Embedded COMPUTING DESIGN

www.embedded-computing.com

Got a Sec...

E-casts
Learn by viewing 46

@embedded_mag
Embedded Computing Design

APRIL 2010
VOLUME 8 • NUMBER 2

Embedded imitates (phone) life

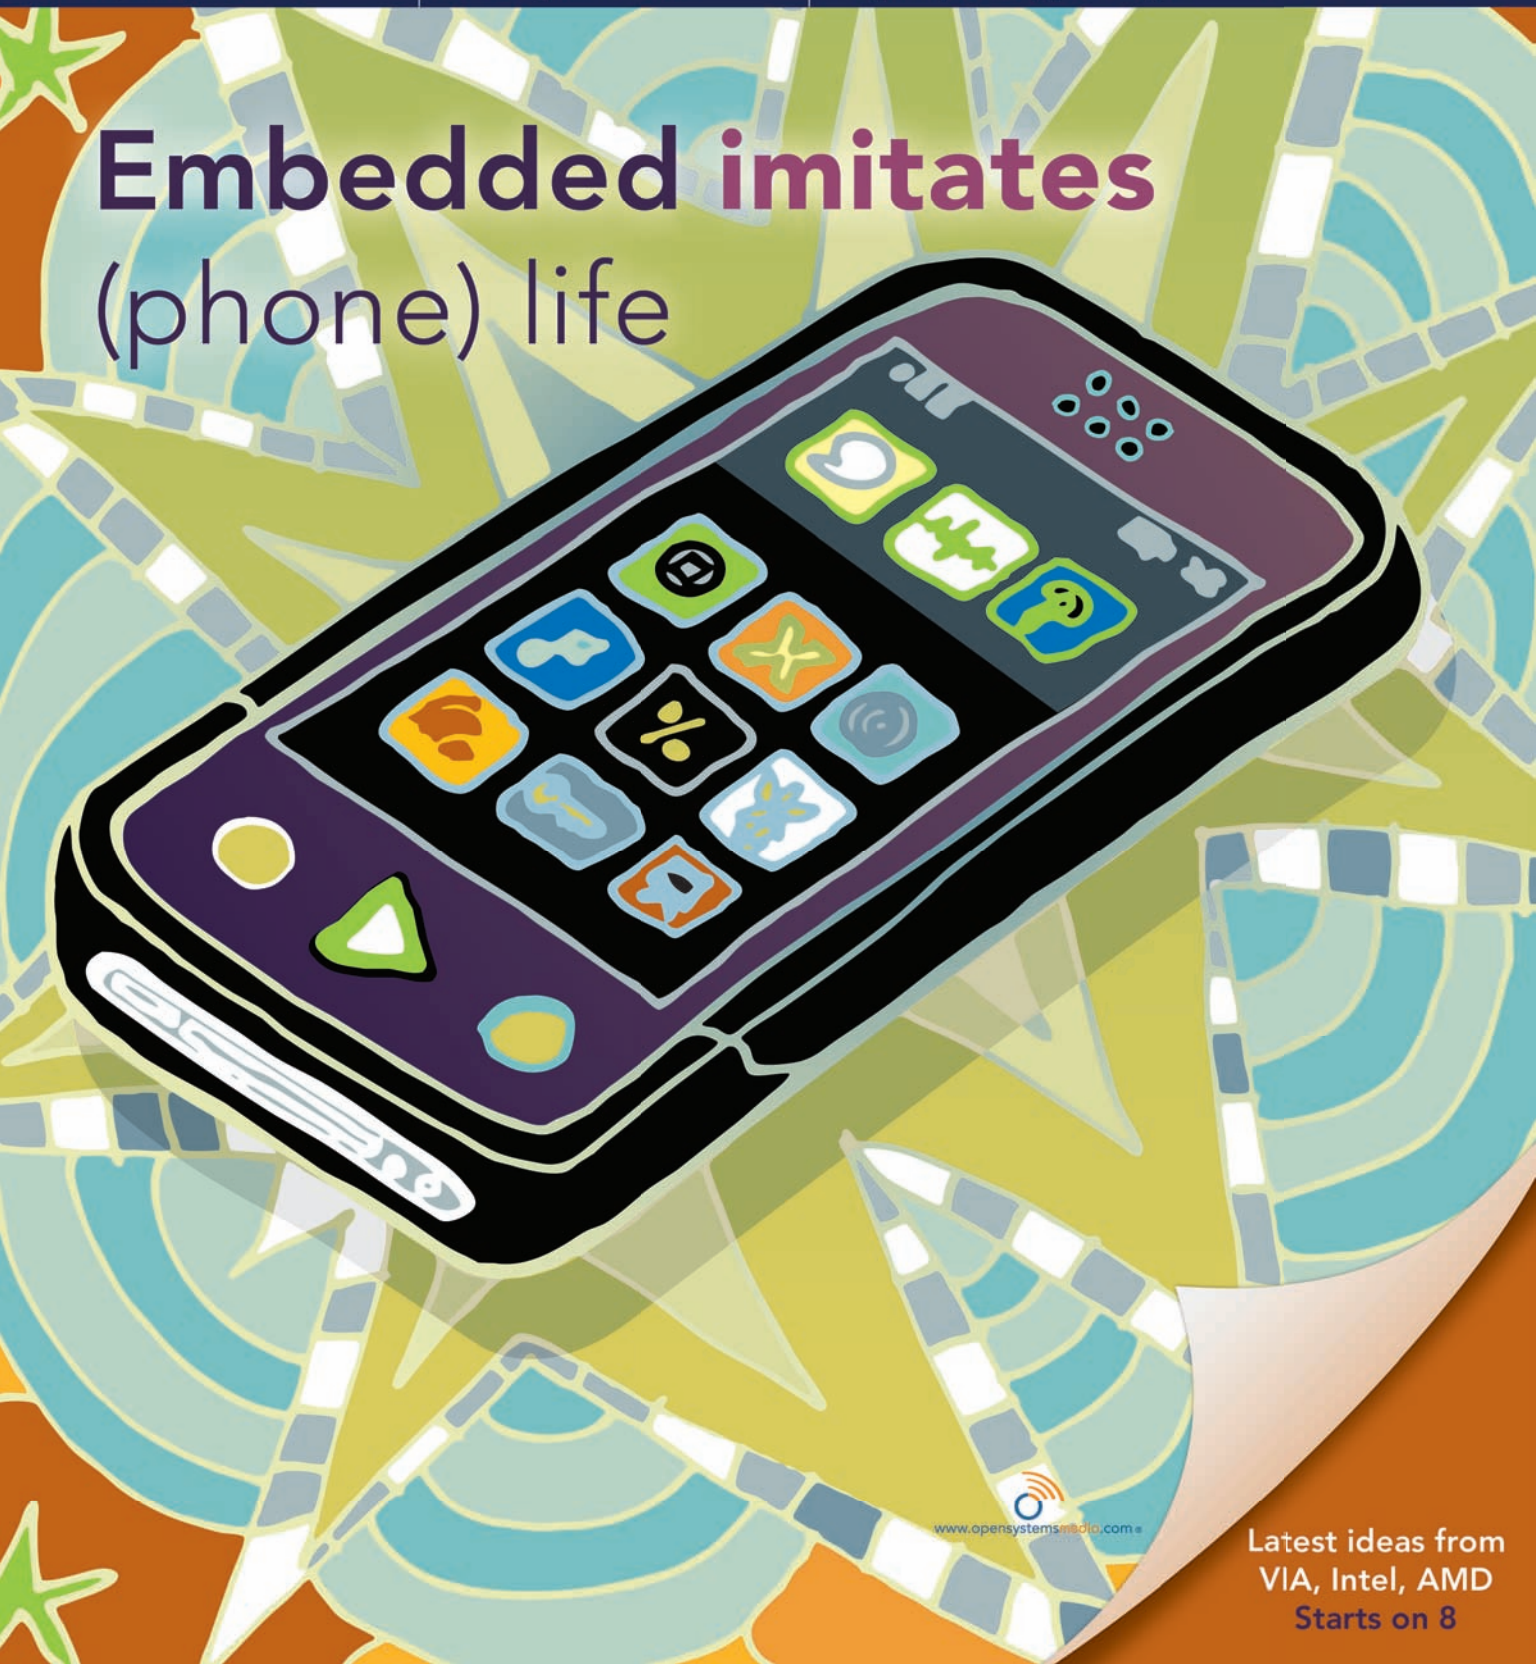

www.opensystemsmedia.com

Latest ideas from VIA, Intel, AMD
Starts on 8


THE EMBEDDED
DEVELOPER
A MALFUNCTION IN THE SYSTEM
COULD COST THE PLANT MILLIONS...
THEY'RE COUNTING
ON ME TO DELIVER.
600
500
400
300
200
100
0


Microsoft
Windows Embedded
BUILT WITH WINDOWS® 7 TECHNOLOGIES
VOL. 5
THE DEVICE HAS TO WORK PERFECTLY. TO THE MICROSECOND.
AND HAVE THE CONNECTIVITY TO TRACK PERFORMANCE IN REAL TIME.
WINDOWS® EMBEDDED OFFERS A HIGHLY RELIABLE PLATFORM, WITH THE LEVEL OF PERFORMANCE YOU NEED TO HELP DELIVER CONNECTED DEVICES THAT STAND OUT.
WHICH WINDOWS EMBEDDED PLATFORM CAN HELP YOU DELIVER STANDOUT DEVICES?
FIND OUT AT WINDOWSEMBEDDED.COM/DEVICESTORIES

Embedded COMPUTING DESIGN

www.embedded-computing.com

April 2010 | Volume 8 • Number 2

On the cover

Phones are now the center of many of our 24/7 lives. In "Left to My Own Devices" on page 7, Don Dingee makes the case that embedded apps will be forced to use a similar interface. There's proof in this month's Strategies department starting on page 32.

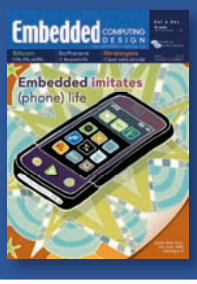

Silicon
CPUs, GPUs, and NPUs

Software
New ground in OSs

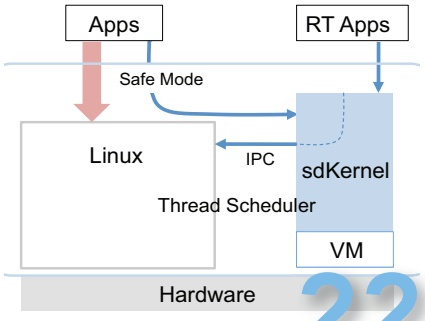

Strategies
Speech, medical, pick-to-light

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

2010 OpenSystems Media ® © 2010 Embedded Computing Design
All registered brands and trademarks in *Embedded Computing Design* are property of their respective owners.

ISSN: Print 1542-6408, Online 1542-6459

D-TA SYSTEMS INC.

SENSOR INTERFACE & PROCESSING COMPANY

10 GIGABIT SENSOR PROCESSING

THE FUTURE IS NOW

DELIVERING 10 GIGABIT NETWORK ATTACHED COMPLETE SOFTWARE RADIO SOLUTIONS

Combining High-Precision RF, High-Performance Data Conversion, High-Speed FPGAs & 10 Gigabit Network Technologies, D-TA Systems is enabling you to connect sensors to your computer. D-TA products come fully integrated. Now there is no need to struggle to integrate COTS boards from multiple vendors. **D-TA PRODUCTS WORK RIGHT OUT OF THE BOX.**

DTA-2300 Synchronized Multi-Antenna Software Radio

A 16-Bit Digital IF Transceiver for up to 16 Antennas. DTA-2300 is supported by multi-channel DTA-3200 RF Front-end and DTA-5000 Record & Playback.

DTA-2210 Single Antenna Software Radio

A 16-Bit Digital IF Transceiver that costs less than typical SDR COTS boards. DTA-2210 is supported by DTA-3200L RF Front-end and DTA-5000L Record & Playback.

CALL 1-877-382-3222 FOR A SPECIAL PACKAGE OFFER FOR DTA-2210 WITH DTA-3200L TUNABLE RF AND DTA-5000L RECORD & PLAYBACK

DTA-3200L Tunable RF Transceiver

20 MHz to 1 GHz with 40 MHz BW & 70 MHz IF

DTA-5000L 10 Gigabit Record & Playback System

with 2.4 Terabyte Storage

SENSOR PROCESSORS THAT DRASTICALLY REDUCE DEPLOYMENT TIME AND COSTS

WWW.D-TA.COM

Reader resources

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

E-casts:
ecast.opensystemsmedia.com

White Papers:
whitepapers.opensystemsmedia.com

Videos:
videos.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

Comment on blogs, articles, news, and more
Twitter: @embedded_mag
LinkedIn: www.linkedin.com/groups?gid=1802681

Advertiser Information

Page #	Advertiser	Ad title
34	ACCES I/O Products, Inc.	USB embedded I/O solutions
19	Actel Corporation	Innovative intelligent integration
14	Advanced Micro Devices, Inc.	Leveraging leading-edge PC graphics
15	American Portwell Technology	Your trusted provider for embedded applications
47	Annapolis Micro Systems, Inc.	The FPGA systems performance leader
17	Apacer	The most reliable storage for industries
38	Axiomtek	Intel Atom low power
21	CMP Techinsights	A new era, a new ESC
12	Concurrent Technologies, Inc.	PP 712/08x 6U CompactPCI board
5	D-TA Systems	10 gigabit sensor processing
11	Extreme Engineering Solutions	Intel Core i7 processor solutions
13	InnoDisk	Top reliability, finest solution
12	Kontron	ETXexpress-AI COM Express COM
39	LynuxWorks, Inc.	We won't leave you high and dry
41	MEN Micro Elektronik GmbH	MEN Micro's rugged COMs
27	Micro Digital, Inc.	What does our 35 years of experience bring to you?
2	Microsoft	The embedded developer
48	Moxa	Tiny and mighty
12	RadiSys Corporation	Procelerant CEQM57XT module
29	Sealevel Systems, Inc.	The right connection creates incredible power
35	Technologic Systems	TS-8100
26	Tri-M Systems Inc.	PC/104 Can-Tainer
33	Tri-M Systems Inc.	100MHz PC/104 Module
23	VersaLogic Corp.	2nd generation Intel Core 2 Duo processor
8	VIA Technologies, Inc.	VIA: Processors, digital signage, and USB 3.0
24	WinSystems, Inc.	PPM-LX800

Embedded COMPUTING DESIGN

OpenSystems media.

CompactPCI Advanced & Micro TCA SYSTEMS | DSP-FPGA.com | INDUSTRIAL EMBEDDED SYSTEMS | Military EMBEDDED SYSTEMS | PC/104 and small form factors | VME and Critical Systems

Embedded Group

Don Dingee, Editorial Director
ddingee@opensystemsmedia.com

Jennifer Hesse, Assistant Managing Editor
jhesse@opensystemsmedia.com

Monique DeVoe, Web Content Editor

Hermann Strass, European Representative
hstrass@opensystemsmedia.com

Konrad Witte, Senior Web Developer

David Diomede, Art Director

Steph Sweet, Creative Director

Joann Toth, Senior Designer

Phyllis Thompson
Circulation/Office Manager
subscriptions@opensystemsmedia.com

Reprints and PDFs

Nan Holliday
800-259-0470
republish@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

Andrea Stabile
Advertising/Marketing Coordinator
astabile@opensystemsmedia.com

Christine Long, Digital Content Manager
clong@opensystemsmedia.com

International Sales
Dan Aronovic, Account Manager – Israel
daronovic@opensystemsmedia.com

Sally Hsiao, Account Manager – Asia
sally@aceforum.com.tw

Regional Sales Managers
Ernest Godsey, Central and Mountain States
egodsey@opensystemsmedia.com

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

Editorial/Business Office

16626 E. Avenue of the Fountains, Ste. 203
Fountain Hills, AZ 85268
Tel: 480-967-5581 ■ Fax: 480-837-6466
Website: www.opensystemsmedia.com

Publishers: John Black, Michael Hopper, Wayne Kristoff
Vice President Editorial: Rosemary Kristoff
Vice President Marketing & Sales: Patrick Hopper
phopper@opensystemsmedia.com
Business Manager: Karen Layman

Left to My Own Devices

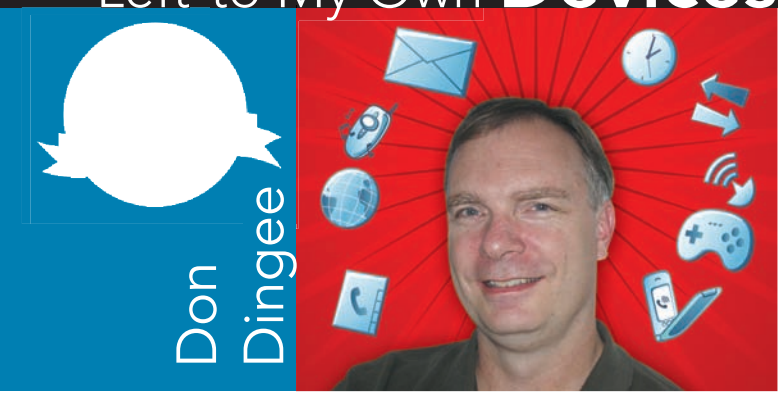

Don Dingee

Phonemance

A couple days ago, someone asked me if *Embedded Computing Design* was all about mobile phones now. I'll say that's not quite true, but it is symptomatic. My job is to spot trends, and I'm pretty sure that "phonemance" (think "bromance," only the object of affection is your phone, not your best buddy) has implications much larger for the embedded computing industry than any phone.

Have you heard or said this recently: "This phone is my life"? It's even a hook in phone advertisements, with messages urging you to bring everything in your life together in one place and promising the ability to do anything, anywhere, even multiple things at the same time. What needs to be recognized is that this notion is driving our industry, for better or worse, and it's a good idea to get your arms around it if you haven't already.

When the business-ready nerds broke out of command lines and assembly language and showed up with Windows, that drove a big slice of the embedded computing industry as we have it right up to today. It seemed almost every embedded system was required to run Windows because that's what the system user was familiar with. So the pendulum swung – look at CompactPCI, PC/104, and several other examples – toward x86 boards running Windows in many forms. Deeply embedded stuff, high-end Linux, and Real-Time Operating System (RTOS) environments survived because Windows simply didn't fit on everything and didn't do every job, but it still required serious nerds for those jobs, both to design and use them.

Now, look at what is happening today and learn from what has happened in the past. People want everything to behave like an iPhone or Android because the experience on their favorite device is what's driving their 24/7 life. (Just those two, you ask? Maybe two, with some chance of three ultimately. BlackBerry better get moving fast; it still has a chance. Palm is in *serious* trouble. So is Symbian. So is Windows Phone 7 unless billions of dollars get it right on the first try. Mark my words; I'm a certified smart mobile device pundit now. Follow the 2D barcode on this page for more on that.) It's not about the technology from the user perspective – it's about what they love and use, and what they will spend their money and time on.

I don't succumb easily, but as a baseball freak armed with a Droid, the first app I actually paid for was MLB At Bat 2010. My tweet after the first few minutes of using it was: "This is a lot cooler than the transistor radio I had as a kid." If a Droid can tap into my emotions and make me feel really good about something, good enough to yap about it on a social network, it's for real. Other Android and iPhone users will show you their favorite app with just as much passion and excitement.

Today's bottom line for your consideration is in two parts:

First, embedded devices have historically been forced to adopt the interface people use the most, and that interface is now the smartphone, which will soon shake out into two, maybe three flavors tops.

Second, the stuff that brings us the most joy is the stuff that we don't have to think much about to enjoy.

twitter @dondingee, @embedded_mag

ddingee@opensystemsmedia.com

Up next:
Design Automation Conference
Anaheim, CA
June 13-18
www.dac.com

Interview

VIA: Processors, digital signage, and USB 3.0

Joseph Chung, VIA Technologies

Editor's note: We asked VIA Technologies to provide some insight into their product offerings, examining where they are vis-à-vis the competition and showing us some of their forward-looking strategies.

ECD: What does VIA plan for a processor features roadmap, looking beyond the current C7 and Nano processor offerings? Where and how will you focus on competing with other vendors in the processor arena?

CHUNG: Back in 2000, VIA perceived a fundamental shift in the x86 computing industry: the demand for a power-efficient x86 processor that can deliver the necessary computing capacity while boasting a low power consumption of sub-10W and below. Our engineers therefore set out to design a generation of processors that will meet that demand. The result of our efforts was the C7/Eden processor released in 2005. VIA quickly established its reputation as an industry leader having the smallest and most power-efficient x86 solution.

In 2008, we released the successor to the C7/Eden processor, the Nano 1000/2000 Series, for the desktop and notebook markets. This earlier version of Nano processor compares favorably against the first versions of the Intel Atom processor, as confirmed in several independent reviews on a few well-known sites including Tweaktown and Wired.

Subsequently in 2009, we released the Nano 3000E (Embedded), which is the latest version of our Nano processor. It offers up to 7 years of availability for embedded system requirements, and it is available at processor speeds from 1GHz to 2GHz. The Nano 3000 Series processors deliver up to 20% higher performance while consuming up to 20% less power than the earlier Nano processors. The Nano 3000E is based on VIA's 64-bit superscalar "Isaiah" architecture, which boasts flawless playback of high bit-rate 1080p HD video. It also supports CPU virtualization technology, SSE4, and security capabilities integrated in the VIA PadLock Security Engine (NIST FIPS-197 Advanced Encryption Standard certified since December 2009).

Later this year, we expect to announce the addition of a dual core Nano to our processor lineup in response to our customers' continuing demand for additional processing power and their desire to leverage existing software infrastructure that takes advantage of multicore architecture. Further out, we will continue developing our x86 processor technology while taking advantage of newer semiconductor manufacturing processes to offer even more capable next-generation x86 solutions with Performance-Per-Watt leadership. This is in keeping with our strategy of improving performance while keeping the TDP to a minimum, a feature we have been known for in the industry for some time.

ECD: How has VIA been affected by the launch of the Intel Atom processor and the company's entrance into the small form factor low-power arena?

CHUNG: Before the Atom arrived, we had little or no competition in the x86 solution space in the small form factor low-power arena. When Intel launched the Atom, it was marketed as a new processor based on an entirely new low-power design built specifically for a new wave of Mobile Internet Devices and simple low-cost PCs. Although the Atom targeted many markets we are already very successful in, Intel's marketing prowess and brand recognition helped validate these markets, and thus greatly increased the overall demand for low-power x86 solutions.

Many critics continue to question Intel's strategy for Atom. The fact that further developing and improving upon the Atom product line may cannibalize their higher-end CPU markets is without dispute. Upon closer examination, Atom is merely a step back from the latest generation of x86 processors. It is only capable of in-order processing with L2 cache of only 512MB. It is primarily a 32-bit processor with only 533MHz FSB support. There is no cryptographic hardware implemented, either.

Fortunately, VIA does not have to limit the feature set of our processors, as we need not to worry about cannibalizing our higher-end markets. We have been focusing on the low-power small form factor market since 2001. Some of our recent successful consumer design wins include the Samsung NC-20 and the Lenovo S12 netbooks, both designed with the Nano processor and the VX800 chipset. These netbooks received rave reviews based on performance, low power consumption, and compact design (12" screen size).

> As we add more features (such as 64-bit, Advanced Cryptography Engine, virtualization, SSE4, and dual core support) into our processor offerings, we can address new markets that are currently underserved by x86 platforms.

As we add more features (such as 64-bit, Advanced Cryptography Engine, virtualization, SSE4, and dual core support) into our processor offerings, we can address new markets that are currently underserved by x86 platforms. Our integrated graphic chipsets such as the VX855 and the VX900 all-in-one multimedia system processors already support H.264 full hardware decode. For 3D graphics-intensive solutions we also have our S3G Embedded GPU offerings along with our CPU and integrated chipsets, enabling us to deliver DirectX 10.1, Hi-Def video, Blue-Ray, and Windows 7 to the SFF desktop, the thin and light notebook, and the embedded markets. The VB8003 Trinity Platform (Figure 1) serves as an example implementation of our CPU and chipset together with our S3G Embedded GPU.

ECD: Digital signage is a fast-growing area. What does VIA consider to be the key needs for signage platforms, and how are you addressing those needs?

CHUNG: Indeed, this is one of the fastest-growing markets not only in North America, but worldwide as well. Frost & Sullivan, Investor's Business Daily, and iSuppli all expect the digital signage market to reach US$5 Billion in North America alone and US$8 Billion worldwide by 2011. VIA believes it has the right platform solutions to address this vast market.

For the first generations of digital signage, the digital signage player platform requirements were very basic, with primarily 2D graphics and MPEG-2 video playback over analog displays (RGB, NSTC, or PAL) or digital displays in standard definition. The content was primarily stored locally on either CompactFlash or some sort of fixed or removable storage. The current generation of digital signage has evolved greatly. Today's requirements often include high bit-rate HD video playback in H.264 and WMV9 (VC1) formats at 1080p resolutions. In addition, HDMI, Dual-Link DVI, and DisplayPort interfaces are needed to drive the latest large digital LCD displays in either single or multiple display configurations. Customers looking for small form factor systems with these requirements will also look for additional features such as:

- VESA mounting (for easy mounting with plasma or LCD panels)
- Extended operating temperature support
- Low-power or fanless design
- Semi or fully sealed chassis design
- Robust and configurable I/O options

VIA has embedded system and board solutions in a variety of form factors that offer long life along with some or all of these features. For the very low-cost semi embedded digital signage system, we have the ARTiGO A1100 system built with the

Figure 1 | VB8003 Trinity Platform

EPIA-P820 Pico-ITX board. For customers that need ruggedized fanless chassis designs, the AMOS-3001 (Figure 2) is designed specifically for the EPIA-P820 board.

In 2009, we announced a new modularized ITX form factor for the embedded market, the Em-ITX and Em-IO. In that same year, we released our first Em-ITX product, the EITX-3000 with three Em-IO modules, the EmIO-3430 Wireless Carrier, EmIO-3210 Datacom Module, and the EmIO-3110 Video Module. Couple this with our AMOS series of ruggedized fanless chassis, and our customers can build a successful line of digital signage systems with different I/O configurations and capabilities, from the very basic, entry-level, single display system to the higher-end multi-display system that has up to six displays with four independent content streams.

For the very high-end solutions with High-Def performance and 3D graphics-intensive applications, we have S3G embedded graphics card solutions with single and dual GPU offerings. The dual GPU graphics card in a single PCI Express slot can run a wall display, the Video Wall, with up to 8 displays. The configuration can be 1 x 8, 2 x 2, 2 x 3, or 2 x 4. The 2 x 2 configuration can support up to 3800 x 2400 resolution natively.

ECD: USB 3.0 is starting to pick up speed. What are the benefits? Where do you see the technology fitting into the strategy, and what areas you are focusing on?

CHUNG: USB is a hugely successful peripheral interconnect not only for the PC market, but for the Consumer Electronics and the mobile/cell phone markets as well. Over the years, we had also seen a great deal of USB penetration into the embedded space as a robust interconnect for adding I/O. A couple of good examples are SUMIT from SFF-SIG and StackableUSB. Both use a stacking approach to add I/O expansion capability by leveraging interconnecting technology such as USB and its existing infrastructure of software and development tools. As technology advances progressed, we saw newer and faster devices demanding faster connection and higher bandwidth. The once-adequate bandwidth of USB 2.0 at 480Mbps nominal had quickly become obsolete, and thus, USB 3.0 was needed. The benefit of USB 3.0 is more than the much-needed improvement in link speed (5Gbps) and data bandwidth (~4.8Gbps). It is backward compatible with USB 2.0 and USB 1.1. In addition, USB 3.0 employs an even more aggressive power management scheme compared to USB 2.0 while doubling the bus power from 500mA to 900mA.

VIA has been successful with USB 2.0. Our USB 2.0 Host Controller ICs, VT6212 and VT6210, account for over 50% of the USB host controller market share. In addition, USB 2.0 is a standard feature in our x86 chipset offerings. To ensure our continuing success with USB, VIA has established VIA Labs Inc., a wholly owned subsidiary to solely focus on developing our own USB 3.0 IP portfolio.

In January this year, we released two brand-new USB 3.0 products, VL700 and VL810. The VL700 is a USB 3.0 to SATA Controller IC while the VL810 is a USB 3.0 Hub Controller IC. In recent years, personal storage devices of all shapes and sizes have proliferated the market. Among the more popular are the 2.5" and 3.5" HDDs with USB 2.0 connectivity. The capacity of these storage devices has progressed from merely 40GB to 1TB and beyond. We believe VL700 will unleash the performance of these devices, as the USB 3.0 5Gbps link speed is a better match with the SATA 3Gbps link speed.

We plan to release a USB 3.0 Host Controller product sometime later this year. This will be a quick and easy way for us to add USB 3.0 support to any of our existing x86 solutions with PCI Express connectivity. Coupled with the VL810, you can build a robust platform solution with multiple high-speed I/O connectivity, which will find ready markets in applications such as Hi-Def AV authoring, home surveillance, and high precision instrumentation. As the next step, we will integrate USB 3.0 into our next generation of x86 chipsets, which we plan to release in 2011.

Joseph Chung *is technology advocate at VIA Technologies Inc., based in Taipei, Taiwan. He has worked at VIA for 12 years and is currently responsible for the company's embedded component business in North America. He also chairs the SFF-SIG Pico-ITXe work group. Joseph received his BS in Electrical Engineering and BS in Computer Engineering from the University of Kansas.*

VIA Technologies
www.viaembedded.com

Figure 2 | AMOS-3001

Intel® Core™ i7
Processor Solutions

Optimized For Embedded Computing Applications.

6U VPX
3U VPX
6U cPCI
3U cPCI
XMC
VME

X-ES Intel® Core™ i7 Processor Solutions: Unmatched Performance

Extreme Engineering Solutions, Inc. (X-ES) unleashes the performance of the Intel Core i7 processor for embedded computing. By utilizing a processor with integrated graphics, PCIe, and ECC DDR3 memory controllers, the X-ES solutions deliver unmatched power savings and processing performance for compute intensive commercial and military applications.

X-ES offers an extensive product portfolio that includes commercial and ruggedized single board computers, high-performance processor modules, multipurpose I/O modules, backplanes, enclosures, and fully integrated systems.

Intel Core i7 processor solutions available now in a variety of form factors. Call or visit our website today.

X-ES
Extreme Engineering Solutions

608-833-1155 • www.x-es.com

RadiSys
THE POWER OF WE

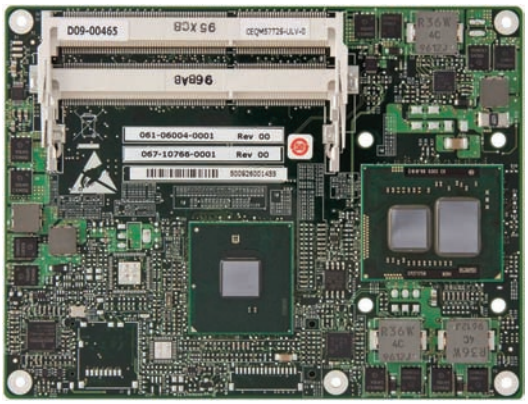

The Procelerant™ CEQM57XT module combines the Intel® Core™ i7 and i5 processors and the Mobile Intel® QM57 Express chipset with extended temperature range from -25 to +70 °C and ruggedized vibration specifications that are required for many mil/aero and industrial applications such as ruggedized computers/tablets, unmanned vehicles/robots and in-vehicle computers. RadiSys designs ruggedness into the CEQM57XT by utilizing its proprietary implementation of the Highly Accelerated Life Test (HALT) during the design process to identify and correct product weaknesses prior to production; Highly Accelerated Sample Screen (HASS) is then used during production to continually monitor supplier quality and component robustness.

800-950-0044 | www.radisys.com

CONCURRENT TECHNOLOGIES

The PP 712/08x 6U CompactPCI® board supports a choice of Core™ i7 processors: the 2.53GHz Intel® Core™ i7-610E, 2.0GHz Intel® Core™ i7-620LE and the 1.06GHz Intel® Core™ i7-620UE. In addition, the board provides up to 8Gbytes of DDR3-1066 ECC SDRAM, dual PMC/XMC sites, three Gigabit Ethernet ports, four SATA300 disk interfaces and dual head graphics. Rear I/O compatible with a previous generation, the board ensures a continuing upgrade path. Extended-temperature variants are available as well as support for many of today's leading operating systems. Further Core™ i7 based boards for release shortly include 3U CompactPCI and VME, including rugged conduction-cooled variants.

781-933-5900 | www.gocct.com

---

kontron

The Kontron ETXexpress®-AI COM Express™ Computer-on-Module offers OEMs a new level of design flexibility thanks to its 32 nm Intel Core i7/Core i5 processor technology, high energy efficiency, wide graphics support, customizable PCI Express configuration and ECC dual channel RAM to ensure data accuracy. The ETXexpress®-AI Computer-on-Module offers improved computing and graphics performance in the high-end performance sector.

High-performance COM with ECC and robust, flexible graphics:

- Processor performance up to Intel® Core™ i7 SV (2.53GHz)
- Support for up to 8GB system memory (DDR3 with ECC support)
- Up to 2.6x graphics performance improvement 3D Mark06 over previous solutions

888-294-4558 | www.kontron.com

iNNODISK®
Beyond your imagination

Top Reliability, Finest Solution

InnoDisk, the leading manufacturer of industrial storage devices, provides highly reliable products for extreme operating conditions. We offer comprehensive solutions for embedded systems and industrial applications.

-40°C~+85°C Wide Temp. Available

FiD 2.5" SATA 20000 Capacities: 16GB~256GB
Read: 230MB/sec.(max.) Write: 160MB/sec.(max.)

The total solution for Memory and Flash Storage

SSD Industrial Solid State Disk -Capacity to 512GB

-40°C~+85°C Wide Temp. Available

2.5" SATA 10000
Capacities: 8GB~256GB
Read:160MB/s(max.)
Write:140MB/s(max.)

1.8" SATA 10000
Capacities: 8GB~64GB
Read:160MB/s(max.)
Write:140MB/s(max.)

2.5" SATA 10000-RS
Capacities: 2GB~32GB
Read: 100MB/s(max.)
Write: 50MB/s(max.)

2.5" ATA 8000-J
Capacities: 8GB~128GB
Read: 85MB/s(max.)
Write: 75MB/s(max.)

SATADOM

i-Series
Capacities: 2GB~16GB
Read: 100MB/s(max.)
Write: 50MB/s(max.)
Dimension: 25x39.4x6.5mm(WxLxH)

SATADOM i-series

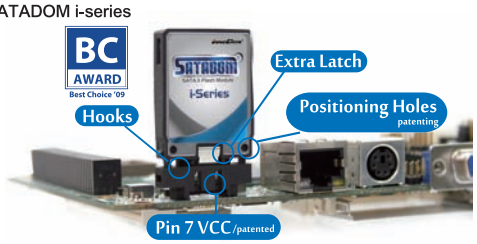

8000 Series-EDC/iCF

-40°C~+85°C Wide Temp. Available

Vertical type 40pin
Capacities: 2GB~16GB

iCF 8000
Read: 80MB/s(max.)

iCF Write Protect Series
Write: 75MB/s(max.)

i-DIMM Series

-40°C~+85°C Wide Temp. Available

Low Profile Available

Embedded Solution:
Low Profile, Wide Temp.
Server/Workstation Solution:
w/ECC, w/ECC w/PLL, w/ECC w/R
Printer Solution

-40°C~+85°C Wide Temp. Available

Low Profile / Vertical Type
Capacities: 128MB~8GB
Read: 24MB/s(max.)
Write: 14MB/s(max.)
Dimension: 0.81x1.54x0.31(inch)
Dimension: 0.90x1.28x0.31(inch)

www.innodisk.com

InnoDisk USA Corporation
43154 Christy Street Fremont, CA 94538 Tel: 510-770-9421 Email: usasales@innodisk.com

Leveraging leading-edge PC graphics technologies for embedded systems

By Peter Mandl and Richard Jaenicke

The pace of innovation and performance increases in the PC graphics industry is tremendous. Instead of just keeping up with Moore's Law, PC graphics double performance approximately every 18 months versus every 2 years (see Figure 1). At the same time, innovative features are added with each new generation. Although embedded systems do not change at nearly the same rate as the PC industry, embedded designers can take advantage of the fast-paced innovation in graphics technology to help design systems capable of productive use in the field for 5, 10, or even 20 years.

One of the most exciting new graphics technologies is the ability to drive up to six independent displays from a single Graphics Processing Unit (GPU). Found on the latest generation of GPUs, ATI Eyefinity technology can be used by a wide variety of applications to provide a large display surface with an area of up to 24 million pixels split across several monitors or independent content on each monitor, depending on the graphics card model. Simulation systems can use multiple monitors to offer peripheral vision and greater spatial awareness. Casino gaming systems can drive the main screen, entertainment screen, and player hospitality screen all from a single GPU. High-end digital signage systems can drive multiple screens without special-purpose display hardware.

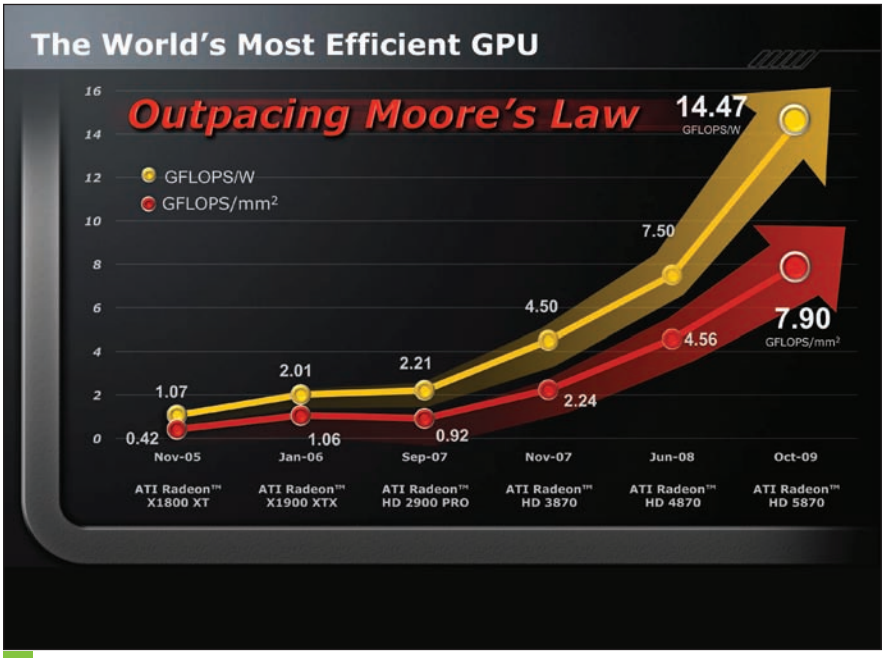

Figure 1 | GPU performance increases historically double every 18 months.

Another display technology of interest to embedded developers is DisplayPort. The latest version of the standard, DisplayPort 1.2, can double the bandwidth of HDMI and support stereoscopic 3D. Embedded system developers can appreciate that DisplayPort is an open standard free of royalties. DisplayPort 1.2 also adds the ability to daisy-chain monitors with multiple independent video streams, thereby cleaning up the swarm of cables needed for multi-monitor systems such as display walls.

The graphics technology likely to have the greatest impact on embedded applications is General-Purpose computing on Graphics Processing Units (GPGPU), also known as stream computing. Today's GPUs have evolved into massively parallel, programmable architectures optimized for performance per watt and per mm². By leveraging both the CPU and GPU for computation, developers have a broader toolbox to optimize and accelerate their applications. With the advent of the OpenCL programming standard for general-purpose computations on heterogeneous systems, software developers now have a way to use the same language for both CPUs and GPUs.

The parallel processing of stream computing is a great match to challenging problems in embedded applications with image and signal processing. The GPU can be used to accelerate encoding and decoding in video conferencing applications, creating the 3D model in medical imaging, the image formation process in radar imaging, pattern recognition in signal analysis, and seismic data processing for oil and gas exploration.

Embedded designers can access the first level of this technology with the Embedded ATI Radeon™ E4690 GPU (Figure 2). With 320 stream processors producing up to 3.88 GFLOPS of peak single-precision performance, the ATI Radeon E4690 GPU consumes only 25W and includes graphics memory in the same package.

The stream processing power is accessed through the ATI Stream SDK, including the compiler, device driver, performance libraries, and performance profiling tools.

The ATI Radeon E4690 GPU matches well with an AMD Turion™ processor to provide balanced serial and parallel computing as

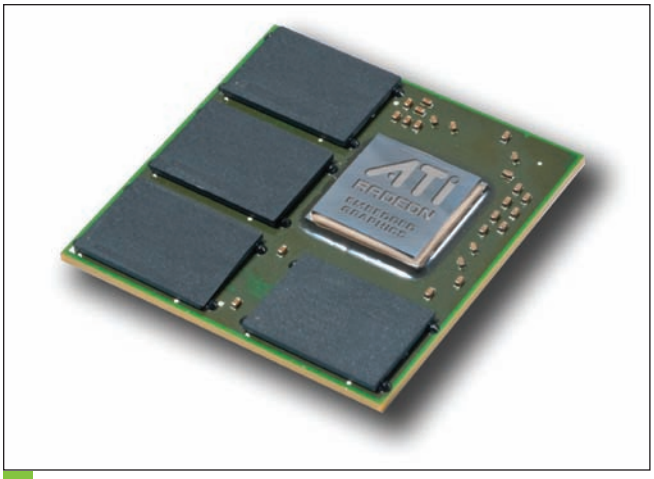

Figure 2 | ATI Radeon™ E4690 GPU with 512MB GDDR3

well as high-performance graphics and high-def video. The ATI Radeon E4690 GPU supports multiple types of display outputs, including DisplayPort, for flexible choices in embedded displays. For the next level of technology, one only has to look at the recently introduced ATI Radeon HD 5000 series of consumer graphics boards to see which technology may make its way into embedded graphics solutions.

Peter Mandl *is senior product manager at AMD, where he is responsible for embedded discrete graphics product management and marketing. His activities include competitive analysis, gathering customer and market feedback, generation of marketing/sales collateral, and definition of overall product strategy. Peter has extensive experience marketing graphics and video processing ICs. He holds a BASc in Electrical Engineering as well as an MBA.*

Richard Jaenicke *is director of the Embedded Graphics Product Division at AMD. His group gathers customer input to define, manufacture, and market graphics and stream computing solutions for embedded applications. Prior to joining AMD, Richard spent 15 years defining and marketing multiprocessor systems for embedded imaging and signal processing. He holds a Master's degree in Computer Engineering.*

AMD
www.amd.com

Your Trusted Provider for Embedded Applications

Single Board Computers

- Latest Intel® technology: Xeon®, Core™ 2 Duo, Pentium®, Celeron®
- Complete PICMG portfolio: PICMG 1.0, PICMG 1.2, PICMG 1.3
- Performance plus backward compatibility
- Long life cycle support

Extended Temp

ECX Boards

- Intel® Atom™ processor (Z510P, Z510PT, Z520PT or Z530P) with industrial temperature
- Dual display (LVDS and SDVO)
- Multiple USB ports
- Gigabit Ethernet
- Low power fanless and small footprint
- Long life cycle support

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

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

American Portwell is both an ISO 9001:2000 and ISO 13485:2003 certified company.

intel Embedded Alliance

Portwell
www.portwell.com
1-877-278-8899

Grok-ability and the multicore processor

By Bruce Edwards

With multicore processors being such a hot topic, it's worth noting that not all multicore solutions are created equal. Some are rooted in legacy with good reason, some are created from scratch for specific jobs, and some are being re-created to try to solve certain architectural limitations. The solution discussed herein addresses this last point.

Hypothetically disregarding money, time, or commercial constraints, and with physics the only limitation, how would you design the ultimate computer processor? Would it be massively parallel, run at insanely high frequencies, and use exotic optical or quantum interconnects? Would it run familiar software, like an x86 or PowerPC, or have a new optimized instruction set? Would it be very large or very small? Would it require intelligent compilers or unique software constructs?

For many years, all designers had to do to make processors go faster was crank up their clock speed. That worked fine until power consumption and the associated heat dissipation caught up with the speed increases. Beyond that point, going faster meant doing something besides simply going faster.

Multicore goes faster, but …

Thus began the multicore era. If two heads are better than one, then four must be twice as good. To some extent, that axiom is true. But today's dual- and quad-core processors aren't running two to four times faster than the previous generation.

There are two reasons for this: hardware and software. The great majority of today's multicore chips don't scale very well, so four cores don't really offer four times the performance of a single-core implementation. On-chip buses can't keep up, cache coherence overhead eats performance, pipelines stall too frequently, and so on. For a variety of reasons, conventional microprocessor architectures don't come close to doubling performance when their core count is doubled.

On the software side, many programmers aren't comfortable or familiar with multicore programming. This is especially true when the multicore chip in question includes different types of processor cores (often called a heterogeneous architecture). Programming one processor is hard enough; programming four different ones with separate tool chains is exponentially more complex.

Heterogeneous, homogeneous, or just humungous?

An argument can be made that different compute problems require different resources, and that microprocessors should therefore include a spectrum of different processing resources. For instance, some tasks might require signal-processing ability, others might require single instruction multiple data vector processing, while still others might involve complex decision trees and massive data movement.

One school of thought is that no one processor architecture can efficiently handle all these different tasks; thus, a mosaic of different architectures is required. In the extreme case, one can envision a processor made up of wildly different compute engines with nothing in common but the package they share. These processors are really cohabitating, not cooperating.

The opposite approach is to choose one instruction set and stick to it. This no doubt simplifies programming but runs the risk of deploying overly generic processors that aren't fine-tuned to a particular task. On the other hand, processors are programmable, and it's easier and cheaper to change software than hardware.

Ease of programming is not a trivial issue, either. Delays are typically caused by software bugs, not hardware problems. Complicating things further, programmers are scared to death of multicore processors. Getting one high-end processor to work reliably is hard enough; how do you program and debug 10 of them? It's easier to program, debug, and design with one core architecture instead of dealing with an amalgam of different cores with different instruction sets, architectures, buses, tools, and debug methods.

Contrasting approaches

Intel and AMD have taken most of this advice to heart and produced dual- and quad-core versions of their legacy x86

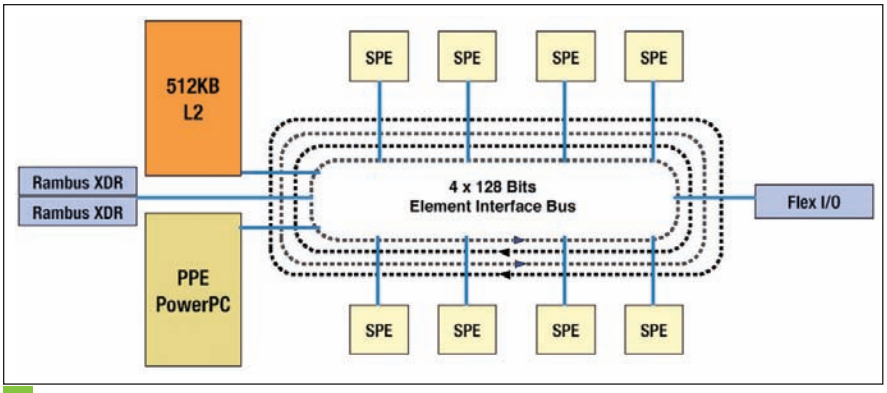

Figure 1 | IBM's Cell Broadband Engine chip includes nine processors, one of which is based on PowerPC. The processors are connected via an element interconnect bus, with a total of 12 masters. It is implemented as four counter-rotating unidirectional rings.

architectures. In part, that's simply making an asset out of a liability; x86 is what they know how to do, and backward compatibility is vitally important to their markets. Existing x86 code runs nicely on these upgraded designs, though it rarely runs much faster than before or makes significant use of the additional cores.

In contrast, many RISC CPU and network processor (NPU) vendors have taken a radically different approach, mixing an assortment of different processor cores and architectures into a variety of Swiss Army knife designs. IBM's famous Cell processor (Figure 1), for example, has one general-purpose processor core plus eight specialty cores, requiring different tools and programming techniques. Several wide buses – some rings, some more conventional – connect the cores in various ways. Cell's performance is impressive, but PlayStation programmers complain that Cell is a tough beast to tame, partly because managing bandwidth, latency, bus transactions, and coherence are all part of the game.

It's one thing to corral all the right hardware resources onto a single chip; it's quite another to make the combination usable. Massively parallel chips with a mixture of architectures combine the worst of both worlds: large-scale multicore complexity with disparate and distinct tools and architectures. It's like holding a United Nations meeting inside your chip.

Meshing together

A better approach is to keep the massively parallel part, which is *de rigueur* for high performance, but ditch the differences and connect lots of the same processor core

together in a two-dimensional mesh. In concept, it's not much different from connecting individual computers over a network, just on a microscopic scale.

Meshing also has "grok-ability" on its side. It's not hard for programmers to wrap their heads around the idea of ten, 100, or 1,000 identical processor cores working the same way and communicating with one another in a simple yet mostly transparent way. Whether each of the 1,000 elements is perfectly tuned for a given job is almost irrelevant; what's important is there are 1,000 processors to throw at a problem.

Such a homogeneous arrangement also aids scalability. While Cell-like combinations are well-suited to their specific tasks, building a larger or smaller version of Cell requires a significant amount of redesign work from the chip maker, and even more work from the programmer on the receiving end. Existing Cell code won't magically scale up or down to a chip with a different mix of resources. It might not run at all. In contrast, adding 25 percent more processors to a mesh of identical processors adds 25 percent more computing power without breaking existing code.

That doesn't mean designing this type of chip is trivial. Bandwidth between and among the cores is the first challenge. If the cores can't talk to each other efficiently, there's not much point in connecting them. An example of this approach is Tilera's TILE-Gx100 processor (Figure 2) packed with 100 identical cores. In this processor, the bandwidth between adjacent cores is 1,100 Gbps. With four connections per core in the north/south/east/west directions, the 100-core processor has an aggregate bandwidth of 200 TBps. Most applications would be hard-pressed to use a fraction of that. Even Tilera's relatively modest Gx16 chip with a 4x4 array of cores boasts 20 TBps of on-chip bandwidth.

Another challenge of this tile-based design is memory latency. If memory isn't close or accessible enough, all those processors can grind to a halt. Here again, Tilera breaks up its device into easily replicated tiles, each with its own local L1 and L2 cache. Interestingly, even though the memory is local to each tile, it can also be part of a larger shared distributed cache that maintains coherence among all of the sharers. In some scenarios, programmers might want to define an arbitrary number of islands of cache coherence, either cooperating with or ignoring neighboring tiles as necessary.

The overall chip architecture is like a fabric of computing. Identical blocks of logic, memory, and interconnect are replicated in rows and columns to make larger or smaller chips. And like an FPGA or fractal Mandelbrot image, a tiled processor looks the same at any scale. Large or small, it's programmed the same way. Scalability squared.

Like a quad-core x86 but unlike Cell or NVIDIA chips, the TILE-Gx mesh interconnect works transparently under the hood. Mesh traffic doesn't need to be massaged manually, nor do transactions need to be hand-tuned to avoid conflicts or arbitration. As central as it is, the mesh is basically invisible, which is the way programmers like it.

Scalability ultimately wins

As with most ecosystems, many different kinds of processors will survive. Some will thrive, while others will barely eke out a living in some specific niche. Outside forces will cull the herd, as has happened with graphics and network processors, winnowing those that don't fit the current environment.

For the past few decades, scalability and programmability have been the keys. Developers want a chip they can understand and stick with for successive generations. They want a roadmap for growth, both up and down the price/performance scale. And making it really, really fast doesn't hurt, either. ECD

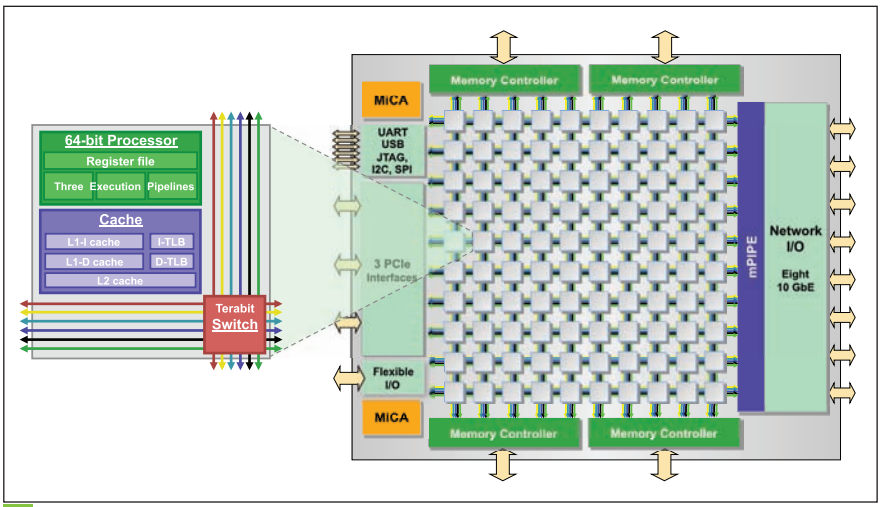

Figure 2 | One of Tilera's TILE-Gx processor cores, where n is 16 to 100, efficiently handles bandwidth between and among cores. Each core tile has its own 64-bit processor, L1 and L2 caches, and network connections to four neighbors in the north/south/east/west directions.

Bruce Edwards *is a Principal Architect at Tilera Corporation and has more than 25 years of experience in processor design. He previously worked on microprocessors at Symbolics and Digital Equipment Corporation and designed multiple generations of wireless networking systems at Broadcom. He has a BS in Electrical Engineering from MIT.*

Tilera
508-616-9300 ext. 242
bedwards@tilera.com
www.tilera.com

Innovative Intelligent Integration

SmartFusion™

FPGA + ARM®Cortex™-M3 + Programmable Analog

© 2010 Actel Corporation. All rights reserved.

Actel POWER MATTERS

Get Smart, visit: www.actel.com/smartfusion

Visit us at ESC Silicon Valley, Booth #930

Q&A: Getting to 100G with NPUs

Niel Viljoen, CEO, Netronome Systems

ECD: A lot of people are talking about 10G, 20G, 40G, and 100G packet processing. Where are you now, and where do you see technology going in the next year?

VILJOEN: Network traffic in enterprise and carrier networks continues to rise with today's bandwidth requirements at 10 Gbps and increasing to 40 Gbps and 100 Gbps in the coming years. It is no longer sufficient for computing network infrastructures to forward packets blindly, even at those incredible throughputs. Instead, our communications network must be capable of intelligently and securely processing every piece of data to ascertain its nature rather than just its destination. This intelligence is required due to the explosion in new IP-based applications, and is needed to counter the security threats that we face.

As the demand for intelligent networking applications at 10 Gbps to 100 Gbps increases, general-purpose CPU architectures struggle to keep pace with these extreme performance requirements. To meet these needs, Netronome's family of Network Flow Processors (NFP-32xx) are powered by 40 programmable networking cores each operating at 1.4 GHz to deliver 2,000 instructions per packet at 30 million packets per second, enabling up to 40 Gbps of L2-L7 processing with line-rate security and I/O virtualization per chip. In addition, Netronome is designing the follow-on to the NFP-3240, which will offer even greater throughputs while retaining the important L2-L7 programmability and integrated hardware-based cryptography engines.

ECD: How does your current hardware architecture, with 40 cores as you mentioned, fit into real-world problems in IP-based wireless networks? Where have you seen other NPU architectures fall short in this application?

VILJOEN: This explosion in bandwidth and the need for increasing content awareness applies not only to fixed LAN/WAN-based networks, but increasingly to mobile networks as well. With the bandwidths that 3G wireless and long-term evolution networks offer, along with the converged data, voice, and video services that users utilize over these networks, wireless infrastructure needs to support all of the same services as fixed networks, which means it also shares vulnerabilities to the same types of threats. Support for IP-based wireless services requires massively parallel, highly programmable systems. Accordingly, the NFP-3240 contains 40 programmable RISC cores optimized for networking and security with unparalleled performance per watt.

In contrast, general-purpose processors excel at application and control plane processing but stall when used for high-throughput networking tasks, security processing, and deep packet inspection. Traditional network processors only support processing traffic at L2-L3, typically support a pipeline rather than parallel architecture, and are configurable rather than programmable.

The optimal architecture for IP-based wireless networks and services utilizes heterogeneous multicore processing coupling the NFP-3240 microengine cores with the performance and scalability of general-purpose multicore x86 systems over a high-speed, virtualized PCI Express data path via enhanced I/O virtualization.

ECD: What is the key technology in your software solution, and how will you compete with NPU vendors buying Linux and middleware vendors?

VILJOEN: Netronome's processors are optimized for high-performance programmable data planes for L2-L7 packet processing, I/O virtualization, and security. Our strategy is to have application and control plane processing execute on x86 cores, the industry's highest-performance general-purpose processing architecture.

The NFP is specifically designed to be a network coprocessor to x86, not replace it. This architecture accelerates applications to 100 Gbps while offering the comfort and familiarity of x86. As such, our software focus is to provide an extensive library of software for network flow processing and cryptography and rely on our partners for operating system and middleware solutions.

Netronome's expansive software library support spans low-level functionality like Rx/Tx code, memory operations, and bulk cryptography to high-level abstractions like PCI Express with I/O virtualization, switching, routing, packet classification, load balancing, stateful flow processing, and TCP-related operations. We also make these libraries available as full software distributions exposed as APIs to host operating systems.

***Niel Viljoen** is CEO of Netronome Systems, with more than 20 years of experience focusing on IP, ATM, security, and system design. His experience includes angel investing in Nujira, Intune, and Azuro; CTO of the Marconi Group and VP with FORE Systems; and CEO of Nemesys Ltd. He attended the University of Stellenbosch and Cambridge University.*

Netronome Systems
724-778-3290 | www.netronome.com

A new era, a new ESC

New tracks. New courses. New focus.

ESC Silicon Valley brings together systems architects, design engineers, suppliers, analysts, and media from across the globe. With cutting edge product demonstrations, visionary keynotes, and hundreds of essential training classes, ESC is the ideal conference for the embedded design community to learn, collaborate, and recognize excellence.

ESC Silicon Valley 2010 features:

- Two "Build Your Own Embedded System" programs featuring the new Intel Multicore Atom and the Freescale Tower development system
- Multicore Expo and Conference co-located with ESC Silicon Valley
- ARM Partner Pavilion
- Display and Connected Devices Emerging Technology Zones
- Keynotes from Dr. Michio Kaku, theoretical physicist and bestselling author and Richard Templeton, CEO & President, Texas Instruments

Start your own personal development at ESC 2010. You can't afford to miss it.

Silicon Valley
McEnery Convention Center, San Jose, April 26 - 29, 2010
Conference: April 26 - 29, 2010
Expo: April 27 – 29, 2010

Register Today.
www. embedded.com/sv

Learn today. Design tomorrow.
ESC

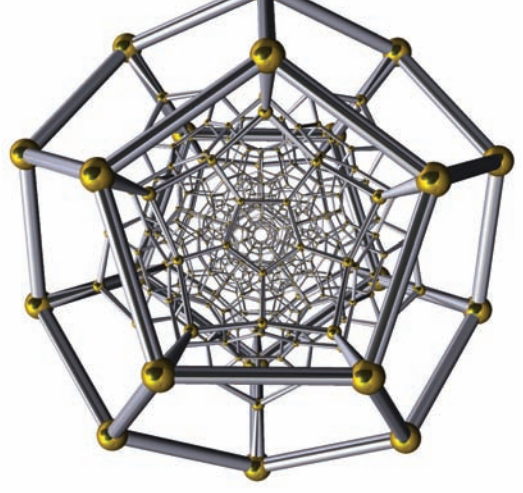

Performance management: A new dimension in operating systems

By Val Popescu

Given the increased complexity of processors and applications, the current generation of Operating Systems (OSs) focuses mostly on software integrity while partially neglecting the need to extract maximum performance out of the existing hardware.

Processors perform as well as OSs allow them to. A computing platform, embedded or otherwise, consists of not only physical resources – memory, CPU cores, peripherals, and buses – managed with some success by resource partitioning (virtualization), but also performance resources such as CPU cycles, clock speed, memory and I/O bandwidth, and main/cache memory space. These resources are managed by ancient methods like priority or time slices or not managed at all. As a result, processors are underutilized and consume too much energy, robbing them of their true performance potential.

Most existing management schemes are fragmented. CPU cycles are managed by priorities and temporal isolation, meaning applications that need to finish in a preset amount of time are reserved that time, whether they actually need it or not. Because execution time is not safely predictable due to cache misses, miss speculation, and I/O blocking, the reserved time is typically longer than it needs to be. To ensure that the modem stack in a smartphone receives enough CPU cycles to carry on a call, other applications might be restricted to not run concurrently. This explains why some users of an unnamed brand handset complain that when the phone rings, GPS drops.

Separate from this, power management has recently received a great deal of interest. Notice the "separate" characterization. Most deployed solutions are good at detecting idle times, use modes with slow system response, or particular applications where the CPU can run at lower clock speeds and thus save energy. For example, Intel came up with Hurry Up and Get Idle (HUGI). To understand HUGI, consider this analogy: Someone can use an Indy car at full speed to reach a destination and then park it, but perhaps using a Prius to get there just in time would be more practical. Which do you think uses less gas? Power management based on use modes has too coarse a granularity to effectively mine all energy reduction opportunities all the time.

Ideally, developers want to vary the clock speed/voltage to match the instantaneous workload, but that cannot be done by merely focusing on the running application. Developers might be able to determine minimum clock speed for an application to finish on time, but can they slow down the clock not knowing how other applications waiting to run will be affected if they are delayed? Managing tasks and clock speed (power) separately cannot lead to optimum energy consumption. The winning method will simultaneously manage/optimize all performance resources, but at a minimum, manage the clock speed and task scheduling. Imagine the task scheduler being the trip planner and the clock manager as the car driver. If the car slows down, the trip has to be re-planned. The driver might have to slow down because of bad road conditions (cache misses) or stop at a railroad barrier (barrier in multithreading, blocked on buffer empty due to insufficiently allocated I/O bandwidth, and so on). Applications that exhibit data-dependent execution time also present a problem, as the timing of when they finish isn't known until they finish. What clock speed should be allocated for these applications in advance?

An advanced performance management solution

One example of managing performance resources is VirtualMetrix Performance Management (PerfMan), which controls all performance resources by a parametrically driven algorithm. The software schedules tasks, changes clock speed, determines idle periods, and allocates I/O bandwidth and cache space based on performance data such as bandwidth consumed and instructions retired. This approach (diagrammed in Figure 1) solves the fragmentation problem and can lead to optimum resource allocation, even accounting for the unpredictability of the execution speed of modern processors and data-dependent applications.

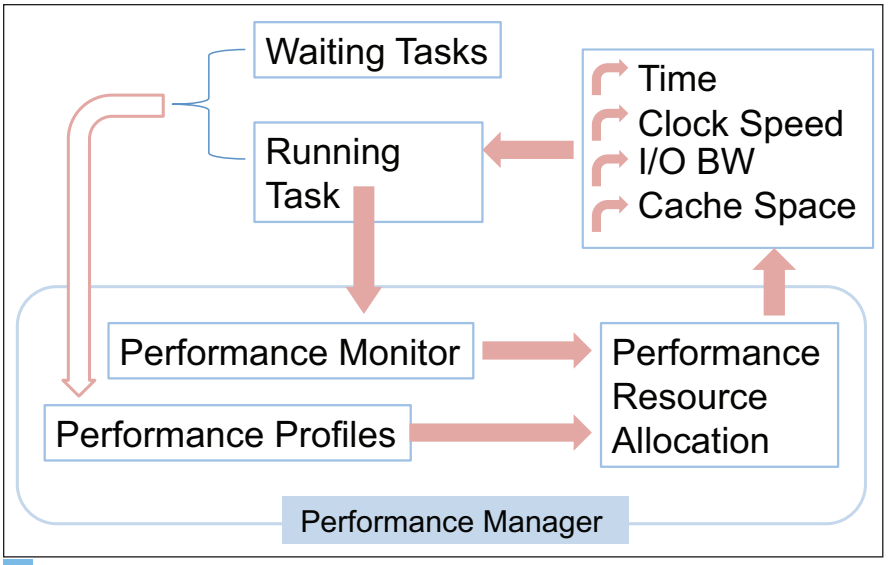

Figure 1 | PerfMan controls all performance resources using a parametrically driven algorithm, leading to optimum resource allocation.

The patent-pending work performed allocation algorithm uses a closed-loop method that makes allocation decisions by comparing work completed with work still to be performed, expressed in any of the measurable performance quantities the system offers. For example, if the application is a video player or communication protocol that fills a buffer, PerfMan can keep track of the buffer fill level and determine the clock speed and time to run so that the buffer is filled just in time. The time to finish will inevitably vary, so the decision is cyclically updated. In many cases, buffers are overfilled to prevent blocking on buffer empty, which can lead to timing violations. PerfMan is capable of precise performance allocation, keeping buffering to a minimum and reducing memory footprint. The algorithm can handle hard, soft, and non-real-time applications mixed together.

If the application execution graph is quantified into simple performance parameters and the deadlines are known when they matter, the algorithm will dynamically schedule to meet deadlines just in time. Even non-real-time applications need some performance allocation to avoid indefinite postponement. Allocating the minimum processor resources an application needs increases system utilization, resulting in a higher possible workload. The method does not rely on strict priorities, although they can be used.

2nd Generation Intel® Core™ 2 Duo processor

Unbeatable performance. Unbeatable expandability. Take a look at the 2nd generation Intel Core 2 Duo processor-based 2.27 GHz embedded computer. The "Komodo" features VersaLogic's legendary quality, extreme performance, and reliability in even tough environments. The industry-standard SUMIT™ expansion site supports 2nd generation speeds and legacy expansion.

- Very high performance
- 2.27 GHz Intel Core 2 Duo processor
- Up to 4 GB RAM
- 3D video accelerator
- Low power: 15–20W typical
- SUMIT-EPIC expansion with PCIe and ISA
- Extended temp. (-40º to +85ºC) version
- Compact EPIC (4.5" x 6.5") footprint

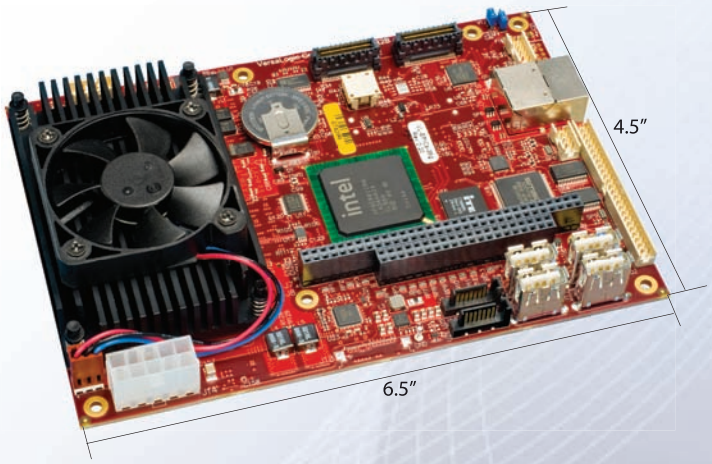

Includes VersaLogic's long-term (5+ year) product availability guarantee. Customization available on quantities as low as 100 pieces!

Delivering extraordinary support and on-time delivery for more than 30 years, VersaLogic has perfected the fine art of service, one customer at a time. Experience it for yourself. Call **800-824-3163** for more information.

Recipient of the VDC Platinum Vendor Award for five years running!

VDC research 2008 PLATINUM VENDOR

VERSALOGIC CORPORATION

800-824-3163 | 541-485-8575 | www.VersaLogic.com/kom

PC/104-*Plus* SBC Supports Industrial Control, HMI, Networking and I/O

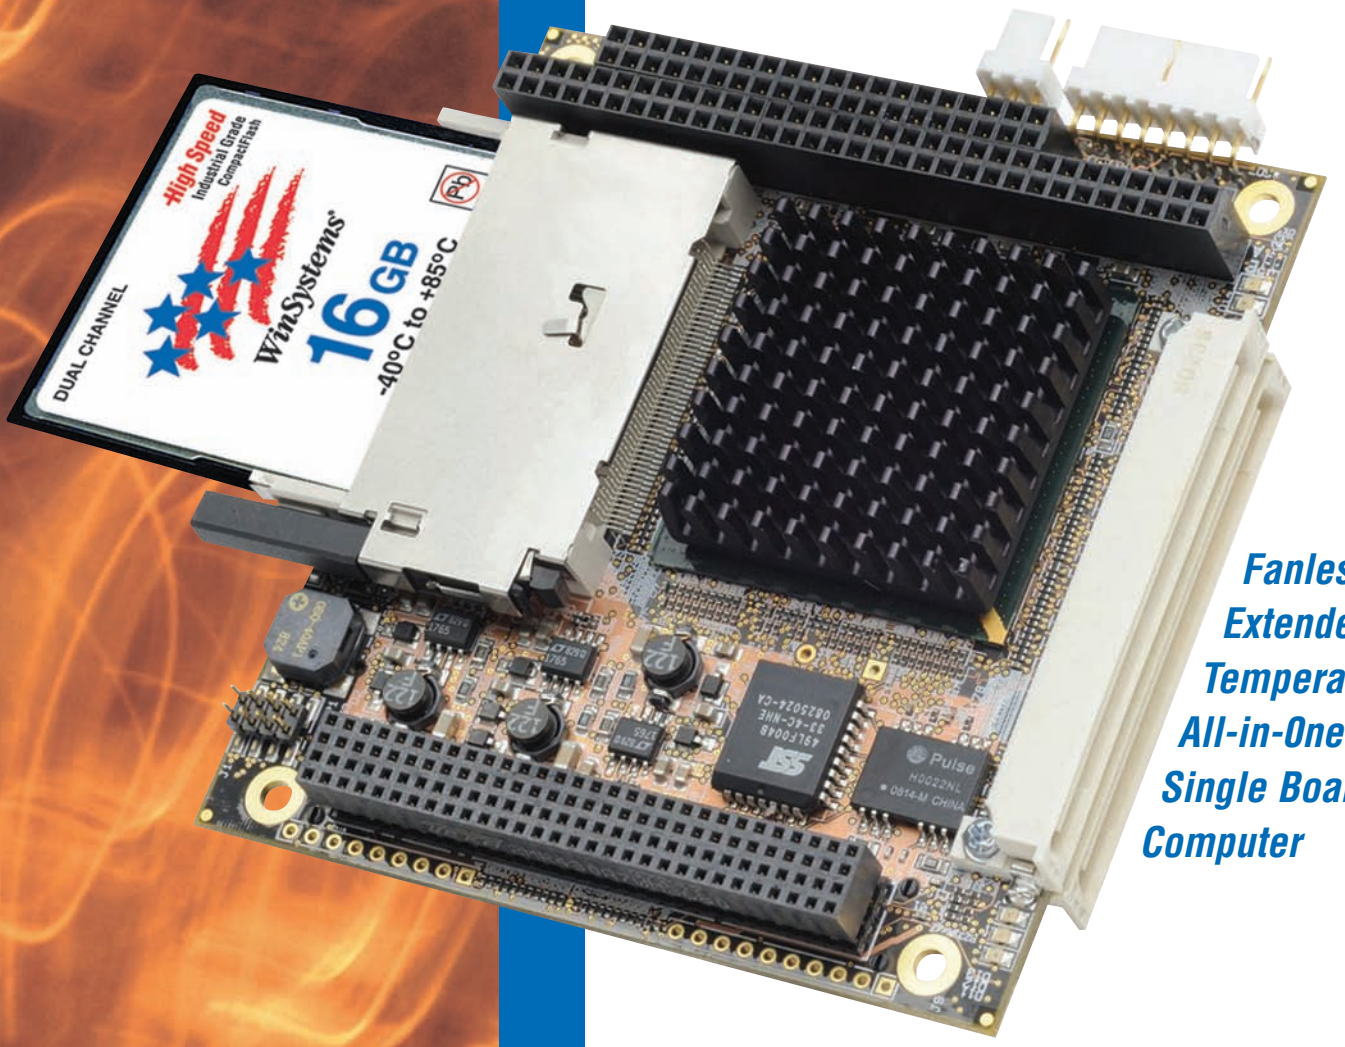

Fanless Extended Temperature All-in-One Single Board Computer

The PPM-LX800 SBC is available on a 6.5-inch Open Frame Panel PC.

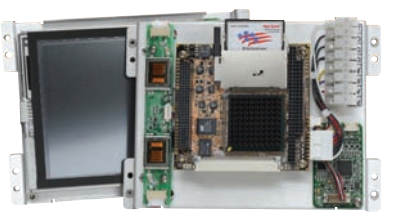

If your application needs a compact, industrial PC with fanless -40°C to +85°C operation, check out WinSystems' PPM-LX800.

Designed for embedded, space-limited, and low-power applications requiring reliability and long-term availability, the PPM-LX800 runs Linux, Windows® XP embedded, and other x86 real-time operating systems. Plus it comes with a wealth of on-board I/O functions and free software drivers.

Features include:

- AMD LX800; x86-compatible CPU
- Compact size, 3.6" x 3.8" (90mm x 96mm)
- Socket for a bootable CompactFlash
- High-resolution video controller provides
  - CRT resolutions up to 1920 x 1440
  - Panel resolutions up to 1600 x 1200
- Custom splash screen on start up
- 10/100 Mbps Intel Ethernet controller
- Four serial ports with FIFO
- PS/2 keyboard and mouse controller
- +5V only operation
- Ultra DMA100 EIDE controller for 1 or 2 devices
- PC/104 and PC/104-*Plus* expansion connectors
- AC97 Audio with MIC, Line In, and Line Out
- -40°C to +85°C extended operational temperature
- Long-term product availability
- Responsive and knowledgeable technical support
- Quick Start Kits offered for easy software development

Contact our factory application engineers for additional product information, custom configurations, and pricing.

Call 817-274-7553 or
Visit www.winsystems.com/PPM-LX800

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

M2M 100 COMPANIES 2010 · Linux · Microsoft Windows Embedded Partner · Embedded PC/104 Modules · Extended Temperature · RoHS compliant

The priority or order in execution is the direct result of the urgency the application exhibits while waiting its turn to run, which is a function of the basic work to be performed/worked completed paradigm.

Extending to more dimensions

If tasks are ready to run in existing OSs, they will run, but do they need to? Can they be delayed (forced idling) if the OS knows it will not affect their operation?

Knowing the timing of every task and whether it is running or waiting to run with respect to its progress toward completion allows the software to automatically determine the minimum clock speed and runtime. Thus everything completes on time under all load conditions. Matching clock speed to the instantaneous workload does not mean the clock speed is always minimized. The goal of low energy consumption sometimes calls for a burst of high speed followed by idle, as in Intel's HUGI. But even then, there is no benefit in running faster than the optimum utilization (executed operations per unit of time) would indicate. Fast clocking while waiting for memory operations to complete does not save energy.

The algorithm's mantra of "highest utilization/workload at the lowest energy consumption" is largely accomplished with a closed-loop algorithm managing all performance resources.

In multicore systems, a balanced load, low multithreading barrier latency, and the lowest overall energy consumption cannot be achieved simultaneously. To resolve this, PerfMan can be configured to optimize one or several performance attributes. If minimum energy consumption is the goal, an unbalanced system with some cores that are highly loaded and others that are empty and thus shut down might offer the lowest energy consumption at the expense of longer execution latency and overall lower performance.

Accelerating threads to reduce barrier latency can also lead to higher energy consumption. However, meeting deadlines (hard or soft) overrides all other considerations. The precise closed-loop-based performance resource allocation algorithm can safely maintain a higher workload level, which in turn, allows pushing the core consolidation further than possible with existing methods and thus achieving higher energy reduction.

Implementation on VMX Linux

PerfMan has been implemented as a thin kernel (sdKernel) running independently of the resident OS. It has been ported to Linux 2.6.29 (VMX Linux), as shown in Figure 2. An Android port is nearing completion. The software takes over

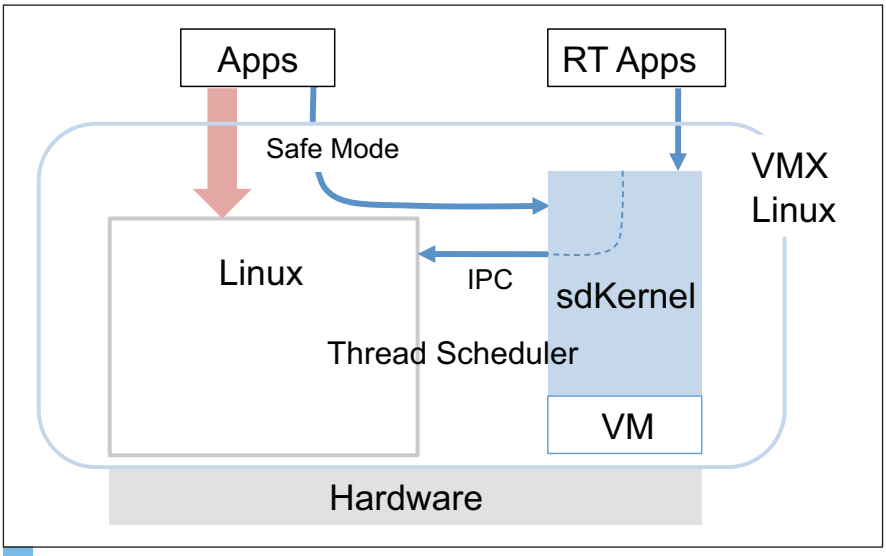

Figure 2 | In a Linux implementation, PerfMan takes over Linux task scheduling and interworks with the existing power management infrastructure.

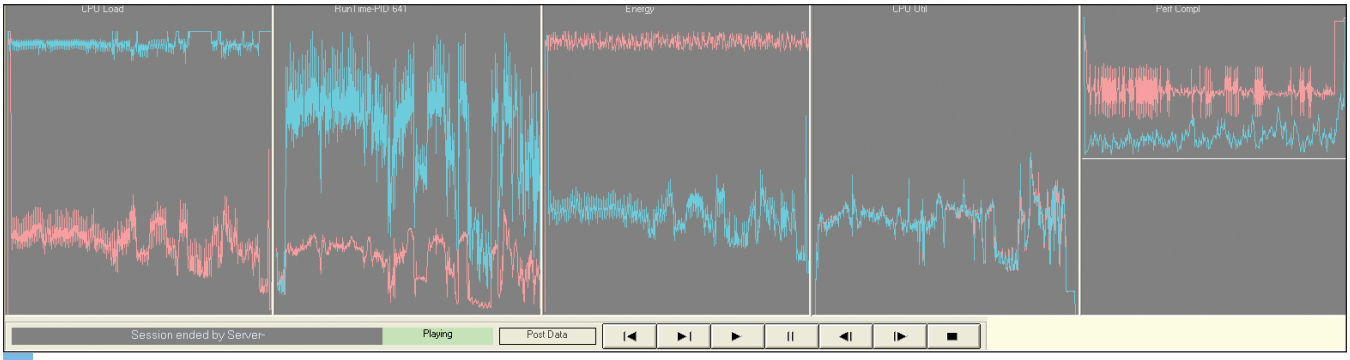

Figure 3 | Using VMX Linux on an OMAP35xx BeagleBoard achieves a 95 percent average load that finishes just in time.

Linux task scheduling and interworks with the existing power management infrastructure. A separate version of the sdKernel provides virtualization and supports hard real-time tasks in a POSIX-compliant environment. Scheduling/context switching is at the submicrosecond level on many platforms, but because most Linux system calls are too slow for hard real-time applications, the sdKernel provides APIs for basic peripherals, timers, and other resources.

By monitoring performance, the software can detect unusual execution patterns that predict an upcoming OS panic and crash. In such cases, the sdKernel will notify mission-critical applications to stop using Linux system calls and temporarily switch over to sdKernel APIs (safe mode) while Linux is being rebooted.

VMX Linux supports a mix of real and non-real-time applications with efficient performance isolation while minimizing energy consumption. It can also provide hardware isolation/security and safe crash landing.

Benchmarks show the results

The energy consumption, measured in real time using a VMX-designed energy meter, was accumulated for the system and correlated to individual applications. A media player application (video and audio) was run on an OMAP35xx BeagleBoard first using standard Linux 2.6.29 (Figure 3 red graph) and then VMX Linux (Figure 3 blue graph).

Performance compliance (Perf Compl graph) shows how close the application tasks come to finish on time (center line). Below the line indicates deadline violations. Notice that with VMX Linux, a 95 percent average load is achieved with no prebuffering and no deadline violations, but it gets close. The total board energy consumption for the 46 seconds of video dropped from 68.7 W*sec to 27.6 W*sec with VMX Linux. The displayed data represents averages over a preset interval. As an additional bonus, when Linux is purposely crashed, the video disappears but the music plays on in safe mode with no audible glitches.

In short, the implementation creates a new approach to performance management with exciting results. ECD

Val Popescu *is president and CEO of VirtualMetrix, which he founded in 2005. He was founder and CEO of Metaflow, a microprocessor design company that pioneered the out-of-order/speculative execution adopted in many high-performance microprocessors including the ARM Cortex-A9. Val holds nine patents on the subject.*

VirtualMetrix
858-395-5793 | val@virtualmetrix.com
www.virtualmetrix.com

- Out of the box solutions
- Extensive product line
- Reliable, proven code
- Free evaluation kits
- No royalties
- Full source code

SMX® RTOS	WiFi	USB Device
BSPs	TCP/IP	USB Host
Device Drivers	FAT File System	USB OTG
Kernel Awareness	Flash File System	USB Class Drivers
Simulator	GUI	Floating Point

Free Evaluation Kits: **www.smxrtos.com/eval**
Free Demos: **www.smxrtos.com/demo**

Embedded Systems Conference ESC Silicon Valley • April 26-29, 2010 BOOTH #939

Micro Digital 35TH ANNIVERSARY
800.366.2491 sales@smxrtos.com

www.smxrtos.com

ARM • ColdFire • Cortex-M3 • PowerPC • x86

Real-time and general-purpose operating systems unite via virtualization

By Chris Main

Virtualization for embedded systems has many implementations in which two or more operating systems coexist to gain the benefits of each. One approach puts Microsoft Windows and a Real-Time Operating System (RTOS) together.

Much is being said about virtualization these days in the software world. Simply stated, virtualization is about getting multiple OSs to run on the same computing platform at the same time. Virtualization has been cited as a key technology for getting the most performance out of the newest multicore processors. But just as not all computing applications are the same, not all virtualization approaches are appropriate for all applications.

Embedded systems have a key requirement that doesn't normally apply to office and server computers: the need for

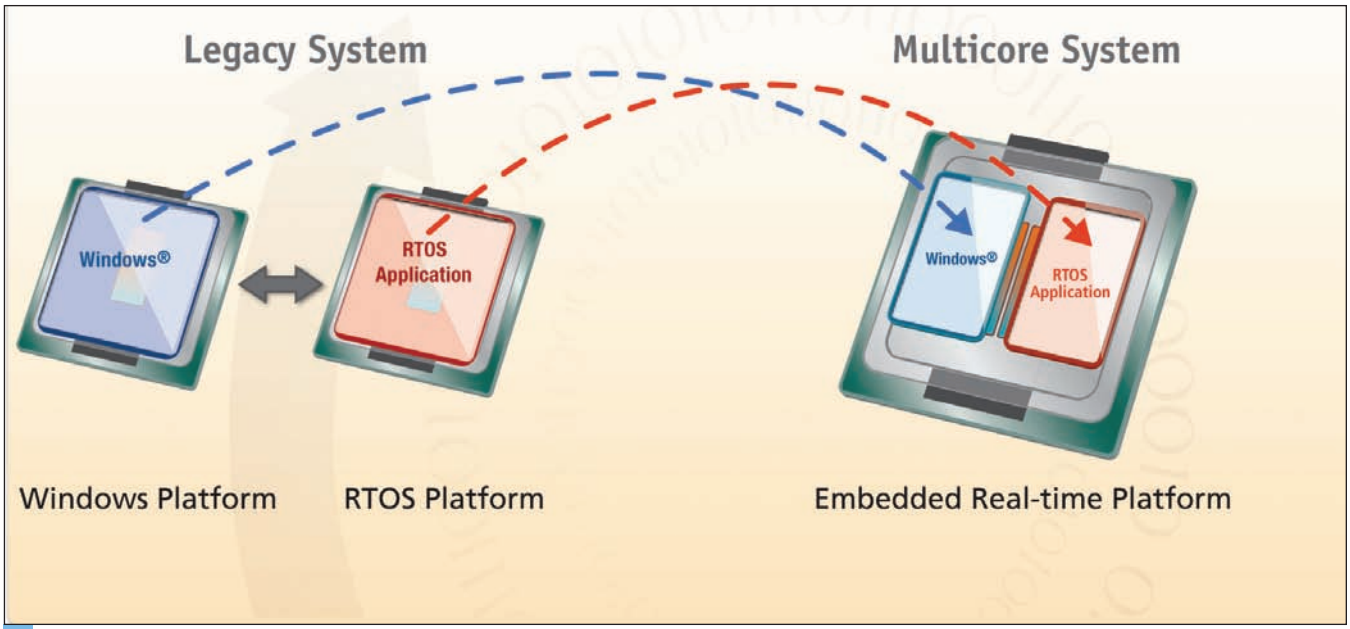

Figure 1 | Embedded virtualization combines real-time and general-purpose processing on a single platform.

deterministic response to real-time events. To support the requirement for determinism, embedded applications typically use RTOSs. Embedded applications also employ general-purpose OSs to handle operator interfaces, databases, and general-purpose computing tasks.

In the past, because OSs couldn't successfully co-reside on computing platforms, system developers employed multiple processing platforms using one or more to support real-time functions and others to handle general-purpose processing. System designers that can combine both types of processing on the same platform can save costs by eliminating redundant computing hardware. The advent of multicore processors supports this premise because it is possible to dedicate processor cores to different computing environments; however, the software issues posed by consolidating such environments require special consideration. Combining real-time and general-purpose operating environments on the same platform (Figure 1) places some stringent requirements on how virtualization is implemented.

Paravirtualization

A common approach to supporting multiple OSs, on an embedded platform is to employ operating software called a hypervisor that boots first on power-up, then loads the OSs. The hypervisor is in charge of the platform and handles the memory partitioning and other processor resources between the OS environments. If the CPU doesn't support Intel's VT-x hardware extensions for virtualization, modifications must be made to the guest OSs, so they know that they're supported by a hypervisor. They must carry out their own address translation to avoid conflict with memory belonging to other OSs.

A situation such as this, where the guest OS is aware that it is a guest, is called paravirtualization. Though the scheme might work, the main downside to this approach is that developers can't use a standard out-of-the-box OS as a guest. This is the virtualization approach that the Xen hypervisor takes. Modifying a third-party OS can lead to major support

SEALEVEL®

Computing/HMI | Serial | I/O

The right connection creates incredible power.

Connectivity and control. Making your interface as reliable as the tides—and just as strong. Sealevel creates hardware and software solutions for both digital and serial interface requirements.
We Listen. Think. And Create.

Experience exceptional computing at low power with the SBC-R9, an application-ready platform for your next product design.

sealevel.com > sales@sealevel.com > 864.843.4343

© 1986-2010, Sealevel Systems, Inc. All rights reserved.

headaches, and the cost of supporting such a product over its long-term life cycle through upgrades and updates is significant.

Server virtualization

To minimize the changes that need to be made to a guest OS so that it can run in a virtual environment, developers can build a hypervisor that emulates the entire machine, providing each guest OS with what it thinks is exclusive access to the processor platform. This approach is being used today on multicore processor systems to run multiple instances of the same application.

While this approach might be satisfactory for server applications that are optimized to keep the processors busy, it is unacceptable for embedded applications that need to ensure the fastest possible response to external events. Currently available software products of this type are not optimized for Intel Architecture processors or a particular class of applications. Every time a new OS configuration needs to be supported, significant work is required to tailor the hypervisor to work with the selected OSs.

Hardware-aided embedded virtualization

The best type of virtualization for embedded applications involves hardware assist features provided by the processor and platform architecture, which the hypervisor uses to ensure that the system meets the performance goals of the applications to be supported. With embedded virtualization, parts of the machine are emulated to provide a standard environment for multiple OSs, and parts of the system are not emulated to ensure that system performance goals are met. Basic functions such as the PCI bus interface and CMOS registers are emulated, but the I/O interfaces that are necessary for real-time responsiveness are assigned for exclusive use by the guest RTOS.

The embedded hypervisor must partition the machine to separate resources for exclusive use by each OS. TenAsys' eVM for Windows accomplishes this by configuring the Microsoft Windows bootloader to limit the number of hardware threads and memory that it uses. Windows boots first and runs normally with the remaining resources allocated to it. Because it is running on the bare machine rather than an emulating software layer, Windows tasks execute with maximum performance.

Next, the RTOS and real-time application software are loaded into memory allocated for the real-time portion of the application, then the RTOS is booted from a Windows driver and begins executing application code in its isolated hardware environment. With access to its own performance-critical I/O devices, the real-time application will run completely independently of Windows. Special drivers are provided so that the real-time environment can

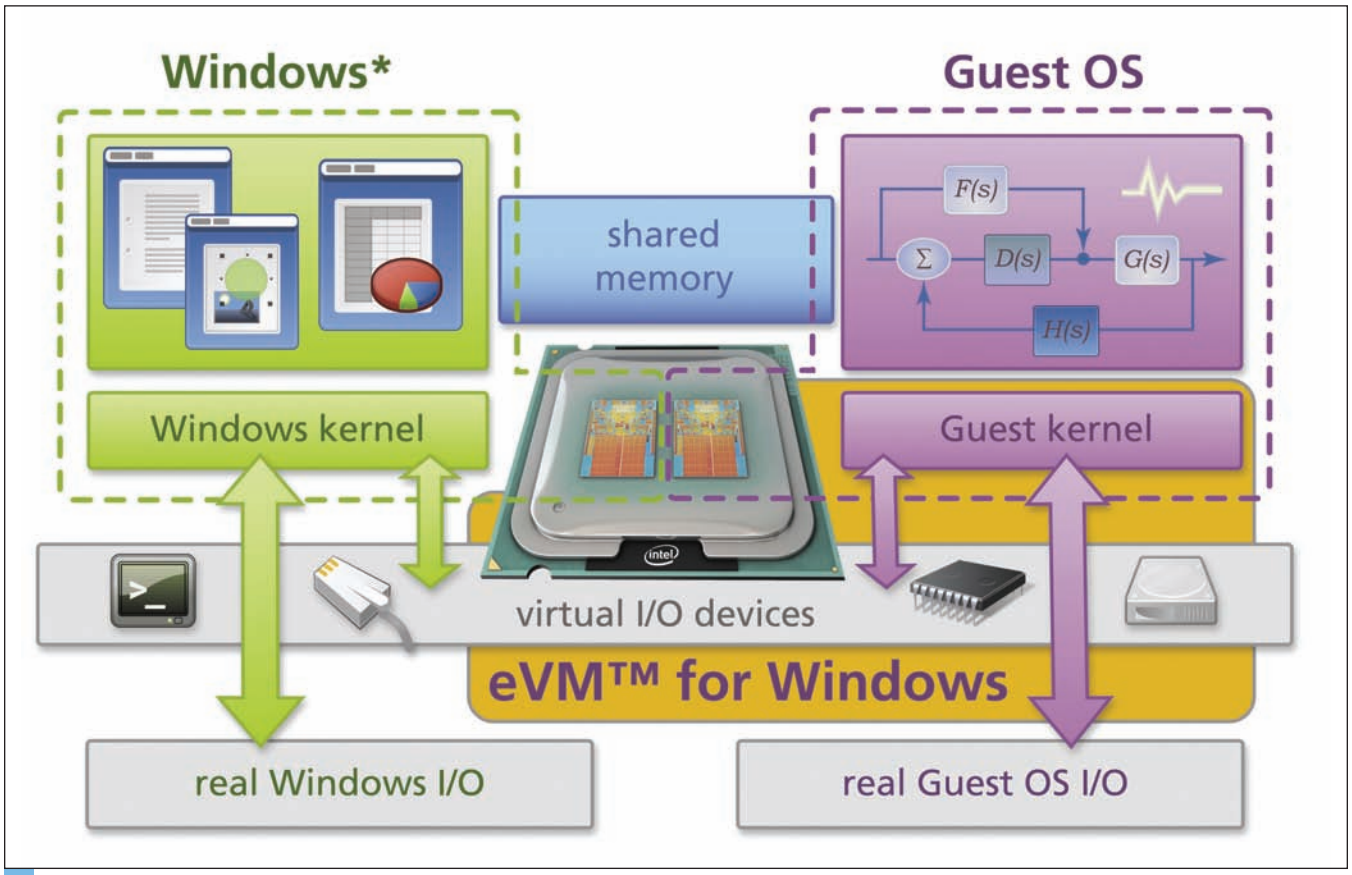

Figure 2 | Windows and real-time tasks communicate via virtual communication channels between OSs.

use Windows resources. When Windows and real-time tasks need to communicate or the real-time application needs to use Windows resources, the two environments communicate via emulated communications links in shared memory such as virtual Ethernet and serial links (see Figure 2). Real-time I/O devices are configured via a generic device driver on the Windows side, and real-time interrupts are configured to be delivered to the virtual machine manager and not to Windows.

Leveraging Intel Architecture

Embedded system developers have made huge investments in intellectual property that they would like to protect. Besides costing money and time to re-implement, changing application code to run in new OS environments can inject errors into applications or cause them to become unreliable.

Efficient embedded virtualization requires hardware support so that a variety of OSs can run without modification. By using the VT-x hardware extensions that Intel provides in its new CPUs and chipsets, an efficient execution platform can be delivered without requiring the guest OS or its drivers to be modified or the system disk to be repartitioned.

New hardware support under development at Intel will enable embedded virtualization with other general-purpose OSs. In addition, Intel is currently developing support for virtualizing the CPU's paging mechanism and interrupt controller. With the new hardware features, guest OSs can be started and stopped independently, and Windows can be rebooted without needing to reboot the entire platform.

Embedded virtualization saves costs

In the embedded market, developers have been accustomed to tailoring operating environments with additional work and expense using multiple processing platforms to guarantee real-time system responsiveness. But the future can be different. Using eVM for Windows to enable system consolidation without requiring software customization can avoid solutions that are not adaptable while enabling solutions that can be implemented quickly at reasonable cost. **ECD**

Chris Main *is CTO of TenAsys Corporation, where he has led development of the company's virtualization technologies. He has worked in real-time systems starting with minicomputers and was involved in the iRMX group at Intel. He was on the original development team for INtime and is a cofounder of TenAsys. He earned a graduate degree in Physics from York University and postgraduate degree in Education from Bath University.*

TenAsys
503-748-4720
www.tenasys.com

Cavium Networks released their latest PureVu video coprocessors to take on the full HD H.264 task in set-top boxes, home media distribution, video conferencing, and other applications.

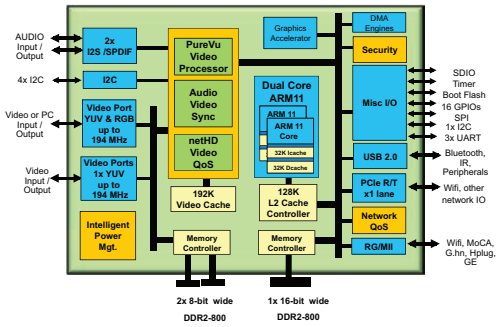

Cavium Networks | **www.caviumnetworks.com**

Using third-generation voice processing technology and sophisticated Computational Auditory Scene Analysis to group and process complex mixtures of sound, the Audience A1026 identifies the primary conversation and pulls it out of a noisy environment to enable clear calls.

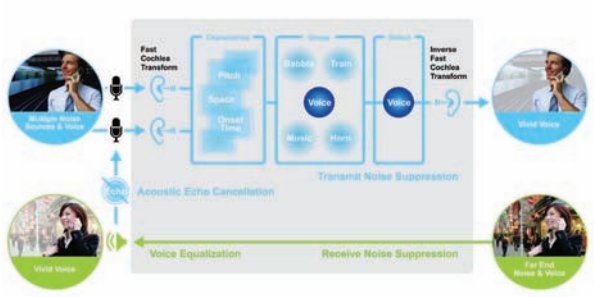

Audience | **www.audience.com**

The Atheros AR1520 and its FYX location core feature high-accuracy GPS tracking at less than 30 mW, low-power tracking at less than 10 mW, and an AlwaysReady mode for "hot start" consumption at less than 1 mW.

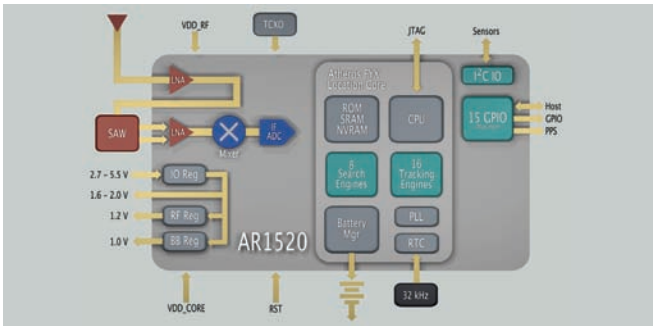

Atheros Communications | **www.atheros.com**

DynaVox Mayer-Johnson advances human expression

By Hilary Tomasson

Designing a small speech-generating device durable enough for children to use called for a small form factor module based on the Intel Atom processor running Microsoft's Windows Embedded Operating System (OS).

DynaVox Mayer-Johnson, a leading provider of Augmentative and Alternative Communication (AAC) products and services, offers complete solutions for individuals with speech, language, and learning challenges. DynaVox is dedicated to helping individuals and families who need alternatives to gain or regain the power of speech. The company is committed to helping individuals challenged by significant speech disabilities make meaningful connections with the world and with those who care for them.

DynaVox speech-generating devices help those living with conditions such as stroke, autism, ALS, brain injury, cerebral palsy, Down Syndrome, or aphasia. The DynaVox Xpress, a battery-powered handheld device (Figure 1), lets those with speech challenges take DynaVox technology with them into the broader society. Powerful communication functionality, small size, and sleek design make mainstreaming possible.

DynaVox's core technology, refined in dozens of products used by tens of thousands of people worldwide, is based on a clear

understanding of an individual user's needs. People can find speech challenging for various reasons, and people of different ages have different requirements.

For example, an adult with autism might have only the most basic understanding of communication, while an adult with ALS might have the cognitive ability to be highly verbal and communicative but lack the muscle control to move lips and vocal cords to form words. An adult with aphasia resulting from a stroke or brain injury might be able to speak words, but not in the appropriate sequence or context for the communication to be functional.

Whatever the speech problem, different ages require different vocabulary. The communication needs of an adult with Down Syndrome and a child with Down Syndrome can be quite dissimilar.

DynaVox approaches this problem by offering an extensive set of content in its devices organized in a matrix called the InterAACT Framework. One axis differentiates communication ability into emergent, context dependent, and independent communicator categories. The other axis is typified by the individual's age-appropriate vocabulary and needs. The InterAACT Framework spans more than 30,000 communication pages, each of which link some form of visual input (such as an icon or text) to some form of speech output. Also, additional logic allows words and images to be assembled into complete sentences and thoughts.

Identifying improvements needed

After refining its core technology for many years, DynaVox faced the challenge of making it more accessible to more people through smaller size and enhanced ergonomics. Shrinking the device to pocket or purse size required more than a mechanical effort.

Figure 1 | DynaVox Xpress is a battery-powered speech-generating device based on technology tailored to an individual user's disabilities, age, and needs.

TRI-M SYSTEMS
proudly distributes
TRI-M ENGINEERING
100Mhz PC/104 Module
MZ104
Featuring the new edition ZFx86 FailSafe® Embedded PC-on-a-Chip
Dual watchdog timers, Phoenix BIOS and FAILSAFE Boot ROM
Extended temperature -40°C to 85°C
TRI-M ENGINEERING
PC/104 VersaTainer
VT-104
The VT104 VersaTainer is a rugged aluminum enclosure that can be used as either a PC/104, PC/104+ or EBX enclosure.
The solid one-piece extruded body provides dual internal shock and vibration protection.
TRI-M ENGINEERING
75 Watt High Efficiency PC/104
HE104-75W
75 Watt output
+5V, +12V, -12V outputs
6V to 40V Dc input range
PC/104 compliant
www.tri-m.com info@tri-m.com
1.800.665.5600
HEAD OFFICE: VANCOUVER
tel: 604.945.9565 fax: 604.945.9566

In addition, DynaVox decided to integrate the latest projected capacitive touch screen technology, which is lighter, thinner, and more scratch resistant than the resistive touch screens used in other communication devices. This technology never needs calibration and responds to an extremely light touch, enabling a gestural interface. DynaVox also wanted to ruggedize the device since a user might drop it. Therefore, the device needed to operate using flash storage instead of a hard disk.

DynaVox planned to incorporate more media technology like MP3 playback and video, which could be used for training as well as entertainment. To assist in mainstreaming, DynaVox identified a wide range of spoken voice options to let users pick their personality and work in many languages. The DynaVox Xpress needed to be Internet-enabled so that the augmented communicator could quickly access new vocabulary and symbols and facilitate maintenance and updates over the Internet.

Technology to meet the challenge

To accomplish these goals, DynaVox needed to create new hardware and software platforms to take advantage of technology advancements. The first decision was to choose an OS.

After considering Linux and other OSs, DynaVox eventually chose Microsoft Windows Embedded Standard 2009, which offers many advantages including access to a wide selection of high-quality text-to-speech voices in many languages. It provides a familiar, robust development environment and offers the scalability necessary to customize the OS size to meet the goal of using smaller, flash-based storage instead of a hard disk.

Early in the Xpress design cycle, DynaVox selected the Intel Atom processor, even before production silicon was available. This decision was made because of the revolutionary claims about the processor's combination of low power operation and good performance, plus support for Windows Embedded Standard 2009. DynaVox had previously met specifications using Intel processors with its V family of products.

While DynaVox could have integrated Intel Atom technology at the chip level, a cost analysis made it evident that it was better to use a packaged Computer-On-Module (COM) such as the Eurotech Catalyst Module (Figure 2) based on the Atom processor and focus the company's design efforts on its device-specific carrier board and basic application/software technology. This modular approach offered DynaVox an easier path to upgrade its products as technology evolves by incorporating the next-generation computer on the module rather than redesigning the boards.

In addition to offering an Intel Atom-based hardware platform, Eurotech worked with DynaVox to customize the BIOS and other supporting software, an important concern for a smaller company like DynaVox.

The result

DynaVox introduced the Xpress to the market in August 2009. The device has a simple user interface that lets users

select words, ideas, and full sentences so they can easily interact with others. It is small, portable, easy to use, and discreet, so it can be used in virtually any environment.

The new system is well received by users and is funded by Medicare, Medicaid, and private insurance. DynaVox has a network of more than 110 sales consultants to help medical professionals, individual users, and families with children select and support a communication device.

Many things have changed at DynaVox since it was founded in 1983, but one thing that remains constant is the company's commitment to bring the gift of communication to people with speech challenges across the globe. ECD

Hilary Tomasson *is VP of marketing for Eurotech, based in Columbia, Maryland, where she leads product management and marketing communications initiatives. She formerly led BroadSoft's corporate marketing activities as the company grew to be one of Deloitte's Top 100 Fastest Growing Companies in North America in 2006. Hilary also held leadership positions at Spirent Communications in the areas of marketing, communications, product marketing, and sales operations. She has a BS from Clarkson University and an MS from the University of Maryland.*

Eurotech
301-490-4007
www.eurotech-inc.com

ECD in 2D:

For more on how DynaVox helps children express themselves – use your smartphone, scan this code, watch a video: http://bit.ly/cCCplO, or visit http://bit.ly/ecdin2d_apr10

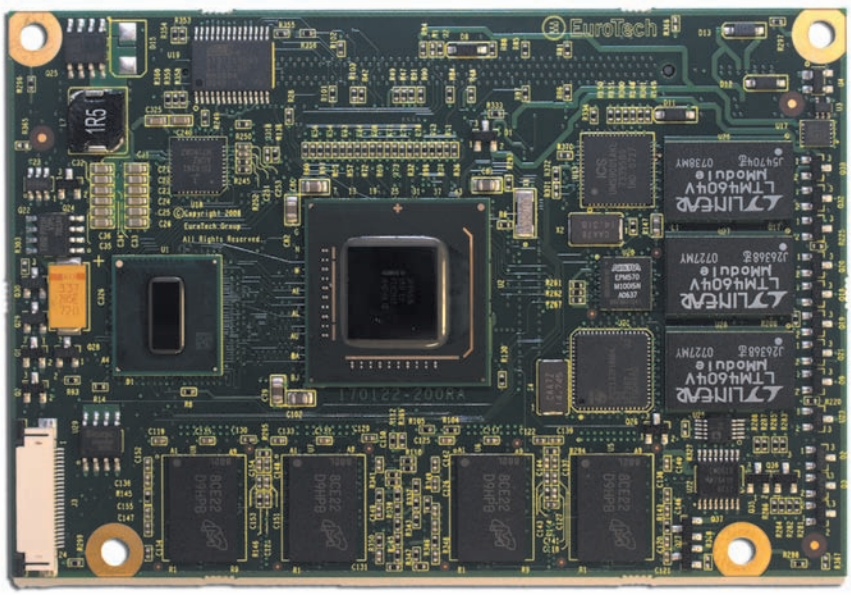

Figure 2 | Using a COM like the Eurotech Catalyst Module gave DynaVox an easier, more cost-effective means to upgrade its products as technology evolves.

TS-8100
Ultra Reliable
w/ 128MB ECC RAM

Qty 1 $269 Qty 100 $229

shown w/ optional SD Card

400 MHz PowerPC
w/ Floating Point Unit

- POE ready
- Dual-execution unit
- Double-precision FPU
- Multifunctional PC/104 connector
- 12K LUT customizable FPGA
- 512MB NAND Flash
- 1 USB Host, 1 USB Device (12 Mb/s)
- Boots Linux 2.6 in < 2 seconds
- Fanless < 4W, sleep mode < 1mW
- Regulated 5-28V power input
- 5 10 bit ADC
- 4 COM ports
- SPI & DIO
- RS485/RS422
- 2 SDHC sockets
- 2 10/100 ethernet
- RTC & WatchDog
- 2 DMX Channels

TS-8150 available w/ extra data acquisition features

- Over 20 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

Android and RTOS together: The dynamic duo for today's medical devices

By Matthew Locke

One strategy getting a lot of play today is Android, and it can go beyond the smartphone if applied correctly. Bringing the strengths of Android together with the benefits of a Real-Time Operating System (RTOS) can yield interesting results for medical devices.

Android has steadily gained momentum in the mobile handset arena, but that's only one side of the story. While the Android software platform is ideal for handsets, it's

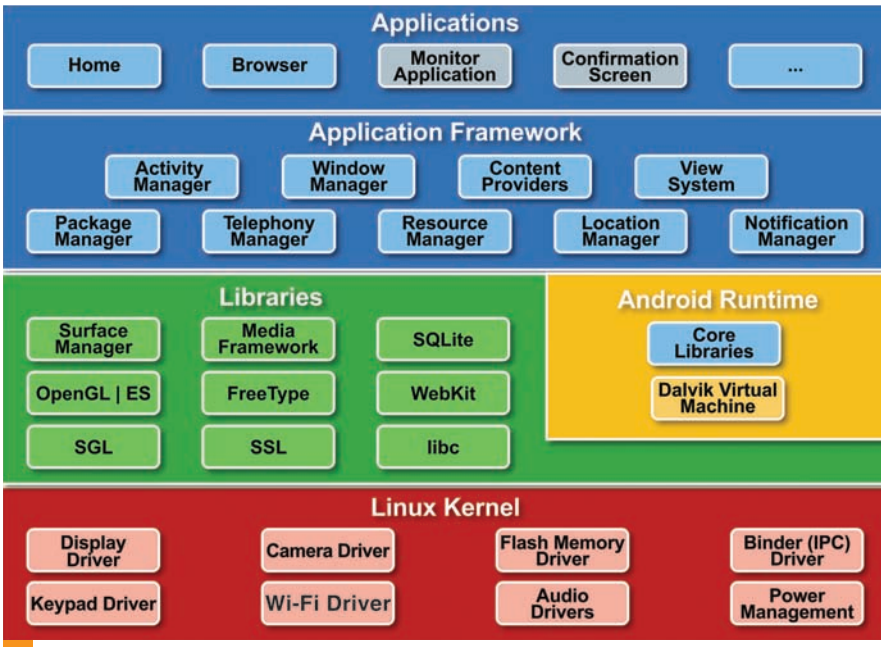

Figure 1 | In the Android system architecture, medical applications for the monitor and confirmation screen can be added at the application layer.

also an extremely good fit for other types of devices that require wireless connectivity, graphical data displays, and intuitive user interactions (see Figure 1).

The emergence of remote patient monitoring devices

One such device category is medical devices, which are becoming more complex by the day. Within this sector, remote patient monitoring is gaining serious traction as a legitimate telemedicine technology. These types of devices are essentially clinical and medical data-gathering systems designed to remotely improve diagnostics, administer health screenings, and deliver therapeutic capabilities. Examples include glucose analyzers, infusion pumps, and cardio monitors. These devices can be used in the home, at the doctor's office, or in the hospital, and each has its own specialized use case and user experience requirements.

Consider the software requirements for this type of device. A remote patient monitoring device typically comprises two components. The first component is connected to the patient for data collection. The second component is used to display the data and configure the system. It is not uncommon for the device to have requirements that go beyond what Android can offer. This is where the capabilities of an RTOS are needed (Figure 2). When working together, Android and an RTOS deliver a powerful combination of embedded efficiencies that make health care specialists and patients both extremely satisfied users.

A flexible, intuitive UI

A health care professional using a patient monitoring device will need to access data quickly and intuitively. The latest trends in User Interface (UI) technology for mobile phones – gesture, cover flow, and widget-based applications – fit nicely with this type of device, as user interaction is extremely important in medical devices. The available Android software developer kit provides a number of options for how the end device and UI can be used. Will the UI have widgets only? Touch-screen buttons? Perhaps a combination of the two? Will the device require a swipe feature as part of the UI? All this can be achieved with the standard APIs available in the Android application framework.

The lightweight determinism of an RTOS

Exactly what is the role of the RTOS? An RTOS is best used in the most safety-critical aspects of devices where real-time data needs to be gathered from the patient. Think of the RTOS as the part of the device that touches the patient. Often, this is achieved by the patient wearing a wristband or belt, or in some cases, noninvasive sensors that attach directly to the skin.

The RTOS can measure functions such as an EKG, heart rate, respiration, and oxygen saturation. The RTOS is also the part of the device that normally requires a small footprint, deterministic latencies, and low overhead. An RTOS not only fits these requirements, but also tends to have fewer lines of code than Android, which makes meeting specific certification or safety regulations a more straightforward process.

Android's many APIs

One of the key reasons to use Android is because its application framework is built around connectivity. Most of the software stack and necessary APIs are already developed. This is especially true with connectivity options such as USB, Bluetooth, Wi-Fi, and others.

Many of today's medical devices demand wireless connectivity, and Android does not disappoint. In fact, Android offers straightforward expansion into new standards such as ZigBee, which is an emerging short-range wireless protocol. The Android software stack can be extended to include a ZigBee driver and stack in the platform layer, all of which gets tied into Android's connection management utility. This connectivity requirement also applies to the RTOS, which, as part of the device connected directly to the patient, often sends the data back to the display unit over a wireless connection. The applications on the RTOS must manage the connection, gather the data, and transmit the data while meeting deadline requirements for each activity. Integrating the connectivity stack with the RTOS is a key piece of the Android/RTOS multi-OS solution.

The ability to extend the I/O is critically important to the development of embedded systems for medical devices. Developers

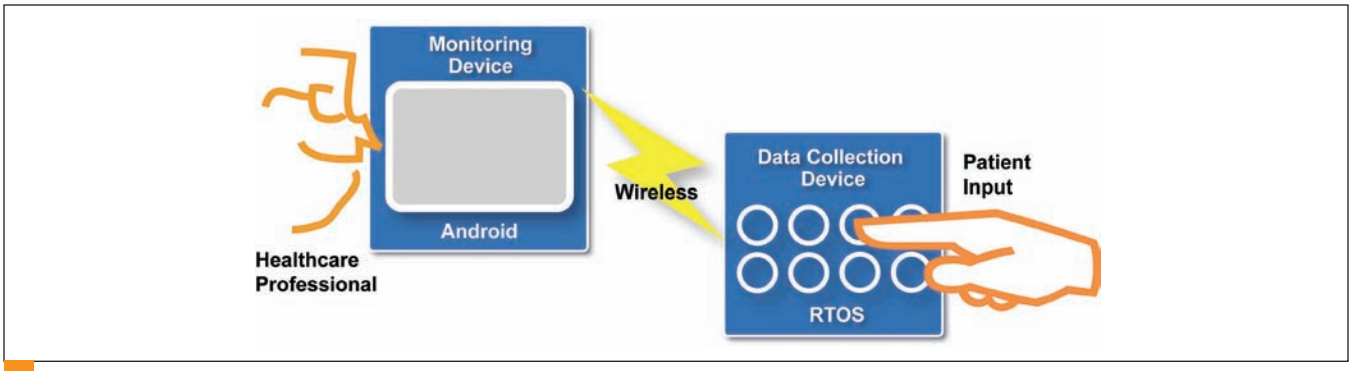

Figure 2 | When developing medical devices for remote patient monitoring, the best approach is to assign Android and user interface capabilities to the system (health care professional) and assign real-time data collection to the RTOS (patient).

must be able to freely add other I/O classes to the framework as new classes become available. Android does a good job of including various drivers, platform code, and the specific application frameworks to make this possible. Granted, the Android application framework might need additional classes as its use spreads to other device types. But as an open source project, developers have all the source code, tools, and freedom to extend Android to meet their target requirements.

Combined to succeed

For a medical device to be successful, it must be able to gather sensitive data in real time and meet all of the necessary safety-critical requirements and certifications. This is no small task.

Having an RTOS and Android work together can alleviate much of the pain. As stated earlier, the RTOS must support all of the same connectivity options as Android. This is a critical requirement for the RTOS. If the connectivity is not already built into the RTOS, it is incumbent upon the developer to make sure it can be added easily when the time comes. It's important to realize that the RTOS not only needs to gather the data quickly, but also support Android in whatever communication/connectivity protocols Android is using. ECD

Matthew Locke *is technical director for Mentor Graphics' Embedded Software Division, with more than 12 years of experience in technology and business leadership roles at Embedded Alley, MontaVista Software, and Lockheed Martin. Prior to Embedded Alley, Matt held CTO positions at several start-ups including NomadGS, which focused on building custom embedded Linux solutions for the consumer device market, and Suunnata, a company developing a system for securely delivering services to networked connected devices. As the product architect for the MontaVista Linux Consumer Electronics Edition, Matt was responsible for the overall product architecture and led the development of key technologies for running Linux on consumer devices. Matthew has a BSEE from the University of Maryland, College Park.*

Mentor Graphics Embedded Software Division
matt_locke@mentor.com
www.mentor.com/embedded

ECD in 2D:

Mentor is working on Android for other segments, too – use your smartphone, scan this code, watch a video: http://bit.ly/9E94VH, or visit http://bit.ly/ecdin2d_apr10

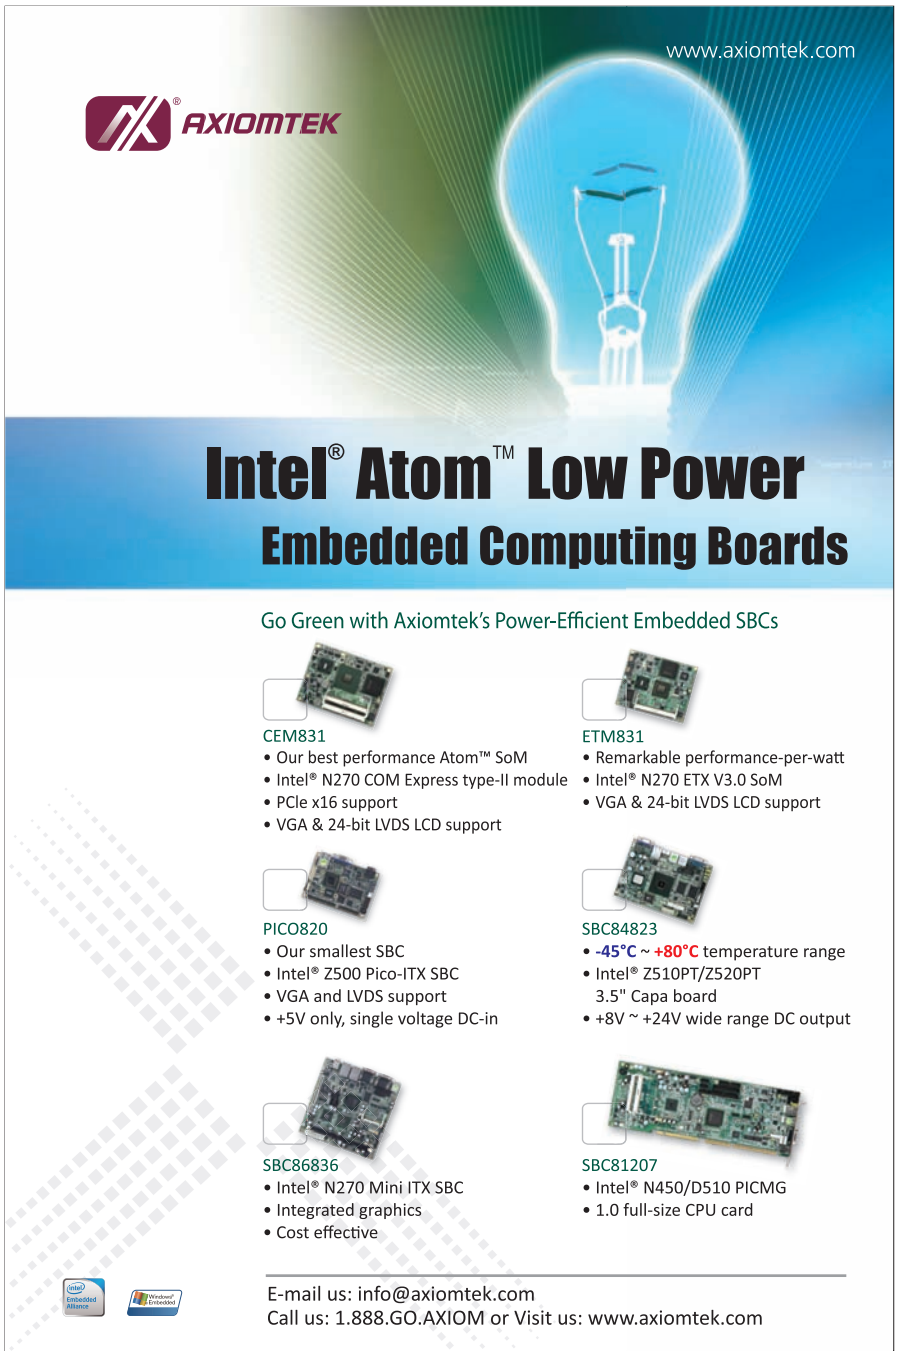

We Won't Leave You High and Dry.

Recent consolidations in the embedded industry have left many PowerPC® users wondering where their next real-time or embedded Linux® operating system will come from.

At LynuxWorks, we are here to help. Our new PowerPC rescue package offers an open-standards based RTOS or embedded Linux, and a migration kit to help PowerPC users stay with their processor of choice.

Visit www.lynuxworks.com/rescue-me to request a migration whitepaper and an emergency hand-crank flashlight-radio*.

LYNUXWORKS™

Open.Reliable.Safe.Secure.

*limited quantities available

Wireless networks provide more with less

By Gary Marrs

802.11 networks are becoming nearly ubiquitous in many enterprise settings and can often be tapped to solve embedded problems in the same location. In this example, a logistics cart used for warehouse operations becomes 802.11 enabled, improving the efficiency of a critical process.

In a tough economy with increased competition from global markets, companies are forced to do more with less. Supply chain management is one area where companies realize this challenge. Organizations are continually struggling to ship more orders, decrease processing time, and increase order accuracy, all the while reducing costs. The nature of order processing has changed as well. With just-in-time and lean manufacturing techniques becoming common practices, distribution warehouses often need to process a high volume of small orders, which often involve a large mix of products, adding further complexity to the process.

The end result is that the order picking process is increasingly challenging. In a typical warehouse or distribution center, studies have shown that order picking can consume as much as two-thirds of the facility's operating cost and time. As a result, companies are continually looking for new technologies and innovations to automate the order picking process.

Pick-to-light goes paperless …

In the first days of order picking, an order list was created, and a fast-footed person would chase down all of the parts, collect the correct quantities of each part, and transport them to the shipping area. This batch process saw many improvements, including techniques to optimize the ideal locations for each part and algorithms to minimize the travel distance of the person picking the order.

Computer and networking technologies facilitated the development of an innovative system known as pick-to-light that significantly improved this process. The *pick-to-light* system provides an automated, paperless, order picking solution. With this system, the order list is sequenced to LED displays that direct the order picker

to a location and indicate the quantity of parts to pick. This system significantly increased order accuracy and reduced order picking times by providing the shortest pick path through the warehouse.

The next step was extending this system to leverage a mobile vehicle or cart for pickers to gather more parts in a highly effective manner. In the mid-1990s, Innovative Picking Technologies (IPTI) developed the first pick-to-light cart called the *RF Batch PicKart* (Figure 1), a patented paperless order fulfillment system that utilizes PICK-MAX pick-to-light displays mounted to a mobile cart.

The first versions of the RF PicKart used proprietary RF radios to communicate. The radios were mostly 900 MHz point-to-point. Although these radios could provide a wireless data link to the cart, they were expensive to deploy, and the data throughput was generally slow.

… then moves to 802.11

With the growing popularity of IEEE 802.11 wireless networking, IPTI found that customers usually had a secure and reliable 802.11 Wireless Local Area Network (WLAN) infrastructure that could be leveraged. IPTI migrated the design of the RF PicKart to 802.11b/g, providing customers with high-speed data transfers, seamless data connections to back-end server processing, and much lower price points.

IPTI performed extensive system testing to make sure that the wireless network could:

- Allow robust communications to multiple mobile carts in a warehouse environment
- Extend to provide complete coverage in large warehouse environments
- Handle the normal shock/vibration issues with robust client radios used on carts

Figure 1 | The RF Batch PicKart was the first pick-to-light cart that enabled a paperless order fulfillment system.

MEN Micro's Rugged COMs for Harsh Environments

XM2 – ESMexpress® with Intel® Core™ 2 Duo

MM1 – ESMini™ with Intel® Atom™

ESMexpress® and ESMini™

- Processor-independent
- Conduction or convection cooled
- Aluminum enclosure
- 100% EMC protection
- -40°C to +85°C operating temperature
- For transportation, commercial vehicles, agricultural machines, avionics, critical industrial and medical control

Rely on MEN Micro for rugged computer boards and systems in harsh, mobile and mission-critical environments.

men micro, inc.

Military & Aerospace Electronics Forum

avionics usa

June 3-4, 2010
Booth 700

MEN Micro, Inc.
24 North Main Street
Ambler, PA 19002
Tel: 215.542.9575
E-mail: sales@menmicro.com
www.ESM-express.com

www.ESM-express.com

To support the pick-to-light functions and WLAN communications, the RF PicKarts have an onboard 16-bit microprocessor that communicates to the WLAN through a UART connection. The microprocessor also supports two serial channels for bar-code scanners and a serial printer. The displays and push buttons interface to the microprocessor through the address/data bus. The carts' pick-to-light displays and push buttons are mounted on a faceplate for easy visibility. All of the electronics are powered by a built-in 12 V lead-acid battery that can provide power for up to 12 hours between charges.

Device servers make design straightforward

For the RF wireless link, IPTI initially chose the Lantronix WiBox, a wireless device server module that creates a transparent data tunnel between an RS-232 serial connection on the RF PicKart microprocessor and the 802.11b/g WLAN. The WiBox was interfaced to the RF PicKart in a couple of days, and a working wireless demonstration was completed in less than one week.

To provide tighter integration of the WLAN interface in the PicKart at a lower price point, IPTI upgraded the RF PicKart design to use the Lantronix MatchPort b/g (Figure 2), an embedded radio module that easily integrates an 802.11b/g link directly onto the main CPU board. The module interfaces to the host processor's UART and transparently tunnels data from the host to the server without requiring any special software driver development.

On power-up, the MatchPort b/g is configured to make a TCP connection directly to the back-end server, which runs four turnkey software programs for a complete solution. The software modules include the Host Interface program, Real-Time Pick-to-Light program, IPTI Statistics program, and the Order Reconciliation program. The Host Interface program is the middleware that processes order information to be used in the pick-to-light system. The Real-Time Pick-to-Light program shows the orders that are active in the system and sends order picking details across the WLAN network to the displays. Communication is event-driven, so when the push buttons are pressed, the system acknowledges those events in real time.

For WLAN security, IPTI uses IEEE 802.11 Wi-Fi Protected Access (WPA) Temporal Key Integrity Protocol Pre-Shared Key security. To set up the security, a few simple configuration settings are established through the serial or network interface. A Web-based configuration manager is also available. For data security from the pick-to-light displays all the way through to the server, the system offers support for optional 256-bit Advanced Encryption Standard, enabling secure end-to-end data transfer. For applications that require enterprise-grade WLAN security, the MatchPort b/g Pro offers enhanced features, including WPA2 Transport Layer Security (TLS), Tunneled TLS, or Protected Extensible Authentication Protocol.

The MatchPort b/g is pretested for European telecommunications regulations and FCC certified, which allowed IPTI to leverage the FCC license grant and bypass 802.11 regulatory testing, speeding time to market and minimizing development and testing costs.

Fitting in the network

One problem identified during application testing was roaming. Larger deployments with more than one Access Point (AP) experienced some delays or dropped TCP connections between the RF PicKart and the enterprise server when the carts traveled between APs in the warehouse. Evidently, roaming is not handled the same by all AP vendors. The MatchPort b/g thus could not maintain a seamless association between the different APs as the carts moved around the warehouse.

Figure 2 | The MatchPort b/g provides an easy way to integrate an 802.11b/g link onto the main CPU board.

To solve this problem, Lantronix implemented SmartRoam technology, which enables the MatchPort b/g to continuously track the signal strength of all APs within range that meet the network and security profile. If a stronger AP is available, the module will make a seamless transition to the better AP. This greatly enhanced the mobility of the RF PicKart within a large warehouse and resolved all of the roaming issues.

For the WLAN infrastructure, IPTI chose the Cisco AP1231 APs, which set the enterprise standard for high-performance, secure, manageable, and reliable WLAN networks. These APs provide robust communications in noisy environments, a particularly important feature in warehouse environments where multiple vehicles and excessive RF traffic make it difficult to obtain good link margins. The Cisco APs also offer extensive debugging facilities and the ability to log debug messages with time stamps. These debug features were helpful when troubleshooting the roaming issues.

WLANs solve embedded challenges

Continuous process improvement is demanding cost-effective ways to solve tough supply chain management issues. Secure and robust WLAN networks are providing solutions to these challenges. The order picking process is just one area where WLAN-enabled carts are allowing companies to remain competitive by doing more with less. On a wider scale, many more creative applications can utilize embedded WLAN technology to solve problems with similar ease. ECD

Gary Marrs *is a senior field application engineer with Lantronix, where he is responsible for helping customers design-in Lantronix device servers for applications that require network connectivity and security. Prior to working at Lantronix, Gary spent 17 years at Grayhill, where he worked in design engineering, marketing, and sales capacities. He holds a BS in Electrical Engineering from Illinois Institute of Technology and an MBA from DePaul University's Kellstadt Graduate School of Business.*

Lantronix
949-453-3990
www.lantronix.com

ECD in 2D:

For more on connecting devices, Lantronix overviews device networking – use your smartphone, scan this code, watch a video: http://bit.ly/drrcin, or visit http://bit.ly/ecdin2d_apr10

S3 Group has the first Continua Certified USB-enabled platform in the form of the Embedded Agent Stack (EAS). EAS is an ISO/IEEE 11073-compliant PAN device (agent) stack. The standard defines how to exchange and evaluate vital signs data between point-of-care devices.

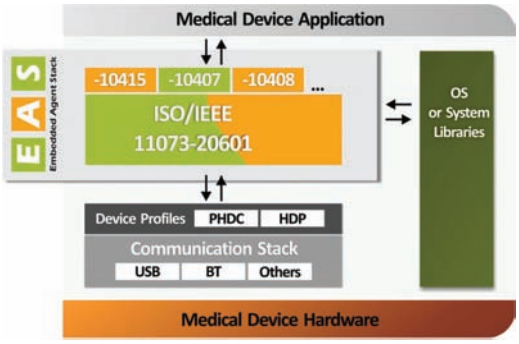

S3 Group | www.s3group.com/telehealth

LensVector has created a breakthrough solid-state autofocus lens, dramatically reducing the size and power of the lens package. The technology is based on a liquid crystal cell, which, when applied with a small control voltage, changes in refractive index.

LensVector | www.lensvector.com

RF Micro Devices PowerSmart platforms provide a multiband front end for mobile devices, with a wide range of mode coverage (GSM/GPRS, EDGE, EDGE Evolution, CDMA, TD-SCDMA, WCDMA/HSPA+, and LTE/4G) and industry-leading band coverage (bands 1-6 and 8-10).

RF Micro Devices | www.rfmd.com

Editor's Choice

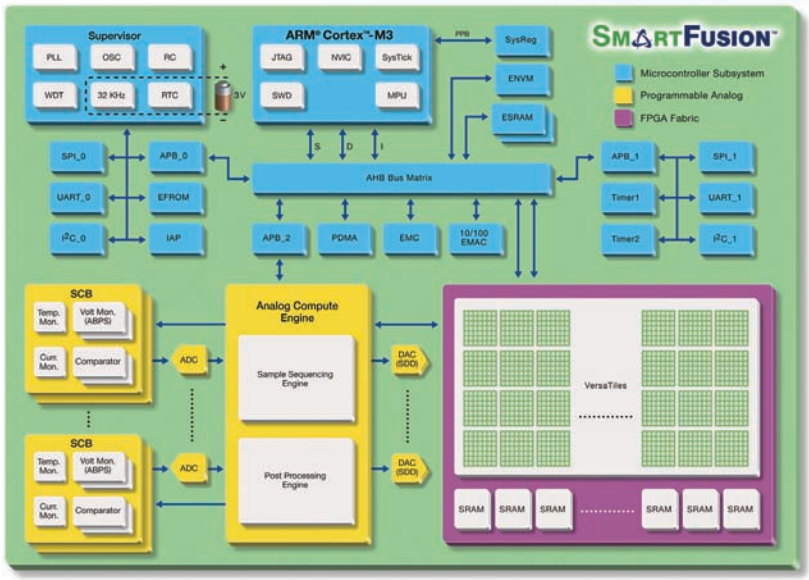

Now, it's possible in an FPGA

A real MCU coupled with real programmable gates combined with real mixed-signal I/O in a single FPGA device – that's what the new Actel SmartFusion FPGA is. As I wrote in my blog, it's an interesting new class of FPGA that's really a multicore device with complete gate programmability, which should allow designers who can't afford to tackle a custom System-on-Chip (SoC) design to do many things not possible before.

There are three key features to note: an ARM Cortex-M3 core running at 100 MHz, a 60K to 500K FPGA gate field with 350 MHz clocking, and an Analog Compute Engine that provides sequencing and computation with up to three 12-bit ADCs and three 12-bit DACs.

Actel | www.actel.com | **RSC#** 44591

Two bits buys 16-bit MCU

Economies of scale are wonderful. The folks at Texas Instruments went after the problem of delivering a 16-bit MCU at an 8-bit price point – starting at "two bits," or 25 cents. For this low price, you get 512 bytes of flash and 128 bytes of RAM, a two-wire debug interface, and 10 GPIO, all in a 14-pin package.

Going a bit higher in the MSP430 Value Line MCU family, a nice feature is a smart ADC with a data transfer controller that manages samples throughout memory, even while the main core is sleeping: an impressive 200 KSps at 0.6 percent CPU load. And there are plenty of other features – ultra low power, code compatibility, and pin compatibility – that make this addition to the MSP430 line very interesting.

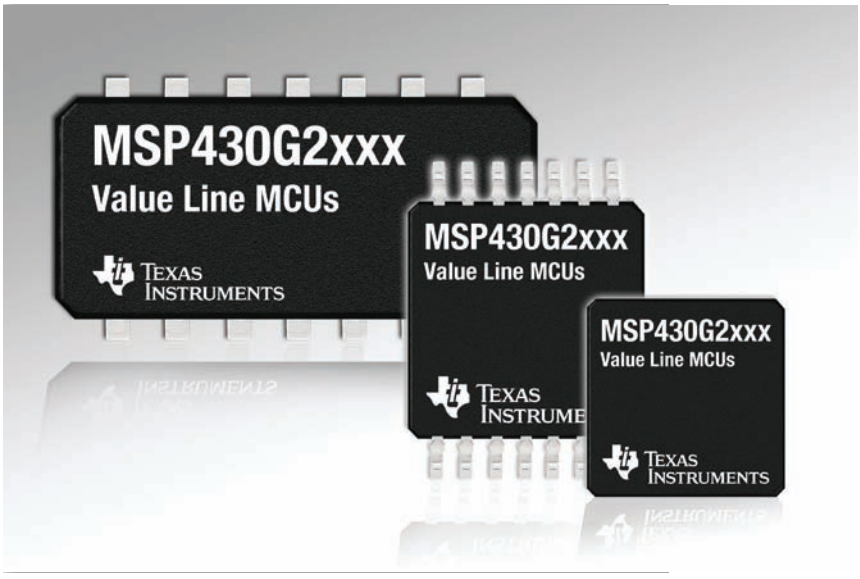

Texas Instruments | www.ti.com | **RSC#** 44590

Editor's Choice

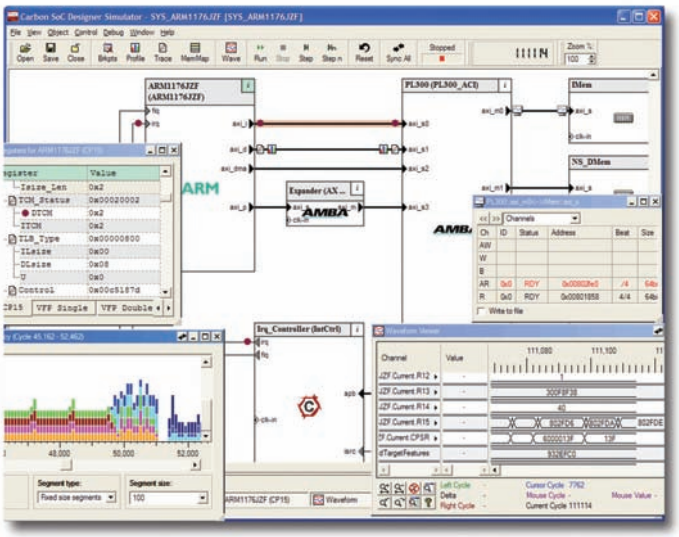

Models for Mali GPU

Graphics Processor Units (GPUs) are becoming increasingly important in multicore devices to accelerate display processing, and ARM's Mali cores are providing designers a path to build SoCs with these high-performance GPUs. Carbon Design Systems has just introduced models for the Mali-200 and Mali-400MP to give designers better tools to work with.

These new models are compiled directly from RTL and are purportedly 100 percent implementation accurate. Each has a TLM-2.0 interface so it can operate with a system prototype and boot an operating system like Linux or Android. This gives teams flexibility to explore architectural options, create SoC firmware, or even develop application software quickly and easily.

Carbon Design Systems | www.carbondesignsystems.com | **RSC#** 44592

More simulation, a lot less waiting

Chip design simulation is critical to working out reliability before committing to silicon, but traditional Monte Carlo analysis often hits a wall in speed and accuracy. Magma Design Automation has developed proprietary dynamic error-controlled algorithms along with sophisticated statistical techniques to provide up to 100 times improvement in speed and substantially better accuracy compared to traditional Monte Carlo statistical analysis.

With this kind of speed, FineSim Fast Monte Carlo will enable statistical analysis in cases where it would have been hard to justify waiting for results. As can be seen from the curve in the graph, accuracy holds up better than traditional Monte Carlo methods as the number of simulations increases – again, providing motivation to do more simulation.

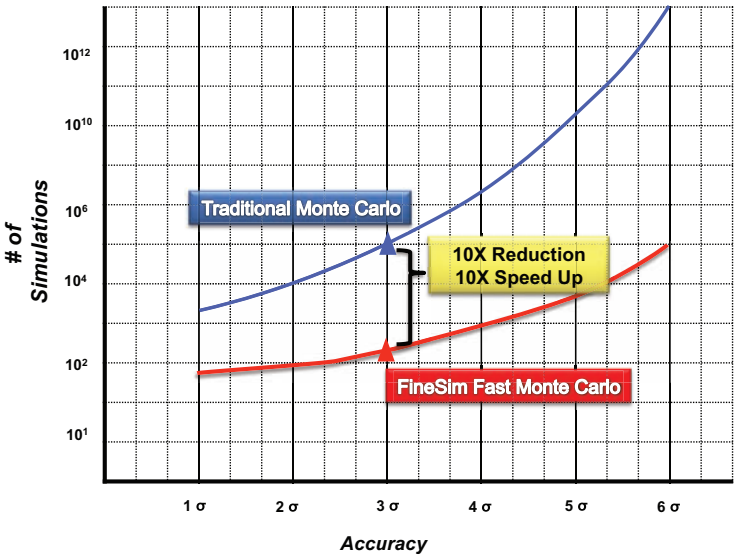

Magma Design Automation | www.magma-da.com | **RSC#** 44593

E-casts and virtual conferences

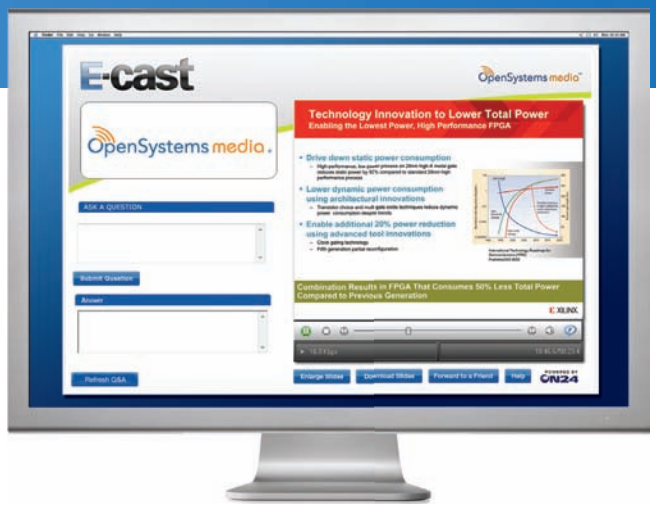

Many topics inside the pages of *Embedded Computing Design* – like multicore processing and software, static code analysis, and FPGAs – come alive in our E-cast programs. These programs, led by a moderator and featuring expert guests from industry-leading firms, are typically one-hour-long live audio presentations followed by Ask the Experts sessions in which the speakers answer questions submitted by the audience.

Last month, we launched an exciting new E-cast format with the FPGA Virtual Summit, which featured four events with technical tracks capped off by a keynote executive roundtable session led by Xilinx and Synopsys, all of which was presented on one day. This unique format showcased viewpoints from nine different firms and multiple experts, providing concentrated information on the latest in FPGA technology straight to engineering desks.

If you're able to attend these events live, you'll find timely, exclusive information and gain the opportunity to ask experts your questions. But if you can't attend, don't worry; these events are archived for one year after the original broadcast date and can be played on demand at your convenience, so you don't miss out on the information.

In addition to developing an extensive calendar of E-casts covering a wide range of topics, we're also making plans for virtual conferences on EDA and AdvancedTCA, so make sure to check back at ecast.opensystemsmedia.com for more information on those and other programs.

ECD in 2D

What you're seeing throughout the magazine are 2D barcodes – our newest innovation to for delivering information quickly and easily. The codes are in QR format, readable by a wide variety of barcode scanner software programs available on popular smartphones. Each barcode in this issue contains a URL, which you can access directly with your smartphone browser and e-mail or text to anyone.

For example, this barcode takes your smartphone to a YouTube video explaining more about how DynaVox Mayer Johnson's products are helping children and others with speech impairments.

QR code scanner software varies, so look for the one that works the best for you. Here are some choices that can be found in your app store or online.

iPhone: App names like BeeTagg, i-Nigma, Kaywa, NeoReader, Optiscan, QuickMark (staff favorite), UpCode, and ZXing Barcode Scanner

Android: App names like BeeTagg and ZXing Barcode Scanner (staff favorite)

BlackBerry: Search for BeeTagg

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

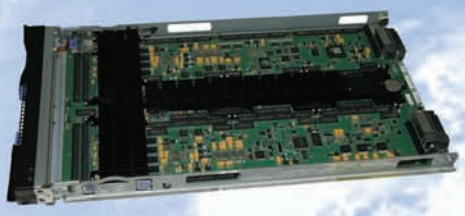

Annapolis
Micro Systems

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

Tiny and Mighty

A World Class Rugged Atom Embedded Computer

V2100 Series

- LVDS connection with industry-leading display configuration options
- Dual independent displays (VGA + LVDS)
- -40 to 85°C wide operating temperature range

V2101

X86-based Atom embedded computer

High Performance

- Intel Menlow XL 1.1 GHz processor 400MHz FSB
- One 200-pin DDR2 SODIMM socket that supports DDR2 400 at up to 2 GB
- Built-in industrial high-speed CompactFlash to store OS

Multiple Interfaces

- 2 RS-232/422/485 serial ports, 2 10/100/1000 Mbps LAN ports for network redundancy
- 2 USB 2.0 hosts and 1 SD socket for storage expansion
- 3 digital input channels and 3 digital output channels

Small Form Factor with Rugged Design

- Highly efficient heat conduction with fanless design
- 2 G anti-vibration and 20 G anti-shock certification for reliable operation on moving objects

Dual Displays

- VGA and LVDS panels for field site monitoring
- 18-bit/24-bit single channel for LVDS connection with multiple resolution selection

Low Power Consumption

- Less than 20 W power consumption for cost-effective industrial application

Wide Temperature

- -40 to 85°C operating temperature models for harsh environments

Moxa Americas, Inc.
Tel: 1-888-669-2872
Fax: 1-714-528-6778
usa@moxa.com www.moxa.com