Input
- Gigabit Ethernet with IEEE-1588™
- USB 2.0 Ports and USB On-The-Go
- Dual FlexCAN Ports
- Multiple Storage Options
- Mini-PCIe and IO60 Expansion
- Linux and Android™ Supported

Learn more at

[www.WinSystems.com/ARME1](http://www.WinSystems.com/ARME1)



Call 817-274-7553

Ask about our product evaluation program.



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



WinSystems® is a registered trademark of WinSystems, Inc.

Freescale and the Freescale logo are trademarks of Freescale Semiconductor, Inc., Reg. U.S. Pat. & Tm. Off.

Android is a trademark of Google Inc. The Android robot is reproduced from work created and shared by Google and used according to terms described in the Creative Commons 3.0 Attribution License.

Scan this tag to  
read more about  
our ARM SBCs.

