Embedded COMPUTING DESIGN

www.embedded-computing.com

@embedded_mag
Embedded Computing Design

MAY 2012
VOLUME 10 • NUMBER 3

Creating VISUAL REALITY

SCAN ME:

www.opensystemsmedia.com

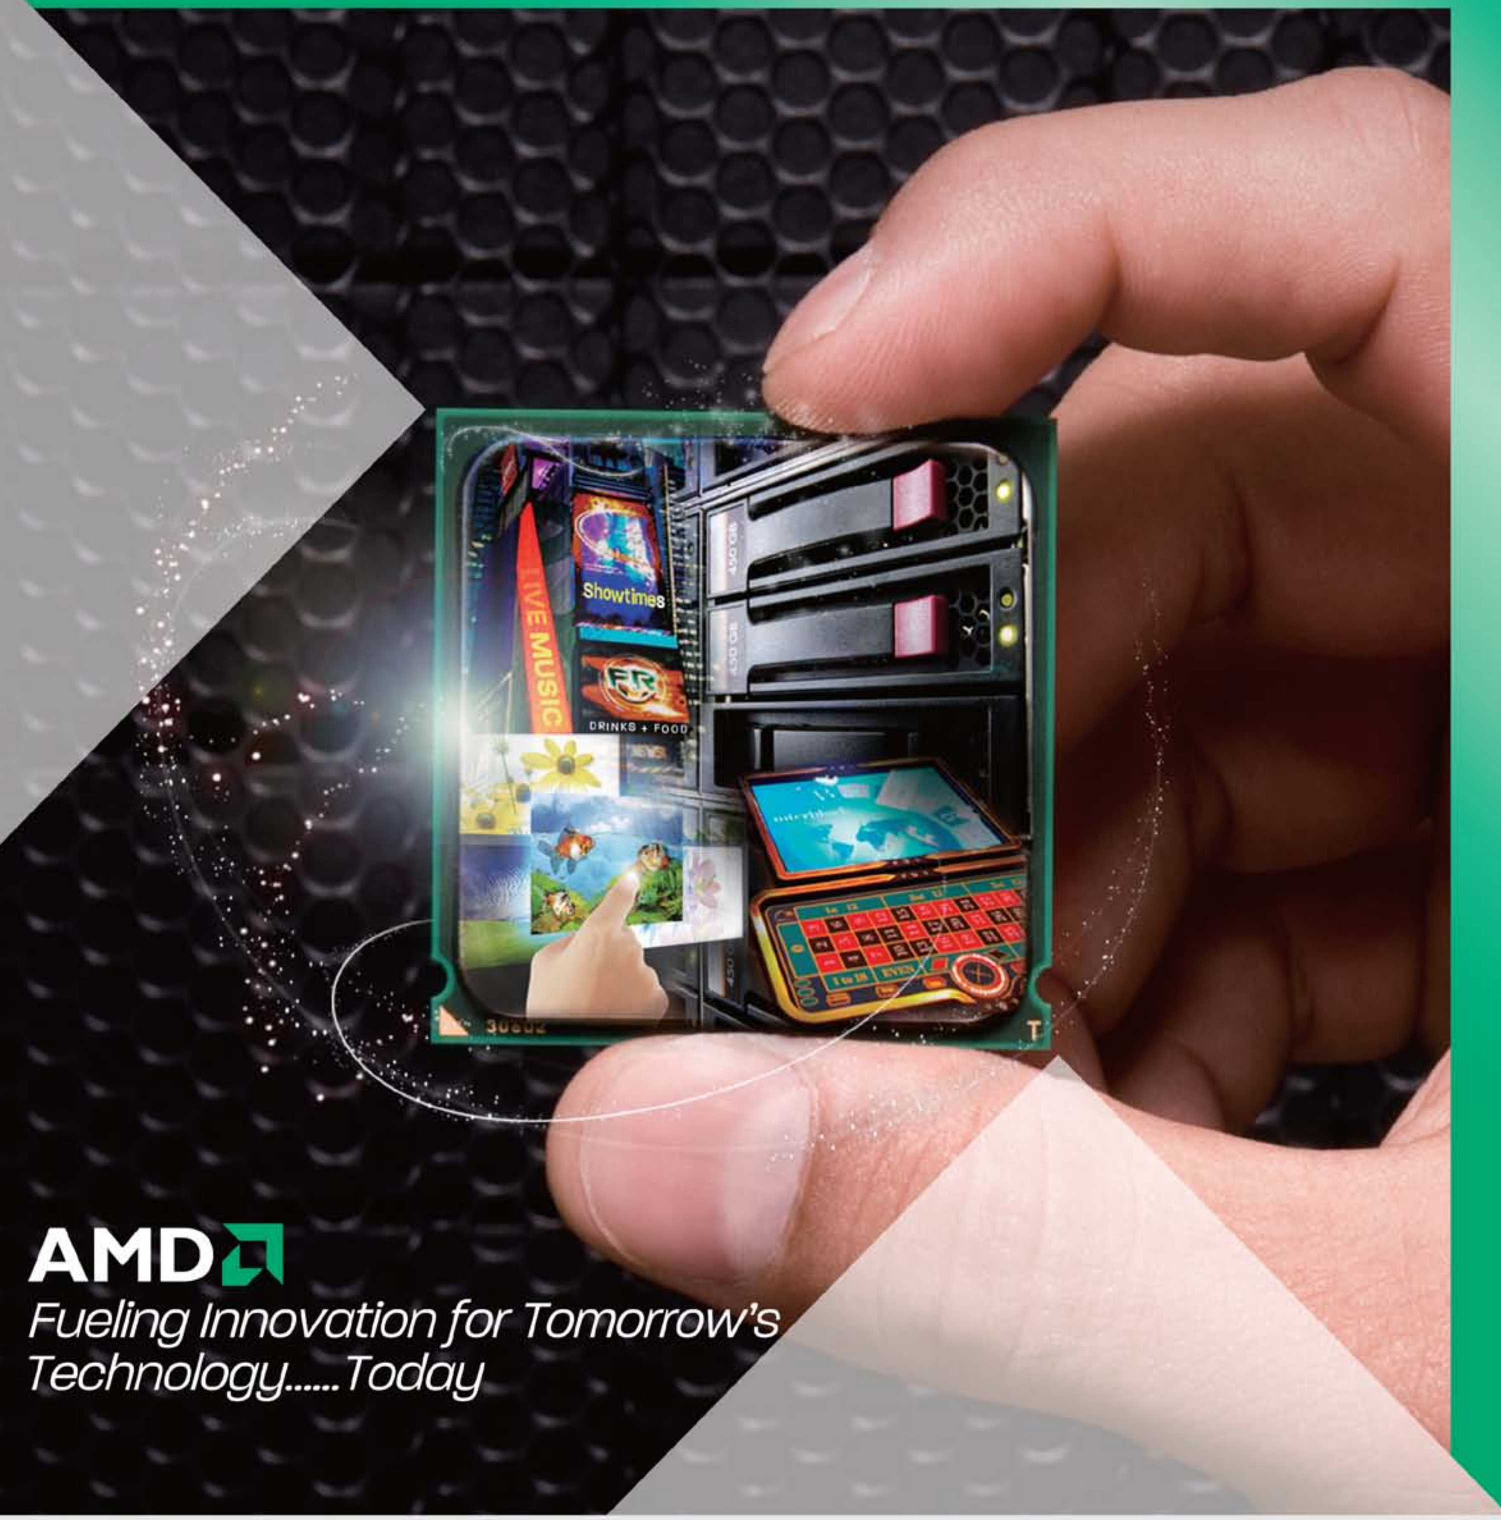

AMD is ushering in a new era of embedded computing. The AMD Embedded G-Series processor is the world's first integrated circuit to combine a low-power CPU and discrete-level GPU into a single embedded Accelerated Processing Unit (APU).

AMD is also proud to offer extended availability of the AMD Geode™ LX processor family until 2015.

AMD Geode

AMD Embedded Solutions

Learn more about new levels of performance in a compact BGA package at: www.amd.com/embedded

© 2011 Advanced Micro Devices, Inc. All rights reserved. AMD, the AMD Arrow logo, ATI, the ATI logo and combinations thereof are trademarks of Advanced Micro Devices, Inc. Other names are for informational purposes only and may be trademarks of their respective owners. Features, performance and specifications may vary by operating environment and are subject to change without notice. Products may not be exactly as shown. PID# 50599C

SUPERMICRO®
Embedded Building Blocks
Modular Computing Solutions for Embedded Applications
High-Performance
Application-Optimized
X8DA6, C7SIM-Q, C7Q67
Compact Form Factor
Short-Depth
SuperServer® 5015A-EHF-D525
SuperServer® 5017C-LF
Small Form Factor
X7SPE-H(F)-D525, X7SPA-H(F)-D525
X9SCV Series
• Supports Intel® Xeon®, Core® i7/i5/i3 and Atom™ Processors
• Energy Efficient, Low Power and Fanless Solutions
• High Performance, Scalable Systems
• High-Density, Compact Form Factors
• Open Standard and Multi-OS Support
• Ruggedized and Industrial Grade Chassis
• 7-Year Life Cycle Support
Versatile, Whisper Quiet and
Highly Configurable
SuperServer® 5037C-i/T
Industrial PC
Short-Depth
SuperServer® 6046T-TUF
ISC
2012 International Supercomputing Conference
JUNE 17-21st
Hamburg, Germany
Congress Center Hamburg
Come visit our booth @ #520
intel inside™
Xeon®
94 93%
21 dB
Whisper-Quiet
intel
Embedded
Alliance
www.supermicro.com/Embedded
© Super Micro Computer, Inc. Specifications subject to change without notice.
Intel®, the Intel® logo, Xeon®, and Xeon® Inside, are trademarks or registered trademarks of Intel Corporation in the US and other countries.
All other brands and names are the property of their respective owners.
SMCI-20120221-1

Embedded COMPUTING DESIGN

www.embedded-computing.com

May 2012 | Volume 10 • Number 3

On the cover

Embedded image processing has advanced to enable realistic, visually stunning media experiences for end users. Our May issue highlights digital imaging and launches new digital watermarks that unlock videos and other content viewable on your smartphone. Download the Digimarc Discover app, scan the phone icon in the lower right corner of the cover, and watch how it works.

Download the Digimarc app and take a picture of this image to get more details

Look for this icon throughout this issue of *Embedded Computing Design* for watermarked images, pull out your smartphone, and start scanning to view the latest videos and other multimedia focused on embedded computing.

2012 OpenSystems Media ®
© 2012 Embedded Computing Design

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

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

Embedded Computing Design Find Us On facebook

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

LOOKING FOR EXTREME SOLUTIONS?

X Marks the Spot.

At Extreme Engineering Solutions, you'll find products that are designed from the ground up to handle any environment. From boards to integrated systems, our embedded solutions are rugged and reliable–ensuring your application is a success, no matter how extreme the conditions.

Embedded solutions that always hit the mark. That's the Extreme way.

X-ES

Extreme Engineering Solutions

608.833.1155 www.xes-inc.com

Advertiser Information

Reader Resources

Subscribe to magazines and E-letters:
www.opensystemsmedia.com/subscriptions

E-casts:
ecast.opensystemsmedia.com

TechChannels:
tech.opensystemsmedia.com

White Papers:
whitepapers.opensystemsmedia.com

Videos:
video.opensystemsmedia.com

Get news via RSS:
www.embedded-computing.com/news/rss.xml

Advertiser Resources

Submit information:
submit.opensystemsmedia.com

Media kits:
www.opensystemsmedia.com/publications

Join the Discussion

Facebook: www.facebook.com/Embedded.Computing.Design

Twitter: @embedded_mag

LinkedIn: www.linkedin.com/groups?gid=1802681

Embedded COMPUTING DESIGN

OpenSystems media

ECD Editorial/Production Staff

Warren Webb, Editorial Director
wwebb@opensystemsmedia.com

Jennifer Hesse, Managing Editor
jhesse@opensystemsmedia.com

Brandon Lewis, Associate Editor
blewis@opensystemsmedia.com

Steph Sweet, Creative Director
ssweet@opensystemsmedia.com

Sales Group

Dennis Doyle
Senior Account Manager
ddoyle@opensystemsmedia.com

Tom Varcie
Senior Account Manager
tvarcie@opensystemsmedia.com

Rebecca Barker
Strategic Account Manager
rbarker@opensystemsmedia.com

Eric Henry
Strategic Account Manager
ehenry@opensystemsmedia.com

Ann Jesse
Strategic Account Manager
ajesse@opensystemsmedia.com

Christine Long
Director of Online Development
clong@opensystemsmedia.com

International Sales
Elvi Lee
Account Manager – Asia
elvi@aceforum.com.tw

Regional Sales Managers
Barbara Quinlan, Midwest/Southwest
bquinlan@opensystemsmedia.com

Denis Seger, Southern California
dseger@opensystemsmedia.com

Sydele Starr, Northern California
sstarr@opensystemsmedia.com

Ron Taylor, East Coast/Mid Atlantic
rtaylor@opensystemsmedia.com

Reprints and PDFs

republish@opensystemsmedia.com

OpenSystems Media Editorial/Production Staff

Mike Demler, Editorial Director
DSP-FPGA.com
mdemler@opensystemsmedia.com

Joe Pavlat, Editorial Director
CompactPCI, AdvancedTCA, & MicroTCA Systems
jpavlat@opensystemsmedia.com

Jerry Gipper, Editorial Director
VME and Critical Systems
jgipper@opensystemsmedia.com

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

Warren Webb, Editorial Director
Embedded Computing Design
wwebb@opensystemsmedia.com

Jennifer Hesse, Managing Editor
Embedded Computing Design
jhesse@opensystemsmedia.com

Sharon Hess, Managing Editor
VME and Critical Systems
sharon_hess@opensystemsmedia.com

Monique DeVoe, Assistant Managing Editor
PC/104 and Small Form Factors
DSP-FPGA.com
mdevoe@opensystemsmedia.com

Brandon Lewis, Associate Editor
CompactPCI, AdvancedTCA, & MicroTCA Systems
blewis@opensystemsmedia.com

Curt Schwaderer, Technology Editor

Steph Sweet, Creative Director

David Diomede, Art Director

Joann Toth, Senior Designer

Konrad Witte, Senior Web Developer

Matt Jones, Web Developer

Editorial/Business Office

Patrick Hopper, Publisher
Tel: 586-415-6500
phopper@opensystemsmedia.com

Subscriptions Updates
Karen Layman, Business Manager
www.opensystemsmedia.com/subscriptions
Tel: 586-415-6500 ■ Fax: 586-415-4882
30233 Jefferson, St. Clair Shores, MI 48082

Rosemary Kristoff, President
rkristoff@opensystemsmedia.com

Wayne Kristoff, CTO

16626 E. Avenue of the Fountains, Ste. 201
Fountain Hills, AZ 85268
Tel: 480-967-5581 ■ Fax: 480-837-6466

By Warren Webb

@warrenwebb wwebb@opensystemsmedia.com

Designing remote health systems reaps rewards

As the world's population continues to age, remote health care or telehealth devices are becoming an enormous growth area for the embedded systems industry. Low-cost, ubiquitously connected hardware and its associated services are rapidly growing in popularity, allowing doctors to monitor, diagnose, and treat specific health conditions remotely. Telehealth devices can also be used to provide remote data collection, patient reminders, multimedia content delivery, and wellness coaching by monitoring and adjusting the user's daily fitness and diet.

For most health care device designs, embedded developers must address a number of common requirements, including a graphical interface and small form factor, patient safety, low-power operation, remote management, secure operation, and universal connectivity. In addition, these devices must be built with long-life support plus the ability to interoperate with other equipment in the field. Although hardware and software development for telehealth can be challenging, the rewards are significant, and manufacturers are rapidly seizing opportunities in the marketplace.

To delve into the latest technologies revolutionizing health care, we interviewed telehealth industry experts in this month's Strategies section. Alan Boucher, Director of Software Architecture and Engineering at Intel-GE Care Innovations, a joint venture between Intel Corporation and GE Healthcare, explains how embedded technologies such as wireless connectivity, software, and sensors can be used to enable and enhance products in a telehealth environment. Also on the forefront of telehealth technology, Dr. Joseph Kvedar, founder and director of the Center for Connected Health, emphasizes the need for programs that move care from the hospital or doctor's office into the day-to-day lives of patients by incorporating engagement strategies such as games, social networking, coaching, reminders, incentives, and punishments.

In this month's Software section we investigate the latest techniques developers can use to optimize code for a variety of embedded applications, including telehealth. John Lockhart from SmartBear Software looks at the advantages of standard and automated peer code review – a process by which team members inspect and review source code to eliminate errors earlier in the development process. Presenting another method to optimize software, Dominic Tavassoli and Jonathon Chard of IBM Rational detail an automated Unified Modeling Language (UML)-based code testing technique that can shorten the development life cycle and reduce costs. In considering how embedded designs are moving to multicore processors, Henk Muller of XMOS presents a unique approach to substitute a parallel software paradigm known as *concurrent real-time programming* for FPGA and even ASIC technology for a range of industrial and consumer applications.

> "Designers must evaluate hardware components to ensure the design will be compatible with changing requirements and the latest standards."

Besides analyzing software, designers must also evaluate hardware components with each new project to ensure that the design will be compatible with changing requirements and the latest standards. For example, in many applications, display connections are transitioning away from traditional interfaces such as LVDS and VGA toward newer technologies such as DisplayPort. In this month's Silicon section, Eurotech's Haritha Treadway describes the cabling and multiple display advantages of DisplayPort I/O-capable interface hardware for embedded projects. As these new display technologies utilize high-resolution images for display and analysis, designers must consider the different approaches to embedded image processing. Allen Rush of Nethra Imaging investigates the use of linear and nonlinear filters to analyze and correct noise, dynamic range, color accuracy, optical artifacts, and other details in high-quality imaging applications.

The telehealth trend promises to influence the technology in many of the Silicon, Software, and Strategy areas covered by *Embedded Computing Design*, and our goal is to report the latest updates as they occur. Feel free to give us your ideas for future technical articles and online updates to support your design efforts. We are always looking for contributed technical articles, which can be an excellent tool to gain exposure in the embedded computing industry. For example, the August issue is our annual Resource Guide and will feature articles and Q&As from industry experts to help simplify and shorten embedded projects. If you have an idea for a technical article that would be of interest to our readers, please send me a short abstract.

Comparing linear versus nonlinear filters in image processing

By Allen Rush

Historically, real-time or embedded image processing was limited in terms of complexity due to cost/power constraints of the underlying silicon. However, with today's sub-90 nm geometries, it is possible to consider complex filtering techniques that up until now could only be completed in offline image data manipulation. Examining the differences between linear and nonlinear filters can help designers implement the most effective filtering technology for detecting and manipulating image information.

Filtering in image processing is a mainstay function that is used to accomplish many things, including interpolation, noise reduction, and resampling. The choice of filter is often determined by the nature of the task and the type and behavior of the data. Noise, dynamic range, color accuracy, optical artifacts, and many more details affect the outcome of filter functions in image processing.

The following discussion will explore the differences between two major categories of filtering – linear and nonlinear – as well as highlight image processing approaches that benefit from these filter types and identify situations where one filter might be preferred or required over the other.

Filter theory background

In image processing, 2D filtering techniques are usually considered an extension of 1D signal processing theory. Almost all contemporary image processing involves discrete or sampled signal processing. This is compared to signal processing that was applied to analog or continuous time domain processing that characterized television and video several generations ago. The two are related, and the foundation for discrete signal processing is derived from continuous time signal processing theory.

Linear filters

To review and compare the two types of filtering, the first step is to briefly describe the attributes that comprise linear filtering.

Several principles define a linear system. The first two are the basic definitions of linearity. If a system is defined to have an input as $x[n] = ax[n1] + bx[n2]$, then the linear system response is $y[n] = ay[n1] + by[n2]$. This is known as the superposition property, and is fundamental to linear system design. The second property is shift invariance. If $y[n]$ is the response to a linear, shift-invariant system with input $x[n]$, then $y[n-n0]$ is the response to the system with input $x[n-n0]$.

In addition, two extra conditions are imposed, causal and stable. The causal condition is needed when considering systems in which future values are not known (for example, in video streaming). It is possible to consider a system that is not causal when looking at captured

images with samples before and after the target location (for example, in a buffered version of an image frame). Stability is imposed to keep a filter's output from exceeding a finite limit, given an input that also does not exceed a finite limit. This is called the Bounded-Input Bounded-Output (BIBO) condition.

For most cases, a system is evaluated in the spatial frequency domain. To accomplish this, the convolution theorem is used, providing the necessary tools to evaluate frequency domain information.

If x[n] and h[n] are two sequences, their convolution is defined as shown in Equation 1.

$$y[n] = \sum_{k=-\infty}^{\infty} x[k]h[n-k]$$

Equation 1

The corresponding frequency response is shown in Equation 2.

$$Y(e^{-i\omega}) = X(e^{-i\omega})H(e^{-i\omega})$$

Equation 2

In Equation 2, $e^{-i\omega}$ is the frequency domain representation and ω is the frequency variable from $-\pi$ to π. This fundamental relationship describes the response of a filter in terms of frequency – low pass, high pass, band pass, and so on. Depending on the nature of the filter kernel h[n], a wide variety of responses can be realized for any image data set.

A typical low-pass filter with 25 taps (h[0..24]) is shown in Figure 1. The idea of a low-pass filter is to preserve low-frequency information and reduce or eliminate high-frequency information in an image. It blurs edges but keeps smooth areas of an image intact. In a similar manner, high-pass filters preserve edges and other high-frequency information but filter low-frequency regions of an image.

Nonlinear filters

Nonlinear filters have quite different behavior compared to linear filters. For nonlinear filters, the filter output or response of the filter does not obey the principles outlined earlier, particularly scaling and shift invariance. Moreover, a nonlinear filter can produce results that vary in a non-intuitive manner.

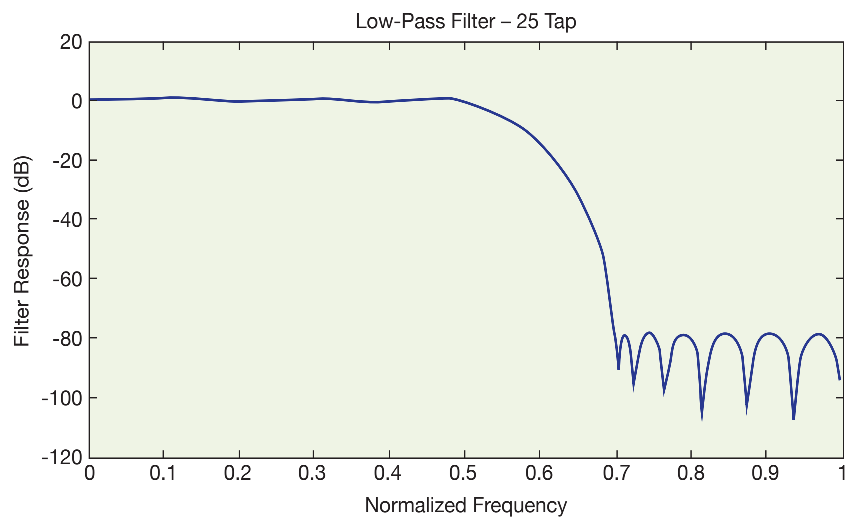

Figure 1 | A typical low-pass filter maintains low-frequency elements and reduces or removes high-frequency elements in an image.

TS-WIFIBOX-2

A Complete Solution for 802.11g WiFi Applications

qty 1 **$185**

Powered by a
250 MHz ARM9 CPU

- Low power (3.2 watts), fanless
- Power via 5-12 VDC, USB, PoE (opt.)
- 64MB DDR-RAM
- 256MB ultra-reliable XNAND drive
- Micro-SD Card slot
- RS-232, RS-485, CAN, RTC
- Header with SPI and 11 DIO
- 480Mbit/s USB, Ethernet, PoE option
- Boots Linux 2.6.24 in < 3 seconds
- Un-brickable, boots from SD or flash
- Customizable FPGA - 5K LUT
- Optional DIN mountable enclosure

Ideal for gateway or firewall, protocol converter, web server, WiFi audio, and unattended remote applications

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

Technologic SYSTEMS

We use our stuff.

visit our TS-7800 powered website at

www.embeddedARM.com

(480) 837-5200

The simplest nonlinear filter to consider is the median or rank-order filter. In the median filter, filter output depends on the ordering of input values, usually ranked from smallest to largest or vice versa. A filter support range with an odd number of values is used, making it easy to select the output.

For example, suppose a filter was based on five values. In the region of interest, x0..x4, the values are ordered from smallest to largest. The value at position 2 is selected as the output. Consider the case at low frequency; all the values are the same or close to it. In this case, the value selected will be the original value ± some small error. In the case of high frequency, such as an edge, the values on one side of the edge will be low and the values on the other side will be high. When the ordering is done, the low values will still be in the low position and the high values will still be in the high position. A selection of the middle value will either be on the low side or the high side, but not in the middle, as would be the case using a linear low-pass filter. The median filter is sometimes called an edge-preserving filter due to this property. It is useful in removing outliers such as impulse noise.

Selecting the right filter

Both filter types have their place in image processing functions. In a typical pipeline for real-time image processing, it is not uncommon to have dozens of both types included to form, shape, detect, and manipulate image information. Moreover, each of these filter types can be parameterized to work one way under certain circumstances and another way under a different set of circumstances using adaptive filter rule generation.

Filtering image data is a standard process used in almost all image processing systems. The goals vary from noise removal to feature abstraction. Linear and nonlinear filters are the two most utilized forms of filter construction. Knowing which type of filter to select depends on the goals and nature of the image data. In cases where the input data contains a large amount of noise but the magnitude is low, a linear low-pass filter may suffice. Conversely, if an image contains a low amount of noise but with relatively high magnitude, then a median filter may be more appropriate. In either case, the filter process changes the overall frequency content of the image. ECD

Allen Rush
is CTO of
Nethra Imaging.

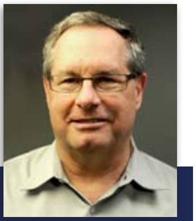

Nethra Imaging
info@nethra.us.com
www.nethra.us.com

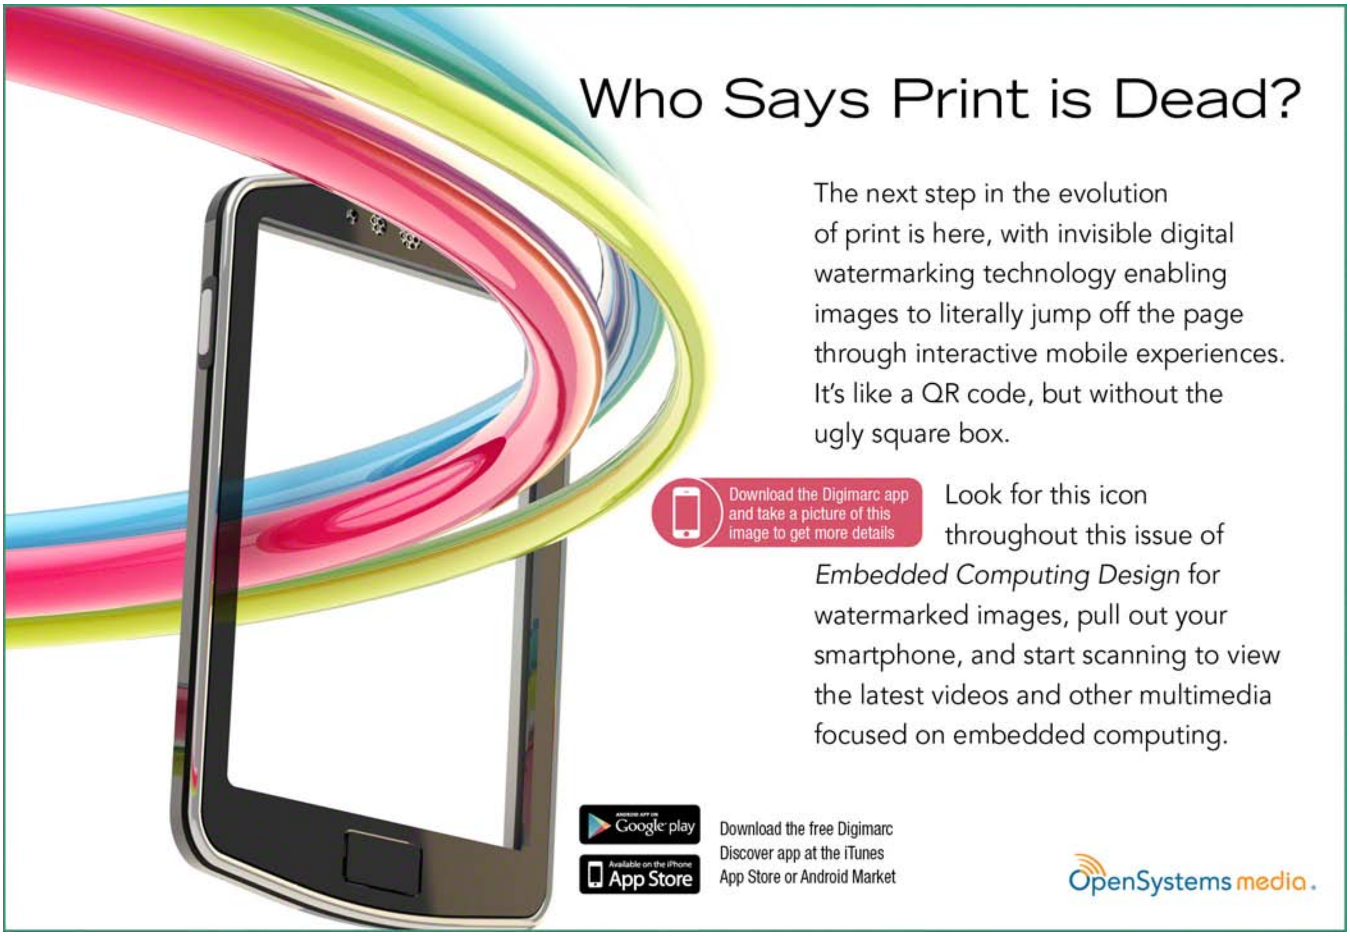

Something old, something new: Marrying and migrating multimedia interfaces

By Haritha Treadway

PCs and mobile computers are rapidly migrating away from traditional LVDS or VGA display interfaces toward newer technologies like DisplayPort. Top players such as Intel and AMD are expected to stop supporting LVDS in the next 3-5 years, making migration to DisplayPort a necessity. Designers must determine how this new display interface fits into new and existing designs while preserving key requirements for embedded systems.

DisplayPort offers several advantages for the embedded market, starting with a clear cost benefit. DisplayPort is a royalty-free, open-standard digital display interface managed by the Video Electronics Standards Association (VESA). For embedded system designs that often involve long and complex Bills Of Materials (BOMs), DisplayPort offers huge cost savings over other standards such as HDMI. According to www.hdmi.org, the annual fee for adopters is $10,000, and royalty fees include 15 cents per end user-licensed product sold. While these fees make sense for commercial volumes, OEMs are often at a loss for how to offer HDMI for smaller-scale embedded designs while still making profits against the annual license fees and royalties.

At the core of embedded system designs are the trade-offs made between power and performance. Embedded applications require low power consumption coupled with optimized performance from single or multicore CPUs. For example, in fields such as biometrics or portable medical devices, OEMs need multicore processing that can support heavy high-resolution visualization and data processing. At the same time, they need a device that can run on batteries for several hours using an embedded module with power consumption below 10 W. DisplayPort was developed to directly address the need for low power, allowing CPUs to drastically reduce total power consumption through a much simpler and efficient layout.

With DisplayPort, a single cable manages several displays, which is an essential requirement in industries such as digital signage where multiple sleek flat-panel displays need to be interlinked and managed through a single controller. Minimizing cabling translates to more flexibility in configurations as well as less maintenance. The simple design also means reduced electromagnetic emissions and improved overall signal integrity, especially over much longer cable lengths than are allowed with LVDS. Whereas LVDS is ideal for applications with shorter cable lengths of up to about 10 meters, DisplayPort can support much longer lengths.

Simplifying multiple display support is certainly an advantage in the embedded market. For example, imagine an embedded device running the displays on the seat backs in an airplane. With DisplayPort, multiple screens can be daisy-chained together and run independent content simultaneously, all without cumbersome additional cabling or central processing complexity. Figure 1 shows examples of different possible display setups with a single DisplayPort connection.

HD resolution is critical for all areas of embedded systems, as today's end users demand more visualization and high-quality human-to-machine interactions and end user experiences. DislayPort can preserve HD graphical resolutions at longer lengths. For example, WQXGA resolutions (2560 x 1600 pixels) are supported at a minimum of 2 meters. Full HD is easily supported at 15 meters. DisplayPort clearly has advantages even over DVI, as demand for higher and higher resolutions in digital displays increases and at some point DVI hits its limitations. By going beyond full HD and enabling fast refresh rates, DisplayPort will help OEMs keep up with future technologies.

Gradual transition to DisplayPort

For now, OEMs can continue using LVDS, VGA, and DVI while keeping DisplayPort on their future technology roadmap. DisplayPort offers backward compatibility with these legacy technologies, which have not been updated or active for several years. In addition, VESA has stated that DisplayPort is not looking to replace HDMI, but rather serves as an alternative for embedded developers.

The ideal embedded system today offers support for both options, allowing developers to leverage new and proven technologies to help speed development. Processors like the Intel Atom N2000 series and embedded boards like Eurotech's Catalyst CV play a role in the migration story by supporting not only traditional interfaces such as LVDS and VGA, but also DisplayPort and Embedded DisplayPort. Supporting both display options gives developers an easy migration path that uses the same board package to extend product life cycles and doesn't necessarily involve reengineering a design down the line. DisplayPort shows potential to be a big boon for the embedded market, reducing the number of standards and interfaces to just one that is actively updated and designed for longevity.

"By going beyond full HD and enabling fast refresh rates, DisplayPort will help OEMs keep up with future technologies."

Eurotech's Catalyst platform, based on the Intel Atom roadmap, is designed for technology migration and flexibility. The Catalyst CV (Figure 2) is available in a 67 mm x 100 mm form factor, and is optimized for the processor and

Figure 1 | With DisplayPort, designers can easily configure multiple display setups with minimal cabling. (Image courtesy of DisplayPort.org.)

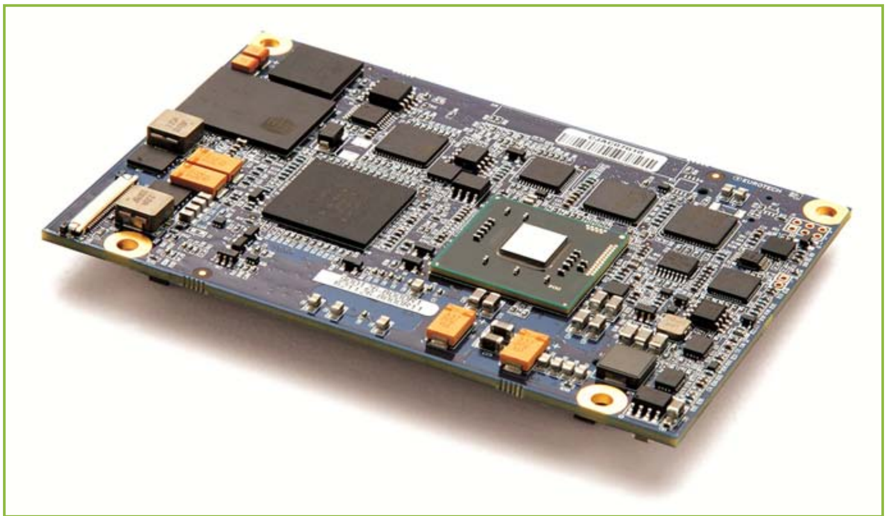

Figure 2 | Eurotech's Catalyst CV supports DisplayPort and all the other multimedia features provided by the Intel Atom N2000 processor.

performance rather than for a standardized form factor. This means that the module brings out all of the capabilities of the Intel Atom N2x00, including the multimedia features such as LVDS, VGA, DVI, HDMI, DisplayPort, and EDP.

The Catalyst CV also offers Eurotech's Everyware Software Framework and Everyware Device Cloud, software programs that help simplify development and manage data through easy-to-use cloud services.

Flexibility to support new and legacy interfaces

The embedded market is an evolving industry that poses the unique challenge of supporting both new and legacy technologies to maintain the balance between performance and product longevity. While many new designs will incorporate DisplayPort, OEMs still need to support LVDS- and VGA-based products. Embedded systems that offer the flexibility of forward and backward compatibility will be critical in making this transition a smooth one. **ECD**

Haritha Treadway is product manager at Eurotech.

Eurotech
haritha.treadway@eurotech.com
www.eurotech.com

Follow:

TechChannels **Application Lifecycle Management**

Anywhere embedded software development takes place, an ALM software platform can help designers and program managers coordinate all of an application's or project's artifacts and activities.

Check out news, white papers, videos, and more on ALM:
tech.opensystemsmedia.com/application-lifecycle-management

OpenSystems media.com

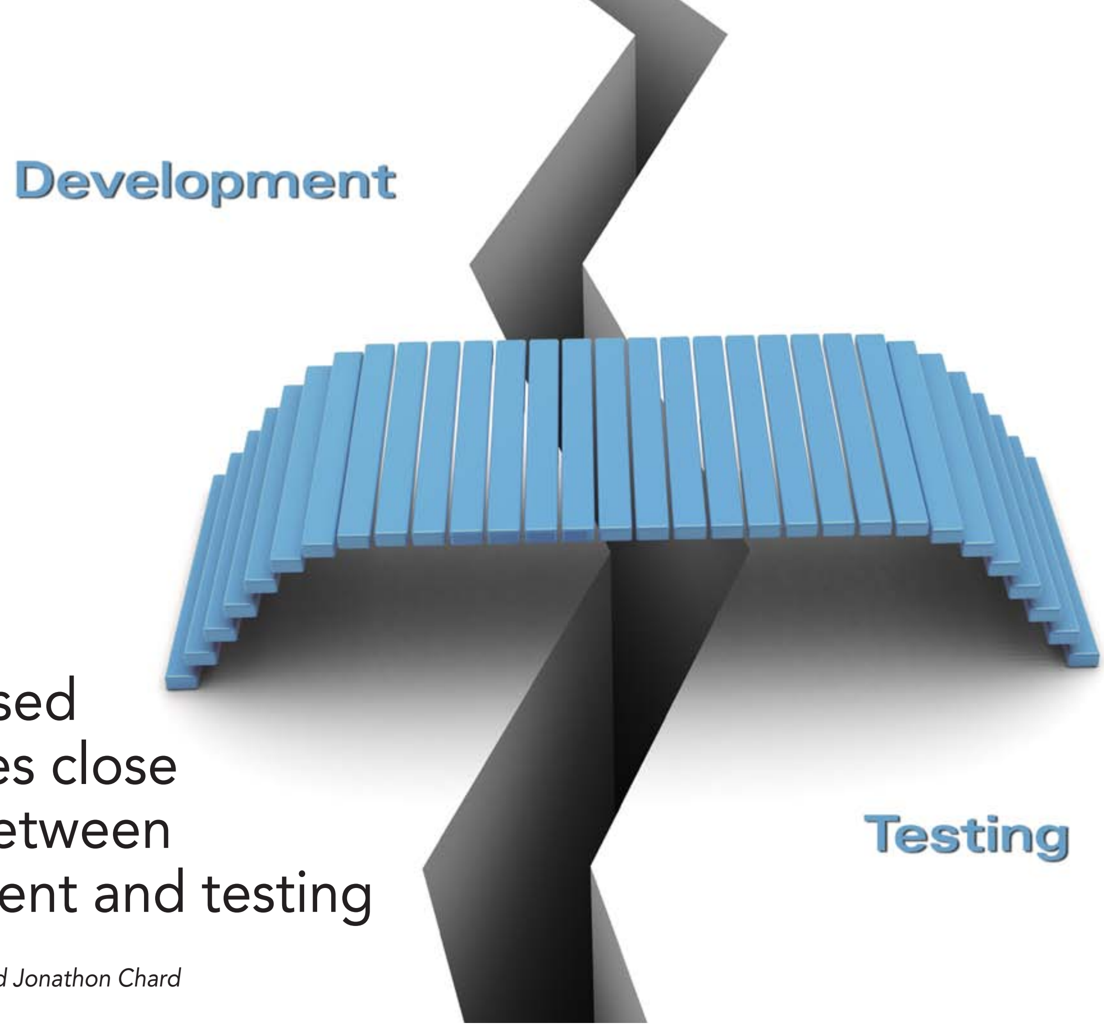

Model-based approaches close the gap between development and testing

By Dominic Tavassoli and Jonathon Chard

While model-based development has helped developers identify defects earlier and cope with increasing design complexity, testing is now the elephant in the room. How can automation, integration, and collaboration around testing deliver the required efficiencies? New approaches are moving the embedded testing challenge from code to models, allowing businesses to gain a competitive edge as a result.

An age of smarter products is ushering in embedded product designs with increased functionality, rising complexity, and compressed delivery windows. These products often need to comply with strict development regulations for use in safety- or mission-critical applications such as aircraft, automobiles, or medical devices.

Model-based development has boosted developer productivity with graphical notations such as Unified Modeling Language (UML) and Systems Modeling Language (SysML) to help manage complexity and uncover design anomalies earlier in the development life cycle. In many cases this has left testing struggling to catch up. It is frequently the biggest time and budget item in projects, and therefore the first to be cut. Despite this, testing costs are increasing; much testing has remained rooted in manual, code-based approaches that do not easily scale to today's demands.

The proven power of model-driven development

Using models, software engineers can more clearly understand and analyze requirements, make architectural trade-offs, define design specifications, validate and verify behavior with simulation, and generate code for direct deployment on target hardware. A key benefit of a model is that consistency and correctness are maintained across the design as it changes. With UML, each diagram can capture different views of the model at different levels of abstraction while remaining consistent across these views. The semantics of the modeling language enforce team consistency and help automate life-cycle tasks, including code generation.

Repeatable processes are the key to improving collaboration and productivity while reducing the cost of demonstrating regulatory compliance. A 2011 embedded development study by Jerry Krasner of Embedded Market Forecasters found that model-driven development reduced development time typically lost to delays by more than 40 percent, delivering typical project savings of $250,000.

Enhancing quality at every step of the life cycle

The next natural step to maximize productivity and agility is to improve the testing process. The "usual suspects" addressing this are the late detection (and resolution) of defects, as well as the communication and consistency issues between teams as changes occur and defects are detected.

Early identification of defects is critical to reduce development costs and meet time-to-market objectives. Many defects are introduced in the design during the early stages of development, but repair costs increase exponentially as defect resolution is delayed. Not only can late defect resolution harm project financial performance, a high intensity of late rework can also derail delivery schedules and delivered quality. Defect resolution must be closely linked to defect discovery, as it is typically much easier to fix a defect that has just been discovered than to attempt a repair after further changes have obscured the original cause.

Successful development projects must deliver products that address customer needs to the required level of quality. Linking unit, integration, validation, and verification tests to requirements is fundamental, and automating traceability is key to efficiently execute projects as changes occur and defects are detected.

Model-based testing and the UML Testing Profile

Model-based testing uses modeling to construct and execute the necessary artifacts to perform software testing. The UML Testing Profile[1] extends the applicability of UML to include model-based testing. Using this profile, test architectures can be automatically created for a system from the definition of its interfaces. Test cases consistent with the latest requirements can be defined graphically using sequence diagrams, state charts, or activity diagrams, providing a common modeling language to define test inputs and expected results (see Figure 1). This approach tightens the linkage between requirements, design elements, and tests, thus facilitating better traceability that can benefit both agility and impact analysis. Test cases can be executed on the developer's desktop and on the target, improving testing productivity.

Download the Digimarc app and take a picture of this image to get more details

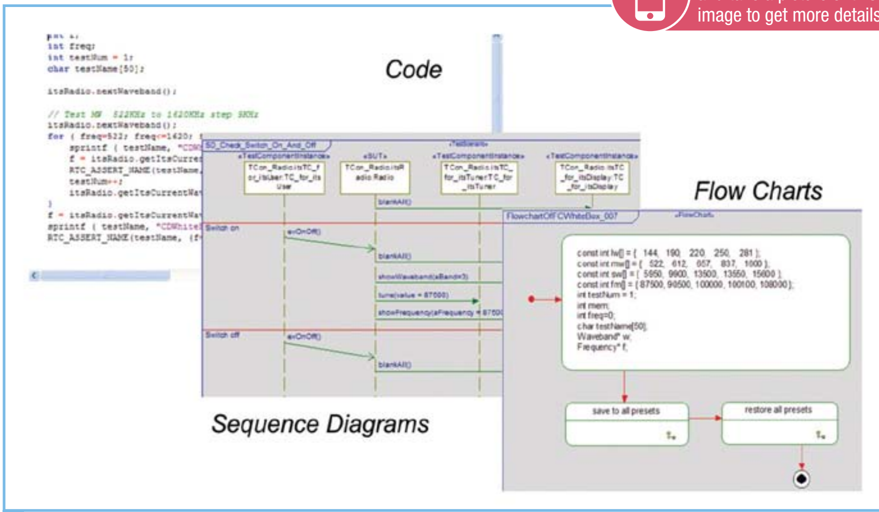

Figure 1 I Model-based testing specifies test cases using diagrams.

Strengthening the bond between development and quality management

This move to model-based developer testing doesn't directly benefit the Quality Assurance (QA) team, as QA engineers typically don't want, nor should they be forced to become UML specialists. What is needed is a way to allow QA engineers to make use of model-based testing resources without the need to author and own them. This can be achieved through model-driven testing tools that link the tests to the QA test management environment. The tools allow tests to be executed in place from the test management tool and the results to be passed automatically to the test management repository. Tests are then managed in a single location, avoiding the inconsistencies and inefficiencies of duplication, and are available for execution at any point in the development process.

Bringing test result data back into the test management environment can automate the defect resolution process, further optimizing the development and testing process. The QA test management environment should also support linkage to requirements management tooling to ensure that tests can be linked back to specific requirements. This will facilitate test coverage automation, enabling test sets to be automatically defined against requirements or changes.

Taking a project to the next level

By implementing a model-based testing tool connected to test management, embedded development teams can take the following tangible steps to make their projects more successful while freeing up valuable time for innovation:

- **Consistency:** A central test repository across the development organization will improve efficiency and lead to higher quality. Rather than individual test engineers creating multiple versions of tests through error-prone, manual replication processes, tests are written once and reused as required throughout the development

We Wrote the Book on Software Radio Again!

Downloaded thousands of times, it's now better than ever

- Expanded
- Fully updated
- Totally new

If you liked it before, you'll like the 9th Edition update even more!

Get your free download from: www.pentek.com/go/ecdsdr9hb

PENTEK
Setting the Standard for Digital Signal Processing

process. A single source of truth for tests can also aid collaboration and ultimately improve delivered quality, as tests are more available throughout the development life cycle, encouraging more frequent testing.

- **Communication:** Model-based test execution within the quality management environment means that QA engineers can run tests and act upon results without having to be modeling specialists. They can navigate from a failed test to the related requirements and to the source of the problem in the design mode or associated code. This can be particularly useful in distributed and offshore development environments where it can help engender collaboration and build empathy between otherwise isolated developers, testers, and quality professionals.
- **Automation:** Improved automation of test creation, execution, and results management will significantly reduce the cost and time of testing. This allows more testing to take place, reducing the risk of regression issues in complex projects. In the same vein, automated defect tracking/resolution connects model-based testing to quality management with a backbone of traceability from requirements through to code.

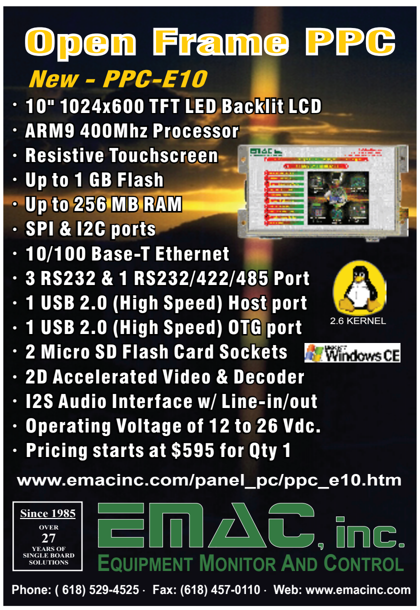

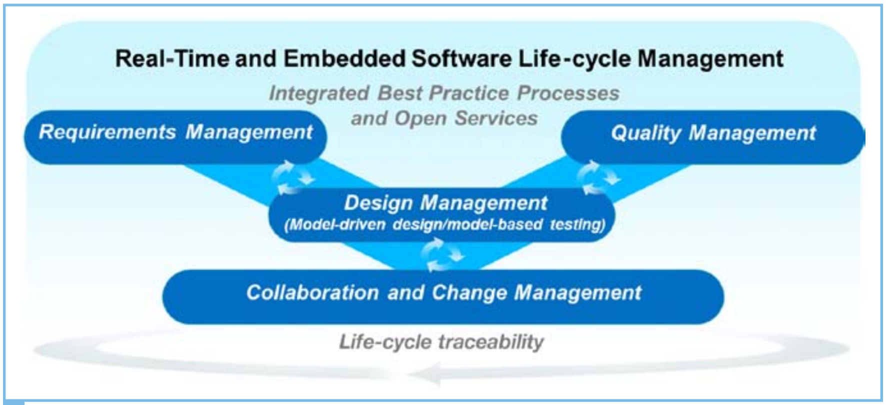

Figure 2 | Linking model-based testing to QA and the software design life cycle enables a high degree of test automation.

This ensures developers have timely and quantitative information to fix defects, and that the effects of defects on delivered functionality are understood. As teams strive to become more agile, it is imperative to prioritize defect resolution over new functionality to avoid accumulating technical debt.

- **Agility:** Model-based testing conducted using the same modeling notation and tools as design activities facilitates test-driven development. Traceability from requirements to testing will be the key to quantifying results, helping answer the critical question in any project: "Are we ready to ship?"

A call to action

Model-based testing can be considered a first step to bring testing efficiency on par with model-driven development. While its immediate effects are in automating the developer testing environment (with correct linkage to the QA environment), it can enable much wider benefits (see Figure 2). By providing tool support for test execution, test version management, and results management within the QA environment, coupled with life-cycle traceability, a greater degree of test automation can be implemented, eliminating a key bottleneck in the development life cycle.

To realize these benefits, embedded development teams should:

- Extend their modeling capabilities from design to testing with tools that support the UML Testing Profile, such as IBM Rational Rhapsody.
- Ensure their modeling and test management environments are closely linked through tools, such as the IBM Rational solution for real-time and embedded software development.
- Adopt, enforce, and continually improve repeatable processes supported through development and testing tools.

In this way, teams can achieve improved collaboration, productivity, and agility, helping them deliver higher-quality products more rapidly at reduced cost. **ECD**

References

[1] UML Testing Profile, OMG, June 2011: www.omg.org/spec/UTP/1.1/Beta2/PDF/

Dominic Tavassoli *is director of industry and systems marketing at IBM Rational.*

Jonathon Chard *is systems market manager at IBM Rational.*

IBM Rational
jon.chard@uk.ibm.com
www.ibm.com/software/rational

Follow: BLOG

Peer review: The best technique embedded developers aren't using

By John Lockhart

Peer review can help embedded hardware designers and software engineers find and fix problems earlier in the development cycle, saving money and time. As opposed to a manual approach, tool-assisted peer review provides reportable data that enables development teams to benchmark and improve their processes.

It's a common scenario in the software development world: A company pins a lot of hope on a new device, hoping to ride the wave of progress in the industry by getting it to market ahead of the competition. The hardware and software teams work long and hard to make it happen. Once the product reaches consumers, though, a problem quickly arises: it doesn't integrate well with a market-leading peripheral.

As engineers think back to where things might have broken down, they realize it couldn't have been during the spec review phase because they routinely review requirements documents, specifications, and test cases. Therefore, the problem most likely occurred during the coding process (see Figure 1). And it probably happened because the software and hardware teams were out of sync and unaware of changes during the design and development phase because they didn't have a simple way to stay connected.

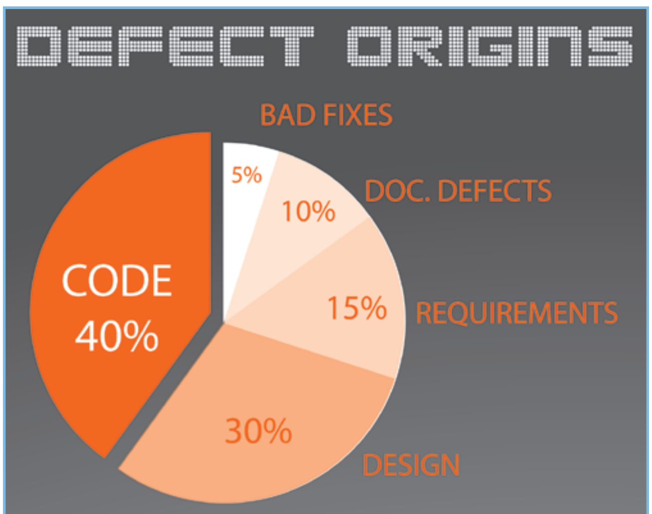

Figure 1 | Defects often arise during the coding process due to a disconnect between the software and hardware teams. (Source: The State of Software Quality in 2011 – Caper's Jones 12/6/2011)

Hardware designers and software engineers who have been working with embedded code for any length of time have probably encountered a situation like this. If not, it's certainly an embedded developer's worst fear: A problem with a product that has already reached the marketplace or is about to be released. It could be an integration issue, a connectivity problem, or a security glitch. Whatever the problem, the outcome is the same; time and money are lost while the code is fixed and the company's reputation suffers.

When so many embedded teams have experienced these problems, why are they still occurring? The reasons are simple: lack of peer and code review and lack of collaboration between software and hardware teams.

Although reviewing artifacts early in the design phase is a common practice and known to be the best way to detect problems early, automating this process between hardware and software engineers has largely been impossible. Although few would disagree that code reviews are always a good idea, other pressures often take precedence. Of course, there are exceptions, mainly in regulated and safety-critical products. However, according to embedded code expert Jack Ganssle, about 98 percent of embedded developers aren't doing peer review. That's a stunning statistic given what's at stake.

While ignoring code review might have worked in a less complicated, less connected world, using that method today leaves companies open to security, interoperability, and connectivity issues. That can lead to expensive recalls requiring time-consuming fixes after products have reached the market.

Peer review – a process by which team members inspect design documents, artifacts, and source code – helps both software and firmware developers, as well as hardware designers, find more bugs or related design errors earlier in the design and prototyping stages, improving product quality and minimizing costly rework later in the development process. It's a process where software and firmware developers as well as the hardware design teams can share and review both technical documents and programming code in a timely manner, keeping the teams in sync when issues are found and changes are made.

Put simply, studies show that peer reviews work. Research by Philip Koopman, an Associate Professor from Carnegie Mellon University, found that peer reviews are the most cost-effective way to find bugs, and 40 to 60 percent of defects are found by such reviews. Koopman also found that reviews cost only about 5 to 10 percent of the project cost.

Peer review helps ensure:

- Higher-quality products in the short term as defects are identified
- Higher-quality products in the long term as technical debt is better managed
- Compliance with applicable regulations
- Interoperability with all potential products, peripherals, and software it may be used with
- Crisper, better documented, and better organized code
- Transfer of knowledge across the entire development team

The peer review process also saves money, as evidenced by a study comparing bug defects before and after code review authored by SmartBear Software in conjunction with Cisco Systems. In both cases, the product had 463 bugs remaining after development. Without code review, getting the bug count down to 194 cost $368,000. The code review process not only fixed more bugs, getting the bug count down to 32, but it did so for $152,000.

Tool-assisted peer review makes sense

On the surface, manual methods of peer review seem like a good way to introduce peer review without spending the money for an automated tool. But these manual methods, while certainly better than no peer review at all, are time-consuming. Furthermore, it can be difficult to collaborate with team members in different locations or time zones. What's more, manual methods like *ad hoc* meetings, water cooler discussions, sending code snippets or PDFs via e-mail, and cutting and pasting code into Word documents tend to be disorganized, and critical points can be lost in the process.

Another consideration is that manual peer review does not produce reportable data. One key to creating support for process improvement is quantifiable results. This is one reason teams wonder if the hours spent in meetings are really worth it.

Automated peer review (also called tool-assisted peer review) solves these problems. With tool-assisted peer review, hardware designers and software engineers can participate in reviews at any time, not on a set schedule. That saves time and increases engineer productivity. Developers also can share and collaborate with team members in different locations. And because all materials are in one place, gathering the right files and design documents is never a problem. In addition, having the review materials and results managed in a reportable database, as well as providing accountability within the review process, helps adhere to multiple regulatory compliance mandates.

Tool-assisted peer review is more efficient and effective than manual peer review. It enables software developers and hardware engineers to catch defects, whether stand-alone or based on changes one team must make that affects the other, earlier in the development process at a time when they are easier and faster to fix. The general rule of thumb is that defects detected later in the process take longer and are more complicated and expensive to fix.

Tool-assisted peer review also provides developers with a host of standard and customizable reports on metrics like defect density, inspection rate, defect detection rate, recent and open defects, lines of code added/modified/deleted, and reviews by change list. These reports and metrics can make a significant difference in the software and hardware development process. With the right metrics, the respective teams can benchmark and improve their processes. For example, it could flag reviews considered trivial or stalled, saving the team valuable time.

Some tools, like SmartBear's PeerReview Complete, also allow development teams to use a variety of review formats such as Word and PDF documents, 2D drawings, schematics, VHDL code and images; develop custom workflows; create custom reports and metrics; integrate with Eclipse and Visual Studio; create customizable fields for tracking and reporting key Capability Maturity Model Integration (CMMI) audit metrics; and implement administrative and

CASE STUDIES

In a recent application, a major heavy equipment manufacturer had a policy that code for any product it ships must be reviewed. After conducting manual in-person code reviews – a process that involved gathering team members in a meeting room with a projector to review code – the company decided to move to tool-assisted code review. Because reviews can be accessed anywhere and anytime, developers can participate in shorter and more focused reviews right from their desks. In addition to increased developer efficiency and higher-quality reviews, tool-assisted review solved the problem of how to include members of a geographically distributed team.

In another application, a major hardware manufacturer wanted to give its internationally distributed development team a way to conduct peer review in a repeatable and verifiable manner. The company wanted a code review process that was easy to use and implement, didn't require explicit scheduling, fostered a close integration between source code management and bug tracking systems, and would create a verifiable record of defect resolution and provide process enforcement. By implementing a peer code review tool, the company measurably improved software quality and built a culture of defect prevention.

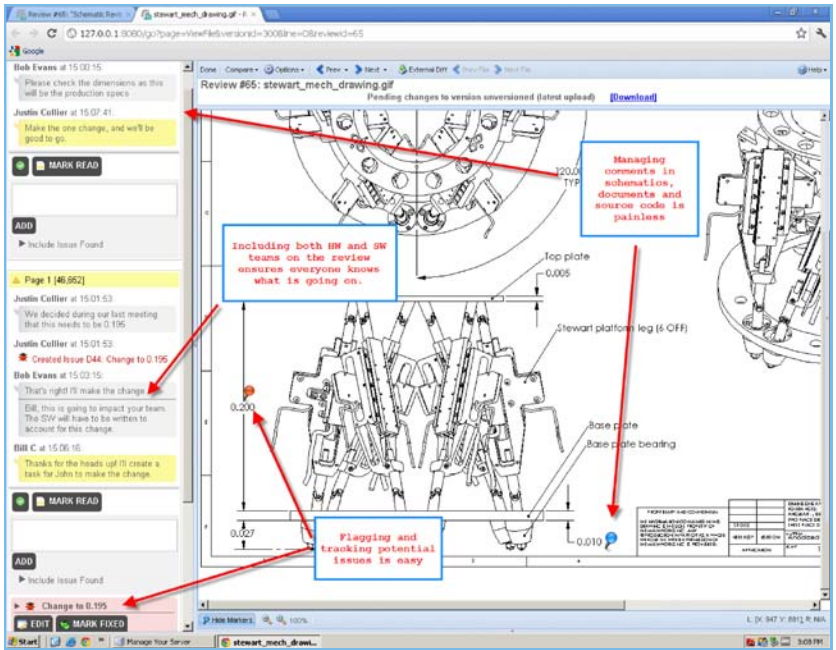

Figure 2 | Using SmartBear PeerReview Complete, an author and reviewers discuss the final aspects of a mechanical drawing before it is used in production.

security controls. It also integrates with a development team's existing issue tracking, development environments, and version control tools. A schematic review is shown in Figure 2.

Tool-assisted peer review for all development artifacts

With these types of capabilities, tool-assisted peer review tools can serve as a comprehensive solution that works with code and all artifacts created during the development process. Requirements documents, hardware and software design documents, schematics, 2D drawings, and test specifications can all be reviewed using the same time-saving tool.

The entire review process comes together in one place, simplifying the existing document review process and extending the review process into code review. The long-established benefits of peer review for design documents can be expanded to the coding process. Code review becomes a normal part of the development cycle, and early detection of defects in code becomes as natural as early detection of defects in design specifications.

A move toward peer review is a positive one for any embedded development team. It eliminates guesswork, improves productivity, saves money, and streamlines workflows. While it generally takes time to implement even a manual review process, tool-assisted peer review provides an immediately impactful peer review process without the headaches of traditional manual approaches. ECD

John Lockhart is the product manager of PeerReview products for SmartBear Software.

Download the Digimarc app and take a picture of this image to get more details

SmartBear Software
john.lockhart@smartbear.com | www.smartbear.com

Follow:

Flexible hardware in software: Enabling customization through concurrent real-time programming

By Henk Muller

Customization and adaptability are important considerations when designing industrial and consumer electronics. Although ASICs and reference designs can be a good starting point for a system, lack of flexibility makes it hard to use ASICs for a complete solution. Using software for customization, even when it involves hardware interfaces, enables unrivalled design flexibility and thus greater product differentiation.

To understand why customization is important, consider the following three examples where a reference design can be used as a starting point, but the designer might want to make changes to the design, even at a late stage:

1. A design implements part of a communications stack (for example, USB audio), but the hardware interface on the audio side can be customized to communicate with a specific type of codec such as I2S master or slave S/PDIF, or it can be customized with an extra control endpoint.
2. An industrial controller implements one of the many standard protocols; for example, a vendor of industrial motors might want to include an Ethernet PHY, but defers the decision whether to run EtherCAT or any other real-time Ethernet stack.
3. An embedded Web server can be used to configure either of the previous examples, provided that the designer can implement the control logic inside the Web server.

The ease with which designs can be customized depends on the technology used to implement the core of the design. Did it use an ASIC, an FPGA, or a real-time processor?

If the design is implemented as an ASIC, then the only feasible way to customize it is to modify parameters built into the ASIC. Considering the first example, a USB audio ASIC will have a method to set the product name of the USB device and might provide a choice between left- and right-aligned I2S. However, it might not offer the options to add an extra control endpoint or to use S/PDIF as an output interface. Indeed, it is not possible to develop an ASIC that supports all plausible interfaces, and it is not economical to develop a family of ASICs that supports well-chosen, small sets of interfaces.

If the design is implemented in an FPGA then, in theory, customization can be as extreme as the user wants it to be. However, a hardware design flow might make customization a not-so-straightforward process. Implementing an additional interface such as S/PDIF is not difficult for hardware designers, but implementing an extra control endpoint requires software to be re-implemented as a piece of hardware. Furthermore, the result must be synthesized, and until the customized design is completely synthesized, placed, and routed it is not clear as to whether the original timing constraints are all still met.

The third option is to implement the design entirely as software. This is often considered difficult to accomplish because of the many conflicting real-time requirements involved. However, a software implementation can be perfectly manageable, provided that the programmer can split the problem into a set of independent real-time tasks that run on a bucket of independent real-time processors.

This gives the designer the flexibility to implement a hardware protocol simply as a software task by dedicating one of the real-time processors to implementing that hardware protocol. For example, one can implement I2S in software by wiggling the clock and data lines in an appropriate manner. Also, USB can be implemented by reading and writing data to the USB PHY. If these two activities are executed as two independent tasks, then all real-time deadlines can be met individually.

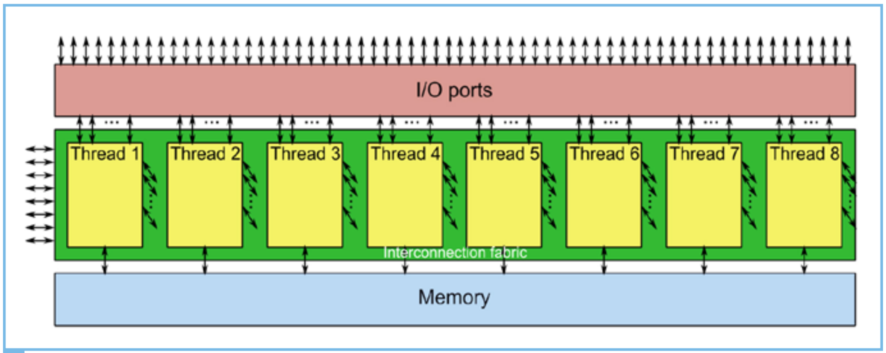

Figure 1 | Concurrent real-time programming executes multiple threads simultaneously.

Download the Digimarc app and take a picture of this image to get more details

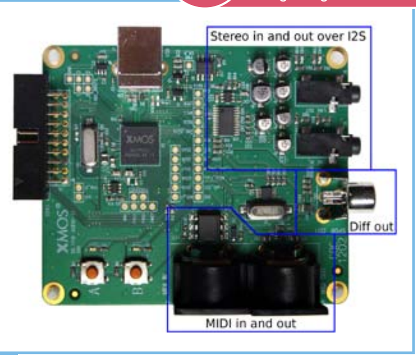

Figure 2 | On a USB board, the same processor can drive any mixture of interfaces simply by changing the software.

The latter is a programming model known as concurrent real-time programming. Figure 1 shows how a single XMOS XCore processor can execute eight threads simultaneously. Each thread can be seen as an individual processor with a guaranteed real-time execution rate. Each instruction will execute within a known amount of time; hence the programmer can predict whether programs will meet timing deadlines or not.

The beauty of this model is that timing prediction is entirely self-contained to the task. As an example, consider the I2S task mentioned before. Once the program has been written in a way that it will execute on, for example, a 50 MIPS processor, then regardless of modifications to other threads, this I2S thread will always meet its timing. The only way to break the timing is by not providing data at the appropriate rate. This is to be expected; if data is not provided at 48 kHz, then the codec cannot be given the data at 48 kHz, and something must break.

Software is not a panacea. Two obvious limitations are interfaces that require too much instantaneous bandwidth and hardware interfaces that need very short turnaround times. For example, it would not be feasible to write a multi-GHz SERDES in software on today's processing hardware, but many commonly found interfaces can be formulated in software.

There are three reasons why a software implementation is preferable over a hardware interface:

1. **Economics:** It is not possible to provide all interfaces available as hardware blocks. A processor can model any type of interface as required.
2. **Adaptability:** An interface built in hardware is set in silicon. If a standard evolves, or if a particular device is not completely compliant or is required to perform some extra operations, then a silicon interface will not do the job. A software-defined interface can be adapted to meet the requirements.
3. **Optimization:** A software-defined interface can be optimized to address the problem at hand. Consider the EtherCAT application mentioned previously. If a hardware Media-Independent Interface (MII) is used, then the packet is stored and forwarded, which does not meet, for example, latency requirements. If the interface is defined in software, then the buffering requirements can be tailored to the problem.

The third reason is important, as many standard blocks dealing with UARTs, I2S, MII, and other standards must have built-in FIFO buffers to decouple the interface, and those buffers in many cases are unwanted because they add delay.

Once the designer chooses a software-based design flow, the possibilities are endless. Figure 2 shows a USB board that simultaneously supports an analog interface using a stereo I/O codec (over I2S), a differential digital interface (using S/PDIF), and a Musical Instrument Digital Interface (MIDI).

The embedded processor implements the USB stack and the audio-over-USB protocol. In this particular configuration, three audio interfaces are driven: MIDI, I2S/analog, and S/PDIF. The latter can be changed to a different protocol such as ADAT by a mere software change. More importantly, variations of the product can be made with different sets of interfaces, such as a multitude of coaxial outputs or I2S-based codecs. The software on the embedded processor can be tailored to the specifics of the interfaces with only minor alterations. There are no predefined numbers of each interface.

Flexible hardware can enable greater product differentiation, allowing developers and product managers to deliver distinguishable value to their customers. Concurrent real-time programming allows tight timing requirements to be guaranteed, with embedded processors such as the XMOS XCore providing the silicon and development tools required to take advantage of this programming method. ECD

Henk Muller *is principal technologist at XMOS Ltd.*

XMOS Ltd.
henk@xmos.com
www.xmos.com

Follow:
Community forum

Q & A

People-centered innovation driving the next generation of telehealth

Q&A with Alan Boucher, Director of Software Architecture and Engineering, Intel-GE Care Innovations

Intel Corporation and GE Healthcare came together in 2011 to form a new company that is transforming the world of technology-enabled health care by rethinking the role of the patient in telehealth system design. Alan describes how Care Innovations is creating a new space for embedded computing technologies in virtual care coordination applications.

ECD: ***What is the mission of Care Innovations, and how does your organization leverage the research conducted by parent companies Intel and GE?***

BOUCHER: Intel-GE Care Innovations creates technology-based solutions that give people confidence to live independently, wherever they are. We are a unique joint venture between Intel Corporation and GE that was formed last year, combining assets and expertise from Intel's Digital Health Group and GE Healthcare's Home Health division.

Both Intel and GE Healthcare have a long history of driving innovation, solving hard problems, and creating new markets, with a massive body of research and expertise to back that up. GE's health care expertise covers medical imaging and information technologies, medical diagnostics, patient monitoring systems, drug discovery, biopharmaceutical manufacturing technologies, and performance improvement/solutions services.

Intel also has a deep understanding of the health care industry. Since Intel began investigating health care in 1999, ethnographic researchers have observed and interacted with more than 1,000 households and 150 hospitals and clinics in 20 countries. Intel also initiated collaborative research projects with the Technology Research for Independent Living Centre, the Center for Aging Services Technologies, and many other organizations.

Today, Care Innovations can cull from this deep body of research in product development and prototypes. Our Care Innovations Connect and Care Innovations Guide products were developed directly out of this research. Current prototypes under investigation are also heavily informed by these extensive studies.

ECD: ***What telehealth products are you currently pushing to market, and how do embedded technologies enable and enhance these products?***

BOUCHER: The Care Innovations Guide (Figure 1) is a next-generation telehealth solution that combines traditional vital

Download the Digimarc app and take a picture of this image to get more details

Figure 1 | The Care Innovations Guide strengthens the connection between health care professionals and patients by enabling vital signs capture and videoconferencing capabilities.

signs capture with advanced videoconferencing and customizable multimedia education. It strengthens the connection between health care professionals and patients, and helps make virtual care coordination possible. We also offer Care Innovations QuietCare (Figure 2), a wireless monitoring technology that keeps caregivers informed about resident activity levels, which can help improve (patient) safety and security while allowing for proactive care.

In these products, embedded technologies such as wireless platform interfaces, ZigBee sensor networks, wireless medical peripherals, and software are designed for both wired and wireless patient use, allowing patients to integrate our products into their lifestyles without imposing restrictions or requiring significant changes. Platforms utilizing untethered devices, radio-based sensors, sensor networks, and a variety of transport media offer patients choices in how they can better manage their health care from a personal preference and lifestyle perspective.

Significant advancements in the Continua Health Alliance Bluetooth specifications, low-power Wi-Fi, Bluetooth profiles, and High-Speed Alternate MAC/PHY (AMP) offer us

Figure 2 | Care Innovations QuietCare utilizes ZigBee network technology to monitor resident activity levels for facilitating independent living.

connectivity options and the ability to make design choices across a variety of radio platforms. However, we still face all the challenges around device integration, data acquisition, data integrity, quality, safety, and design assurance with either commercial or embedded platforms and Operating Systems (OSs). Even with the abundance of Bluetooth medical devices available today versus 3-5 years ago, many of the design challenges persist in medical systems development.

ECD: ***Why does Care Innovations emphasize the need to deliver "human-centered" products and services? What other important technical considerations drive the design of a telehealth device/system?***

BOUCHER: At Care Innovations, perhaps surprisingly, we don't start with technology. Instead, we start with the people who use our products, which makes all of our solutions profoundly human-centered. We've spent nearly 12 years living with, studying, observing, and listening to people at all levels of health care and independent living. And we're including them as active participants in the systems we deliver. From our UX/UI formative modeling to our early design formative testing through to validation testing, we are committed to our patients and institutional customers who use and continue to influence our product design. External focus, quality, and discipline are the cornerstones of our products.

We look closely at user lifestyles when designing our products. For instance, chronically ill patients are not always tethered to their homes. Many of our patients live a more active lifestyle and need to take their devices out of their homes to remain adherent to care protocols and clinical instructions. This requires us to think differently from the core platform out to edge devices, device interactions, data acquisition, and data privacy and security. Concurrently, our clinical customers want to help their patients manage their disease states, which means effectively responding to measurement data while proactively managing patients' needs, regardless of their location.

These are all inputs to our UX/UI and systems engineering use cases and resulting product requirements. One size, model, or type rarely fits all, and that's certainly the case with the products that we build for our patients, clinicians, and independent living institutional customers.

ECD: ***What embedded technology advancements are needed to meet future goals of providing a unified technology platform with interoperable components?***

BOUCHER: Building blocks simply need to be better coordinated. With embedded platforms, OS vendors like Wind River have done a good job of building the requisite core software subsystem capabilities to enable innovative embedded medical device development. To a somewhat lesser degree, we've seen this emerge with Android, iOS, and Windows Mobile as well. However, it's still a complex undertaking. Companies like Care Innovations build systems that inevitably land on both embedded and commercial platforms, which patients interact with directly.

Regardless, the demands on resources, platform usage analysis, mobile/embedded platform roadmaps, and design/development trade-offs remain a complex puzzle for platforms built on embedded OSs and more commercially available operating platforms. Trying to navigate primary OS vendor core capabilities, feature/API exposure, embedded radio, and sensor and network interface availability can be the difference between success and failure in the marketplace. Bluetooth radio testing and qualification, device driver and I/O integration with platform middleware, design assurance, and data acquisition and integrity all play a role and are significant factors in development decisions.

Challenges can still arise with radio firmware, radio interface testing, API implementation/granularity, OS support for drivers, radio shims, profile implementations, and protocol and middleware integration at a platform level. Additionally, medical device development teams have compounded responsibilities in the areas of intended use validation, design assurance, quality, system test, verification, and data integrity that are expected of them according to product and medical regulations. Completing missing building blocks for platform developers, fully implementing profiles, reducing integration complexity, and improving design assurance will give medical device vendors an opportunity to focus on new capabilities and design advancements instead of solving other vendors' core OS, I/O, and subsystem challenges. ECD

__Alan Boucher__ is director of software architecture and engineering at Intel-GE Care Innovations.

Intel-GE Care Innovations
www.careinnovations.com

Follow:

Download the Digimarc app and take a picture of this image to get more details

Q & A

Harnessing technology to transform care

Q&A with Joseph Kvedar, MD, Founder and Director, Center for Connected Health

The Center for Connected Health is creating a new model for health care delivery by developing programs to move care from the hospital or doctor's office into the day-to-day lives of patients. Dr. Kvedar explains how an emphasis on patient engagement affects the design of personal connected health devices that can help providers and patients manage chronic conditions, maintain wellness, and improve adherence and clinical outcomes.

ECD: ***How is the Center for Connected Health currently working with technology developers to foster innovation and evaluate new connected health products?***

KVEDAR: Our in-depth experience, combined with the resources available at the Center for Connected Health and within the Partners HealthCare network, enables us to deliver expert opinions and help organizations prepare their products and services for integration into the health care delivery system. The Center's experienced clinicians and technologists are skilled at working seamlessly with patients and providers to determine potential use cases, identify usability challenges and opportunities, and evaluate technology and workflow issues in medical settings as well as in the home. Our work continues to demonstrate how the right technology, in the right patients' hands, in the right setting can have a profound impact on care and quality of life.

In one example, working with a company developing a medication adherence device, the Connected Health team designed the protocol, recruited more than 120 subjects, and completed the data analysis to evaluate clinical outcomes. This study is currently in publication review. In another case, we assessed a prototype Bluetooth USB device and software platform enabling data upload from the home. Our team provided a thorough and realistic evaluation of the technology's reliability and usability and identified key factors for further product design.

ECD: ***What challenges do telehealth system developers face today, both on the technical side in terms of design requirements, as well as the clinical side in regards to care delivery and technology acceptance?***

KVEDAR: One challenge system developers face is how to test their product or service in a real-world environment. In other words, how does the developer of an activity monitor, wireless weight scale, or blood pressure cuff have individuals use these products on a test basis in their home or connected with their provider? Working within a large provider network, the Center can bridge that need and put these devices and systems in the patients' and providers' hands.

A second significant challenge is the myriad of platforms and Operating Systems (OSs) currently available in the market. Designing a device or program that is suitable for the iPhone, for example, will need to be redesigned for Android, Blackberry, and on and on. Obviously, this requires significant development time and expense. Making all personal health devices plug-and-play so that any sensor or app will work on all platforms, as well as allowing patients and providers to easily and securely share data, are essential requirements for the widespread success of connected health systems. The Continua Health Alliance (www.continuaalliance.org) is creating interoperability standards and guidelines to help streamline the development of these personal health devices and make it seamless for patients and providers to use these technologies.

ECD: ***Where does security fit in the connected health picture, and how must technology advance to address this and any other area of vulnerability and risk?***

KVEDAR: Health care providers are obligated to protect patient privacy and an individual's health information. E-mail communication is one example of a potential vulnerability. However, there are a number of ways to safeguard against potential security breaches. First, e-mail messages can be encrypted. Providers can also communicate with patients using a secure software platform that has the same functionality as e-mail, but is specific to communication with a health care provider. Secured messaging applications are often part of a patient portal or electronic medical records systems offered by physician practices and hospitals.

ECD: ***The Center for Connected Health emphasizes patient engagement and supports initiatives promoting disease prevention and management. How should these objectives influence the design of a connected health device?***

> The success of connected health programs is all about the psychology of engaging participants to motivate them to improve their health.

KVEDAR: That is a very important question. First, the device must be easy for the individual to use. In a recent remote monitoring trial conducted at the Center, we found that for a surprising percentage of diabetes patients, the step of plugging a device into the glucometer and the phone line and then pushing a single button to upload glucose readings was more work than they were willing to do. The technology must be simple and easy to use.

Moreover, we have learned that the most successful technology or product can be personalized to the patient's experience, goals, or motivation. Patients are far more engaged in their care plan or wellness program when their own personal data is presented back to them – the feedback loop – in an easy-to-understand format. They can track their progress, see how their lifestyle choices are affecting their health, and learn how to best manage their health and wellness.

However, personal connected health data alone is not enough, except for a very small group of highly motivated individuals. Objective data is an important part of the solution, but the success of connected health programs is all about the psychology of engaging participants to motivate them to improve their health. We're seeing many attempts at engagement strategies, including gamification, social networking, coaching, reminders, incentives, and punishments.

Our friends at Healthrageous (www.healthrageous.com), a health engagement company, offer a good example of a solution specific to an individual's data, personal habits, and preferences. The company has developed a platform based on dynamic personalization, meaning that each intervention is tailored to meet an individual's needs (see Figure 1). The goal is to know as much as possible about each individual using machine learning to anticipate the engagement experience that motivates each person to stay on the right track. ECD

Joseph Kvedar, MD, is the founder and director of the Center for Connected Health.

Center for Connected Health
jkvedar@partners.org
www.connected-health.org

Follow:

Figure 1 | The Healthrageous platform uses machine learning technology to analyze personal health data and promote strategies tailored to achieve successful outcomes for each individual user.

DESIGN West highlights constant change in embedded marketplace

Embedded systems vendors were in full force at the recent 2012 DESIGN West/Embedded Systems Conference showcasing a wide range of new products and services designed to simplify and shorten your design process. I met with many industry experts to hear about their latest approaches to deal with emerging technology challenges in the embedded world. They covered topics such as the ins and outs of transitioning your design to multicore for higher performance or to lower component count. Cloud connectivity hardware along with the required software security was another popular subject. Many of these challenges are being addressed with off-the-shelf modules and software products displayed by more than 200 exhibitors at this year's conference. One thing is for certain: The embedded systems industry is in a perpetual state of change as designers react to customer demands for higher performance and all the latest features.

– Warren Webb

Tiny SOM includes wireless connectivity

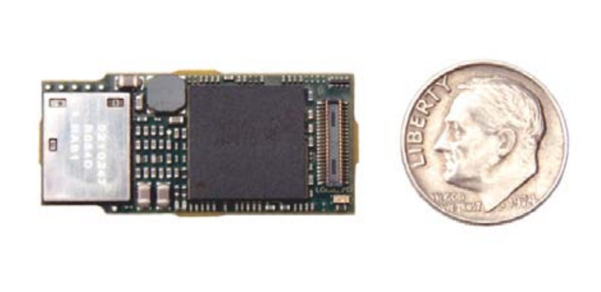

As intelligent embedded devices become ubiquitous, designers are searching for small, low-power techniques to provide continuous, autonomous connectivity and effective remote management. To meet this growing demand, Logic PD recently introduced a small form factor, low-power System-On-Module (SOM) with built-in wireless capabilities. The new Logic PD DM3730 Torpedo + Wireless SOM is based on the Texas Instruments DaVinci DM3730 digital video processor and WiLink 7.0 connectivity products.

Less than 1 square inch in size, the Torpedo + Wireless SOM provides PC-like speeds up to 1 GHz, long battery life, and a WL1283 wireless chipset that includes 802.11a/b/g/n and Bluetooth with a roadmap to support GPS. Additionally, the DM3730 Torpedo + Wireless SOM is FCC certified, saving customers time and money by reducing design risk and shortening the final electromagnetic compatibility testing phase. Logic PD provides Android and Linux board support packages for use with the DM3730 Torpedo + Wireless SOM.

Logic PD | www.logicpd.com | www.embedded-computing.com/p367812

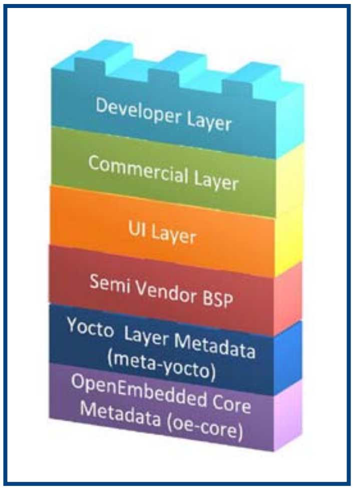

Embedded Linux platform supports open-source tools

Built with contributions from hundreds of software experts, the Linux operating system has captured the attention of embedded developers worldwide and is available in commercial distributions throughout the embedded market. To help eliminate the requirement to match a specific Linux configuration with a new design's hardware architecture, Mentor Graphics recently released the next-generation Mentor Embedded Linux platform with support for the open-source Yocto Project. This project is a collaborative venture established by the Linux Foundation to provide standardized open-source tools and resources to support embedded development.

The new Mentor Embedded Linux platform allows developers to easily select the best Linux kernel for their needs, whether that kernel was developed by Mentor Graphics, a semiconductor company, or another third party. Reference board support packages for leading hardware platforms, including those from Broadcom, Freescale Semiconductor, Intel, and Texas Instruments are included in the Mentor Embedded Linux platform. The Linux platform also supports the QEMU emulator, allowing users to simulate Linux system development without using actual hardware.

Mentor Graphics | www.mentor.com | www.embedded-computing.com/p367797

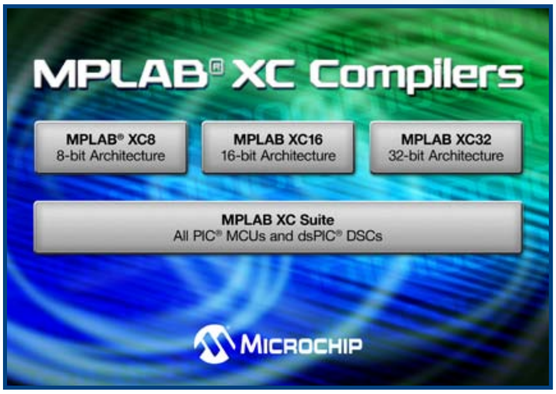

New C compilers improve code performance

With each new project, embedded developers are often challenged to enhance software execution speed and reduce overall code size with fewer resources and a shorter schedule. Offering support for all of their PIC microcontrollers and dsPIC Digital Signal Controllers (DSCs), Microchip Technology is delivering three new C compiler options that improve code execution speed by about 30 percent and reduce code size by roughly 35 percent.

The MPLAB XC8, XC16, and XC32 compilers offer reduced complexity for 8-, 16-, and 32-bit designs and come in three optimization levels: Free, Standard, and Pro. The free editions offer many optimizations, are fully functional, and have no license restrictions for commercial use. Designers can also test their code for 60 days with the Pro optimization levels, which are approximately 50 percent better than the free editions. MPLAB XC compilers also support Linux, Mac, and Windows operating systems, enabling designers to use their platform of choice for embedded development.

Microchip Technology | www.microchip.com | www.embedded-computing.com/p367798

Evaluation kit simplifies certified touch-key designs

Smartphone technology has become commonplace with the general public, and customers expect similar performance and touch-screen interaction features with all of their electronic devices. Supporting the certification of the human interface portion of embedded design, Renesas Electronics America announced a new capacitive touch evaluation kit. The kit is based on Renesas' R8C/3xT-A MCUs that feature a specialized hardware block called the Touch Sensor Control Unit (SCU), which handles touch scanning and processing functions completely independent from the CPU.

The evaluation kit is certified by TÜV Rheinland, a multinational testing facility, to pass the IEC/EN 61000-4-6 standard, which defines compliance to RF-based conducted noise immunity. The system was designed to pass up to 15 V noise injection level over the full frequency range of 150 KHz to 230 MHz for all keys. The touch kit enables engineers to evaluate keypad, slide, and wheel configurations (including a 3 x 4 keypad) for touch sensitivity and noise performance on a single-board system.

Renesas Electronics America | http://am.renesas.com | www.embedded-computing.com/p367799

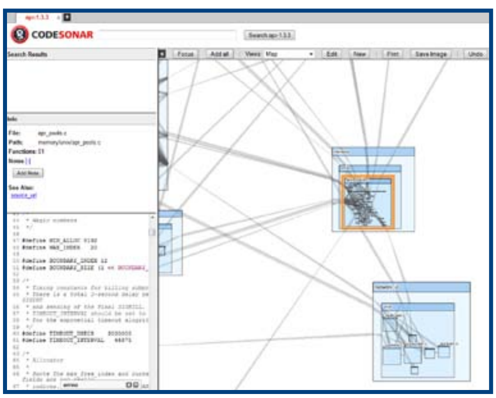

Visualization system guides code optimization and analysis

Software projects continue to grow in size and complexity, making it extremely difficult for developers to test and analyze the relationships between elements of the source code. At this year's DESIGN West conference, GrammaTech introduced a new software architecture visualization system capability designed to perform well on millions of lines of code. The system is integrated with CodeSonar, GrammaTech's static analysis tool for detecting defects.

The visualization system displays the software's call graph organized according to the program's module structure. Graph layouts can be changed in real time and present data in tree, map, circuit, cluster, flow, radial, and other layouts. Data is featured in a layered way that provides a high-level view with drill-down capability. Other features include the ability to search the graph for functions of interest, view the graph in a variety of different layouts, and navigate to and from the source code. CodeSonar visualization runs through a standard Web client such as Microsoft Internet Explorer, Firefox, or Chrome browsers.

GrammaTech | www.grammatech.com | www.embedded-computing.com/p367800

Signal processors target high-performance, portable applications

With the performance and portability requirements of critical applications in markets such as industrial automation, embedded vision, video surveillance, and medical imaging on the rise, embedded designers are looking for low-power, cost-optimized components. With these applications in mind, Texas Instruments (TI) unveiled three new devices based on its KeyStone multicore architecture utilizing the TMS320C66x DSP.

The C665x processors consist of three pin-compatible versions for developers migrating from single core to multicore. The C6657 features two 1.25 GHz DSP cores, delivering up to 80 GMACs and 40 GFLOPs, while the C6655 and C6654 single-core processors deliver up to 40 GMACs and 20 GFLOPS and 27.2 GMACs and 13.6 GFLOPS, respectively. Under normal operating conditions, the C6657, C6655, and C6654 power numbers are at 3.5 W, 2.5 W, and 2 W, respectively. The C665x DSPs also feature large on-chip memory coupled with an external memory controller for low-latency applications. TI offers a series of evaluation modules with development tools and support to simplify the transition to the new DSPs.

Texas Instruments | www.ti.com | www.embedded-computing.com/p367801

Miniature microcontroller fits space-constrained applications

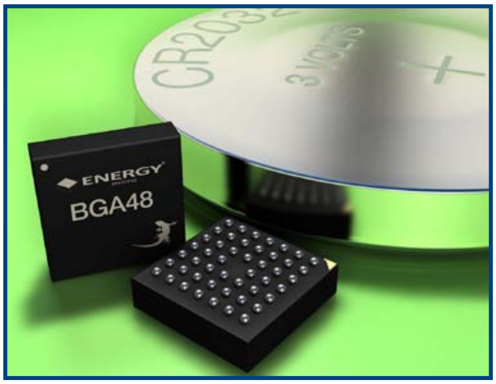

When designers undertake applications with extreme space limitations, such as wireless sensing nodes, home automation systems, and portable health and fitness products, package size becomes a critical component selection factor. Extending its EFM32 Gecko range of microcontrollers to meet these requirements, Energy Micro added products supplied in thin, fine-pitch BGA48 packaging.

Based on the ARM Cortex-M3 processor, the new Tiny Gecko devices require a current of just 150 microamps/MHz when active. Despite their small size (4 mm x 4 mm x 1 mm), the VFBGA48 Tiny Gecko devices include several energy-friendly peripherals, including the company's low-power LCD driver technology and the LESENSE low-energy sensor interface. The LESENSE function enables monitoring of a mix of up to 16 capacitive, inductive, or resistive sensors independently from the processor core. The Cortex-M3 Tiny Gecko microcontrollers run at clock speeds up to 32 MHz and are equipped with up to 32 KB of flash memory and 4 KB of RAM. The RoHS-compliant VFBGA48 devices will be available in sample quantities in Q3 2012.

Energy Micro | www.energymicro.com | www.embedded-computing.com/p367802

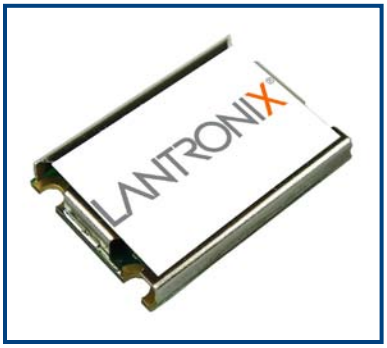

Device server simplifies M2M network connectivity

Machine-to-Machine (M2M) communications is a growing trend in embedded design and promises to transform connectivity and remote management across many market segments, including industrial, transportation, health care, utilities, retail, and consumer electronics. Offering designers lower development costs and faster time to market for M2M applications, Lantronix introduced the xPico device server at the 2012 Embedded Systems Conference.

Touted as one of the smallest embedded device servers available today, xPico can be used in designs typically intended for chip solutions, although there is no need for software development. The xPico comes in a 24 mm x 16.5 mm chip-sized footprint complete with a device server application that includes a full IP stack and Web server. Other features include a serial data rate of up to 921 Kbps, 256-bit AES encryption, and an extended temperatures range of -40 °C to +85 °C. Lantronix expects to begin volume shipping of xPico in the early second half of 2012.

Lantronix | www.Lantronix.com | www.embedded-computing.com/p36780

ELMA's FPGA reference platform has got your back (and your front)

ELMA's VPX-300 FPGA reference platform bundle provides front-end and back-end processing for compute-intensive systems such as managing radar signals or high-resolution images. With a dual-cluster architecture, a front-end Virtex-6 FPGA processor card performs the incoming digital signal processing coupled via fabric to a back-end data cluster. The nine-slot OpenVPX™ backplane provides two star sections, handling both PCIe and Gigabit Ethernet links.

Our VPX-300 reference platform is just one bundled option available. Let ELMA engineers customize a FPGA platform perfect for your organization's needs incorporating the chassis, backplane, power, video, computer, storage, I/O, software and more. Let ELMA show you why we're your solution partner where the smallest detail matters.

ELMA
Your Solution Partner

Learn why the smallest detail matters with FPGA solutions at http://goELMA.com/frap.
Or call toll free (855) GO-ELMA1 or (855) 463-5621.

E-community Beat

Keeping the embedded conversation going

By Jennifer Hesse

www.embedded-computing.com

BLOGS | MARKET STATS | SURVEYS | VIDEOS | SOCIAL NETWORKING

ESC/DESIGN West demo: Microcontroller-powered medical marvels

Mobile health devices such as wearable sensors are getting more sophisticated and powerful, which means they need small, high-performance microcontrollers to manage their functionality. At ESC/DESIGN West 2012, Editor Warren Webb interviewed Steve Kennelly, senior manager of the Medical Products Group at Microchip Technology, to discover how Microchip's PIC MCUs are being used in a variety of medical applications, including a glucose meter for iPod Touch/iPhone/iPad, a drug delivery patch for treating migraines, and a fertility monitor that measures a woman's temperature shifts. http://opsy.st/HXGbxZ

See more videos from ESC/DESIGN West in our library: http://video.opensystemsmedia.com.

Image recognition

Download the Digimarc app and take a picture of this image to get more details

Use your smartphone to scan the Microchip medical solutions photo and other images throughout the magazine with the digital watermark icon to view the latest videos and multimedia focused on embedded computing.

Application Lifecycle Management TechChannel

An Application Lifecycle Management (ALM) platform provides holistic coordination of all software design processes as a whole. Check out our latest TechChannel for ideas on how to manage and coordinate all of an application's or project's artifacts and activities.

http://tech.opensystemsmedia.com/application-lifecycle-management/

Wireless entertainment takes to the air

Airline passengers who dread the customary "please turn off your cell phones and laptops" speech by flight attendants will be glad to hear that the skies are expected to get more wireless-friendly in the near future. Wireless In-Flight Entertainment (IFE) is forecasted to be installed in roughly 9,000 aircraft by 2021, according to a new report by IMS Research, which was recently required by IHS.

While onboard Wi-Fi and cellular communication are already offered by several airlines today, wireless IFE is emerging to offer media such as video, TV, games, and audio on passengers' tablets, smartphones, or airlines' own handheld devices. The technology not only has the potential to revolutionize how passengers are entertained, communicate beyond the plane, and manage their ongoing travel, but also could help airlines reduce weight and improve overall economy by eliminating the need for fixed seat-back systems and cabling.

Read more: http://opsy.st/I8RXsG

Roving Reporter blog: Connectivity gets patients hooked up to care and fired up to stay well

Intel® Embedded Community
Embedded with bright ideas

By Jennifer Hesse

Forgetting to take your meds or not having time to get to the gym no longer cut it as excuses for not maintaining your health. With all the mobile health (mHealth) products available today – from wearable devices that measure vital signs to pills that send texts about missed doses to a host of gaming apps to track fitness or disease management activities – consumers have an arsenal of resources for promoting their well-being.

Members of the Intel Intelligent Systems Alliance are helping enable patient engagement in mHealth applications by offering products that ensure reliable real-time connectivity. For example, wireless modules from Telit Wireless Solutions are deployed in a wide range of mHealth applications, including Vitality GlowCaps pill bottle reminders, pomdevices' Sonamba senior well-being monitor, and MedApps' HealthAir personal health monitor, which uses Telit's GE865-QUAD module to automatically send medical data to a patient's Electronic Health Record (EHR), allowing users to manage and record their health regimen.

Read more: http://opsy.st/IgbYii

A 1.66GHz Intel® Atom™ Industrial SBC of *EPIC* Proportions

The EPX-C380 is a rugged single board computer that provides an open powerful platform for industrial applications. Powered with either a single or dual core processor, it has a wealth of onboard I/O plus expansion options. Also it supports Linux, Windows® XP embedded, and other x86 real-time operating systems.

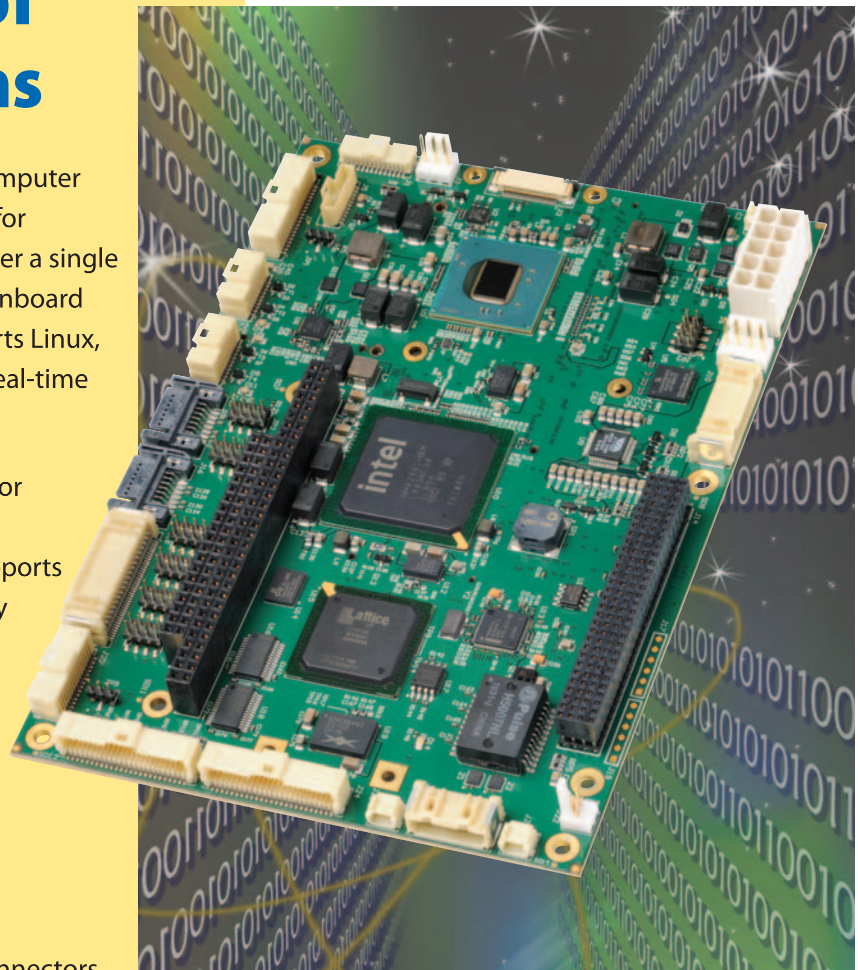

- 1.66GHz Intel® Atom™ N450 single core or D510 dual core processor available
- Embedded Gen 3.5+ GFX video core supports CRT and LVDS flat panels simultaneously
- Custom splash screen on start up
- Optional 1MB of battery backed SRAM
- Two Gigabit Ethernet ports
- Two SATA channels
- Eight USB 2.0 ports
- Four serial RS-232/422/485 channels
- 48 bi-directional digital I/O lines
- CompactFlash (CF) card supported
- MiniPCIe and PC/104-*Plus* expansion connectors
- High Definition audio supported (7.1)
- WDT, RTC, LPT, status LEDs, and beeper
- +5 volt only operation
- EPIC sized: 4.5" x 6.5" and RoHS compliant
- EBX form factor SBC (5.75" x 8.0") also available
- Extended temperature operation
- Responsive and knowledgeable technical support
- Long-term product availability
- Quick Start Kits for software development

Contact us for additional information, pricing, and custom configurations. Our factory application engineers look forward to working with you.

Call 817-274-7553 or Visit www.winsystems.com/EPX-C380E1

Ask about our 30-day product evaluation

WinSystems®
An Employee-Owned Company

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

Annapolis Micro Systems

The FPGA Systems Performance Leader

High Performance Signal and Data Processing in Scalable FPGA Computing Fabric

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

Direct Seamless Connections with no Data Reduction
Between External Sensors and FPGAs
Between FPGAs and Processors over IB or 10GE
Between FPGAs and Standard Output Modules
Between FPGAs and Storage Arrays

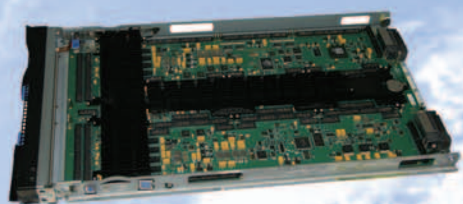

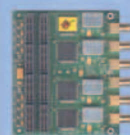

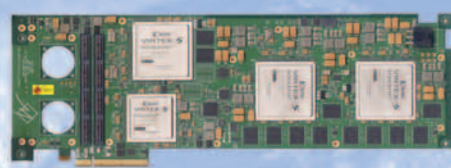

Annapolis Micro Systems

Ultimate Modularity
From 1 to 8 Virtex 4, 5 or 6 FPGA/Memory Modules
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
Infiniband, 10G, 40G or 100G Ethernet or SFPDP

VME/VXS/VPX, IBM Blade, PCI-X/PCI Express, PMC/XMC, MicroTCA

No Other FPGA Board Vendor Streams This Volume of Data Real Time Straight Into the Heart of the Processing Elements and Then Straight Back Out Again

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