Embedded COMPUTING DESIGN

www.embedded-computing.com

MAY 2008 VOLUME 6 NUMBER 4

Thermal designs take the heat

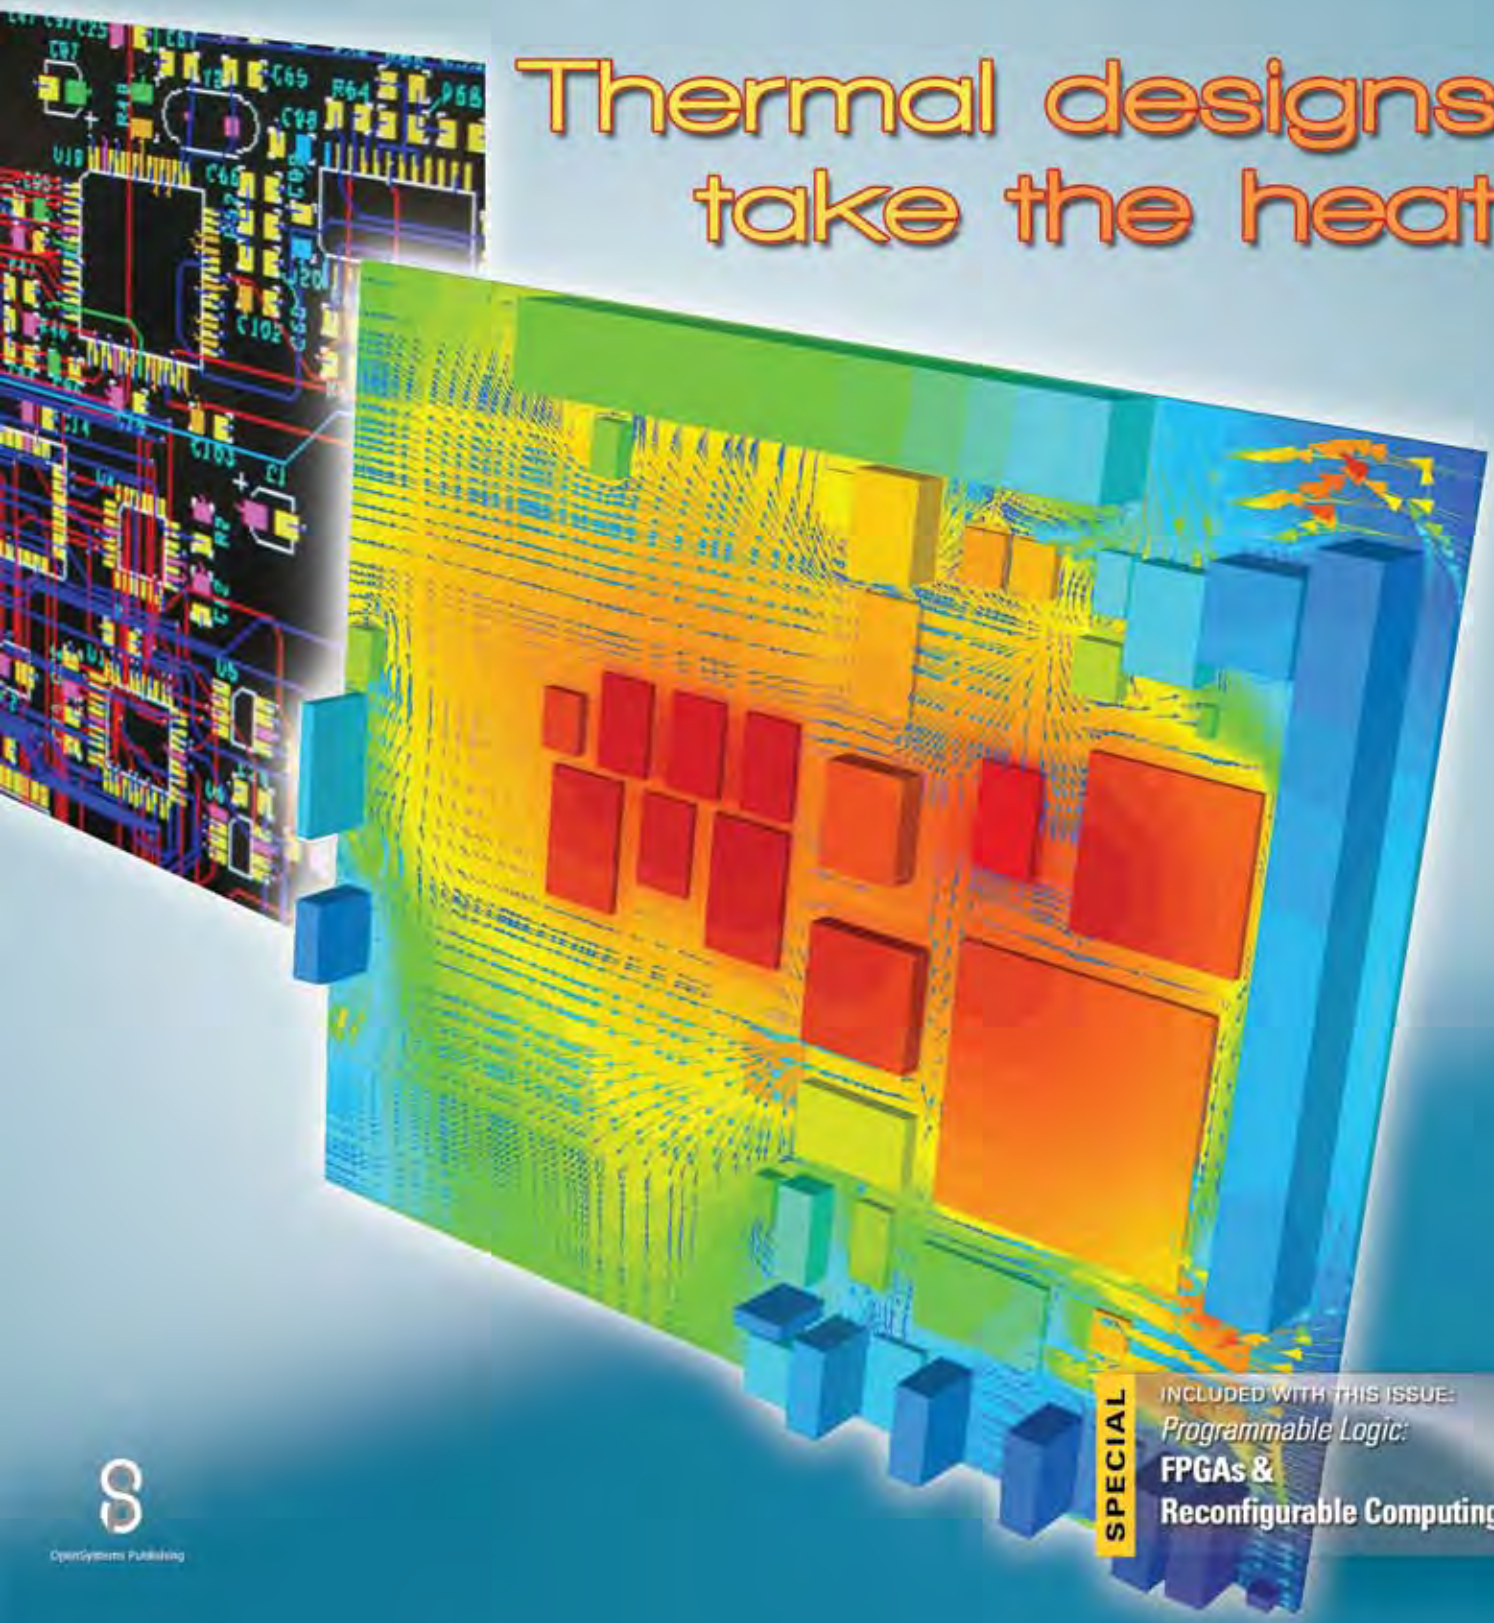

SPECIAL

INCLUDED WITH THIS ISSUE:
Programmable Logic:
FPGAs & Reconfigurable Computing

OpenSystems Publishing

-40º to +70ºC Fanless 1GHz Industrial SBC. List Price $595

WinSystems' EBC-855 is rugged, reliable, and ready for harsh demanding applications.

Windows® and Linux Quick Start Kits available.

If you need a powerful PC with fanless operation, choose WinSystems' EBC-855.

This extended temperature, high-performance, x86-based SBC operates Linux, Windows XP embedded, and other RTOSes to provide a seamless development platform. It also supports popular video and wired/wireless network standands.

EBC-855 features include:

- Intel® 1 GHz ZCD CPU or 1.8 GHz Pentium® M
- Intel® Extreme Graphics 2 technology supports CRT & LVDS flat panels simultaneously with dual independent displays
- Custom splash screen on start up
- 10/100 Mbps Intel® Ethernet controller
- 802.11a/b/g wireless supported
- 4 serial COM ports and 4 USB 2.0 ports
- 48 bi-directional TTL digital I/O lines
- Bi-directional LPT port
- Two EIDE ports (UDMA100) for hard disk
- 3.5-in. floppy disk drive and CompactFlash supported
- PC/104 and PC/104-*Plus* connectors
- Onboard AT keyboard and FDC controller
- AC97 six channel 5.1 surround sound
- +5 volt only operation
- EBX-size: 5.75" x 8.0" (146mm x 203mm)
- Industrial temperature operation
- Long-term product availability
- Off-the-shelf delivery

Contact us for additional information or OEM pricing. Our factory application engineers look forward to working with you.

Call 817-274-7553 or
Visit www.winsystems.com/EBC-855
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

epic™ Linux PC/104 Modules Microsoft Windows Embedded Partner M2M 100 Companies 2008 Extended Temperature RoHS compliant

Embedded COMPUTING DESIGN
Volume 6 • Number 4 www.embedded-computing.com

Columns

Departments

UNIT UNDER TEST

Cover/Web Resources/Events

On the cover

Innovative thermal management methods described in this issue are improving heat dissipation while meeting performance demands in embedded designs. PCB thermal analysis photos on cover and page 24 courtesy of Flomerics, www.flomerics.com

E-casts

www.embedded-computing com/ecast

SDR is more than JTRS: Think military wireless networks and civilian mobile infrastructure
May 21, 2 p.m. EDT

E-letter

www.embedded-computing.com/eletter

SoC microcontrollers power portable medical device innovation
By Kevin Belnap, Texas Instruments

Web Resources

Subscribe to the magazine or E-letter at:
www.opensystems-publishing.com/subscriptions

Industry news:

Read: www.embedded-computing.com/news
Submit: www.opensystems-publishing.com/news/submit

Submit new products at:

www.opensystems-publishing.com/np

Events

NXTcomm
June 16-19 • Las Vegas, NV
www.nxtcommshow.com

Features

Special Feature

Hardware

Software

© 2008 OpenSystems Publishing © 2008 Embedded Computing Design
All registered brands and trademarks in *Embedded Computing Design* are property of their respective owners.

We've Slashed Slow Boot-Up Times

We don't pussyfoot around. Your embedded application needs to come up and execute quickly. Micro/sys CPU boards are optimized to cut through the boot code, which slashes the time from start-up to application readiness for maximum performance over a wide range of hardware platforms.

When you use Micro/sys boards, you gain the value of more than 30 years experience in embedded applications. Our sales team will help you match your boot time requirements with the best board, BIOS, and operating system combination. If you have a tough spec, give us a call.... We welcome the challenge.

MCB58: Compact microcontroller with CAN, 32 DIO, 4PWM D/A, Temp. Sensor, RTC, 16K Serial EEPROM, 60K Flash, and 4K RAM. *Boots in 160 microseconds.*

SBC1586: Headless, Low-Power PC/104 Pentium with Compact Flash, 4 COM ports, USB, and Ethernet. *Boots in under 4 seconds.*

SBC1625: Low-Power ARM processor with Dual Ethernet, 4 COM ports and 24 Digital I/O on PC/104. *Boots in under 14 seconds to Linux Prompt.*

SBC4495: 486/586 EPIC CPU with GPS, Ethernet, 24 bit I/O, 14 bit A/D and D/A. *Boots in under 4 seconds.*

MICRO/SYS

3730 Park Place, Montrose, CA 91020
Voice (818) 244-4600 Fax (818) 244-4246
info@embeddedsys.com

www.embeddedsys.com

©2006 Micro/sys. All brand and/or product names listed are registered trademarks of their respective owners.

TS-7800 High-End Performance with Embedded Ruggedness

$269 qty 1 $229 qty 100

500 MHz ARM9

- New unbrickable design- 3x faster
- Backward compatible w/ TS-72xx
- Low power - 4W at 5V
- 128MB DDR RAM
- 512MB high-speed onboard Flash
- 12K LUT user-programmable FPGA
- Internal PCI Bus, PC/104 connector
- 2 USB 2.0 480 Mbps
- Gigabit ethernet
- 2 SD sockets
- 10 serial ports
- 110 GPIO
- 5 10-bit ADC
- 2 SATA ports
- Sleep mode uses 200 microamps
- Boots Linux in < 2 seconds
- Linux 2.6 and Debian by default

Design your solution with one of our engineers

- 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 stocked and available for next day shipping

See our website for options, peripherals and x86 SBCs

Technologic SYSTEMS

We use our stuff.

visit our TS-7200 powered website at

www.embeddedARM.com

(480) 837-5200

Embedded COMPUTING DESIGN

OpenSystems Publishing™

Embedded and Test & Analysis Group

- Embedded Computing Design
- Embedded Computing Design E-letter
- Embedded Computing Design Resource Guide
- Industrial Embedded Systems
- Industrial Embedded Systems E-letter
- Industrial Embedded Systems Resource Guide
- Unit Under Test
- Unit Under Test E-letter

Editorial Director	Jerry Gipper jgipper@opensystems-publishing.com
Contributing Editor	Don Dingee
Senior Associate Editor	Jennifer Hesse jhesse@opensystems-publishing.com
Assistant Editor	Robin DiPerna
European Representative	Hermann Strass hstrass@opensystems-publishing.com
Special Projects Editor	Bob Stasonis
Art Director	David Diomede
Senior Designer	Joann Toth
Graphic Coordinat r	Sandy Dionisio
Senior Web Developer	Konrad Witte
Web Content Specialist	Matt Avella
Circulation/Office Manager	Phyllis Thompson subscriptions@opensystems-publishing.com

OpenSystems Publishing

Editorial/Production office:
16872 E. Avenue of the Fountains, Ste 203, Fountain Hills, AZ 85268
Tel: 480-967-5581 ■ Fax: 480-837-6466
Website: www.opensystems-publishing.com

Publishers	John Black, Michael Hopper, Wayne Kristoff
Vice President Editorial	Rosemary Kristoff

Communications Group

Editorial Director	Joe Pavlat
Managing Editor	Anne Fisher
Senior Editor (columns)	Terri Thorson
Technology Editor	Curt Schwaderer
European Representative	Hermann Strass

Military & Aerospace Group

Group Editorial Director	Chris Ciufo
Associate Editor	Sharon Schnakenburg
Senior Editor (columns)	Terri Thorson
Senior Associate Editor	Jennifer Hesse
European Representative	Hermann Strass

ISSN: Print 1542-6408, Online 1542-6459

Embedded Computing Design is published 8 times a year by OpenSystems Publishing LLC., 30233 Jefferson Ave., St. Clair Shores, MI 48082.

Subscriptions are free to persons interested in the design or promotion of embedded computing systems. For others inside the US and Canada, subscriptions are $56/year. For 1st class delivery outside the US and Canada, subscriptions are $80/year (advance payment in US funds required).

Canada: Publication agreement number 40048627
Return address: WDS, Station A, PO Box 54, Windsor, ON N9A 615

POSTMASTER: Send address changes to ***Embedded Computing Design***
16872 E. Avenue of the Fountains, Ste 203, Fountain Hills, AZ 85268

Annapolis Micro Systems

The FPGA Systems Performance Leader

WILDSTAR 5 for IBM Blade

The Perfect Blend of Processors and FPGAs

Fully Integrated into IBM Blade Management System
Abundant Power and Cooling Ensure Maximum Performance

Made in the USA

Ultimate Modularity

From 2 to 8 Virtex 5 FPGA/Memory Modules
Input / Output Modules Include:
Quad 130 MSps thru Quad 500 MSps A/D
1.5 GSps thru 2.2 GSps, Quad 600 MSps A/D
Dual 1.5 GSps thru 4.0 GSps D/A
Infiniband, 10 G Ethernet, FC4, SFPDP

Direct Seamless Connections with no Data Reduction

Between External Sensors and FPGAs
Between FPGAs and Processors over IB or 10GE Backplane
Between FPGAs and Standard Output Modules

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

Editor's Foreword

Jerry Gipper

Multi to the core

Multiple processor systems have been around for decades. System architects have designed single computers with thousands of processors, and multicore processors have continually showed up in various niche applications like DSPs.

Today, though, it is different. Multicore processors are the mainstream method for processor development. Given that high-performance processor designers have hit the wall on power consumption, moving to multiple core processors is the most practical way to continue improving processor performance while maintaining a manageable power budget. Multicore is a way past the performance plateau. Embedded platform designers have no choice but to learn to use these multicore processors, and they will need to modify their applications to take full advantage of multiple cores.

One resource that can help engineers is the Multicore Association. I recently attended the group's Multicore Expo, an event that showcased a mix of sponsors from both the high-performance and embedded computing segments. While the high-performance segment has years of experience using multiprocessors, it is a new frontier for many in the embedded segment.

During a panel session on the future of multicore, Jon Peddie of Jon Peddie Research (www.jonpeddie.com) commented that not much has changed in the past 30 years and predicted that not much more will change in the next 5 years. Most of the presenters supported this position in their sessions.

The high-performance segment has implemented Symmetrical Multiprocessing (SMP) solutions where the multicore architecture is a natural fit. In that segment, everyone agrees that the best way to get the most out of every clock cycle in a multiprocessor or multicore system is to structure the application and data set to run in parallel so that every processor or core is busy all the time. Unfortunately, only certain applications lend themselves to a parallel structure, most notably image processing and simulation modeling. This has not changed since the conception of parallel computing. What has changed is that the size of processing elements has expanded and GFLOP numbers have increased two or three orders of magnitude higher.

Despite these challenges, designers are making progress in two areas: Asymmetrical Multiprocessing (AMP) and virtualization. In these environments, multiple cores are running multiple operating systems or multiple threads within an application. This is the hot spot for typical embedded applications that have many dependent and independent tasks running at the same time on a single platform.

AMP embedded applications often use combinations of homogeneous and heterogeneous cores. Many Systems-on-Chip (SoCs) are multicore but have heterogeneous cores for specific functions. One core may be the general-purpose processing core and other cores may be optimized to process I/O or provide DSP functionality. Some newer parts are even introducing dual or quad homogeneous multicore for general-purpose processing with an extra heterogeneous core for specialized functions.

As more players enter the multicore processor market, we will observe more divergence with no convergence in sight despite efforts like the Multicore Association's Multicore Communications API (MCAPI) announced at the Multicore Expo. This will pose a challenge for system developers because their choices will not be compatible, complicating design changes and forcing developers to make decisions that will prevent platforms from being as portable between iterations as they desire.

Numerous similar but distinct topics are covered under the word "multicore." Designers shouldn't be frightened off by the term. It is forecasted that by 2010 nearly 100 percent of processors shipped will be multicore. Kerry Johnson of QNX summed up this trend as he pointed out the variety of options available to embedded systems designers. Most applications will work well with multicore processors, requiring no changes to the software and enabling designers to start moving applications now. Later, they can work out more optimized changes to their system architectures to truly gain from multicore processors.

Feel free to share your comments via e-mail or visit our blog at www.embedded-computing.com to add your comments.

Jerry Gipper, Editorial Director
jgipper@opensystems-publishing.com

Check out the recent E-cast, "Maximizing Performance on Intel Multicore Processors," cosponsored by QNX and Intel. Download it now at www.opensystems-publishing.com/ecast

Embedded Computer Solutions For Harsh Environments

-40°C to +85°C Operating Temperature

Five Year Product Availability Guarantee

ECM401

World's Fastest Embedded Computer Module

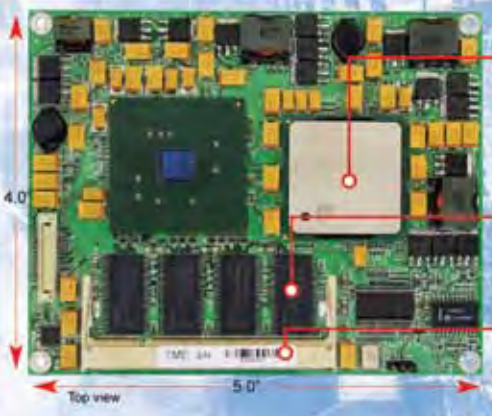

CPU, up to
2.8 GHz Pentium 4
(Low power 1.6Ghz and 2.0Ghz Pentium 4 also available)

512Mbytes soldered on-board DDR memory expandable to 1.5Gbytes using SODIMM socket

SODIMM socket

Features:

- (6) Serial ports
- (4) USB 2.0 ports
- (1) 10/100Base-T
- Compact Flash Socket
- E-IDE supports 2 devices
- Video I/F (1600X1200, 32 Mb)
- AC'97 2.2 Audio
- Less than 5 seconds boot up time
- Intelligent thermal management with independent microcontroller
- Power requirement: +12V @ 3A (2.0Ghz P4, 256MB)
- Over 200,000 hours MTBF

Over 15 years in business of designing and manufacturing embedded computer products for OEMs and System Integrators

We can design your Embedded Computer Board in 30 days.

Time-to-market is very critical to a company's success.
TME has a proven development and manufacturing process that allows new embedded computer products, using TME's core technologies, to be developed in 30 days.
OEMs and System Integrators can now bring their products to the market without compromising cost and features at the expense of time.

Typical Applications

- ▲ Robotic
- ▲ Medical
- ▲ Avionics
- ▲ e-Kiosks
- ▲ Transportation
- ▲ Military/Aerospace
- ▲ Industrial Automation
- ▲ Inventory management
- ▲ Point Of Sale Terminal
- ▲ Test & Measurement

ECM401 Embedded Computer Module provides all functions and features of a high performance computer on a very small Module.

Embedded Computer Designers can easily design an I/O board with proper connectors and form factor that is most suitable for their applications.

The ECM401 supports up to 1.5Gbytes DDR memory, 2.8GHz CPU, with built-in Audio, Video, USB, Network, serial, LPT interfaces and PCI BUS extension capability.

tme

Toronto MicroElectronics Inc.
5149 Bradco Boulevard, Mississauga, ON., Canada L4W 2A6
Tel: (905) 625 - 3203 ~ **Fax:** (905) 625 - 3717
sales@tme-inc.com ~ www.tme-inc.com

Technology Passport

Hermann Strass

Electronic firefighting

Fireproof data center

A modern fire engine carries more than just a pump, ladders, and fire hoses. It also contains a distributed multiprocessor data center. The Z8 Airport Crash Tender from Albert Ziegler GmbH, which has manufactured firefighting equipment since 1890, is used at commercial and military airports in Stuttgart (as shown in Figure 1), Zurich, and other places around world.

Figure 1

The fire engine is based on an 8 x 8 chassis with a V12 diesel engine delivering 1,000 metric horsepower. The Z8 weighs 40 metric tons and can rush to the site of a fire at almost 140 km/h. Two tanks carry 15,000 liters of foam and water. A separate pump motor can propel water or a water-foam mixture more than 90 m from the hose nozzle. The fire engine has nozzles all around its body for spraying water and foam to defend itself against approaching fire.

Figure 2

Two Multiprocessor Control Panels (MCP6) with CAN bus networking control the extinguishing unit. The MCP6 units from Graf-Syteco GmbH are located in the fire engine, one inside the driver's cabin (see Figure 2) and one at the rear of the fire engine. Each unit is equipped with an AMD processor, 16 programmable push buttons, two digital potentiometers with push-button functions, four video input channels, two USB 2.0 ports, two CAN ports, eight digital inputs, 10/100 Ethernet, two Pulse Width Modulated (PWM) outputs, several counters, and a TFT display.

The fire engine also has three keyboards with programmable buttons and LEDs in various locations. All electronic items are protected against environmental influences (Enclosure Codes IP65-IP68), including night-vision design and displays that are easily readable in poor visibility environments. The fire engine includes 10 multichip modules that include various digital and analog I/O channels, including PWM variants connected via CAN interfaces.

European events and awards

The embeddedworld2008 exhibition and conference in Nürnberg (February 26-28) is the world's largest gathering of embedded computing experts. The exhibition set new records for all key figures: 675 exhibitors (+14 percent), 17,341 trade visitors from all over the world (+27 percent) with 56 percent more international attendance, and 18 percent more space, thanks to a new hall. CEOs and CTOs from Green Hills Software, Infineon, and Kontron provided high-level information related to embeddedworld's three focal points: hardware, software, and tools.

Organizers presented the embedded AWARDs 2008 the day before the exhibition. F&S Elektronik Systeme, Germany, received the software AWARD for the Failsafe Flash File System, which continues operating even in the event of a power outage during software updating. Atmel, France, received the hardware AWARD for the first 32-bit microprocessor core with single-cycle SRAM read/write and direct interface to the three-stage pipeline; it does not need a system bus. Hitex, Germany, received the tools AWARD for the STM32-PerformanceStick. The USB stick with peripheral connector and tools software can look deep inside applications, not just the processor core.

Volker Wunsch of the University of Mannheim, Germany, received the university AWARD for a small, self-contained, MEMS-based 6D navigation system for location-based services.

SYSTEMS THAT CAN GO ANYWHERE, DO ANYTHING.

Extreme cold. Intense heat. Moisture. Vibration. Impact. Dust. Dirt. Grime. No matter what the operating conditions, we pride ourselves on our ability to build high-quality, ruggedized solutions that virtually do whatever it is you want them to, regardless of the conditions you need them to do it under. It's all just part of our dedication to producing the best-in-class computing platforms and displays available on the market. And it's why you need to call us. Because we're not just building durable solutions. We're building a reputation.

POWERED BY iEi®

TCS

UNIQUE COMPUTERS & LCDS

www.totalcontrolsolutions.com / 763.416.1610

Consortia Connection

Expanding multicore, RapidIO, and mobile platform innovation

Multicore Association www.multicore-association.org

The Multicore Association provides a neutral forum for vendors interested in, working with, and/or proliferating multicore-related products, including processors, infrastructure, devices, software, and applications. The organization currently has three working groups: Multicore Communications API (MCAPI), Hypervisors, and Multicore Resource Management.

On April 1, the Multicore Association announced that version 1.0 of the MCAPI specification was finished and would soon be available to the public. MCAPI captures the basic elements of inter-core communications and scales to support hundreds of processor cores. This API's principal use will be in embedded multicore systems with tight memory constraints and task execution times.

"MCAPI … has been developed by a wide variety of member companies that range from processor vendors such as Freescale, Intel, MIPS Technologies, and Tilera to software and development tool suppliers such as Codeplay, eSOL, Imperas, PolyCore Software, and Wind River," remarks Multicore Association president Markus Levy. "Such a variety of contributors helps ensure that the specification will be applicable to many different types of multicore systems."

THE Multicore ASSOCIATION™

RapidIO Trade Association www.rapidio.org

The RapidIO Trade Association enables, supports, and drives development of the RapidIO embedded interconnect and its supporting ecosystem. The group will hold the first of its RapidIO Global Design Summits 2008, "Converge, Learn, Innovate, and Design," on May 21 at the Novotel following WiMAX World EMEA in Munich, Germany. This event will showcase technology experts representing CommAgility, Fabric Embedded Tools, IDT, VMETRO, Mercury Computers, RIOLAB, Texas Instruments, Freescale, Tundra Semiconductor, and Xilinx.

"RapidIO technology is the most efficient multihost/processor available and delivers the lowest cost per bit in the industry," asserts RapidIO Trade Association executive director Tom Cox. "These design summits provide a place for the world's best designers to converge, learn the latest about the RapidIO specification and its expansive ecosystem, and discover innovative ways to speed high-performance, highly reliable solutions to market."

RapidIO Global Design Summits 2008 attendees will gain practical information on the RapidIO standard, including the latest RapidIO roadmaps and technology updates as well as news from the RapidIO ecosystem.

RapidIO®

Open Handset Alliance www.openhandsetalliance.com

A multinational alliance of technology and mobile industry leaders, the Open Handset Alliance is committed to fostering mobile device innovation and bringing consumers new mobile experiences. The organization's first product, Android, is a truly open and comprehensive platform for mobile devices. Using Android, developers, wireless operators, and handset manufacturers will be able to bring new products to market faster and at much lower cost, giving consumers flexible and more personal mobile experiences.

Built from the ground up, Android allows developers to create novel mobile applications that take advantage of Internet-capable handsets. The mobile platform is built on the Linux 2.6 kernel, offering a robust operating system, comprehensive set of libraries, rich multimedia user interface, and complete set of phone applications. Android's inventive application model makes it easy for developers to extend, replace, and reuse existing software components to produce rich and integrated mobile services.

Developers can download the Android Software Development Kit to create applications for what promises to be an interesting handset platform.

The Flash Storage Leader

F r Sing

pqi

PQI Corporation
Tel: (510) 651-7281
Fax: (510) 651-7240
www.pqimemory.com

Ingenuity @ work

Parallel programming for multicore processing

problem

Developers cannot always take advantage of multicore processors without understanding the impact on existing code. Various system architecture combinations make optimization challenging, and keeping code portable can be difficult in heterogeneous designs.

solution

By implementing a code partitioning system, programmers can write complex software scalable for use on parallel processors. Separating data allows multiple memory spaces to improve multicore software performance by allocating memory space for each processor.

PlayStation 3 photo courtesy of Sony

Filtering out software side effects

Placing multiple processor cores on a single die can improve system performance while reducing power consumption. Embedded designs can benefit from multicore because designers can spread out an embedded device's functions to different cores so that processing units can perform separate tasks concurrently.

Developers want the ability to take a single-core C++ program, pass it through an autoparallelizing compiler, and automatically generate a multicore program. However, single-core C++ software contains many dependencies in which one part of the software must be executed after another part. It is impossible for the compiler to automatically change the program's order of execution to execute several parts of the program at the same time on different processors.

Codeplay's Sieve C++ parallel programming system aims to minimize the amount of programmer intervention required when writing software for multicore processors. Using a *sieve marker*, programmers mark a section of the program called a *sieve block*. Inside sieve blocks, side effects are delayed, meaning that program code inside sieve blocks can be automatically split into three parts: reading data from memory, pure computation, and writing data back to memory. Because of this division, memory operations can be separated, parallelized, and performed by a DMA system. The system is aptly named "Sieve" because it filters out side effects from software and then lets programmers apply them later. Sieve has been tested successfully on the Cell BE-based PlayStation 3 and AGEIA's PhysX processor.

While Sieve represents a step forward, more innovation is needed to make multicore architectures easier to use. The software community must continue to develop better tools and processes that make multicore architectures more practical.

Quick facts

Codeplay
Founded: 1999
Management: Andrew Richards, founder and managing director
Headquarters: Edinburgh, Scotland
URL: www.codeplay.com

AMAX®

Imagination to Innovation

For Sing[illegible]nt Only

The right manufacturing partner makes all the difference.

AMAX offers comprehensive prototype development, custom manufacturing, and logistics services that deliver robust quality and performance, rapid time to market, and reliable support for your products.

From custom chassis, bezels, and faceplates to your uniquely branded labels and packaging, AMAX provides quality products with a look and feel that reflects your company identity. With AMAX's global logistics and support services, you can deploy your system anywhere in the world with the necessary support to make your product a success.

Contact us **(510) 651-8886 / ext. 8800**
or visit us: **www.amax.com.**

Voted #1 2006, 2007 CRN Award
Server & Storage Leader

© 2008 AMAX Information Technologies. 1565 Reliance way, Fremont, California 94539. All rights reserved. Performance may vary depending on the system configuration and software applications used. All rights reserved.

Packing heat ... into reusable energy

By John Lin

Excessive heat is a significant concern in any computer design. Though designers have implemented varying methods of heat dissipation, all share one common problem. Heat dissipation relies on the transferring medium's thermal conductivity. Many ingenious contraptions have been invented to achieve faster and more efficient heat dissipation. But instead of throwing heat away, why not recycle it? John presents a potentially energy-salvaging thermal management strategy.

Ask most laptop or notebook computer users about their main annoyance, and their responses will likely have something to do with the heat their computers produce. Laptop and notebook computers use many of the same components as embedded computer designs – components such as the central and graphics processing units, Northbridges, Southbridges, and so on.

Components play a significant role in a thermally friendly computer design. Using such low-heat processors like the VIA C7 CPU help reduce overall heat[1]. But generally, as processors become more powerful, they generate more heat, making thermal management an important consideration.

Herein lies the problem: Computers require power to operate, and that power inevitably produces waste energy in the form of heat. If the accumulated heat within a system becomes too great, the system cannot operate properly. Designers' objective, then, is to get rid of the generated heat and use it productively, if possible.

Current thermal management techniques

At present, most if not all thermal management methods concentrate on heat dissipation, either passive (for example, fanless heat sinks) or active (for example, fan-forced cooling or water cooling). An ideal thermal management system based on heat dissipation would dissipate all the generated heat. Though existing thermal management devices are adequate, they are limited in how much heat they can dissipate. Heat dissipation is only satisfactory until thermal equilibrium is reached.

How much heat needs to be dissipated? Generally, the amount of energy used is proportional to the amount of heat produced. So if a component uses 30 W of power over 60 seconds, it will output 1,800 joules of heat.

$$Q = Pt = VIt$$

In the equation above, Q represents heat in joules (J), P represents power in watts (W), t represents time in seconds (s), V represents voltage (V), and I represents current in amps (A). The delta T degree Celsius can be calculated using the following equation, where m represents mass in kilograms (kg) and c represents specific heat (J/gK).

$$\Delta T = Q/mc = VIt/mc$$

A heat-sink system is based on transferring heat from the heat sink to ambient air. For a heat sink to properly work, thermal equilibrium must never be reached while the system is running (that is, ambient air must always be cooler than the heat sink when the system is on). Otherwise, once thermal equilibrium is reached, the heat will no longer transfer from the heat sink to ambient air.

Once thermal equilibrium is reached, the heat sink cannot do much to relieve the processor of its heat burdens. Even assuming that thermal equilibrium is never reached, it is evident that much energy is wasted. Such potential should not go unused. An alternative thermal management technique involves converting heat energy and storing it for later use.

We Simplify Technology

CORE SYSTEMS

Reliable operation in severe environments

- Integrates hardware mechanically, electrically and thermally
- Fanless, total thermal solution
- Mini PCI & PC/104 expansion
- MIL–810F, EN50155 available
- Installed Windows® XPe or Linux available
- Virtually indestructible mobile power supply
- NEMA 4X
- –40° to 85° C

- Transportation
- Mining vehicles
- Outdoor kiosks
- Industrial servers

- Digital signage
- Military and police
- Traffic management
- Digital video surveillance
- Asset management

303–430–1500

OCTAGON SYSTEMS

www.octagonsystems.com

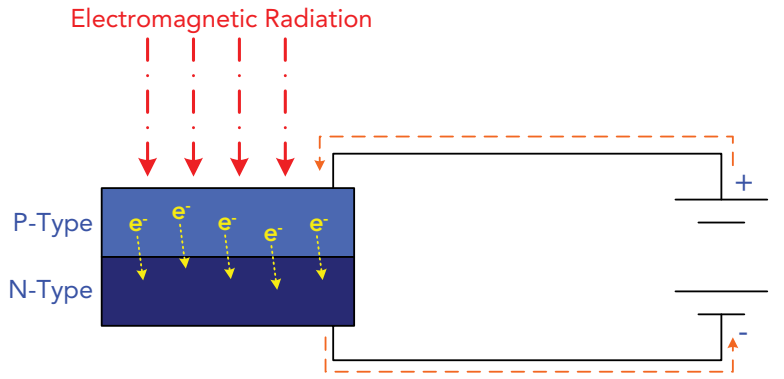

Figure 1

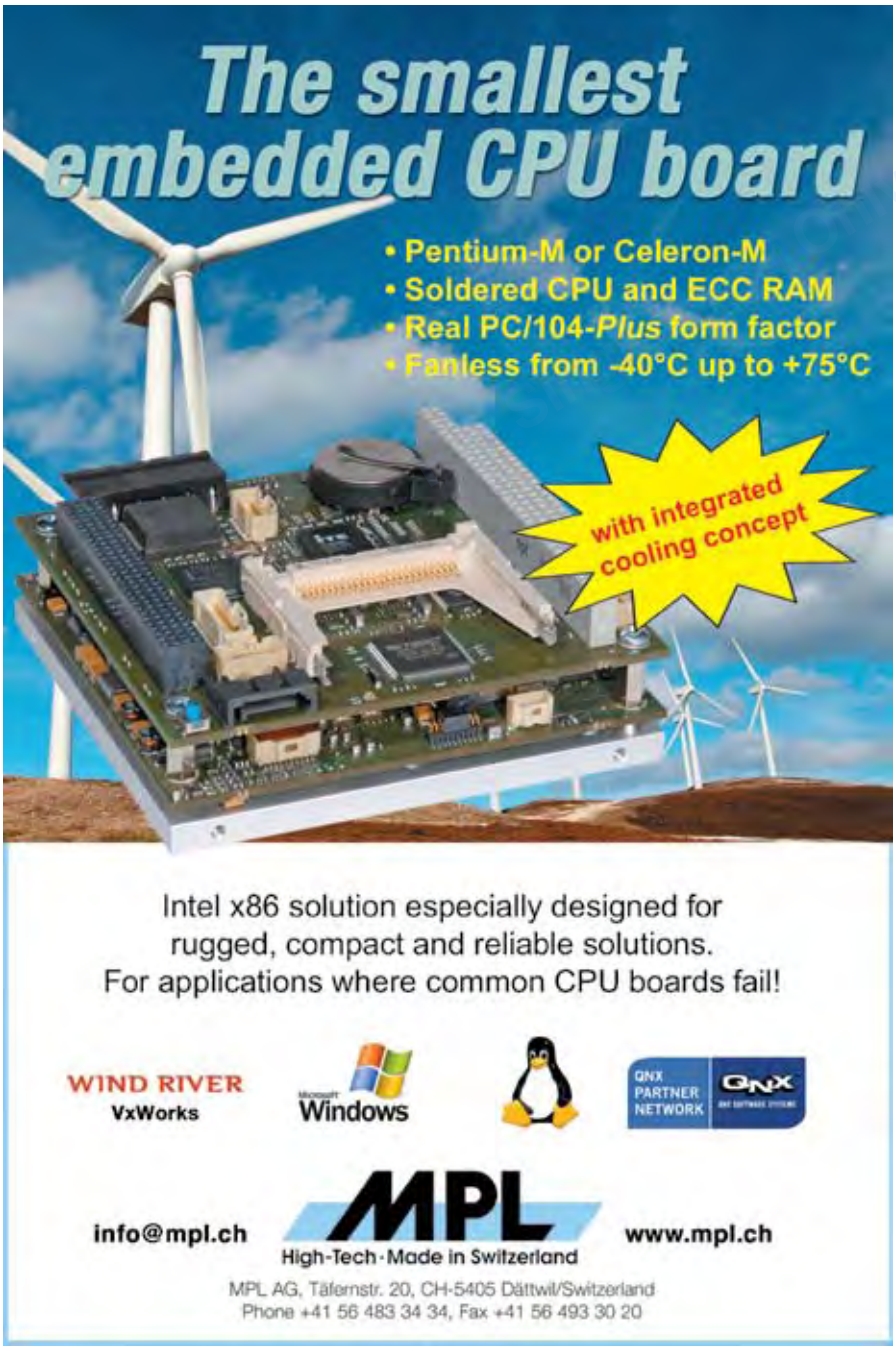

Making heat productive

Converting heat energy to electricity is not a new idea. Because heat is infrared radiation and thus part of the electromagnetic spectrum, photovoltaic-based thermal management is possible[2]. The term *thermophotovoltaic* (TPV) has been used to describe photovoltaic devices for converting heat energy[3]. Currently, TPV devices are not practical enough to consider as primary power sources. However, it is important to remember that thermal management systems' chief objective is heat removal, not power generation. Any power that is gleaned from the process is a bonus and not the main goal.

Photovoltaic devices require two layers of photosensitive material: one p-type layer and one n-type layer. The p-type comprises a material whose atoms have an extra electron that prevents the material from being completely nonconductive. The n-type consists of a material whose atoms lack one electron. When the p-type layer is exposed to light, the photons in the light source cause the extra electrons to be released, resulting in a flow of electricity (see Figure 1).

Implementation for a fully enclosed system

A thermal management device that recycles waste energy is ideal for airtight sealed systems. In such systems, passive cooling and fan-forced cooling have little if any effect. When waste energy is recyclable, rugged airtight systems can be designed without sacrificing performance. High-performance components generally produce more heat, but for a thermal management device fueled by heat, it is no longer a problem.

To implement the thermal management device, the chassis interior must be lined with TPV cells to absorb as much heat as possible. In addition, a double-sided TPV cell should be placed directly over the heat sink (see Figure 2). Having a double-sided design enables the TPV cell to draw heat not only from the heat sink but also from ambient air. The chassis interior wall is lined with single-sided TPV cells to capture any remaining ambient heat.

This thermal management strategy is particularly suitable for portable applications or external applications where the environment may adversely affect the computer (for example, a street kiosk, GPS unit on a motorcycle, or navigation unit on marine craft). The technique also would work in partially enclosed systems such as fixed applications where the environment will not have adverse effects on the computer (for example, a kiosk inside a building).

Perpetuating battery life

By using TPV cells in thermal management devices and choosing the right components, designers can create winning low-heat systems. As photovoltaic cell technology advances, such designs may result in products with ridiculously long battery lives – and perhaps even a nearly perpetual energy source. ECD

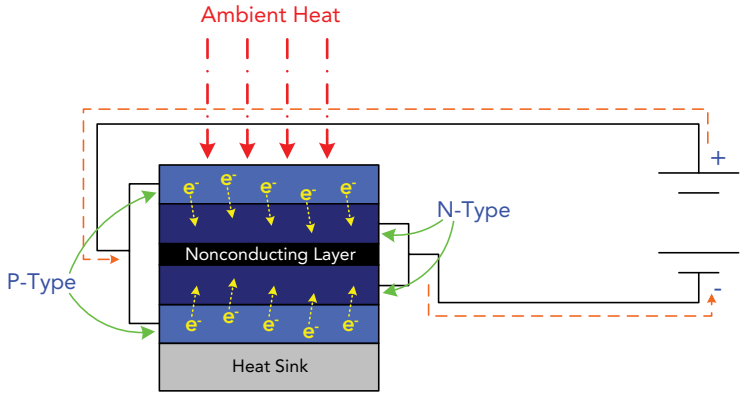

Figure 2

References

[1] "VIA C7-M ULV Processor," VIA Technologies website. 2007. VIA Technologies, Inc. 6 Dec, 2007. www.via.com.tw/en/products/processors/c7-m_ulv/index.jsp

[2] "Solar Cells that Harness Infrared Light," Environmental Science & Technology. 2 Mar. 2005. Konarka Technologies, Inc. 14 Nov. 2007 www.konarka.com/news_and_events/konarka_articles/2005/3_march/environmental_sci_and_tech/solar_cells_that_harness_mar2.php

[3] "Thermo-photo-voltaic Cell Based on GaSb," TPV Network. 2007. 12 Nov. 2007 tpvcell.hmg.inpg.fr/WT.html

John Lin *is a senior technical writer at VIA Technologies, Inc., based in Taipei, Taiwan. He has seven years of experience in the computer industry, with five-plus years at VIA. John graduated with a Bachelor's degree in Business Economics from UCSB.*

VIA Technologies, Inc.
+886-2-2218-5452
JohnLin@via.com.tw
www.via.com.tw

Switching strategies for increasing processor power efficiency

By Brian Law and Greg Ferrell

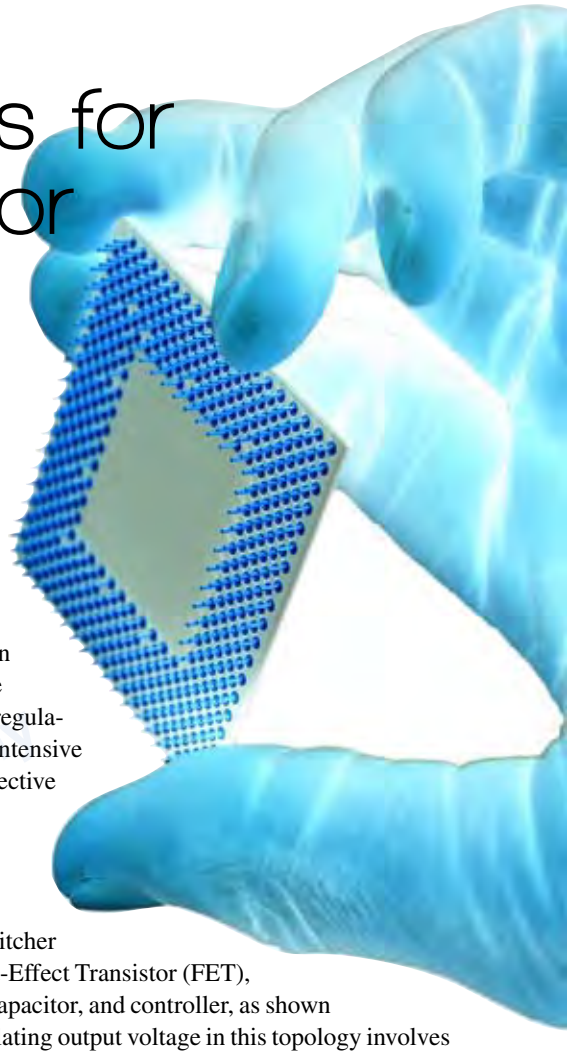

Continual advancements in smaller silicon geometries have enabled lower operating voltages in embedded devices, particularly microprocessors. Operating voltages for processor cores are now as low as 1.2 V and quickly moving to 0.8 V, presenting a challenge as to how designers can efficiently power these devices. As opposed to the traditional method of using linear regulators to power low-voltage devices, using switching regulators in these applications can help solve this problem. Depending on system requirements, designers can use various architectures to maximize switching regulators' efficiency.

A typical low-power embedded processor draws 300-600 mA. Using a switching regulator in older-generation processors where core voltages were 3.3 V provided little benefit. However, decreasing core voltages are presenting a significant opportunity for efficiency improvements, especially when powered by a lithium ion cell (4.2 V) or 5 V rail. For example, a linear regulator in a 4.2 V system will waste 1.8 W [(4.2-1.2 V) x 600 mA] when regulating down to 1.2 V. In contrast, a switching regulator can perform with up to 95 percent efficiency under the same conditions, which can add a significant amount of operation time to the system.

The potential noise created by a switching regulator, its inefficiency at light loads, and the need for more complex control have traditionally discouraged designers from using this type of regulator for an embedded processor. Additionally, linear regulators' larger footprint has made linear regulation the preferred method. Nonetheless, designers can employ several techniques that greatly improve efficiency in switching power supplies and make switching regulators feasible in different types of designs. Single ICs with an integrated controller, pass devices, and compensation components have made switching regulators less design-intensive and more cost-effective to implement.

Standard buck topology

A simple buck switcher comprises a Field-Effect Transistor (FET), diode, inductor, capacitor, and controller, as shown in Figure 1. Regulating output voltage in this topology involves varying the duty cycle on the FET's gate to increase or decrease current through the inductor, a method referred to as *Pulse Width Modulation* (PWM). A PWM switching regulator's efficiency can be upwards of 95 percent when operated at full load. However, when operated at light load, a switching regulator's efficiency drops off significantly, making it undesirable for systems that operate under changing load conditions or require low current or sleep modes.

To overcome a switching regulator's inefficiency during light load conditions, designers can put the regulator in pulse skipping or *Pulse Frequency Modulation* (PFM) mode. When in PFM mode, the FET in the switching regulator only operates when output voltage falls below the low limit. This reduces the number of pulses, which in turn reduces switching losses through the FET, inductor, and diode, thus improving efficiency under light load conditions.

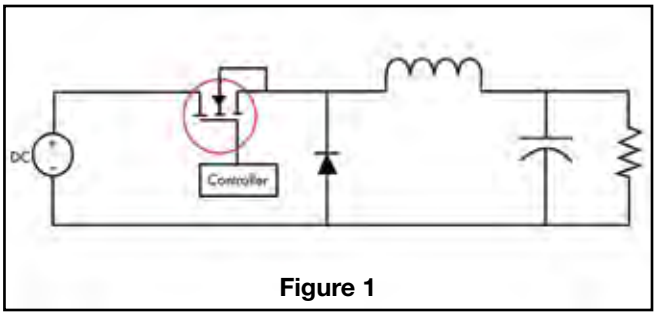

Figure 1

SUPER ★ TALENT®
THE BEST MEMORY

Reinventing Storage

Rugged, Reliable, Unparalleled

Industrial Temp

SATA / IDE SSD

Solid State Drive

SATA/ IDE SSD

• Performance

Operating Temp: -40 to +85C
Sequential Read Rate: 60 MB/sec (min)
Sequential Write Rate: 45 MB/sec (min)
Sequential Read IOPS: 58,675 IOPS
Access Time: 0.1ms
Power Consumption: 85% less than Average HD

• Reliability

Shock: 1500G (operating)
Vibration: 16G (operating)
MTBF: > 1,000,000

• Endurance

Read: Unlimited
Write: >140 years @ 50GB write-erase/day

"AWARD WINNING PRODUCT"

CeBIT
"Innovation of the year" Award

DISKCON USA
Best of Show Award

- **100% compatible with standard hard disk drives**
- **IDE form factors in 1.3" / 1.8" / 2.5"**
- **Designed and manufactured in USA**
- **Capacities range from 4GB to 256GB**
- **SATA form factors in 1.8" / 2.5" / 3.5"**

Super Talent USA Headquarters
www.supertalent.com
Tel: 1.408.934.2560 | Email: sales@supertalent.com
2077 North Capitol Avenue San Jose, CA 95132 U.S.A

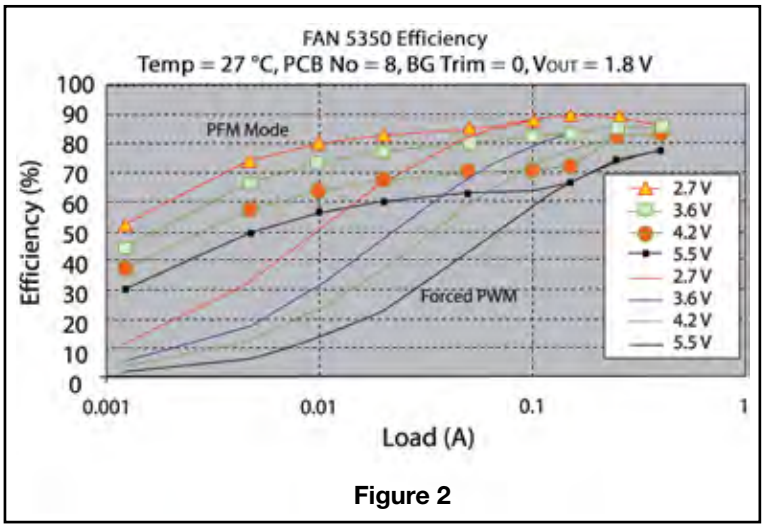

Figure 2

For the device shown in Figure 2, the regulator transitions from PFM to PWM at about 100 mA, thereby maximizing efficiency throughout the entire range of output currents.

When designing with a PFM regulator, designers must take into account the applications in which it will be used. Since the switcher's frequency varies with the output load, it is possible that switching frequency can fall as low as the audio band, which can produce undesirable noise problems. Fortunately, certain tools available in the market today can ensure that frequency never enters the audio band. Although these tools may cause a slight decrease in efficiency under light loads, they can save designers countless hours spent eliminating a noise issue.

Synchronous topology

Designers can make additional improvements to a switcher's efficiency using a synchronous topology, as shown in Figure 3. This topology can reduce switching losses and reverse recovery losses through the diode by replacing the diode with a low $R_{ds(on)}$ FET.

In this topology, the timing of synchronous switching is critical. If both FETs are conducting, power will be lost on every cycle and efficiency will be compromised. In addition, designers must consider which devices to choose. For example, the FET's gate capacitance is an important variable because it can keep the low side FET in conduction after the controller has switched off the FET. During the short period of time that the gate remains charged, the input power is shunted directly to ground. The $R_{ds(on)}$ and FETs' gate capacitance will have the greatest impact on efficiency in this topology; thus, it is important to optimize both.

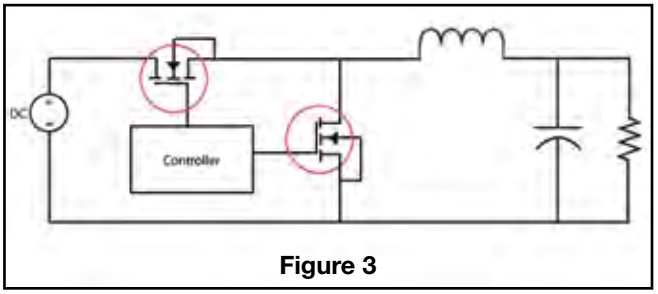

Figure 3

In the aforementioned topologies, using higher switching frequencies can be advantageous when board space is very limited. Higher switching frequencies allow designers to use smaller passive components such as the inductor and output capacitor, which can decrease the design's cost and overall footprint. However, this comes at the price of decreased efficiency. When frequency increases, the number of times that FETs switch increases, which in turn increases losses. In addition, the smaller inductor and capacitor will likely create higher ripple on output voltage.

Multiphase topology

To overcome today's stringent ripple requirements while maintaining efficiency in embedded designs, designers can use multiphase DC-DC switchers. This topology can reduce switching losses while effectively increasing the regulator's switching frequency.

For example, consider a single-phase DC-DC synchronous switcher connected to a load, I_{load}. Switching losses through each FET are I^2R or $I_{load}{}^{2}* R_{ds(on)}$. In a multiphase design, switching losses are the same through each FET. However, the current in each phase is divided by the number of phases. Therefore, switching losses in a two-phase design can be reduced by:

$$\left(\frac{I_{load}}{2}\right)^2 * R_{ds(on)}$$

or, in general:

$$\left(\frac{I_{load}}{N}\right)^2 * R_{ds(on)}$$

where N is the number of phases in the design.

Additionally, a multiphase design improves the regulator's ripple currents and transient response time. This comes at an increase in cost and footprint because an inductor and two FETs must be added for each phase and because the controller becomes larger and more complex.

Narrowing the options

Power efficiency is the key to meeting the demand for extended battery life in portable products. Processor manufacturers have aided this effort by reducing operating voltages, but power supplies must adapt to maximize efficiency. When deciding on the most efficient power supply, it is important to examine the variables to assure the power supply will meet processor requirements. Power supply cost must be part of the decision as well.

The standard buck switcher provides efficiency when operating in a PWM mode and is less complex than the synchronous buck, making it a less expensive option. However, the diode creates a voltage drop, which wastes some power. The synchronous buck reduces this voltage drop using a FET in place of the diode, which increases efficiency but at higher cost.

The multimode regulator provides efficiency improvements over the full load range by switching from PWM to PFM when the processor is in sleep mode. This adds some output ripple voltage, but as long as it stays within the processor's power specification, it will significantly improve battery life. The multiphase regulator maintains efficiency while eliminating much of the output ripple voltage, also at higher cost.

Designers can follow these guidelines when designing power supplies for an embedded processor. All designs require trade-offs, and power supplies are no exception. Given the budget constraints, power requirements, and efficiency target for any design, these strategies will help narrow the options to identify a power supply that provides the best compromise of all three. ECD

Brian Law *is a field applications engineer at Fairchild Semiconductor, based in San Jose, California. Prior to Fairchild, he worked as an account executive for ON Semiconductor and a district manager for Toko. He holds a BS in Electrical Engineering from Iowa State University and is working toward an MBA from the University of Phoenix.*

Greg Ferrell *is a field applications engineer at Fairchild Semiconductor. Greg has 17 years of experience in engineering design and field applications engineering, with specific focus on design, test, and field applications for the cellular handset market. He holds a BSEE from Ohio State University and an MSEE from the Illinois Institute of Technology.*

Fairchild Semiconductor
207-775-8100
Brian.Law@fairchildsemi.com
Greg.Ferrell@fairchildsemi.com
www.fairchildsemi.com

Embedded Computer Solutions For Harsh Environments
-40°C to 85°C Operating Temperature
Five Year Product Availability Guarantee

EMBEDDED COMPUTER

Fanless Dust Tight High Performance Embedded Computer for Outdoor, POS, In-Vehicle, Marine or Applications.
- Pentium M / Pentium 4 / Celeron CPU from 1.3GHz to 2.1GHz
- Over 200,000 hours MTBF

DIGITAL VIDEO RECORDER

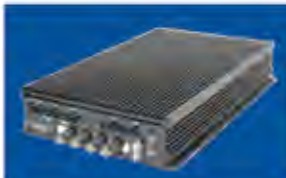

In-Vehicle Fanless Rugged Digital Video Recorder
- Up to 8 Video and Audio Channels @ 704x480 resolution (D1), 30fps per channel
- Operated from 7V ~ 28V DC-IN
- +12V DC Out for cameras
- Anti-vibration and anti-shock mounting kit

POINT OF SALE SBC

Ultra low power Embedded Computer designed for POS terminals
- Fanless operation from 200MHz to 900MHz Intel Tualatin CPU, on-board memory, LAN, CRT/LCD, Touch Screen, AUDIO, 6 Serial, 2 USB ports
- Over 250,000 hours MTBF

SYSTEM ON MODULE (SOM)

Design your own or have TME design a custom I/O board to be used with TME's SOM
- From 800MHz to 3.4Ghz CPU
- Small form factor 5.5" X 4"
- Up to 1.5 GBytes soldered on-board DDR memory
- ISA, PCI, PCI-Express expansion

EBX FORM FACTOR SBC

Pentium III 1.2GHz EBX SBC designed for **Mobile** and **Outdoor Applications.**
- Ultra low power, Passive Heat Sink for CPU up to 900 MHz, soldered on-board 256MB SDRAM
- CRT/LCD, Audio, LAN, 6 serial, 4 USB 1.1, Intelligent thermal management, PC104+ expansion

PC/104-*PLUS* SBC

High Performance PC/104-*Plus* Embedded Computer (1.2GHz Pentium III)
- **Fanless** up to 900MHz
- CRT/LCD, LAN, PCI on PC104+, ISA on PC104 connectors, 2 serial ports, dual USB, 256 Bytes EEPROM, 64-bit unique electronic ID

tme

Toronto MicroElectronics Inc.
5149 Bradco Boulevard, Mississauga, ON., Canada L4W 2A6
Tel: (905) 625 - 3203 ~ **Fax:** (905) 625 - 3717
sales@tme-inc.com ~ www.tme-inc.com

Thermal simulation optimizes cooling in high performance systems

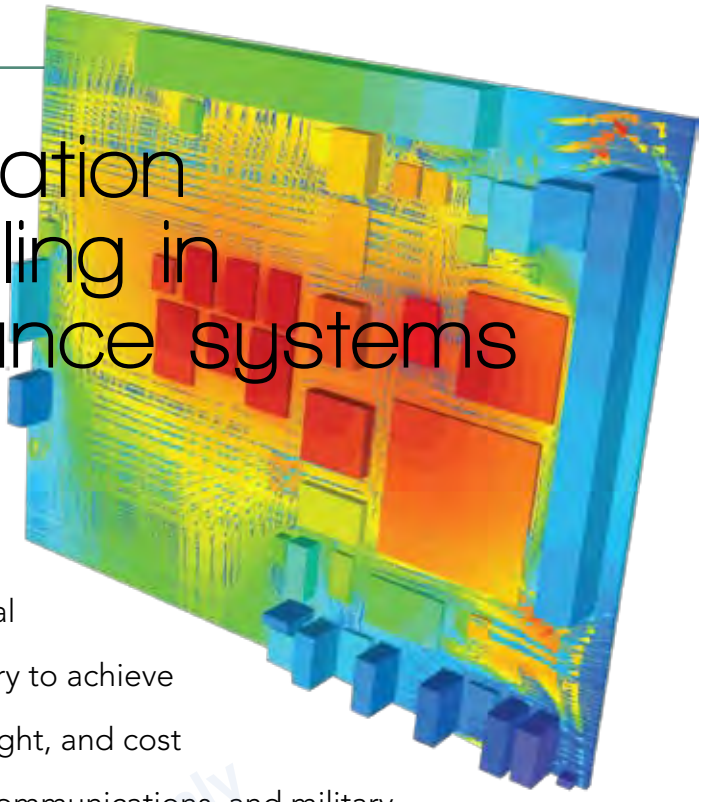

By Kevin Rhoades

Thermal issues present some of the most critical challenges embedded designers face as they try to achieve maximum performance while keeping size, weight, and cost down. Applications in medical, industrial, telecommunications, and military environments have huge performance demands, thereby requiring faster processors. As operating frequencies increase and process geometries continue to shrink, high-performance processors are becom ng the biggest culprit of heat generation, giving off as much as 120 W per processor. With power levels rising, designers need effective tools to validate thermal management early in the design process.

For Reprints, Not Only

Importance of thermal design

The lifespan of a component as well as the entire system is directly tied to the heat it generates and the system's internal temperature. Excessive temperatures can result in performance loss and, ultimately, component failures. For embedded systems that require high availability and longevity, this can be disastrous. Consequently, cooling has become an important design consideration. Where designers choose to place components, airflow vents, pressurization, and heat sinks has a measurable effect on thermal management.

To further complicate today's embedded system designs, certain applications have additional requirements for sealed enclosure designs. For example, military equipment must be able to operate in extreme temperature fluctuations and high altitudes where air is thin and ineffective for cooling – all while undergoing extreme vibration and shock. This often calls for sealed designs that rarely or never exchange enclosure air with outside air, limiting which thermal management techniques designers can employ. These types of rugged designs are often encased with outside protection using high-density rubber that retains heat. Engineers must be mindful of application requirements when determining how to dissipate heat in these designs, particularly rugged handheld computers that cannot be too hot to handle in the field.

Thermal management techniques

Embedded system engineers have many options available for addressing heat issues. Heat sinks come in a vast assortment of types and sizes. Several special epoxies allow a heat sink to be glued to a component, providing increased surface area to dissipate heat away from the IC. Mechanically modifying the embedded system enclosure and mounting custom-designed heat sinks can remove heat from individual or multiple components. The back corners of the chassis can act as heat traps, allowing special baffles to deflect air to a particular location.

Internal component-mounted as well as enclosure-mounted fans can be incorporated to provide a constant airflow directly to the components and the internal system circuit board. When system design restrictions prohibit using fans or advanced cooling systems and require hermetically sealed enclosures, designers must often resort to more creative methods. Utilizing a different board with decreased features, reduced clock speeds or an alternate CPU, and more power-efficient storage is sometimes the only way to resolve thermal problems in these designs.

However, predicting and adjusting a system's thermal footprint is only part science. Thermal management requires

When You Design for the Extreme Choose Your Partner Wisely

Did we mention "Reliability" is our middle name?

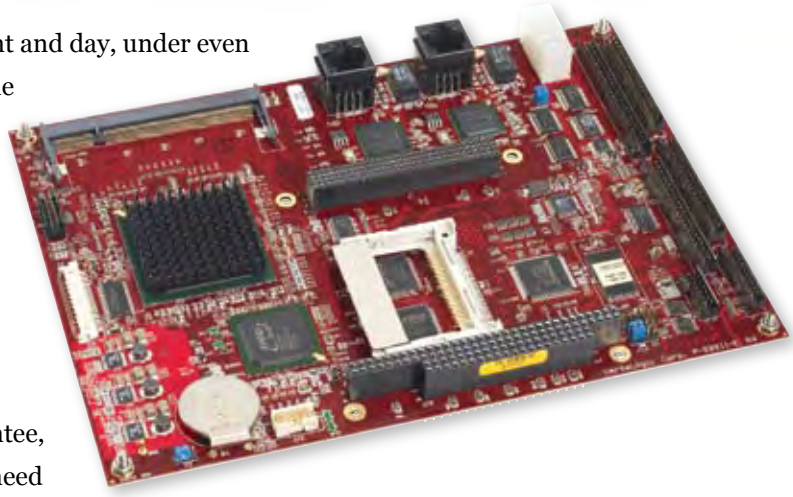

Industrial equipment needs to perform flawlessly, night and day, under even the most extreme conditions. Whether designing for the manufacturing floor, clean room or the field, you can depend on VersaLogic to deliver the highest quality embedded computer products, from prototyping and design-in, through years of product production. We design our boards for high reliability and long-term availability, then run them through exhaustive quality tests, ensuring that we deliver only the best. And with our world class service and five year availability guarantee, we'll always be there when you need us. Whether you need one of our standard products or a customized version, our skilled technical staff will work to meet your exact requirements.

Contact us to find out how for more than 30 years we've been perfecting the fine art of extra-ordinary support and on-time delivery: One customer at a time.

1.800.824.3163
1.541.485.8575
www.VersaLogic.com/ind

VERSALOGIC CORPORATION

some guesswork or tribal knowledge based on an engineer's previous experience. Yet there is no guarantee that the engineer will pick the right combination of values to investigate which one will lead down a time-consuming path of trial-and-error refinement techniques.

New technologies eliminate guesswork

To help improve the design's overall reliability and avoid costly redesigns, thermal management must be validated early in the design process so that any problems can be addressed before the hardware is built. New technologies like Computational Fluid Dynamics (CFD) software are making it possible for engineers to effectively deliver thermal management expertise early in the design process long before physical prototyping and testing.

Prior to CFD software, an engineer would attack thermal issues through a cycle of exercises. Each approach required building a prototype, then taking the thermal measurements and making modifications until the desired thermal results were achieved. This often took months of building and testing before the correct design was identified, increasing time to market.

Today, simulation tools can enable engineers to create virtual models of electronic equipment, perform thermal analysis, and test design modifications quickly (often in less than 24 hours) and easily in the early stages of the design process. With a few mouse clicks, designers can rapidly develop "what-if" models, calculate the results, and compare performances.

Flomerics' FLOTHERM software is an example of an application designed to address the challenges of modeling thermal management in electronic and electrical systems. The software enables users to predict airflow and heat transfer in and around electronic equipment, including the coupled effects of conduction, convection, and radiation.

Figure 1 illustrates the active heat-sink temperature plot of an ETXexpress module using this software. The stream in blue to red flowing through the heat-sink fins represents the airflow blowing from the fan on top. The right side shows the temperature reading starting at +50 °C from blue to red.

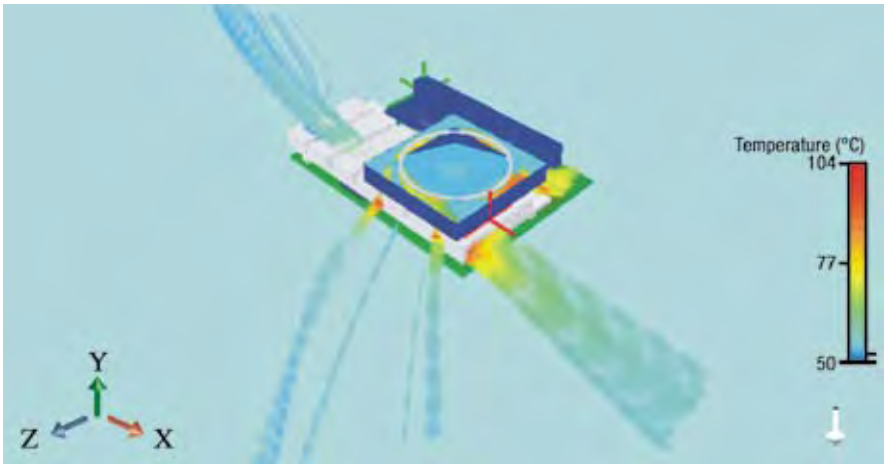

Figure 1

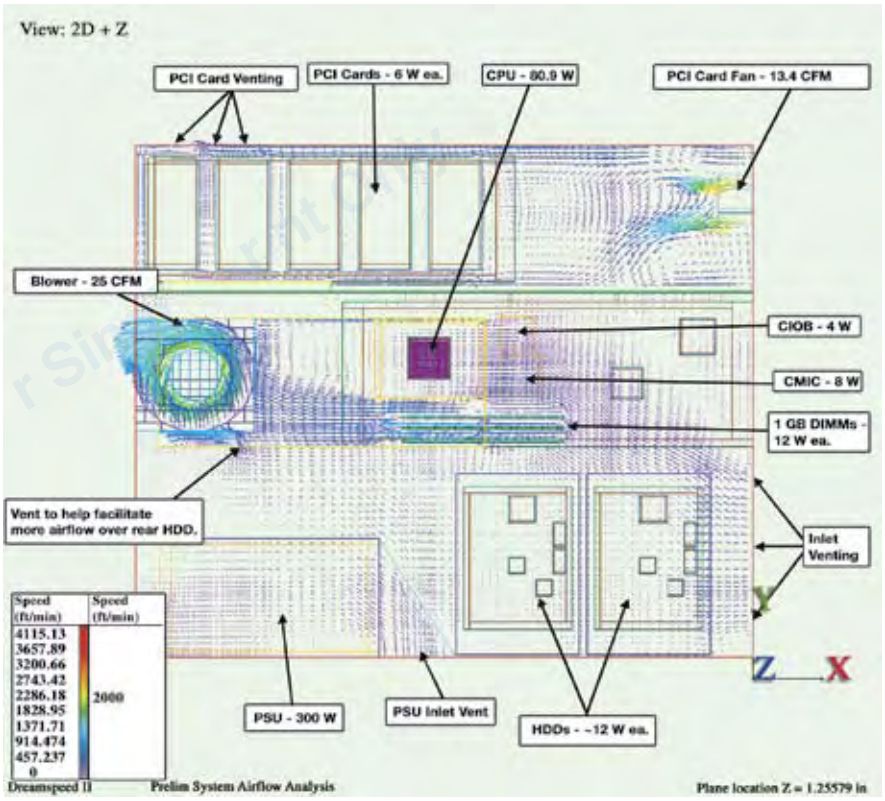

Figure 2

Thermal simulation can be helpful in designing an optimal chassis cooling system. The ability to create a model environment that predicts airflow, temperature, and heat transfer and validates the thermal design can cut design time significantly. An independent survey by the Aberdeen Group found that thermal analysis software users completed thermal design verification 3x faster and had more than 40 percent fewer re-spins on average per PCB design than nonusers.

Thermal analysis in action

Figure 2 illustrates the preliminary system airflow analysis showing airflow direction, magnitude of speed, and fluid field strength on a 1U chassis temperature vector plot. Using this analysis, engineers can determine if individual commodities on the board need to be moved or replaced or if a particular component or its placement is causing a temperature anomaly.

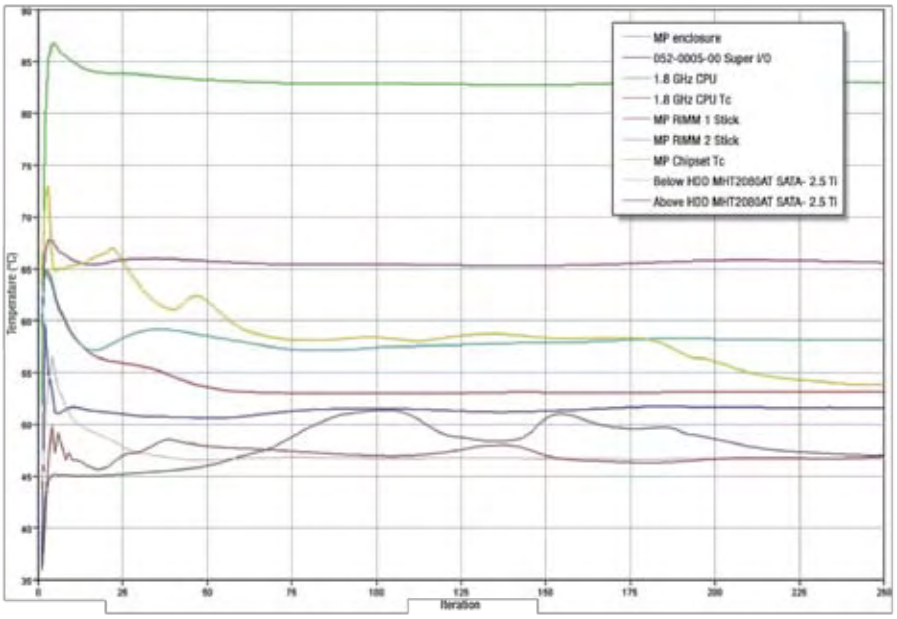

Figure 3

Kontron used FLOTHERM in a network appliance wireless server design because of its ability to reliably simulate the final system's performance and predict worst-case scenarios early in the design process. This particular system had several unknown variables and was based on proprietary cards that were available only as a specification at the time. To complicate matters, the spec had a wide range of power variables from 60-120 W.

Kontron's engineering team took the specification and assigned different values to the CPU, HDD, memory, and I/O to generate a number of scenarios. Figure 3 shows an example of a small form factor monitor point graph that illustrates the temperature of different variables, such as the enclosure, CPU, and other components within the chassis. This simulation provided designers a view into the chassis to see thermal characteristics and anomalies and allowed them to predict worst-case scenarios, including what would happen in the event of a fan or other key component failure.

Once designers validated the model using simulation tools, they built a prototype and tested it in the customer's real-world environment. Because analysis was completed early on and the results were virtually validated using simulation tools, the design held up to the customer's needs.

Tools ensure reliability

Thermal requirements must be addressed at the beginning of the design cycle in today's high-performance embedded systems. With the advent of CFD, engineers can virtually e iminate prototyping time and cost, software can pick the direction in which to move the parameters, and users can find the optimum design. Using this conceptual design process, thermal simulation can provide engineers with a key tool to reduce design time and the number of re-spins. Perhaps even more importantly, making good use of these tools enables designers to predict and plan for failures – reducing design risks and improving reliability for mission-critical applications. ECD

Kevin Rhoades *is VP of business development at Kontron America, where he is responsible for managing the North American Systems Division in Poway, California. Kevin has 25 years of experience in global supply chain and contract manufacturing and holds a BS in Business from Gonzaga University in Spokane, Washington.*

Kontron America
888-294-4558
sales@us.kontron.com
www.kontron.com

TRI-M SYSTEMS
proudly distributes

TRI-M ENGINEERING
100Mhz PC/104 Module

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

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

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

Thermal concepts enhance DRAM memory subsystem designs

By Phan Hoang

Controlling thermal issues related to the CPU is often a top priority for designers; however, the memory module is not necessarily less important. Thermal management presents challenging design considerations within embedded environments, requiring knowledge, precision, and creativity to diagnose and overcome memory subsystem design parameters.

In the past, memory was not as complex and did not warrant the same kind of thermal attention that designers paid the CPU. Because the CPU required cooling, chipsets came equipped with heat sinks as a production standard. By comparison, the memory module needed only minor airflow adjustments to keep temperatures in check. But with DDR3 and DDR4 technology increasing speeds in today's embedded designs, memory module design is complex and requires thermal attention as well.

Clock speed is just one reason why memory is running hotter than ever. The customer environment, overall design choices such as memory modules' location on the board, horizontal or vertical module orientation, and the amount of airflow over the system also can influence memory modules' thermal condition.

Embedded system designers typically work with tight board layouts requiring near-perfect engineering to achieve flawless signal integrity and extreme performance. Though other design issues exist, successful system designers consider memory thermal management as a higher-level design issue, keeping abreast of evolving memory technologies and thermal management techniques for reducing heat in memory modules.

Memory designers can mitigate heat and design better memory subsystems using a range of simple but powerful thermal concepts. Similarly, system designers can enhance products by incorporating those concepts when creating their designs.

The heat is on

Memory designers begin the process by choosing memory modules that mitigate heat and deliver the best overall heat reduction scheme. Incorporating modules that use the least amount of DRAM within the greatest number of module ranks can achieve the desired module density and manage power effectively. The more DRAM in standby mode, the less power the module consumes – frequently achieved by using DRAM with the widest data bus, as shown in Table 1. For example, a 36-chip four-rank x8 DIMM uses less power than a 36-chip two-rank x4 DIMM.

For an additional example, a 512 MB error-correcting code DIMM can be made using five 64x16 DRAM chips versus nine 64x8 DRAM, resulting in 44 percent heat reduction. The actual reduction may be slightly less due to differences in IDD

	ECC (x72)			nonECC (x64)		
DRAM	1 rank	2 rank	4 rank	1 rank	2 rank	4 rank
x16	5	10	20	4	8	16
x8	9	18	36	8	16	32
x4	18	36	72	16	32	64

Table 1

values specified in the datasheets for 64x16 and 64x8 DRAM. Memory designers would typically explore whether or not the memory controller chipset can support the wider DRAM data bus width.

Overall, memory modules appropriately spaced between DRAM either non-stacked or without large, hot semiconductors will have better thermal characteristics. Small form factor memory such as stacked very low profile or stacked SODIMMs have higher power density (watts/area) and need special considerations for cooling. Fully buffered DIMMs also have high power densities because of onboard advanced memory buffer and may require additional cooling aids or airflow.

System versus memory

Thermal sensors are critical tools for memory designers. JEDEC's standard specifies that memory modules have thermal sensors to give users monitoring and triggering mechanisms that adjust system performance according to fluctuations in temperature.

Depending on defined parameters, the system can issue an extended mode register set command, which would double the internal refresh rate on the DDR2 DRAM to a 32 millisecond period (tREFI = 3.9 microseconds) at a trigger temperature of +85 °C to extend the DRAM operating temperature to +95 °C. If that feature is not available, designers can incorporate special programming on the memory module for extended temperature operation. Alternatively, the system can use closed loop dynamic temperature throttling and fan speed control to optimize memory performance.

The key point here is that the CPU manages the memory board's thermal sensors, demonstrating that system-level and board-level thermal issues are closely related. The system's BIOS reads output from the sensor and evaluates performance options based on preprogrammed thresholds identifying acceptable temperature ranges. For example, if the memory runs over the limited temperature, the system thermal monitor alerts administrators about temperatures

Embedded Computer Solutions
For Harsh Environments
-40°C to 85°C Operating Temperature
Five year product availability guarantee

Micro-P3
1.26GHz Pentium III
System On Module

$US 275.00/each
(OEM Qty. 900Mhz CPU, 256K L2 cache)

- Choice of ultra low power Tualatin Pentium III / Celeron CPU from 500MHz to 1.26GHz with 256/512 KB L2 cache
- Up to 512MB SDRAM
- Intelligent thermal management with independent microcontroller
- Over 200,000 hours MTBF
- 5V @ 2.2A, 3.3V @ .7A with 900 MHz CPU and 256 MB memory

Back View

PCI/LPC/SMB Bus •
CRT/LCD video •
Dual channel enhanced IDE •
10/100Base-T •
AC'97 2.2 Audio •
2 RS232 Ports •
1 LPT Port •
4 USB 1.1 Ports •
8 GPIO Ports •
Keyboard and Mouse •
+3.3V •
High Density Connector

5831
Pentium III EBX SBC designed for Mobile and Outdoor Applications

- 500MHz to 1.26GHz Pentium III CPU
- Ultra low power
- Passive Heat Sink for CPU up to 900MHz
- 14W (using 900Mhz CPU, 256K L2 cache)
- -40°C to +85°C operating temperature

- Soldered Onboard 256MB SDRAM (Optional 768MB with SODIMM)
- CompactFlash Disk (up to 6 Gbytes)
- LCD/CRT Video, 10/100Base-T Ethernet interface, AC'97 Audio, TV Out
- 6 serial ports, dual USB, 256 Bytes EEPROM, 64-bit unique electronic ID
- 128/512Kbytes battery backed SRAM with 10 years data retention
- Intelligent thermal management with independent microcontroller
- Less than 4 seconds boot up time
- 8" x 5.75" standard EBX form factor
- Supports DOS, Windows 98, NT, 2000, XP, CE, QNX, pSOS, Linux, VxWorks

tme
Toronto MicroElectronics Inc.
5149 Bradco Boulevard, Mississauga, ON., Canada L4W 2A6
Tel: (905) 625 - 3203 – Fax: (905) 625 - 3717
sales@tme-inc.com – www.tme-inc.com

above defined thresholds, prompting them to take the necessary steps to lower the temperature, such as checking processors and chassis fans, addressing any chassis airflow vents that may be blocked, or adding another chassis fan.

Airflow matters

Airflow is a simple but critical issue for memory; the primary goal is to avoid blowing preheated air directly over the memory subsystem. Whenever possible, designers should place the memory subsystem on the sides of the processor and outside the flow of warm air generated by the processor's heat sink or other hot components such as the power supply or chipset. The ambient intake air should flow evenly over the memory subsystem and other hot components such as the processor.

Too small an air gap between modules may create airflow back pressure from the physically obstructed DIMM modules within the airflow path. This could result in an airflow pressure drop along the side of the DIMMs, generating a decrease in airflow, or could divert the airflow away or around the entire memory subsystem. DIMM socket spacing should be 10 mm or greater from center to center.

In general, maximizing airflow extracts heat away from the memory. Designers should use a blower or dual fans to optimize airflow if acoustical noise is not an issue. Airflow with a minimum pressure drop is best achieved by extracting the hot air at the exhaust point but also can be improved by pushing air in at the intake point. Plenums, ducts, or shrouds can be used to direct and contain the airflow through the memory subsystem, flowing parallel to the longest sides of the DIMMs and on both sides. These enclosures may allow for slower fan speeds with less acoustical noise and not affect airflow.

Memory modules can be designed to allow airflow across the short side of the DIMM, eliminating heat from being dragged across the long side of the DIMM. This type of mezzanine connector technique does not expose as much DRAM to preheated air from upstream DRAM.

If the motherboard or system board is mounted flat and perpendicular to the line of gravity, the best orientation for the memory would be a vertical mount since hot air rises up along the line of gravity. A vertical DIMM orientation prevents heat from being trapped under the lower bottom side of the memory modules. If a vertical mount is not possible, then an angled-mount DIMM orientation would benefit from one-sided DIMMs with the DRAM components mounted on the top side. This would hold true for memory DIMMs placed flat over the system board as well.

Designers should choose a module with a DRAM placement that does not allow all the DRAM devices to be active on the same side at the same time. Modules with alternating DRAM placements on each side of the memory module per rank will evenly disperse the heat surrounding the DIMM. If airflow is restricted on one side of a DIMM, memory modules with DRAM placed only on the side with the maximum airflow will perform better at

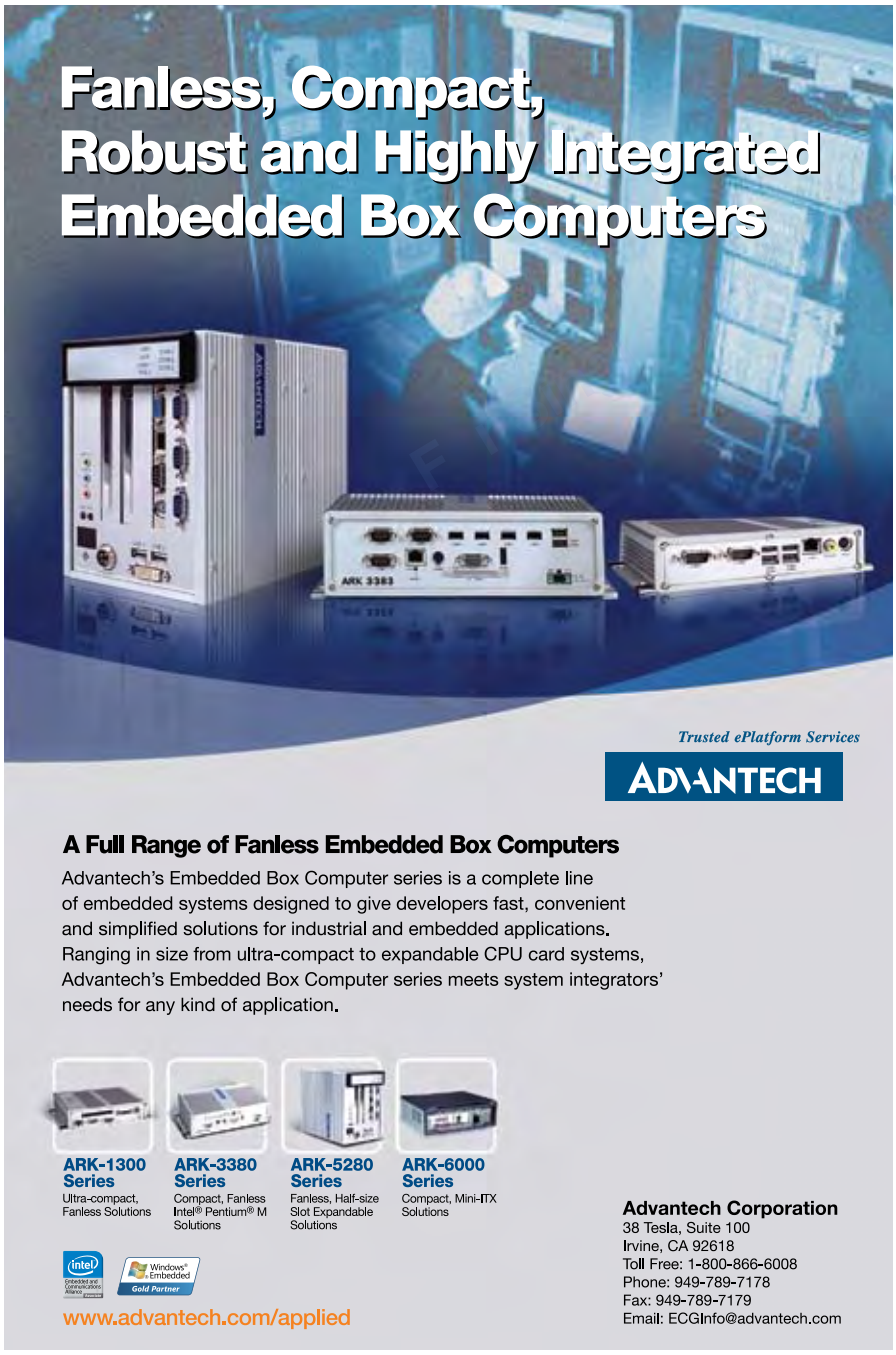

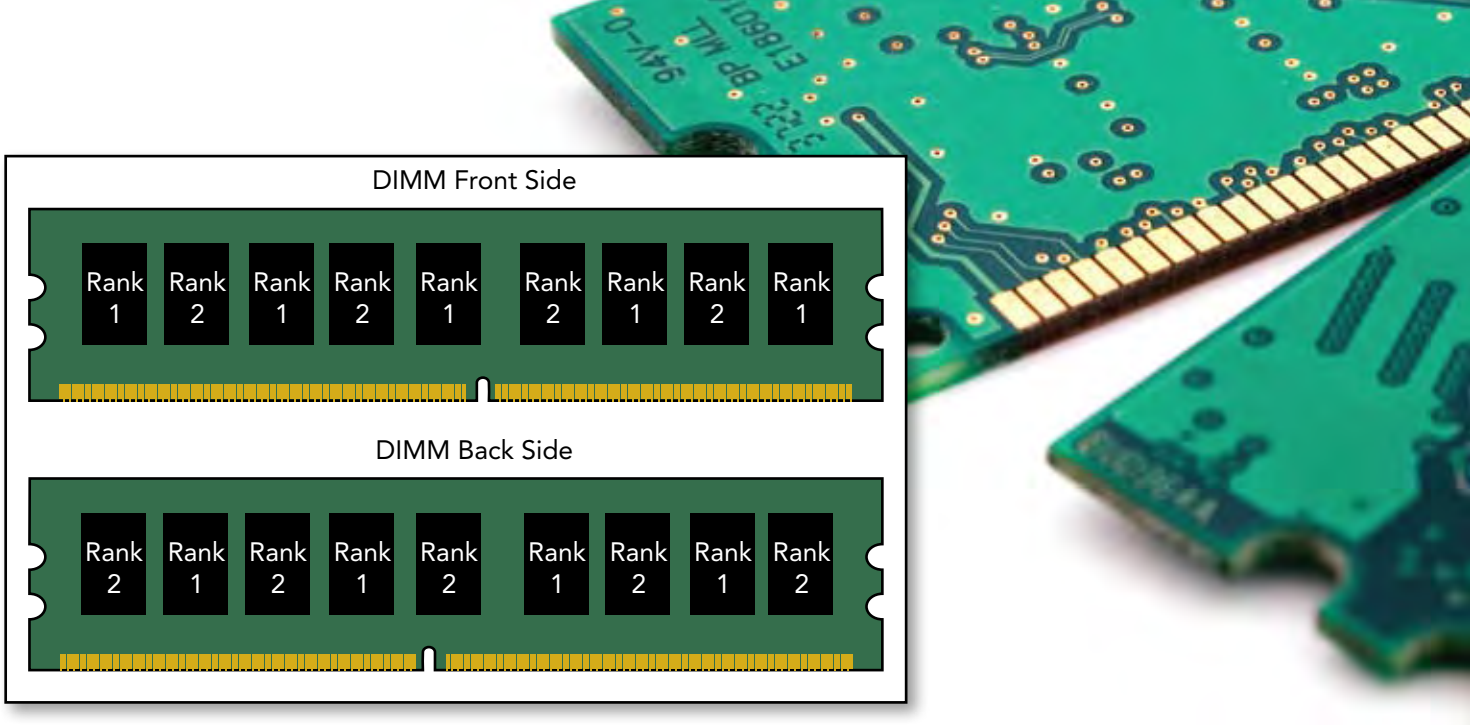

Figure 1

higher temperatures. Figure 1 illustrates how the technique of alternating DRAM ranks can reduce thermal impact.

Heat spreaders and more

Heat spreaders are metal covers placed on the surface of a memory module to disperse heat evenly across the surface and equalize surface temperature by removing localized hot spots. A heat spreader is made of a thermoconductive material such as copper or aluminum in the shape of a clamshell wrapped around the memory module.

If space permits, a heat sink placed on the side surfaces of the memory and/or top edge of the memory module will maximize heat extraction from the module. The additional surface area the heat sink adds to the memory module without affecting airflow determines its overall effectiveness.

Heat-conducting PCBs and PCB cores are also effective options. These metal or carbon composite laminated layers are embedded into the structure of the memory PCB to allow it to operate cooler than standard FR-4. The layers also equalize component temperatures by removing localized hot spots like the phase-locked loop. It is not uncommon to see numerous hot spots created via holes under hot devices to conduct heat into the core. These cores in turn conduct heat into the edge fingers of the module and can be brought to the top edge of the PCB to expose it to heat spreaders or heat sinks. The top edge of this type of PCB has the inner thermal core of the DIMM connected to an integrated heat sink at the top of the module, adding additional height to the DIMM.

During the manufacturing process, memory modules can be tested at elevated temperatures in customers' systems running customers' diagnostic software. This *active* burn-in will screen out potentially weak modules. *Passive* burn-in (on unpowered modules) has no effect on screening out DRAM with weak cells because DRAM cells are semiconductor-based capacitors that need to be continually recharged or refreshed to retain binary information. Some memory modules are available using DRAM screened for an extended operating temperature range of -40 °C ≤ Tcase ≤ +95 °C. This is a specialty item, and not all DRAM suppliers offer industrial temperature DRAM as an option for commercial temperature (0 °C ≤ Tcase ≤ +85°C).

Thermal matters across the board

Thermal management issues are evolving with memory technology and becoming critical to embedded systems' reliability and performance. The design dynamic between system designers and memory subsystem designers is also evolving and can impact designs built for endurance and performance. Trusted system-level and board-level partnerships plus greater understanding of current thermal concepts associated with DRAM memory modules can make all the difference in the final product's success.

Understanding DRAM memory module thermal considerations as part of proven system design tenets can equip designers with a new level of understanding regarding ways to improve thermal performance. General design considerations and alternative thermal options can create a winning memory subsystem design, effectively meeting system requirements for the high memory bandwidth, large memory densities, small physical space, and low cost imperative in embedded environments. ECD

Phan Hoang *is chief engineer and director of R&D at Virtium Technology in Rancho Santa Margarita, California. Phan began his career as a design engineer specializing in memory module design and holds a stacked memory package patent. He has Associate and Bachelor degrees in Computer Electrical Engineering, Business, and Management/Marketing from the University of Phoenix.*

Virtium Technology
949-888-2444 ext. 141
phan.hoang@virtium.com
www.virtium.com

Virtual channels accelerate traffic in Serial RapidIO 2.0

By David MacAdam and Robert Bishop

Further enhancing the RapidIO Interconnect Architecture, Serial RapidIO 2.0 offers system designers improved abilities to control, optimize, and accelerate traffic using Virtual Channels (VCs). By implementing techniques explored in this article, future switches can enable VC features even in Serial RapidIO 1.3 systems and facilitate the evolution to Serial RapidIO 2.0 systems.

Serial RapidIO is a high-performance packet-based technology that can be found in an increasing array of applications, including the wireless infrastructure, storage, medical imaging, and military markets. Serial RapidIO's strength lies in its efficient application-driven features and the breadth of its adoption. The switched fabric continues to offer embedded system designers more compelling features with the Serial RapidIO 2.0 specification.

With the Serial RapidIO 2.0 standard, system designers can select link rates from 1.25-6.25 Gbps and port widths from 1-16x, providing high granularity to select a port's data rate best suited to individual applications. Beyond the physical layer enhancements, Serial RapidIO 2.0 offers several higher-level features designed to provide greater control over the switch fabric traffic flow.

VCs grant the ability to control how different types of traffic flow through the system. VCs allow system designers to control packet flow by dividing the link into separate channels and assigning packets to a particular channel. The first VC, VC0, is a backward-compatible VC provided in Serial RapidIO 2.0 that operates like a Serial RapidIO 1.3 compliant link. Beyond this, Serial RapidIO 2.0 supports up to eight more VCs (VC1-VC8).

Controlling traffic

VCs offer a variety of enhancements to control data flow in the fabric. Each VC can be guaranteed a portion of the link bandwid h System designers have control over how multiple traffic types interact and can, in effect, insulate them from one another through bandwidth allocation. Latency-sensitive traffic, such as streaming video, can be allocated a guaranteed portion of bandwidth throughout the switch fabric. This enables system designers to ensure a high-quality experience for consumers because a minimum level of performance can be guaranteed regardless of other traffic present in the switch fabric.

When a VC demands less than its guaranteed bandwidth, Serial RapidIO 2.0 allows other VCs to use that available bandwidth, thus maximizing link utilization. In essence, bandwidth allocation is intelligent, simultaneously ensuring that a greedy VC demanding more than its allocation cannot rob bandwidth from another VC and that bandwidth does not go unused when packets need to be sent.

Figure 1 shows the benefits of VC bandwidth reservation. Three packet streams share a link and are allocated bandwidth such that 10 percent is dedicated to VC_A, 60 percent to VC_B, and 30 percent to VC_C.

In the first part of the simulation, only VC_A and VC_C need packets transmitted, so their bandwidths increase beyond their allocations to take advantage of the unused portion allocated to VC_B. As shown, VC_C occupies 75 percent of the available link bandwidth while VC_A uses the remaining 25 percent. Serial RapidIO 2.0 allows VCs with packets present to access the unused part of the link in proportion to their bandwidth allocations.

Later in the simulation, VC_B needs traffic sent via the shared link. Serial RapidIO 2.0 allows the switch to rapidly respond to the change in traffic and alter bandwidth usage as needed to match the programmed allocations. In this case, the switch quickly gives an intermittent but latency-sensitive stream such as VC_B its 60 percent allocation. Once the packets from VC_B are transmitted, the switch reapportions link usage to the remaining VCs with packets awaiting transmission.

Changing channels

Serial RapidIO 2.0 provides an additional level of control over link partitioning by offering a function unique to VC0. VC0 can be configured to obey bandwidth reservations. Alternatively, it can be configured to automatically obtain any bandwidth it requires while any remaining bandwidth is divided among all other VCs with packets

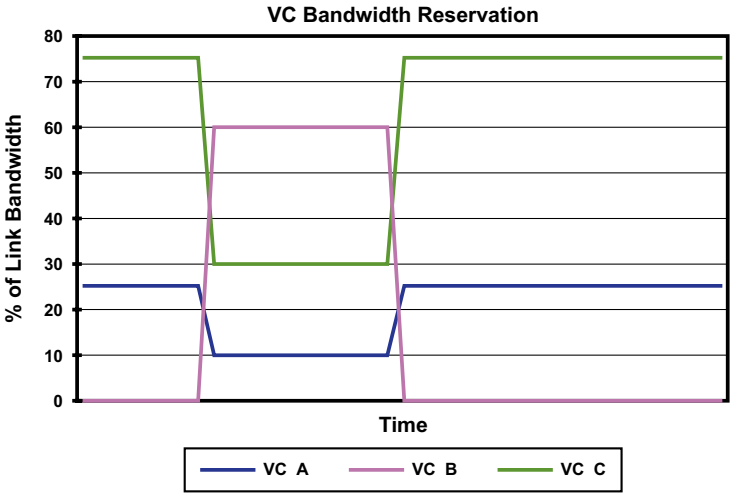

Figure 1

awaiting transmission. This allows control plane traffic transmitted via VC0 to operate completely independent of data plane traffic and only be constrained by Serial RapidIO 1.3 compliant priority rules.

VCs offer two modes of packet transmission – Continuous Transmission (CT) and Reliable Transmission (RT). RT operates like earlier versions of Serial RapidIO in that retransmitting a packet when it cannot be received makes packet transmission lossless. CT is optimized to achieve low latency for traffic flows that can accommodate packet loss by not performing retransmissions. VC0 supports all defined priorities and operates exclusively in RT mode.

Higher VCs (1-8) can operate in CT or RT mode, allowing customers to optimize the transportation method for different types of data. For example, whereas control plane traffic may require the responses and guaranteed delivery that RT mode provides, data plane traffic (such as from an audio stream) may benefit from CT mode's reduced latency and possibly suffer if retransmissions are performed.

Switches take on the challenge

Integrating a specification's new features into present and next-generation systems presents a challenge. Backward compatibility is a cornerstone of the Serial RapidIO specification development process, but promoting progression toward the new standard is equally important. System developers will begin to see more Serial RapidIO 2.0 compliant products in the near future. However, the ecosystem cannot completely transition at once.

Switch vendors will likely be among the first to adopt Serial RapidIO 2.0. Switches serve as the foundation of embedded fabric ecosystems and, in essence, validate new specifications. Processing end points, such as DSPs and FPGAs, may lag in producing new devices compliant to a specification's latest revision. Switch vendors are thus presented wi h a unique challenge to produce next-generation solutions and, at least initially, market them into sys ems limited to current-generation technology based on available end points.

Designers will migrate to end points that support Serial RapidIO 2.0 as existing systems evolve or new ones are initiated. Inevitably, Serial RapidIO 1.3 and Serial RapidIO 2.0 availability will overlap for an extended period of time. Systems that require Serial RapidIO 1.3 subsystems to communicate with Serial RapidIO 2.0 compliant subsystems will likely exist for an even longer period of time.

However, switch vendors can enable many of Serial RapidIO 2.0's benefits independent of available end points because these benefits are focused on the switching fabric. The advantages of using VCs are so compelling that system designers will desire support for adopting the technology. Thus the challenge for switch vendors is to implement transitional support for VC functionality within the Serial RapidIO fabric in systems that may be dominated by Serial RapidIO 1.3 compliant end points. Switch vendors also must provide a simple transition mechanism so that as systems evolve to contain Serial RapidIO 2.0 devices exclusively, the fabric and its use of VCs will evolve as well.

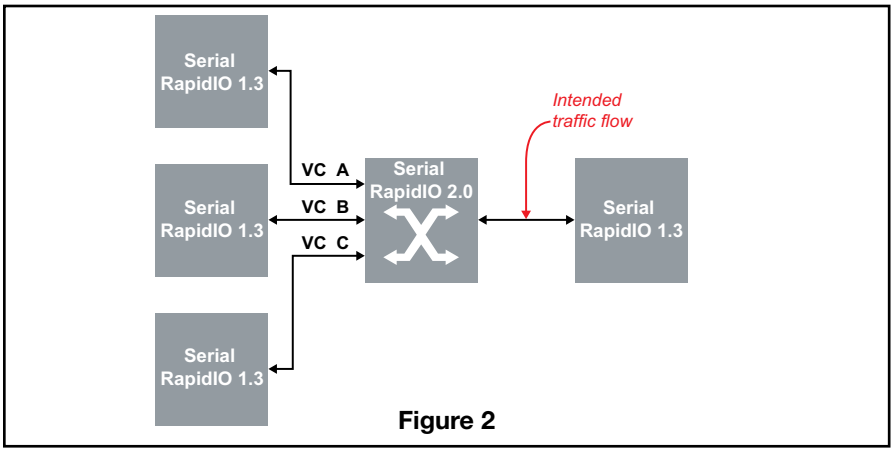

Figure 2

Taking advantage of VCs in a Serial RapidIO 1.3 system

Switches that support VCs must provide internal circuitry dedicated to routing VC packets and supporting Serial RapidIO 2.0's bandwidth allocation requirements. Serial RapidIO 1.3 compliant packets can leverage these paths when the switch routes an incoming Serial RapidIO 1.3 packet as if it was a Serial RapidIO 2.0 packet with a higher VC value.

Switch vendors can allow customers to program a mapping protocol to treat incoming Serial RapidIO 1.3 packets as Serial RapidIO 2.0 packets in terms of buffer utilization, switching algorithm decisions, load balancing, and bandwidth reservations. This function gives system designers more control over the fabric defined in Serial RapidIO 2.0, even if they cannot generate higher VC packets, and allows them to dictate that different data flows be treated as though the entire system is using Serial RapidIO 2.0 devices. Figure 2 illustrates this concept.

Although the devices only generate Serial RapidIO 1.3 compliant traffic, the switch can be configured to treat certain inbound traffic as if it has a higher VC, essentially mapping a Serial RapidIO 1.3 packet to a Serial RapidIO 2.0 packet with a programmed VC value.

One method of telling the switch how to treat a particular packet is using a unique destination identifier. Such an enhancement enables the switch to route the packet to the correct end point while restricting the packet's outbound flow based on the bandwidth allocation assigned to the VC for which it is mapped. Using this feature, the switch can treat three different end points as VC_A, VC_B, and VC_C generators, and the intended traffic flow to the destination can mimic the flexibility and responsiveness as shown in Figure 1, all without using Serial RapidIO 2.0 end points. Although potentially limiting transmitter-based flow control, this VC mapping capability allows system designers to more closely approximate the intended traffic flow in the system rather than restrict the user to obey only the Serial RapidIO 1.3 packet ordering rules.

Another key element of VCs in Serial RapidIO 2.0 is the ability to operate in CT mode, allowing packets to be dropped in favor of reduced latency. Any migration effort to facilitate VC adoption must address this crucial feature. Unlike Serial RapidIO 2.0, Serial RapidIO 1.3 devices will not accept a packet if buffer space is not available.

Switch vendors can approximate Serial RapidIO 2.0 CT mode operation with Serial RapidIO 1.3 end points by offering a pseudo-CT mode wherein a new packet replaces a packet that needs to be retransmitted. This mode allows customers in a Serial RapidIO 1.3 system to leverage the latency benefit that results from VCs operating in CT mode.

Communicating between Serial RapidIO 1.3 and Serial RapidIO 2.0

As new Serial RapidIO 2.0 end points enter the market, system designers will need a technology that can bridge the two specifications so that subsystems designed for Serial RapidIO 2.0 can leverage their advantages while still communicating with legacy Serial RapidIO 1.3 subsystems. Serial RapidIO switch providers can help system designers in this regard by allowing them to migrate their systems as the ecosystem evolves.

Switch providers can offer another powerful feature – a switch that operates as a translator between Serial RapidIO 1.3 legacy traffic and Serial RapidIO 2.0 VC traffic. Serial RapidIO provides a built-in backward-compatible translation of VC values since the portion of the header that specifies the VC value reuses the priority and critical request flow fields. Although this allows Serial RapidIO 1.3 devices to pass Serial RapidIO 2.0 packets that specify a VC value, it is a static translation that may not meet the needs of system designers who must support both versions of Serial RapidIO in the same system.

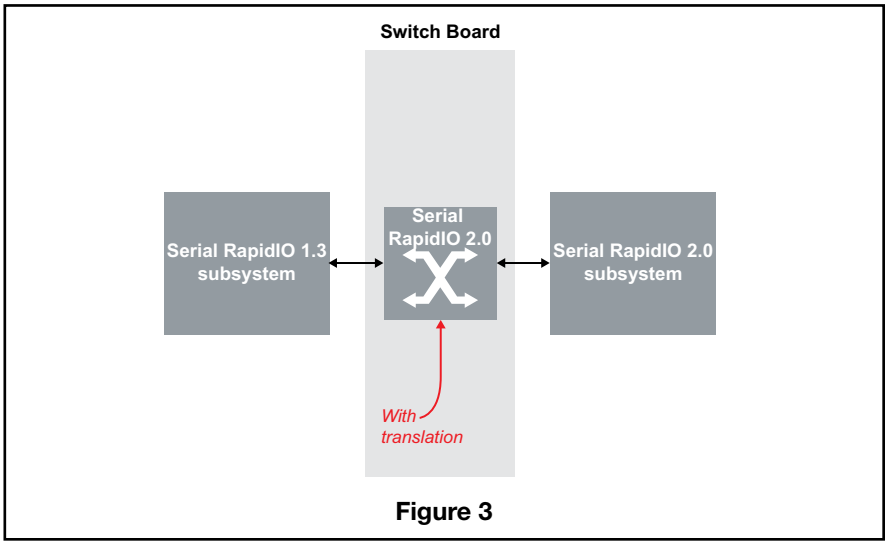

Figure 3

However, a switch can operate as an intelligent interconnect between two subsystems that use different versions of Serial RapidIO by offering a programmable translation capability. Consider a system that includes a legacy Serial RapidIO 1.3 system board, a Serial RapidIO 2.0 system board, and a switching board with translation capability, as shown in Figure 3.

The switch can have a programmable translation function to optionally convert packets between the two systems, enabling system designers to grow into the new specification as subsystems are revised, new subsystems are introduced, and the overall Serial RapidIO 2.0 ecosystem expands to include all necessary components. Rather than simply mapping packets to a VC and changing how the switch handles each packet, this system changes the inbound packet's header, recalculates the cyclic redundancy check, and essentially generates a new packet.

Accelerating Serial RapidIO 2.0's success

Serial RapidIO 2.0 promises to be a powerful embedded fabric with robust features supporting all manner of data traffic. VCs are major components that will change how system designers define and control data flow through their fabric. The challenge for switch vendors is to provide customers with powerful Serial RapidIO 2.0 systems that leverage the compelling features of VCs while offering an easy path to evolve into a complete Serial RapidIO 2.0 system as the ecosystem continues to grow. ECD

David MacAdam *is an architectural engineer with Integrated Device Technology (IDT), based in San Jose, California, where he focuses on next-generation product concept modeling, analysis, and specification. Prior to this role, David spent two years as an RTL team lead and four years in various aspects of full custom digital design. David received his BAS in Electrical Engineering from the University of Waterloo in 1999. David holds four patents in multiple areas of semiconductor design, with 11 patents pending.*

IDT
678-775-2961
David.MacAdam@idt.com
www.idt.com

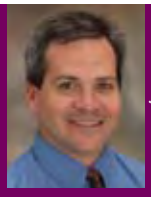

Robert Bishop *is a director of technical staff with IDT. He leads a cross-functional team responsible for architecture and algorithm research to support next-generation IDT product definition. During the past 13 years at IDT, Robert has led a broad scope of projects, including global teams developing network search engines and local teams developing dual-port and FIFO memories and logic devices. Prior to IDT, Robert worked at IBM as a custom circuit designer for six years. Robert received his BSEE from the University of Pittsburgh in 1988 and his MSEE from the University of Vermont in 1992. Robert has been awarded five patents, with five additional patents pending.*

678-775-2928
Robert.Bishop@idt.com

Virtual hardware platforms: Productivity proven for software development

By James Clough

Software developers have a love-hate relationship with hardware development boards. On the one hand, they make software come to life. On the other, they can cause headaches, especially since development schedules depend on their availability. Limited visibility in development boards' behavior for software debugging often forces developers to use more complex approaches. And sometimes they simply do not work the way developers expect. Virtual hardware platforms can help resolve developers' frustrations by executing software quickly and thus shortening development cycles.

Complex multicore platforms such as those used to design cellular base stations are increasingly presenting challenges for software developers, compounding the reasons why they tend to despise hardware development boards. A recent RadioFrame Networks design project demonstrates this dilemma, which leads to developers' double-edged relationship with hardware.

An alternative development methodology

After the initial hardware specification, the development team considered using a reference development board that slightly resembled the preliminary hardware design. None of the design peripherals or the DDR RAM controller would have matched the development board; however, there was nothing else available at the time. Once the final hardware became available, developers would have had to blindly write and then modify the software because they did not have all the hardware components on the reference board. It was the only method they thought they could use to get the job done.

Then the development team came across virtual platforms for software development, a type of technology that simulates hardware models and executes software at speeds close to real time. Although developers were interested in this concept, they had some reservations stemming from a lack of understanding about the modeling technologies and methodologies used to create a virtual hardware platform. Developers wondered if the flash model would be common flash interface/scalable command set compliant and if the new MAC controller would be functional enough to provide a head start on developing the driver. Since the ARM core selected for the ASIC did not use the full ARM instruction set, developers wanted the virtual hardware platform to capture illegal instructions attempted by the kernel as well as trap and report illegal register writes.

As the team started the project, it became clear that some of the concerns were not warranted. They quickly learned that modeling uses a standard language called SystemC, a subset of C++ specifically designed for modeling hardware, and a methodology called transaction-level modeling. The concepts in SystemC were very natural to the developers; thus, modeling the flash in this project was not an issue.

Using this technology saved modifications from session to session. Although the specific core in this case was not modeled, developers still gained control over simulation and stopped execution when illegal instructions occurred, which they accomplished by using a Tool command language (Tcl) scripting feature.

During initial modeling, developers discovered that communication between software and modeling teams was essential to understand what could be achieved with a virtual hardware platform. They learned that they needed to consider virtual hardware platform modeling itself and pay attention to peripheral models and the functionality they support. Developing the virtual hardware platform involved interaction between IP and tool suppliers. After a few weeks of modeling, developers started using the virtual hardware platform to commence software development.

Fast feedback, simulation

The initial software development task was to develop a Linux support package and U-Boot monitor. From the outset, the virtual hardware platform provided valuable feedback, giving developers the ability to determine if they were on the right track. Oversights in the initial assumption for the board support package development were quickly captured and solved.

One specific issue involved the advanced high-performance bus controller. Support had to be included for swapping flash and DDR RAM during the initial boot. The virtual platform quickly helped developers identify and correct how the jump was set up – a simple feature available in the virtual platform, but one that would have required a JTAG tool to catch it in the physical hardware. The modeling done in the platform enabled developers to not only instrument the code but also instrument the platform. The virtual hardware platform provided a view of any and all peripheral states, if desired, without affecting the operation.

During the software development process, the Virtual Platform Analyzer from CoWare (Figure 1) allowed developers to observe and control the virtual hardware platform and use it to efficiently trace peripheral block accesses by initiators. In particular, breakpoints could be placed on peripheral block accesses and specific debug messages could be used through a Tcl application programming interface. The same Tcl scripting capability also enabled developers to adapt the virtual hardware platform to their development needs and thus validate programming of the hardware configuration in the firmware by simulating timing-related configurations without requiring models to be time-accurate. As a result, the team enjoyed fast simulation speed and avoided having to wait for the physical hardware.

Valuable insight into code

In this project, using a virtual hardware platform shortened the development cycle by 33 percent compared to using the physical hardware. The virtual hardware platform provided a *pre-silicon* software development testing environment.

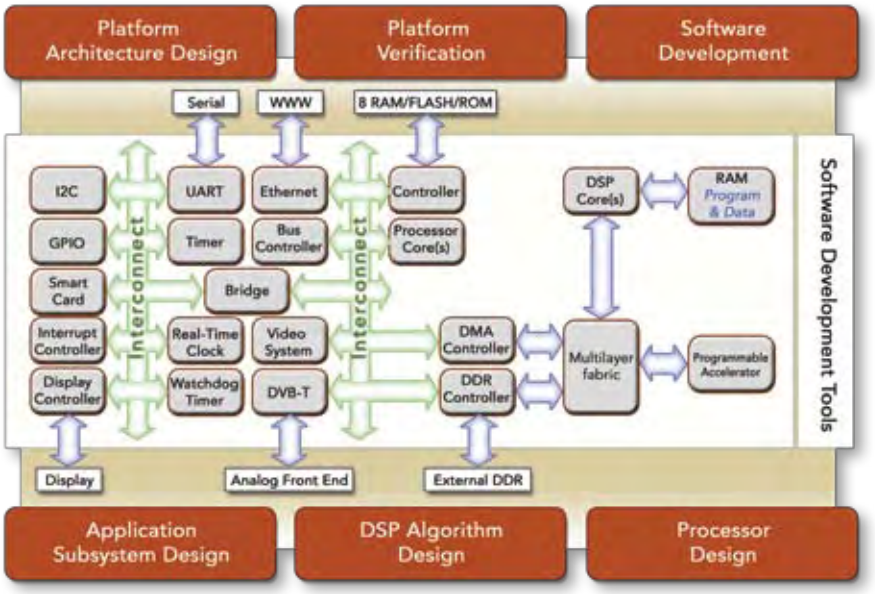

Figure 1

In addition, its unique debug and analysis capability made it superior to similar debug and analysis capabilities provided on the physical hardware.

Given the success of this project, it is evident that virtual hardware platforms can significantly improve productivity in software development teams. With a rapid increase n multicore platform development, the level of visibility provided by a virtual hardware platform can give application developers insight into code they could not see before without specialized equipment.

Therefore, developers should consider using a virtual hardware platform to correct code. The value of doing so far outweighs the initial modeling investment. Communication, education, expertise, and other advantages gained from virtual hardware platform technology suppliers such as CoWare can alleviate any concerns.

To appreciate the technology, consider the benefits of white-box testing versus black-box testing. Making the hardware set logging levels allows developers to record various accesses performed by the operating system and application programs. Designers also can place hardware breakpoints on register accesses down to the bit level, visually verify changing states on interrupts and other discrete signals, and extend the virtual hardware platform's capability through Tcl script procedures such as setting hardware watchpoints and breakpoints. Most importantly, developers can accomplish all this on their workstations without complex hardware setup, cables, and unstable hardware boards.

No more waiting for hardware

Virtual hardware platforms promise a bright future for developers, sparing them the pain of waiting for hardware availability and providing debugging capabilities that would not be possible with hardware development boards. Virtual hardware platforms are definitely productivity-proven, production-ready tools for software developers in this decade and beyond. ECD

James Clough *is currently working as platform software engineer at RadioFrame Networks, a complete radio access system provider based in Redmond, Washington. James has more than 20 years of experience in embedded systems development, covering a wide range of embedded operating systems, including VxWorks, Linux, pSOS, and Windows CE as well as various hardware architectures, including PowerPC, ARM, MIPS, XScale, and 68K cores. He attended the University of Washington.*

RadioFrame Networks
JClough@radioframenetworks.com
www.radioframe.com

Don't blow your stack: Static stack analysis for high integrity systems

By José F. Ruiz, PhD

Stack overflows can have catastrophic consequences in embedded software controlling safety- or security-critical systems. Static stack analysis tools provide the best choice to ensure reliable execution.

Every high-integrity system and, in general, any system that requires reliability must provide evidence demonstrating the impossibility of stack overflow. Static verification is mandatory for high-integrity systems because it is the only way to ensure that worst-case scenarios are found. Safety-critical standards such as DO-178B require worst-case memory utilization to be known, and stack memory is no exception.

The stack is the memory area that stores local information for the subprograms being executed. A stack overflow occurs when there is an attempt to store more data on the stack than what can fit. The consequences of a stack overflow depend on the operating environment and execution context. If this event is properly detected and handled, execution can proceed in error-recovery mode. If this situation remains unnoticed, memory can be corrupted, leading to unpredictable execution.

Stack overflow situations are very difficult to detect and debug, and even when they are detected, recovery is complex and limited. Its effects are random and difficult to reproduce because usually the code that overflows the stack does not go wrong, but instead corrupts the data associated with a different part of the software. Hence, it is of paramount importance to avoid stack overflows in the first place by means of analyzing worst-case stack requirements.

Approaches to analysis

The two approaches to stack requirement analysis are based on either dynamic testing or static analysis techniques. Dynamic testing-based approaches usually involve measuring the maximum amount of memory used while running or simulating the application. Static analysis techniques entail computing per subprogram stack consumption combined with control flow analysis.

The major weakness of dynamic testing-based approaches is they cannot guarantee that the worst-case execution path has been covered during the testing campaign, which is why this method is not valid for high in egrity systems.

When fully applied, static stack analysis techniques provide worst-case results that can be safely used for dimension stacks, ensuring reliable stack memory usage and thus guaranteeing safe execution. Unaffected by the limitations of dynamic testing, this technique rapidly provides precise results without much effort early in the development process.

It is worth noting that static stack analysis is no substitute for clean and precise programming. While static analysis is useful, it will not magically turn carelessly written code into good code. Stack analysis needs to know the stack requirements for each subprogram as well as all feasible control flow during execution. Hence, several constructs typically excluded for high-integrity systems threaten this analysis, including:

- **Cycles in the control flow graph:** In the presence of control flow cycles involving stack consumption, the worst-case boundary is a function of the maximum number of cycle iterations.
- **Indirect or dispatching subprogram calls:** Statically determining the target subprogram when dereferencing an access to the subprogram requires data flow analysis, which is not always possible. In Ada, these occurrences can be forbidden using the restriction *No_Access_Subprograms*. Dispatching calls in object-oriented programming are another instance of indirect calls. In Ada, dispatching calls can be eliminated using the restriction *No_Dispatch*, which prohibits class-wide constructs. However, runtime dispatching calls are problematic when runtime dispatching is inherent to the language, as in Java.
- **Dynamic stack allocations:** The presence of dynamically sized local objects, such as those whose size depends on an input argument, introduces variable, nonstatically bound amounts of stack usage on the paths where they appear.

One approach to static stack analysis is parsing the object or assembly code to extract the code that performs stack allocations and calls to subprograms. This technique can provide accurate information; however, it requires a complex machine code parser and is very sensitive to compiler changes and optimization options.

Leveraging the compiler

A more advantageous approach is leveraging the compiler itself, which is well positioned to know everything about stack allocations and subprogram calls for the code it processes. Three steps are needed to analyze stack usage:

1. Generating the basic stack consumption and call graph information
2. Computing the complete call tree using control flow analysis, decorated with local stack requirements
3. Analyzing worst-case stacks for any entry point in the application (see Figure 1)

The compiler knows the generated code and where it comes from, which is valuable when developers want to point directly to some user-level construct that jeopardizes static analysis. It is also important to note that a compiler has visibility in semantic information that can help tackle some of the previously identified challenging constructs.

Consider an Ada *Integer* subtype with range 1..5, for example. If a variable of this subtype is used to size a local array, the range may be used to compute a bound on the array size and corresponding stack allocation. Furthermore, using subprogram profiles and actual references to subprograms, the compiler can provide a limited yet in-depth list of subprograms possibly reached by an indirect call.

Detecting problematic constructs

AdaCore's GNATstack is a static stack analysis tool that uses a specialized extension to the GNU compiler collection back end to produce the required call graph and stack usage information. It constructs the full call graph annotated with local stack requirements and then performs a depth-first traversal of the tree to extract the worst-case analysis in terms of stack usage (see Figure 2).

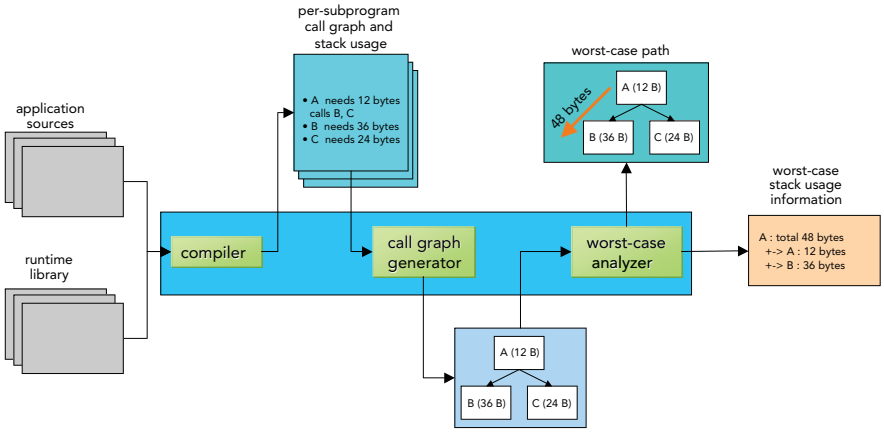

Figure 1

This static stack analysis tool also can detect and report the problematic constructs mentioned previously (cycles, indirect calls, dynamic frames, and missing stack information), providing an nfrastructure that allows users to specify the required information, such as potential targets for indirect calls, stack requirements for ex ernals calls, and user-defined bounds for unbounded frames. Avoiding these program constructs is often part of already established coding guidelines in environments where preventing stack overflows is a real concern. GNATstack ensures that these harmful situations do not go unnoticed.

Relying on the compiler back end to obtain stack information works smoothly at any optimization level for any target machine and allows mixing different programming languages (Ada, C, C++, and so on).

Stack overflow prevention

During the development phase, static stack analysis provides early feedback that helps tackle undesired constructs and hot spots with respect to stack usage. In the validation phase, any possible execution path is analyzed, thus guaranteeing that the computed worst-case stack can never overflow. This characteristic makes static analysis the only viable approach for high-integrity systems and for any system that requires reliability. The generated information can be used as evidence for certifying high-integrity and high-reliability applications. ECD

```
1 package Pck is
2   Size : constant := 5;
3   type Data is array (1 .. Size) of Integer;
4   type Operations is (Shift, Swap);
5   procedure Process (Val : in out Data; Op : Operations);
6 end Pck;
```

```
1 package body Pkg is
2   procedure Shift (Val : in out Data) is
3     Tmp : Integer := Val (Size);
4   begin
5     for Index in 1 .. Size - 1 loop
6       Val (Index + 1) := Val (Index);
7     end loop;
8     Val (1) := Tmp;
9   end Shift;
10  procedure Swap (Val : in out Data) is
11    Result : Data;
12  begin
13    for Index in Data'Range loop
14      Result (Index) := Val (Size – (Index - 1));
15    end loop;
16    Val := Result;
17  end Swap;
18  procedure Process (Val : in out Data; Op : Operations)is
19  begin
20    case Op is
21      when Shift => Shift (Val);
22      when Swap => Swap (Val);
23    end case ;
24  end Process;
25 end Pkg;
```

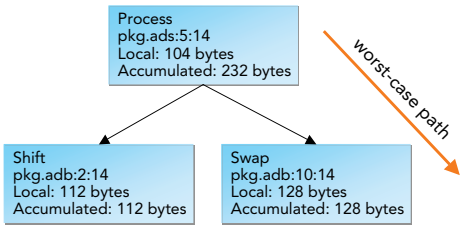


```
Accumulated stack usage information for entry points

pkg.process : total 232 bytes
 +-> pkg.process at Process:pkg.ads:5:14 : 104 bytes
 +-> pkg.swap at Swap:pkg.adb:10:14 : 128 bytes
```

Figure 2

José F. Ruiz *is a senior software engineer at AdaCore, based at the company's European headquarters in Paris, where he specializes in real-time, embedded, and certifiable systems in Ada. He has 15-plus years of software development experience and has authored/coauthored more than 20 technical papers. José holds a PhD from the Technical University of Madrid for his work in the field of real-time and multimedia systems.*

AdaCore
+33 1 49 70 67 16
ruiz@adacore.com
www.adacore.com

UNIT UNDER TEST

Real-world Certified Wireless USB testing

By Onno Harms and Rob Brownstein

As wireless USB makes its debut, designers are questioning if development and production testing align with realistic cost expectations. Using a test system built to capture and analyze wireless USB signals provides comprehensive test results that meet development and production needs within budget and time constraints.

The stars are now aligning for Certified Wireless USB. Chipset makers are finally providing silicon for band group 6 (the only band group universally approved), and more reference designs are starting to appear.

It's only a matter of time before Certified Wireless USB will show up in real-world applications. Soon, users will simply plop their digital cameras down next to their desktops or laptops and transfer photos with the press of a button. Printers will turn a document or spreadsheet into a hard copy with nary a wire between laptop and printer. All the ubiquity that wired USB now offers will be further enhanced by wireless connectivity.

But before designers break out the champagne, they must weigh the realities of wireless USB pricing versus the costs of calibration and testing. As an offshoot of WiMedia Ultra-Wideband (UWB), Certified Wireless USB is being asked to do a lot of things: reflect the utility of wired USB, have low power consumption, coexist with other wireless connectivity modalities, and be low cost. At the same time, Certified Wireless USB must meet a challenging set of parametric specifications, which is where calibration and testing enter the equation.

A closer look at the technology

First, designers should take a brief step back and determine which of the two incompatible versions of wireless USB – WirelessUSB and Certified Wireless USB – they will use.

Backed by Cypress Semiconductor, WirelessUSB is a protocol that uses the 2.4 GHz band with a maximum data rate of 1 Mbps at a distance of 10 m and 62.5 Kbps at 50 m. This article will focus instead on Certified Wireless USB, sponsored by the USB Implementers Forum.

Certified Wireless USB spans 3.1-10.6 GHz, which is divided into 14 contiguous 528 MHz bands. Signals "hop" across three adjacent bands (called a *band group*, covering a spectrum of 1,584 MHz), characterized by fast hopping and short symbol times (242.5 nanoseconds). This technology uses Multiband Orthogonal Frequency Division Multiplexing (MB-OFDM), and depending on the intended data rate, uses Quadrature Phase Shift Keying (QPSK) or Dual Carrier Modulation (DCM), as shown in Table 1. Certified Wireless USB can pump out 480 Mbps at up to 3 m and 110 Mbps at 10 m.

Data rate (Mbps)	Modulation	Coding rate (R)
53.3	QPSK	1/3
80	QPSK	1/2
106.7	QPSK	1/3
160	QPSK	1/2
200	QPSK	5/8
320	DCM	1/2
400	DCM	5/8
480	DCM	3/4

Table 1

The band groups of three 528 MHz bands have different regional regulatory constraints (see Figure 1), and only band group 6 is globally authorized. Thus, the newest chipsets are likely to converge around group 6, whereas the first chipset generation was group 1 oriented because of its allowance in the United States.

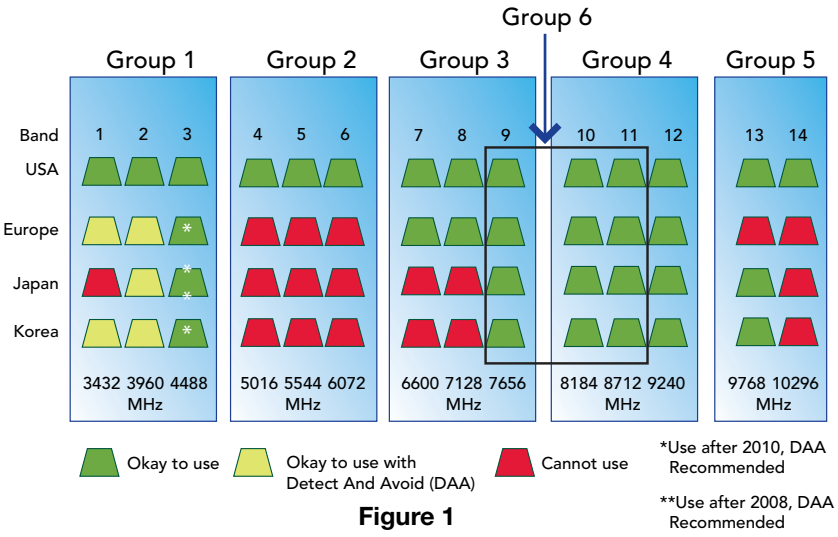

Figure 1

Standardize your PXI platform

with a Mass InterConnect Solution from VPC!

- Quickly change from one fixture or UUT to another
- Digital, analog, RF, power, pneumatic, vacuum, fiber optic and thermocouple connections
- Allows for modularity and expansion
- Systems range from 100 to 14,000 connections
- Reduces and eliminates false failures as well as incorrect wiring
- Reliable connections to 20,000 cycles

With over 47 years of engineering excellence, VPC provides **Mass InterConnect Solutions** for test and measurement applications in defense, aerospace, automotive, medical, contract manufacturing, and other high-volume testing environments. VPC's Mass InterConnect solutions are available for all major instrument platforms, GPIB, VXI, PXI, LXI, etc., delivering **rapid and reliable unit under test changeover, and reduced test costs.**

VPC Virginia Panel Corporation

1400 New Hope Rd. • Waynesboro, VA 22980-2647 • +1.540.932.3300

vpc.com/pxi2

© 2008 Virginia Panel Corporation. All rights reserved.

TFC	Band group 1 band sequence					
1	1	2	3	1	2	3
2	1	3	2	1	3	2
3	1	1	2	2	3	3
4	1	1	3	3	2	2
5	1	1	1	1	1	1
6	2	2	2	2	2	2
7	3	3	3	3	3	3
8	1	2	1	2	1	2
9	1	3	1	3	1	3
10	2	3	2	3	2	3

Table 2

The UWB common radio platform supports a variety of band-hopping sequences, called *Time Frequency Code* (TFC), allowing designers to use only one band or hop between two or three bands (see example in Table 2).

Development and production testing

A fundamental difference exists between testing for development and testing for production. In development, designers are creating a capability that meets all aspects of a standard using components whose worst-case tolerance contributions produce an end result that still fits within specifications. Once that design is verified, however, designers no longer have to test all the same design aspects as before. Now, they are essentially trying to find production errors or faulty components.

It can be argued that for development, it makes no difference whether designers use a bunch of general-purpose instruments hooked together to accurately measure the necessary parameters or a purpose-built system designed specifically for testing wireless USB.

However, in today's commoditized, hypercompetitive electronics market, time is of the essence both for development and production. Designs must be production ready within weeks rather than months, and high-volume production capabilities must be in place as soon as the design is verified.

Under these circumstances, designers don't have the luxury of using *ad hoc* instrumentation clusters for production testing. But it also can be argued that waiting for a purpose-built test system could slow development.

Development test requirements

A Certified Wireless USB device transmits and receives, and both functions need to be tested. On the transmit side, tests are needed for power level and setting accuracy, Error Vector Magnitude (EVM), power spectral density (against a spectral mask), and frequency offset. But that's just the beginning. Designers also will want to test symbol clock offset, phase noise, compression, amplitude and phase balance, amplitude variation over time, and local oscillator leakage.

The reason for more exhaustive testing is to make sure a device meets all specifications. For example, it is very possible for a device to work functionally in all aspects – range, speeds, and error rates – yet produce interfering radiation above a prescribed level.

On the receive side, designers will want to test minimum power sensitivity at a specified Packet Error Rate (PER) and maximum power sensitivity at a specified PER. This will reveal if the device's dynamic range is within specification. But, here again, more testing is needed for the receiver's adjacent channel rejection, alternate channel rejection, and maximum tolerable signal to make sure it is not susceptible to interference at signal levels in other channels within specification.

Using general-purpose instruments for this level of development testing usually involves a lengthy process. These instruments must be capable of accurately measuring signals spanning over 7.5 GHz and responding quickly enough to measure fast hopping and short symbols.

Most general-purpose instruments available today fall short of the 7.5 GHz bandwidth requirement. Those that have sufficient bandwidth are typically high-end systems with prices to match.

HALF COOLED SYSTEMS?
RADIAN HEATSINKS WILL FULLY SUPPORT YOUR APPLICATIONS
radian
ISO 9001:2000 Certified
- We optimize your heatsink designs
- Complimentary designs, from CFD simulations to final product!
- Custom and standard solutions to solve your thermal needs.
web www.radianheatsinks.com tel 408-988-6200 fax 408-988-0683
Radian Heatsinks is a division of Intricast Company Inc.

On the other hand, single box testers with software-controlled vector signal analyzers can be designed for capturing and analyzing wireless USB signals and providing comprehensive test results that meet development needs. Developing such a tester that meets Certified Wireless USB's bandwidth and speed requirements has challenges of its own.

Production test requirements

A test system built to handle wireless USB development testing would unquestionably be capable of production testing as well. In fact, the total number of tests would be significantly relaxed since designers would not be trying to verify the design but rather trying to catch production problems.

In this case, transmitter testing would involve calibrating the output power, measuring the EVM, checking that power spectral density falls within the mask, and (optionally) checking the frequency offset.

On the receiver side, designers must check minimum power sensitivity at a specified PER and maximum power sensitivity at a specified PER. These tests would expose any production shortcomings.

A prototype wireless USB tester

At this time, there are no single box EVM testers designed for wireless USB development and production testing However, that is about to change.

At the Consumer Electronics Show earlier this year in Las Vegas, LitePoint demonstrated a prototype wireless USB test system paired with the latest Certified Wireless USB chipset, as shown in Figure 2. This prototype captures continuous packets transmitted by the chipset's reference design and displays a rich set of data plus informative graphics. The Unit Under Test (UUT) was controlled via its own driver software, allowing band group, TFC, and continuous packet output selection.

The test system is controlled by its GUI software, providing a user-selectable array of graphic displays and data listings. Designers can select the band group and capture mode (single signal or continuous). In single signal mode, the tester captures one packet and displays the instantaneous readings. In continuous capture, a succession of packet "snapshots" shows if the readings are relatively stable or punctuated by large, random deviations.

Figure 2

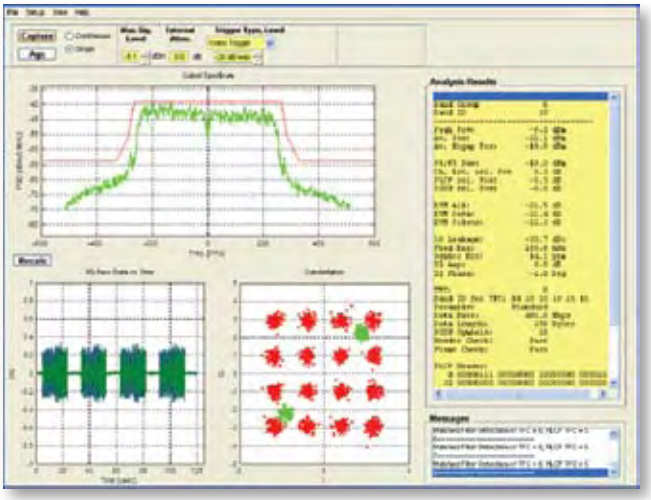

Figure 3

An important test is power spectral density (shown in Figure 3), ensuring that the transmitted signal is well within the spectrum mask limits. Signals that exceed spectral mask limits are more likely to interfere with other signals.

All the development and production tests alluded to earlier can be accomplished within a short timeframe. The setup remains the same, and users simply select different parameters and graphical views. Signal capture, subsequent analysis, and display each take a few seconds to complete, with a single click to produce all the analytical results.

In the near future, at least one single box EVM-based test system – the LitePoint IQultra – will be finished with beta testing and available for rapid Certified Wireless USB-based device testing.

Onno Harms *is product manager for WiMAX test solutions at LitePoint Corporation in Sunnyvale, California. Onno has more than 15 years of experience in wireless systems, ICs, and volume manufacturing and has developed products for voice and data applications involving several popular wireless standards. Onno holds an MS in Electrical Engineering from Twente University in The Netherlands.*

Rob Brownstein *is VP of market relations and communications at LitePoint. He began his career as a test engineer more than 30 years ago and worked as a journalist for three electronic trade media. He holds a BS in Physics from Queens College, City University of New York, and performed graduate work in Electrical Engineering at San Diego State University.*

LitePoint Corporation
408-456-5000
onno.harms@litepoint.com
rob.brownstein@litepoint.com
www.litepoint.com

Editor's Choice

Ultra-small embedded storage combo

Designers are constantly trying to find smaller devices that can be squeezed into the nooks and crannies of their designs. Mass storage is no exception to the usual space limitations.

SiliconSystems, Inc. and Samtec have collaborated to offer a turnkey embedded storage solution – the SiliconDrive II USB Blade and associated SiliconBlade Socket. The combination of innovative storage technology in the SiliconDrive II USB Blade and SiliconBlade Socket enables OEMs to easily incorporate advanced storage into space-constrained designs. SiliconDrive II USB Blade is an ultra-small form factor that complements SiliconBlade Socket's robustness and locking mechanism. The postage stamp-sized USB solid-state drive is an ideal alternative to SD and MMC cards in telecommunication, embedded, industrial, military, and medical applications.

Available in capacities of 512 MB, 1 GB, and 2 GB, SiliconDrive II USB Blade is designed for applications where board space, shock, vibration, temperature, and multiyear product life cycles are mandatory design considerations.

SiliconSystems
www.siliconsystems.com
RSC# 36483

Samtec
www.samtec.com
RSC# 36482

Multithreaded multiprocessor enables multitasking

With frequency scaling limited by power constraints, using multicore processors traditionally was one of the best ways to improve processor performance and manage a reasonable power consumption budget. But when designers made their own chipsets, they had no other options for multicore processor IP – until now.

MIPS Technologies, Inc. has introduced the MIPS32R 1004KT coherent processing system, purportedly the industry's first embedded multithreaded mul iprocessor licensable IP core. The new multicore offering provides performance efficiency and configurability in a multiprocessing system – up to four single or multithreaded processors integrated with advanced system coherency

The 1004KT core optimizes CPU performance on a shared memory system, enabling multiple functions to be implemented in a single product running concurrently under symmetric multiprocessing-based operating systems. The multicore Coherence Manager serves as the foundation block for intelligent system coherency, with an I/O Coherence Unit that provides optional hardware coherence for I/O peripherals. Several vertical applications, including digital home entertainment, home networking, and office automation can benefit from coherent multiprocessing using multithreading.

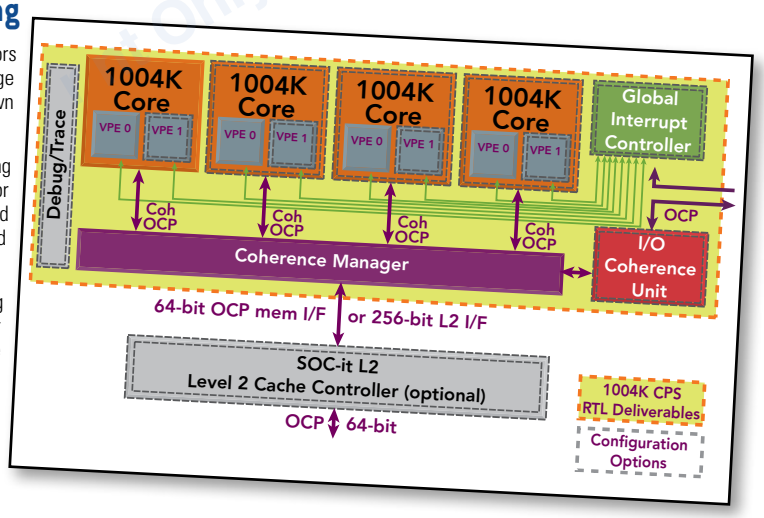

MIPS Technologies
www.mips.com
RSC# 36480

Rolling out low-cost OLEDs

Organic Light-Emitting Diodes (OLEDs) represent the next evolution in lighting products. Their widespread design capabilities can deliver significantly improved levels of efficiency and environmental performance with the same quality of illumination found in traditional products.

GE Global Research, General Electric's centralized research organization, has successfully demonstrated roll-to-roll manufactured OLED lighting devices, a key step toward making OLEDs and other high-performance organic electronics products at lower costs than possible today. The demonstration leads the way to a possible cost-effective system for mass-producing organic electronics products, such as flexible electronic paper displays, portable TV screens the size of posters, solar-powered cells, and high-efficiency lighting devices.

The few organic electronics products available today are relatively high cost and made with more conventional batch processes. A roll-to-roll manufacturing infrastructure for creating high-performance, low-cost devices would allow more widespread adoption of organic electronics products. GE's research program aims to introduce OLED lighting products to the market by 2010.

GE Global Research
www.ge.com/research
RSC# 36481

Editor's Choice Products are drawn from OSP's product database and press releases. Vendors may add their new products to our website at www.opensystems-publishing.com/vendors/submissions/np/ and submit press releases at www.opensystems-publishing.com/news/submit. OSP reserves the right to publish products based on editors' discretion alone, and does not guarantee publication of any product entries.

Choose your path – our AdvancedMC™ solutions will take you there.

Whatever direction you take your application, Performance Technologies has the compute module to get you there. Available in cost-optimized or high-performance configurations, our AdvancedMC™ modules will speed time-to-market with reduced developments costs, and are well suited for the high-end packet processing and multi-threaded software applications found in today's MicroTCA™ and AdvancedTCA® systems.

Design your way to market with our flexible options that include:

- PowerPC® or x86 chipsets
- Single- or dual-core configurations
- Support for up to 4 GB RAM
- User-upgradeable Flash memory
- Integrated NexusWare® CGL OS
- Support for Linux®, Solaris™, or Windows® operating systems

Contact us today to discover how our AMC modules can provide you with a short cut to success.

AdvancedMC™ Modules
that take you there.

PERFORMANCE TECHNOLOGIES

Phone: 585.256.0200 • www.pt.com/AMC • E-mail: info-request@pt.com

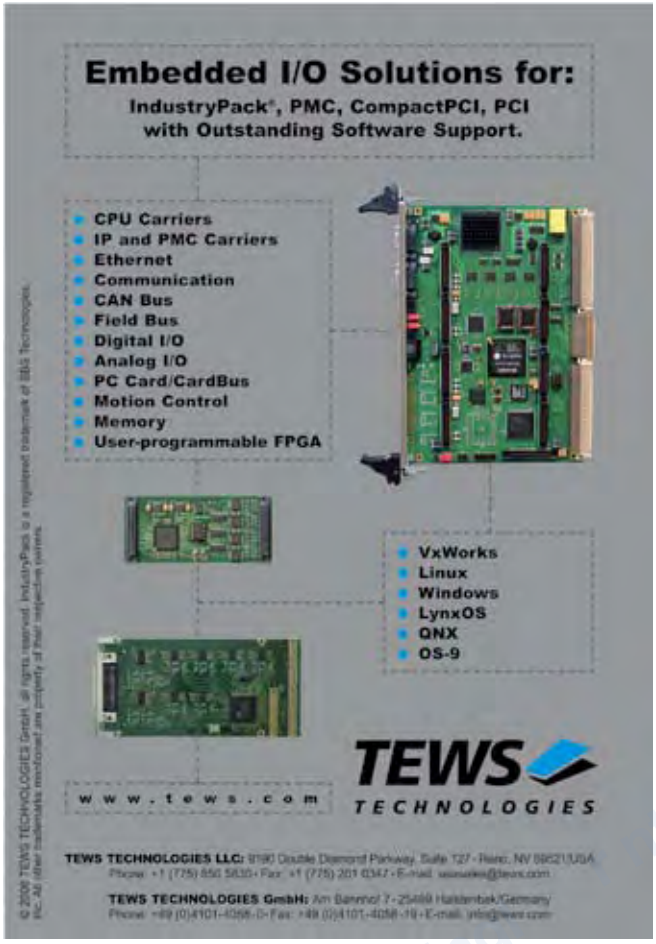

Advertiser Information

Page #	Advertiser	Ad title
19	ACCES I/O Products, Inc.	Boldly go
30	Advantech Corporation	Embedded box computers
15	AMAX	When every second counts
7	Annapolis Micro Systems, Inc.	WILDSTAR 5
47	Intel	Rethink cool
5	Micro/sys, Inc.	We've slashed slow boot-up
48	Microsoft	Fast-forward your masterpiece
18	MPL AG	The smallest embedded CPU board
34	North Atlantic Industries	Power source
17	Octagon Systems	Core systems
45	Performance Technologies	Choose your path
13	PQI Corporation	The flash storage leader
42	Radian Heatsinks	Half cooled systems
21	Super Talent Technologies	Reinventing storage
6	Technologic Systems	TS-7800
46	TEWS Technologies LLC	Embedded I/O solutions
9	Toronto MicroElectronics, Inc.	ECM401
23	Toronto MicroElectronics, Inc.	Embedded computer
29	Toronto MicroElectronics, Inc.	Micro-P3
11	Total Control Solutions	Systems that can go anywhere
27	Tri-M Systems Inc.	100MHz PC/104 module
33	Tri-M Systems Inc.	PC/104 Can-Tainer
25	VersaLogic Corp.	When you design
41	Virginia Panel Corporation	Standardize your PXI platform
2,3	WinSystems, Inc.	-40 °C to +70 °C Fanless

Single Copy Print Only

OpenSystems Publishing

Advertising/Business Office

30233 Jefferson Avenue
St. Clair Shores, MI 48082
Tel: 586-415-6500 ■ Fax: 586-415-4882

Vice President Marketing & Sales
Patrick Hopper
phopper@opensystems-publishing.com

Business Manager
Karen Layman

Sales Group

Dennis Doyle
Senior Account Manager
ddoyle@opensystems-publishing.com

Tom Varcie
Senior Account Manager
tvarcie@opensystems-publishing.com

Doug Cordier
Account Manager
dcordier@opensystems-publishing.com

Andrea Stabile
Advertising/Marketing Coordinator
astabile@opensystems-publishing.com

Christine Long
E-marketing Manager
clong@opensystems-publishing.com

Regional Sales Managers

Jerry Bleich
New England
jbleich@opensystems-publishing.com

Barbara Quinlan
Midwest/Southwest
bquinlan@opensystems-publishing.com

Ron Taylor
East Coast/Mid Atlantic
rtaylor@opensystems-publishing.com

Ernest Godsey
Central and Mountain States
egodsey@opensystems-publishing.com

Denis Seger
Southern California
dseger@opensystems-publishing.com

Sydele Starr
Nothern California
sstarr@opensystems-publishing.com

International Sales

Dan Aronovic
Account Manager – Israel
daronovic@opensystems-publishing.com

Sam Fan
Account Manager – Asia
sfan@opensystems-publishing.com

Reprints and PDFs

Nan Lamade: 800-259-0470
ecdreprints@opensystems-publishing.com

intel Leap ahead™

Rethink cool.

Intel embedded processor/chipset with TDP of less than 5 watts*. Way cool.
For your fanless applications, design in Intel. Our new 45nm microarchitecture, built from the ground up, delivers robust performance while keeping its cool. How? Power savings optimized by an in-order execution core and C6 enhanced sleep states. How cool is that? Go to **intel.com/go/rethink**

*The TDP specification should be used to design the processor thermal solution. TDP is not the maximum theoretical power the processor can generate. Intel and the Intel logo are trademarks of Intel Corporation in the U.S. and other countries. © 2008 Intel Corporation. All rights reserved.

FAST-FORWARD YOUR MASTERPIECE.

Get your next big idea to market—faster. The built-in tools and technologies in Windows® Embedded help to accelerate development, so your teams get to put their efforts where they really count: creating the next generation of smart, connected devices. You also get Windows Embedded features that enable interoperability with industry standards and Microsoft® technologies, so you can build devices that easily connect to PCs, services, servers, and other devices. Learn more about how to fast-forward device development at: **windowsembedded.com/fastforward**

Windows® Embedded

DSP-FPGA.com
Supplemental
www.DSP-FPGA.com

2008

Programmable Logic:

FPGAs & Reconfigurable Computing

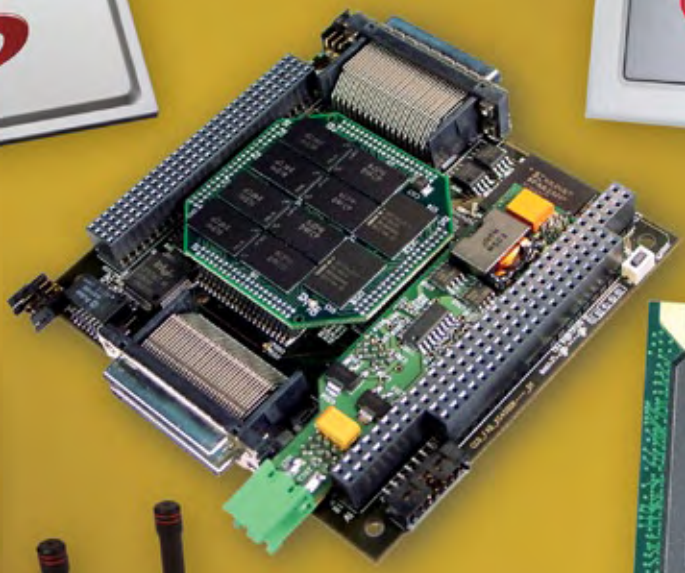

Interviews

Features

OpenSystems Publishing

Annapolis Micro Systems

The FPGA Systems Performance Leader

WILDSTAR 5 for IBM Blade

The Perfect Blend of Processors and FPGAs

Fully Integrated into IBM Blade Management System
Abundant Power and Cooling Ensure Maximum Performance

Made in the USA

Ultimate Modularity

From 2 to 8 Virtex 5 FPGA/Memory Modules
Input / Output Modules Include:
Quad 130 MSps thru Quad 500 MSps A/D
1.5 GSps thru 2.2 GSps, Quad 600 MSps A/D
Dual 1.5 GSps thru 4.0 GSps D/A
Infiniband, 10 G Ethernet, FC4, SFPDP

Direct Seamless Connections with no Data Reduction

Between External Sensors and FPGAs
Between FPGAs and Processors over IB or 10GE Backplane
Between FPGAs and Standard Output Modules

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

For Systems
that Demand
the Most

XG-5000K

Our most advanced FPGA Based PC/104 Plus Circuit Board

Features

- Xilinx Spartan 3, 5 Million Gate FPGA
- 264 User Configurable I/O
- 256 MByte of SRAM
- 32 MByte of Flash
- 6 MByte of Serial Page DataFlash
- CompactFlash Type 1 Connector
- Secondary Xilinx Spartan 3, 400,000 Gate FPGA
- 10/100Base-T Ethernet
- 2 RS232 Ports
- 25MHz FPGA Master Clock Source
- Secondary 0 - 25MHz user programmable DDS FPGA Master Clock Source
- Incorporates Xilinx's design revisioning technology and can retain onboard as many as 16 partial or 4 complete FPGA design BIT files
- Can be powered from the PC/104 connector or an external DC power supply
- PC/104 Plus Form Factor; ISA and PCI 33MHz Bus Interfaces
- Extended temperature -40C to +85C

Jacyl Technology is a Team Member under Lear Siegler Services, Inc for the US Governments CECOM Rapid Response Contract

JACYL TECHNOLOGY

Advancing Today's Technology into Tomorrow

www.jacyltechnology.com

Power consumption as the competitive advantage

Q&A with John East,

President and CEO, Actel Corporation

Our exclusive interview sheds light on how Actel's strategy and technology are reshaping programmable logic for the latest power-sensitive applications.

For Single Print Only

PL: Actel's strategy has Design for Energy Efficiency written all over it. How did you arrive at that as a core strategy for the company, and how is it changing your business?

JE: I guess the real question is why a company *wouldn't* arrive at energy efficiency as a core strategy. The industry is ripe with discussion of "green design" and industry leaders are called upon every day to act as socially responsible citizens.

Power matters. With the world's increasing use of electronic devices, Actel has the opportunity to play a major role in resolving the world's global warming problems with further technological change.

We're changing FPGAs from "power hogs" into "power savers." We've designed nonvolatile, reprogrammable flash-based devices offering 200 times less static power and more than 10 times the battery life than competitive SRAM-based programmable logic solutions.

We are in the business of enabling designers to save power – from low-power FPGAs and power-efficient programmable system chips to power-optimized tools and power-smart IP.

PL: A good example is the IGLOO. What makes the IGLOO architecture a low power solution?

JE: There are three main points of interest in IGLOO.

The differentiating factor for the majority of our product line is the use of a true flash-based architecture. True nonvolatile flash-based FPGAs contain a nonvolatile FPGA array, and do not suffer from the leakage current issues associated with SRAM-based solutions.

Secondly, we felt it paramount that we select a foundry partner with a proven low-power flash process technology. We chose UMC's leading-edge 0.13-micron low-power and e-Flash processes. Combined with IGLOO's power mode options and Flash*Freeze technology, UMC's processes enable power consumption as low as 5 microwatts.

Finally, the IGLOO family has several power modes to optimize power consumption – the Flash*Freeze mode, a low-power active mode, and a sleep mode. While in Flash*Freeze mode, the device conserves power while maintaining FPGA content. Device I/Os are tri-stated and SRAM and register content is maintained, but not clocked. Further, the Flash*Freeze pin enables designers to quickly and easily enter or exit the Flash*Freeze mode within 1 microsecond. Alternatively, the low-power active mode allows the IGLOO device to directly control when the system goes into low-power mode. While in the low-power active mode, the FPGA core, clocks, and all I/O are functional.

PL: What have you done to aid designers in power consumption analysis?

JE: The Libero 8.1 IDE has specific features to further optimize FPGA designs for low power.

The Libero IDE enables creation of realistic power profiles based on design parameters or metrics. An example is the percent of time the FPGA will be in a combination of custom or functional modes, such as Active, Standby, or Flash*Freeze. The tool then displays a realistic and accurate report of power consumption based on the true power profile of the target FPGA.

The design analysis data is then utilized to engage in power-driven layout. Use of this tool enables users to quickly realize

LOWEST TOTAL COST... PERIOD.

GET UP TO 50% LOWER COST

Xilinx Spartan™-3 Generation devices are the only low-cost FPGA family to cut your total cost by half... and we can prove it! We offer the industry's widest range of devices, most comprehensive IP library (8x the nearest competitor!), a huge selection of development boards and we've minimized the need for external components.

Visit us at **www.xilinx.com/spartan** to download our free ISE™ WebPACK™ design tools and start your Spartan-3 design today. You too can experience the lowest total cost. Period.

XILINX®

www.xilinx.com/spartan

The World's Most Widely Adopted Low-Cost FPGAs

©2003 Xilinx, Inc. All rights reserved. XILINX, the Xilinx logo, and other designated brands included herein are trademarks of Xilinx, Inc. All other trademarks are the property of their respective owners.

For Single Point Only

dynamic power savings through the reduction of the capacitive loading of the nets. While average IGLOO power consumption is reduced by 13 percent, some designs can realize a 30 percent reduction in power consumption.

PL: Where do you see processor cores in FPGAs headed?

JE: We see real value there and have several cores – CoreABC, Core8051, Core8051s, and CoreMP7, as well as the FPGA-optimized ARM Cortex-M1 and the Leon3 SPARC processor.

These offerings bring the flexibility, fast time to market, and low implementation cost of FPGAs to users who can't justify the expense of ASICs for their application. We'll continue to explore processor core technologies with more choices, better tools, and improved power optimization.

PL: I see you're betting on MicroTCA specifically – why?

JE: We're really interested in system management. At the high end, there are standards-based and proprietary solutions for telecommunications such as AdvancedTCA and MicroTCA. At the low end embedded applications need low cost, less complex solutions. We're doing both.

Our mixed-signal Fusion Programmable System Chip (PSC) can perform system management functions, including power and thermal management, data logging, and system diagnostics, within high-end and low-end electronic systems. And it's all in a single chip, with a cost and space savings of 50 percent or greater relative to current implementations while also improving reliability.

PL: How can designers get an advantage in their next FPGA-based design?

JE: Nearly every vendor claims low-power leadership. Ignore the claims and check the datasheets – static power, dynamic power, inrush, configuration power, and low-power modes.

Comparing one-million gate devices, for example, SRAM-based FPGAs marketed as the "lowest power FPGAs" state static power consumption between 40 milliwatts and 150 milliwatts, which is 800 to 3000 times higher than Actel's 0.05 milliwatt flash-based IGLOO device. Similarly, when comparing "zero power" CPLD solutions against IGLOO, these devices consume 10 times more static power.

Then look at tools to optimize power. How does the development environment like Libero help the designer? What power-smart IP, like FPGA-optimized 32-bit ARM processors, is available? Tools are just as important as architecture.

PL: What's next for the FPGA industry?

JE: Today, FPGAs are capable of handling so much more within a design – whether it be the interfacing and control within a smart phone or in a system-critical automotive or medical application. FPGAs are expanding into so many different applications that can leverage their flexibility, low power, and low cost.

Looking ahead, I believe a continued focus on power reduction will drive even more applications. There continues to be tremendous demand for low-power programmable solutions at lower densities to handle the control and bridging functions in portable designs. We're now hitting the $1 price point with these solutions, particularly for price-sensitive consumer portable applications. The technology will just continue to improve.

PJ

John East *is President and CEO of Actel Corporation, with nearly 40 years experience in the semiconductor industry. Prior to joining Actel in 1988, East held positions with AMD, Raytheon Semiconductor, and Fairchild Semiconductor. He has a bachelor's of science degree in electrical engineering and a master's degree in business administration from the University of California, Berkeley.*

Actel

650-318-4200 • www.actel.com

Programmable Logic: FPGAs & Reconfigurable Computing

DSP-FPGA.com

Editorial

Don Dingee
Editorial Director
ddingee@opensystems-publishing.com

Chris Ciufo
Group Editorial Director
cciufo@opensystems-publishing.com

Robin DiPerna
Assistant Editor

Rosemary Kristoff
Vice President Editorial

Joann Toth
Senior Designer

David Diomede
Art Director

Konrad Witte
Senior Web Developer

Sandy Dionisio
Graphic Coordinator

Advertising/Business Office

OpenSystems Publishing
30233 Jefferson Avenue
St. Clair Shores, MI 48082
586-415-6500

Patrick Hopper
Vice President Marketing & Sales
phopper@opensystems-publishing.com

Karen Layman
Business Manager

Sales Group

Dennis Doyle
ddoyle@opensystems-publishing.com

Tom Varcie
tvarcie@opensystems-publishing.com

Doug Cordier
dcordier@opensystems-publishing.com

Andrea Stabile
astabile@opensystems-publishing.com

Christine Long
clong@opensystems-publishing.com

Ernest Godsey
egodsey@opensystems-publishing.com

Barbara Quinlan
bquinlan@opensystems-publishing.com

Denis Seger
dseger@opensystems-publishing.com

Sydele Starr
sydele@pacbell.net

Ron Taylor
rtaylor@opensystems-publishing.com

International

Dan Aronovic
Account Manager – Israel
daronovic@opensystems-publishing.com

Sam Fan
Account Manager – Asia
sfan@opensystems-publishing.com

Reprints and PDFs

Nan Lamade: 800-259-0470
dspreprints@opensystems-publishing.com

For Single Print Only

Reconfigurable FPGA Processing Cards

Ultra Performance X5 Modules

Xilinx Virtex5 FPGA logic

4MB QDR SRAM

8 Rocket IO Private Links
2.5 Gbps each

>1GB/s
8-lane PCI Express Host

Available I/O Solutions:

X5-400M
400 MSPS 14-bit A/D (x2)
500 MSPS 16-bit D/A (x2)
1GB DDR2 DRAM

X5-210M
210 MSPS
14-bit A/D (x4), 512MB DDR2 DRAM

eInstruments

User-customizable, embedded instrumentation!
COTS solution combining two XMC modules plus intelligent x86-based SBC in rugged enclosure

Capabilities
- Embedded software radio and IF processors; DSP within FPGA
- FPGA-based real-time (servo) control/analog I/O at rates to 400 MHz
- Perfect for portable or data loggers or waveform generators

Distributed Data Acquisition
- Co-locate eInstrument and unit-under-test for optimized signal quality/minimized cabling
- Limitless expansion via multiple nodes

Single Board Computer (SBC)
- Scalable CPU up to Core 2 Duo w/ 4 GB memory
- USB2.0 x6, Gigabit Ethernet, SATA x4, VGA, AC '97 audio
- GPS timebase support
- Cabled PCI Express expansion to other eInstruments
- Windows or Linux from USB FLASH (diskless) or SATA HDD

Rugged & Portable
- 12V DC-only operation via Battery or AC Adapter
- Conductive cooling plus optional fan
- 240 x 170 x 50 mm
- Optional aluminum instrument enclosure

Low Cost X3 Modules

Xilinx Spartan 3A DSP 1.4M

Two 2MB SRAMs, 48-bit DIO

PCI Express with >200 MB/s data rates

Available I/O Solutions:

X3-25M: 16-bit, 130MHz, A/D (x4), D/A (x4)

X3-A4D4: 16-bit, 4 MHz A/D+ D/A (x4)

X3-DIO: 64-bit, 66 MHz LVDS

X3-SD: 216 kHz, 16-bit, A/D (x16)

X3-SDF: 20 MHz, 24-bit, A/D (x4)

X3-Servo: 16-bit, 250 KSPS, A/D (x12)
16-bit, 1 KSPS D/A (x12)
(ultra low-latency & glitch energy)

Using Innovative Integration's FrameWork Logic VHDL source-code or Simulink/MATLAB board support package you can readily customize FPGA functionality to include real-time processing such as independent FIR and IIR filters on each channel, real-time FFT processing, ultra-fast feedback and control loops and much more.

For more information visit www.innovative-dsp.com/e

For Single Use Only

www.innovative-dsp.com • sales@innovative-dsp.com • 805-578-4260

Innovative Integration
... real time solutions!

Stretching the boundaries for FPGAs and people

Q&A with Jeff Kodosky,

Co-Founder and Technology Fellow, National Instruments

Jeff K, as he's known in the LabVIEW community, has a unique vision for how people should program and who should program. He shares some exclusive thoughts with us on where he sees NI's technology today and in the future.

For Single Print Only

PL: We've heard from Dr. Truchard about graphical system design. How else has the FPGA capability of NI CompactRIO and LabVIEW changed the way people think about designs?

JK: The combination of an intuitive graphical programming language and a heterogeneous hardware platform (MPU + FPGA) has really changed not only how people think about design but also who is thinking about design.

Hardware and software design engineers see this tightly integrated, embedded platform as a solution to their time-to-market pressures. Regardless of their deployment platforms, these experts use NI LabVIEW software and CompactRIO hardware to prototype and iterate on their designs and algorithms at a fraction of the time even imagined in the past.

But even more exciting to me is how this marriage of graphical programming and the latest programmable logic technology is changing the "who" of the design world. Now domain experts – the physicists, robotics engineers, mechatronics and mechanical engineers, and scientists – can actually implement their innovative ideas in hardware themselves, instead of hiring an embedded design expert or giving up seeing their concepts come to fruition.

PL: Where are the boundaries today, and where would you like them to be soon, in choosing to implement on an FPGA or a DSP?

JK: The characteristics of an application typically suggest which alternative might be more appropriate. If the application is computation-bound and involves a great deal of complicated mathematical computations, a DSP may be the best implementation target. An FPGA works better if the application is highly parallel or pipelined, or logic-bound with many bit-level operations – for example, when implementing digital protocols. An FPGA also works better if the application uses custom I/O requiring high-speed timing and triggering logic or special "glue" logic. Finally, an FPGA may be the best choice in some safety-critical situations because the safety-critical portion can run in parallel without interference from anything else; there is no operating system, device driver, critical section, or interrupt to delay or interfere with the execution of the safety-critical actions.

In the long term, I would like to see a convergence between FPGA and DSP targets where highly optimized ALU components are generously sprinkled around in an FPGA connection and logic fabric. Such a hybrid target would be ideal for complex embedded applications that combine lots of parallel computation with high-speed timing and triggering.

PL: At NIWeek 2007, OEM board versions of CompactRIO were shown – give us an example of how that is impacting fielded applications.

JK: The ability to move down what we call the *RIO deployment curve* is very exciting. Engineers and scientists can use powerful platforms like PCs, PXI, and plug-in FPGA devices for rapid prototyping. Then they can easily move to smaller, rugged systems like CompactRIO, and even integrated systems combining a real-time processor and reconfigurable FPGA within the same chassis.

Sanarus, a medical device start-up company, has developed plans for a potentially revolutionary product that could change the way doctors treat benign tumors. With this device, based on LabVIEW and CompactRIO, doctors can eliminate tumors by freezing and killing them in an outpatient procedure, a dramatic change from in-patient surgery or the "wait and see" approach used previously.

The Visica2 Treatment System is an instrument for use in a doctor's office or clinic. The procedure is performed with local anesthesia and uses a real-time, ultrasound-guided probe that is virtually painless. The treatment, which lasts 10 to 20 minutes, freezes and destroys targeted tissue through an incision so small that it does not require stitches.

National Instruments, continued on page 10

Signal processing in FPGAs goes mainstream

Q&A with Jeffrey Milrod,

President and CEO, BittWare

Jeff gives us further insight in this exclusive interview as to how signal processing using FPGAs has changed, and how customers are benefitting from the improvements.

PL: What's different in 2008, from a couple years ago, in how people design a signal processing system with FPGAs?

JM: A few years ago, FPGAs were used mostly, but not exclusively, for communications, interface, and glue-type logic. Only the brave, lunatic, or desperate few implemented processing algorithms in the FPGA itself. Even then, these were mostly very simple, straight-forward, and repetitive types of algorithms.

Today, things are very different and the pendulum has swung the other way. It now seems that everyone wants to use FPGAs as a signal processing resource, but clearly not everyone understands the implications of that. Implementation is still quite challenging and not for the faint of heart. While FPGA vendors and third party tools suppliers have made great strides, there is still a gap between the high-level algorithm implementation and the HW.

PL: I see FPGAs, DSPs, and even GPPs in your diagrams. Where are the boundaries in partitioning a system?

JM: There really are no strict boundaries other than those imposed by hard performance requirements. However, partitioning can have a huge impact on development time and effort.

Traditional DSPs are still much easier for implementing complex algorithms and algorithms that are likely to be frequently modified, since development can be done in higher level languages like C. This also facilitates code reuse – many users already have a great deal of code working on DSPs, with no compelling reason to port to FPGAs.

If code reuse is not an issue, and the algorithm is fairly well-defined, standard, or straight-forward, FPGAs are very attractive since they can offer compelling performance advantages – both in terms of size and speed. Often, signal processing algorithms force the use of FPGAs because they can't be reasonably implemented in DSPs.

FPGAs provide some future-proofing – as they become bigger, faster, lower power, and easier to use, it's likely that they will dominate signal processing. FPGA development investments today will be the code reuse of tomorrow.

As for GPPs, we only use them to do command and control processing. This can greatly ease the processing and development burden of the target DSPs and FPGAs, so that these valuable resources only do what they're best at.

PL: What does your FPGA design toolset look like, and why are those pieces important?

JM: The generic block template in the toolset consists of a function implemented with a standard data interface for sourcing and sinking, along with a standard memory mapped control interface. Many of the specific blocks provide board-level physical interfacing, but some provide data switching and routing functions and others provide resource arbitration, DMA engines, control block, and utility functions.

These pieces are not overly exciting in and of themselves; the real value lies in the fact that they are constructed to be building blocks in a framework – BittWare's ATLANTiS – allowing users to more quickly and easily get their special algorithms and applications up and running in the FPGA on a real board.

PL: What makes the ATLANTiS framework special?

JM: In ATLANTiS, we're using standard interconnects and an orthogonal control plane, and we've done all the low-level physical interfaces and data movement structures that are standard in microprocessors or DSPs. (Figure 1.)

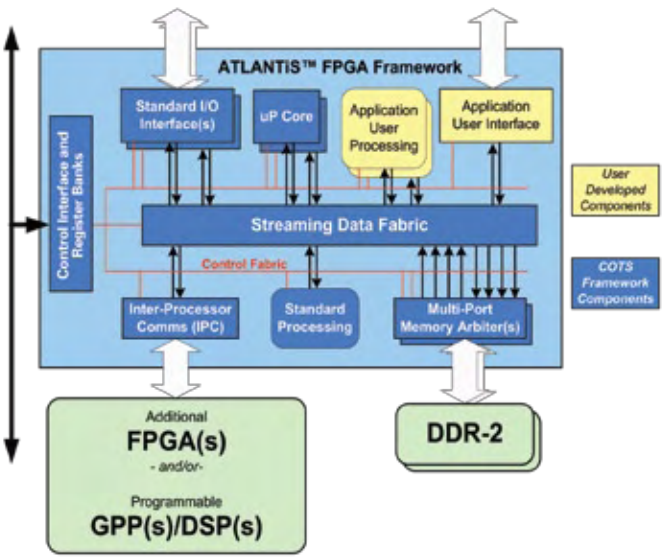

Figure 1

Rather than starting with a blank slate, as it were, and having to deal with all the low level interfacing, data movement, and control, our ATLANTiS modular framework enables users to focus on adding their unique value to the FPGA – much like one

would focus on writing algorithmic code and setting-up memory structures in a processor without worrying about creating a data transport layer, DRAM controller, or arbitrating for resources.

PL: How are people succeeding at shrinking their systems with FPGAs? Please give us an example.

JM: We have seen many examples of this in military, instrumentation, and communication applications where an FPGA is used to implement an algorithm that is particularly difficult to do in standard DSPs or GPPs.

One specific example of this is a high-end cytometry (cell-sorting) instrument from iCyt. Their flow cytometer uses a laser to detect and sort the cells, with optical sensors sampled at over 100 MHz. Using DSPs to detect the cells required several DSPs to perform a weighted threshold on every sample from every sensor. We consulted with them and helped implement the cell detection algorithm in a pre-processing FPGA, thereby only requiring DSPs for the cell analysis – cutting the number of DSP boards in half.

PL: What's the next wave look like, and where can competitive advantage be gained?

JM: The next generation of FPGAs that we've seen through our partnership with Altera is simply amazing. The speed, feature set, and capabilities are astounding, and the power is lower than I'd hoped. The challenge is in harnessing all that tremendous signal processing potential in a practical, timely, and cost effective way.

Competitive advantage will be gained by those who facilitate code reuse and can reduce the effort required to get real signals in and processed in the FPGA. There is a great deal of work being done to improve the algorithmic implementation process, but often the harder part is integrating the algorithms into the real world of signal processing.

PJ

Jeffrey Milrod *is BittWare's president and CEO. He holds a bachelor's degree in Physics from the University of Maryland and an MSEE from John Hopkins University.*

BittWare
603-226-0404 • www.BittWare.com

National Instruments, continued from page 8

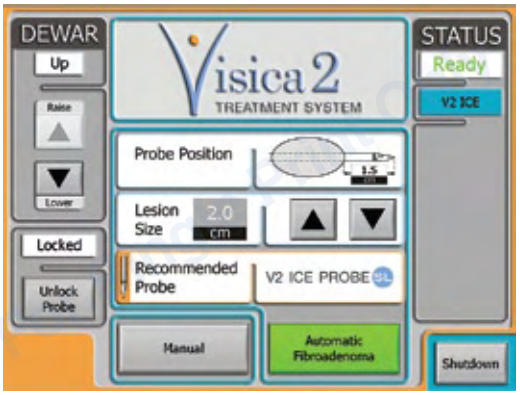

Because this device will be used in offices and clinics around the country, the machine has to be cost-effective. The lower-cost hardware options from NI, still based on the unique reconfigurable (RIO) architecture, enable Sanarus to meet its volume, cost, and technology challenges.

PL: Also at NIWeek 2007 a vision was outlined to get LabVIEW on a chip. Where is the vision is going, and what is it going to take to get there?

JK: It has always been my goal to ensure LabVIEW can reach the lowest levels of hardware and the smallest and lowest-cost programmable targets. This goal is captured nicely by the phrase "LabVIEW on a chip."

In one sense, LabVIEW FPGA already achieves this goal. If an ASIC were made from the FPGA then a LabVIEW application would really be on a chip, but we haven't seen a high enough volume application to justify making an ASIC.

A related interesting topic is whether LabVIEW on a processor could be accelerated by having some code in the run-time engine actually programmed into an FPGA sitting alongside the processor, or maybe directly in the firmware of the processor itself. This is something we continue to brainstorm, and we are beginning to talk about it informally with some processor vendors. We are also developing relationships with several universities to explore these and related ideas.

Another way to think about LabVIEW on a chip is to envision a processor chip that runs LabVIEW "natively," where software modules (virtual instruments) would be loaded and run analogously to the way a typical operating system loads and runs applications. This could be done with a typical Von Neumann architecture processor today, but it is interesting to envision what the ideal underlying hardware might be to fully realize the potential of the LabVIEW parallel language.

In my view, the ideal hardware would be a "super" FPGA containing lots of optimized embedded components sprinkled throughout the interconnect and logic fabric, such as integer and floating-point MACs and ALUs, FIFOs, blocks of memory, and so on. This super FPGA would consist of a large number of separately reconfigurable regions where the configuration memory itself would be multi-buffered so that the contents of a region could be changed in a single clock cycle. I think LabVIEW would be an excellent tool to program such architecture, and, who knows, maybe something like this will eventually be developed.

PL: What should designers be focused on doing differently now to get an advantage in their next FPGA design?

JK: To get the most out of their next multicore or FPGA applications, designers should expose more of the parallelism by drawing more computations in parallel and pipelining computations in a loop. This is already fairly easy and safe to do in LabVIEW, and it will be increasingly important in achieving the highest performance on an FPGA when using newer, higher-speed I/O modules.

PJ

Jeff Kodosky *cofounded National Instruments with Dr. James Truchard and William Nowlin in 1976. Known as the "father of LabVIEW," he invented the graphical programming language that defines the software. Kodosky was named an NI business and technology fellow in 2000. He received his bachelor's degree in physics from Rensselaer Polytechnic Institute in Troy, NY.*

National Instruments
512-683-0100 • www.ni.com

LYRTECH
Infinite possibilities...

MIMO advanced development solution

SDR development platform

Most versatile programmable logic for telecommunications

- Complete baseband-to-RF integration
- Highly modular and scalable
- Complete board support package, including examples for fast kick-off prototyping
- Unsurpassed model-based design support package — Lyrtech products offer the best and most complete I/O and interface integration within configurable Simulink blocks.

SIMULINK Enabled
MathWorks Partner

For Single Print Only

What's more...

The small form factor SDR development platform is:

- Power efficient
- Capable of component-level measurement and repartitioning
- Real-time power monitoring

GSM layer 1/2 software IP on DSP-FPGA mixed architecture

Also features an SCA version:

- All-CORBA enabled (FPGA, DSP, and GPP)
- All-SCA enabled
- Predefined device managers

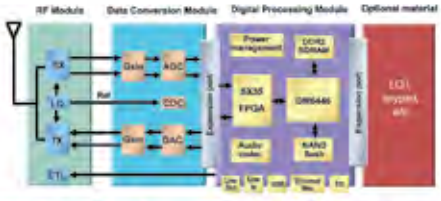

The MIMO advanced development solution is:

- The **only** complete 802.11n solution on the market today
- Capable of multichannel vector signals generation and analysis
- Equipped with real-time host data exchange tools
- Capable of 1-GBps data bandwidth throughout its system

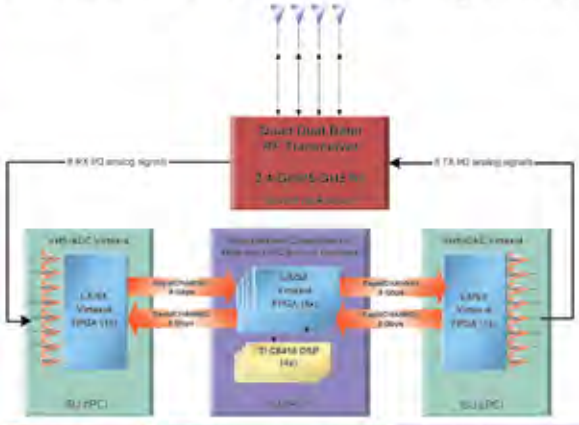

www.lyrtech.com

Driving both technology and cost/performance

Q&A with Chris Fanning and David Lee Rutledge,

Lattice Semiconductor Corporation

Two senior execs at Lattice give us their views on the Lattice strategy, with an in-depth look at non-volatile FPGA technology, in this exclusive interview and short article.

For Single-Point Only

PL: Lattice is making pushes in a lot of areas. Why such diversity in the offering, and how does it tie together from a strategy standpoint?

DLR: When Lattice entered the FPGA market, we chose to focus on two key differentiators – Non-Volatility (NV) and Low-Cost – and set about developing technology and products in each area.

NV is the technology of choice for programmable logic, providing system-level benefits such as single-chip solutions, instant-on capabilities, and in-system programmability. We knew that we could leverage our NV expertise to deliver these benefits to FPGA users.

Lattice saw an opportunity to fill a void in the market by developing a Low-Cost, Feature-Rich FPGA and CPLD product line. Our strategy has been to develop products that work very well for most applications, allowing us to substantially reduce cost with little impact on performance.

So, while our products may seem diverse, they have been developed based on a consistent strategy to deliver value-added products, with NV and Low-Cost, Feature-Rich as our key differentiators.

PL: What's the biggest technology trend driving your FPGA implementations today, and how do you see that affecting what you can do in the near future?

DLR: Achieving higher levels of functional integration and performance while maintaining acceptable power consumption is a major challenge. Historically, the scaling of supply voltages and transistor feature sizes has provided great benefits in functional integration levels and performance, while lowering power consumption. In the future, these techniques will still provide substantially higher levels of functional density and performance, but at the expense of increased levels of static power consumption.

At 45 nm there will be a more direct trade-off between speed and power. For example, a high-density FPGA (~500 K LUT) built on 45 nm technology could easily consume over 10 W of standby power, with no clocks running, at the commercial +85 °C junction temperature limit.

Innovation in power management to optimize this trade-off is a high priority. Innovation is required concurrently across multiple disciplines, ranging from optimizing the basic process technology through the development of new "power-optimized" product architectures and also through the development of "power-aware" design tools.

Also, there is increased use of embedded SERDES channels as high-bandwidth chip-to-chip communication links. An FPGA with 40 channels of soon-to-be-available 10 G SERDES will have an incredible processing capacity of 400 Gbps. This trend will accelerate the development of radically new FPGA-based High-Performance Computing (HPC) platforms.

PL: On the tools side, I see some additions of synthesis and simulation capability, and I'm guessing you're expanding further. What's the latest technology you are working on, and what's the impact?

CF: Lattice has made a very significant investment in its design tool, ispLEVER. Two initiatives we see having great impact are physical synthesis and incremental design.

Physical synthesis should enable more optimal results in improved device performance and help accelerate the debug process. Incremental design is a collection of technologies that ultimately provides

more rapid turnaround time for achieving timing closure. Both these enable customers to meet timing requirements more easily and with fewer design iterations, even as FPGA devices become even larger and more complex.

Other Lattice-driven design tool advancements include:

- Power Calculator allows specifying parameters such as voltage, temperature, process variations, airflow, heat sink, resource utilization, activity, and frequency, and then calculates static (DC) and dynamic (AC) power consumption.
- Reveal uses a signal-centric model for embedded logic debug. Signals of interest are user-defined, and the tool adds instrumentation along with the proper connections to enable the required in-system analysis can then be performed. Users can specify complex, multi-event triggering sequences that make system-level design debug smoother and faster.

PL: What about embedding operating systems on an FPGA core? How is this changing the way people design?

CF: uClinux support expands our commitment to the open source model, and it consistently appears at the top of designer surveys we've seen as the preferred RTOS for embedded design. uClinux and the LatticeMico32 core allow designers to implement control systems in a design flow that builds on Lattice's open source, embedded solutions approach.

The adoption of embedded products has increased among FPGA designers in the last two to three years as the processors provided by FPGA vendors have dramatically improved in functionality, increasing designer productivity and lowering design risk. Embedded processors increasingly include a robust assortment of middleware such as DDR, DDR2, SDRAM memory controllers, Tri-Speed Ethernet Media Access Controller and PCI 33 MHz Target, which automatically integrate into Lattice's Mico System Builder. Middleware has enabled designers to quickly and confidently configure a microprocessor in their design at very low cost, or no cost at all.

PL: What should designers be doing differently now to get an advantage with their next FPGA-based design?

DLR: We think in terms of two fundamental types of FPGA applications: *control-oriented* applications and *data path* applications. Control-oriented applications really have not evolved much over the years, implementing numerous small finite state machines and utilize many parallel I/Os for system monitoring and control. There are not too many new issues to deal with in this area.

Data path applications, however, have continued to grow and evolve. It is more important than ever for a system designer to consider how to best architect systems to effectively leverage the new high-speed SERDES-based capabilities of FPGAs. This increased data processing bandwidth will allow for radically new system-level architectures that can provide dramatically higher levels of cost/performance.

CF: And there's increasingly sophisticated functionality in ispLEVER to help the architect. Also, FPGA IP is particularly important, offering customers time to market and risk management advantages by providing proven, hardware-validated solutions that are typically parameterizable. These pre-engineered IP cores are cost-effective, and can save countless hours of development and validation effort.

PJ

Chris Fanning *is Corporate Vice President of Enterprise Solutions and manages Lattice's software, intellectual property, and technical support businesses. He previously worked at The Boston Consulting Group. He holds an MS in Computer Science from Worcester Polytechnic Institute, and an MBA from The University of Chicago.*

David Lee Rutledge *is Corporate Vice President of FPGA Product Development, with two decades of programmable logic experience including development of the first in-system programmable PLD. After working for Harris Semiconductor, Rutledge joined Lattice in 1983 as its 31st employee. He holds a BSEE and MSEE from Purdue University.*

Lattice Semiconductor

503-268-8000

www.latticesemi.com

Single Print Only

Non-volatility: the choice is clear

By David Lee Rutledge

Non-Volatility (NV) offers the same compelling benefits to FPGA users as it always has to CPLD users: a single-chip solution, instant-on capabilities, design security, and more. In fact, these features are arguably more valuable to FPGA applications than to CPLD applications.

So, why isn't everyone making NV FPGAs?

NV FPGAs require much more sophisticated process technology and circuit design techniques than their SRAM-based cousins. The key to NV success is providing sophistication without significantly impacting cost or performance, compared to an equivalent SRAM FPGA.

Anatomy of a true NV FPGA

A true NV FPGA must employ high-performance, embedded flash memory technology (Figure 1). Historically, this approach had relatively low performance due to the additional wafer processing steps required to embed flash memory in a logic process. Lattice has worked closely with our foundry partner Fujitsu to develop proprietary embedded Flash technologies that enable exceptionally high performance for minimal additional cost.

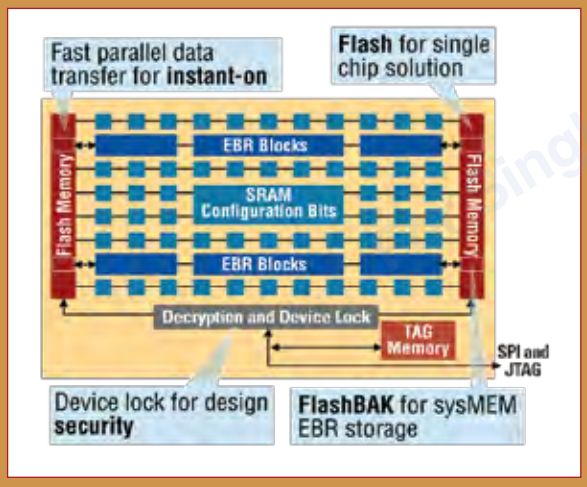

Figure 1

Instant-on, a feature that allows an FPGA to be used immediately upon startup, is essential in many applications. Typical SRAM FPGAs require 10s to 100s of milliseconds at startup using external boot PROMs. This prevents the use of SRAM FPGAs in certain critical control logic sections of systems that must be "live" immediately upon the application of power. Lattice has developed proprietary circuit design techniques that allow its NV FPGAs to become active virtually instantly after startup (within 1-2 mS).

Design security is also a key advantage for our NV FPGAs, in two distinct ways. First, our embedded flash memory allows the entire user design to be stored in on-chip Flash memory. The configuration bitstream is never exposed, as it is when using SRAM FPGAs with boot PROMs. It is virtually impossible for a hacker to access the configuration data of a NV FPGA.

Second, our embedded flash memory allows support of true AES encryption of the configuration bitstreams (Figure 2). This is a new feature that has been added to the latest 90 nm LatticeXP2 family. It supports the storage of a user-programmable, on-chip 128-bit encryption key. Once this key is stored on-chip, the user can encrypt the configuration bitstream for a given design and transmit it to the FPGA (over unsecured transmission lines). Once on-chip, the bitstream is decrypted using the stored AES key, and is never exposed off-chip.

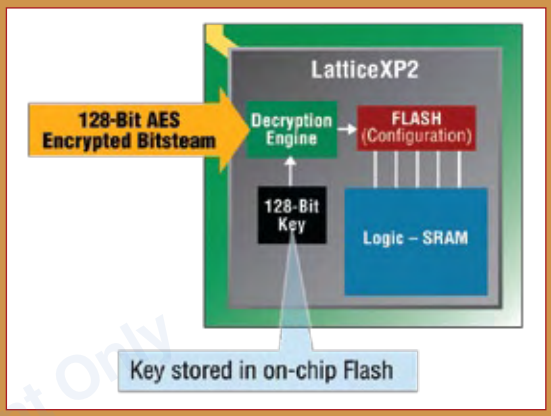

Figure 2

Alternative NV technologies

FPGAs have been developed that are called NV, but they are not true single-chip, embedded flash solutions. Instead, they combine two independent chips (a flash memory and an SRAM-based FPGA) in a single package using a stacked die technology. Compared to true NV FPGAs, these hybrid products have severe limitations.

For example, stacked die technology cannot support instant-on because essentially it is identical to two-chip, SRAM-based technology, and so has the same lengthy boot time as an SRAM FPGA (Figure 3).

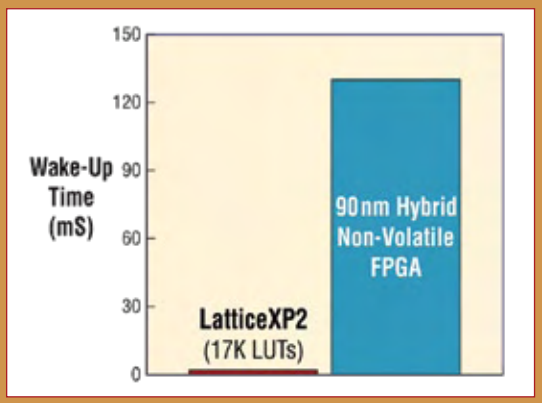

Figure 3

Stacked die technology also does not enable ultimate design security because any stacked-die implementation must transmit the unencrypted configuration data from the boot PROM to the SRAM FPGA over bond wires within the package. A hacker need only perform minor surgery to gain access to the wires that transmit the configuration data. PL

Single Print Only

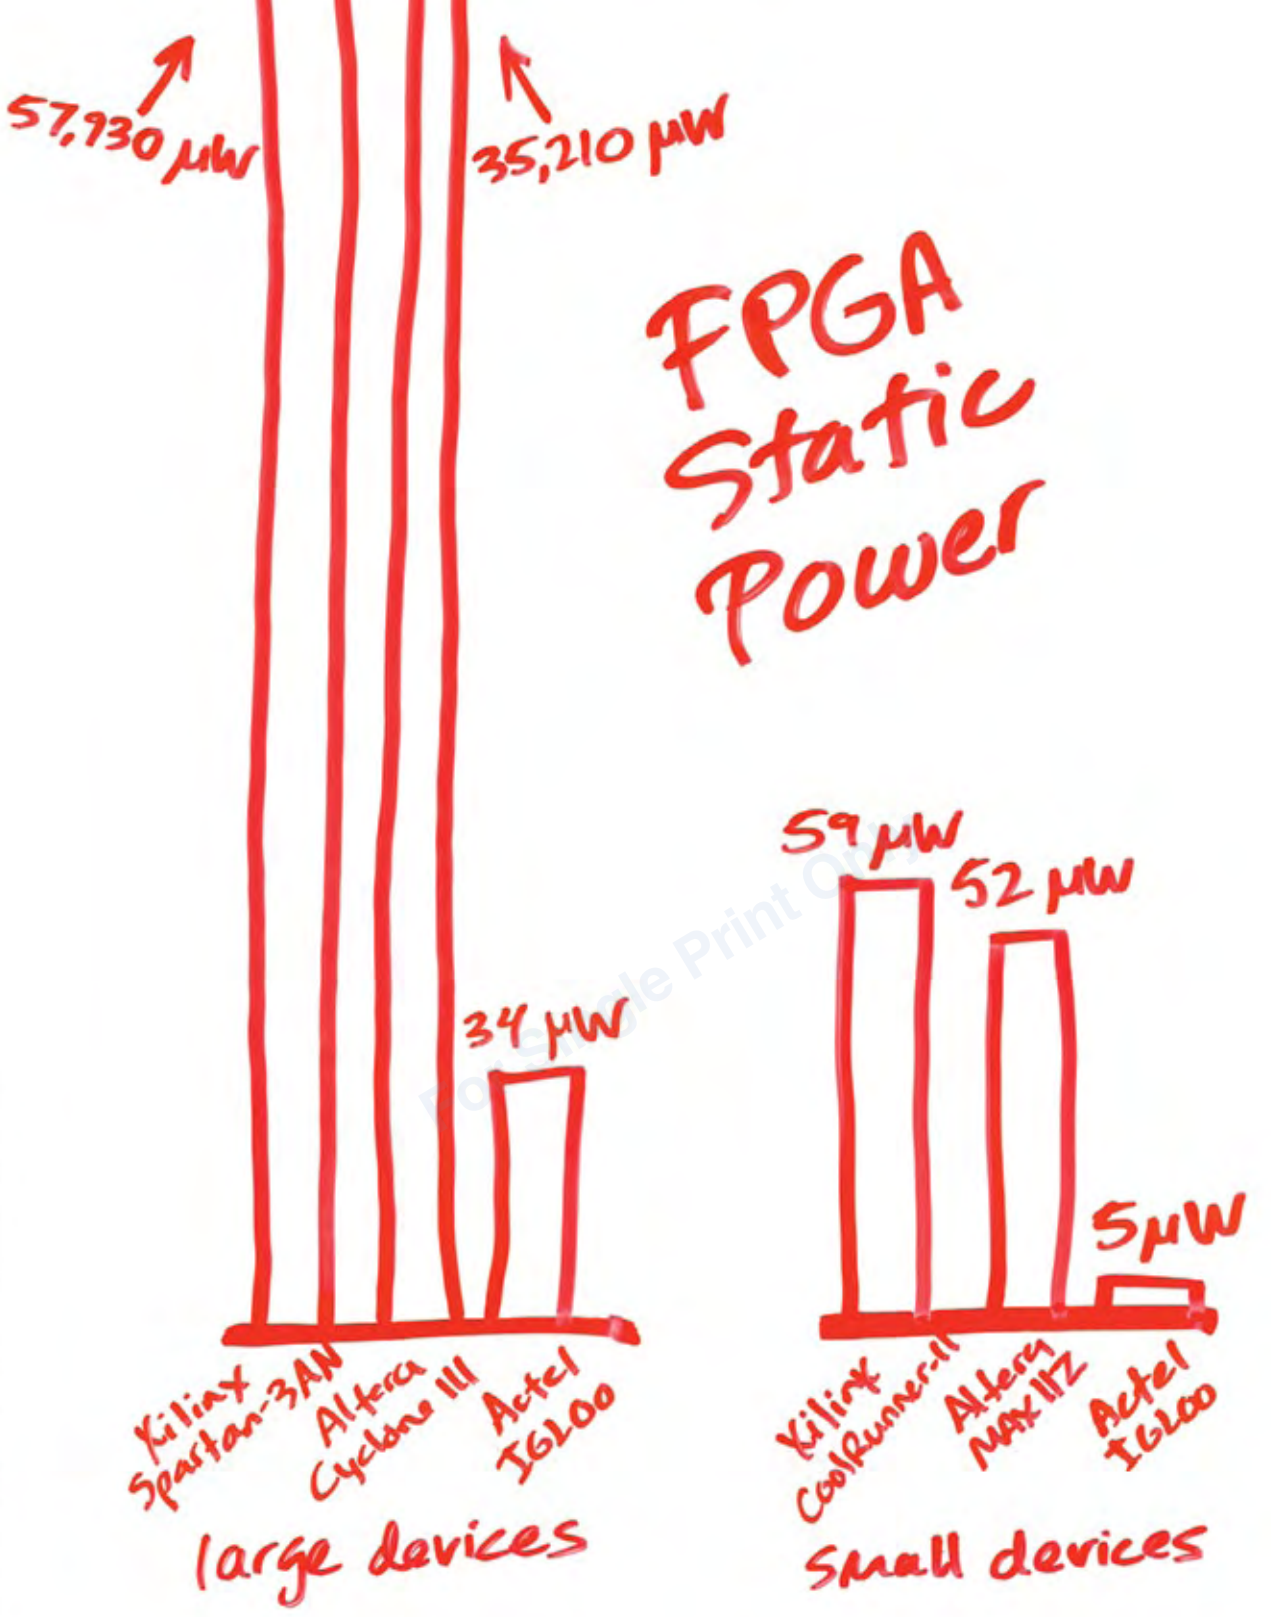

WHAT DO WE HAVE TO DO, DRAW YOU A PICTURE?

Only Actel gets you this close to zero. Any other claims of low power superiority are just that. According to their own data, Altera® and Xilinx® use between 10 and 1700 times the power of Actel IGLOO® FPGAs, depending on device and mode. Want specifics? Visit us to get the whole picture, including a video of actual measurements.

more proof and pictures at **actel.com/power**

Actel®
POWER MATTERS

© 2008 Actel Corporation. All rights reserved. Spartan and CoolRunner are registered trademarks of Xilinx, Inc. Cyclone and MAX are registered trademarks of Altera Corporation.

"Projecting" images in radar and medical applications

By David Pointer

Examples of computationally intense systems being shrunk by use of FPGAs are of keen interest to us. Filtered back-projection is finding its way into both radar and medical imaging applications, and is well served by FPGAs to handle a portion of the algorithm. The results are outstanding.

Searching for an enemy vehicle from a high-flying Unmanned Aerial Vehicle (UAV). Taking a child after a tumble from a bike to the hospital for a CT scan of their broken arm. The imaging techniques in these two seemingly unrelated scenarios rely on the same technology to produce accurate results: the Filtered Backprojection (FBP) algorithm.

The FBP algorithm reconstructs an image of an object by calculating an exact solution for each pixel from a series of planar image data sets (projections) of the object. Both radar and medical imaging require a high degree of imaging accuracy. The FBP algorithm provides excellent imaging accuracy but at a large computational cost.

Heterogeneous reconfigurable systems provide the required computational power for image reconstruction to both Synthetic Aperture Radar (SAR) Backprojection imaging and Computed Tomography (CT) imaging while delivering such significant benefits as smaller form factors and reduced power consumption.

Partitioning the FBP

The general implementation of the FBP algorithm is quite straightforward. A two-dimensional set of floating point data representing the energy detected from illuminating an object (using radar energy for SAR, or X-Rays for CT) is filtered to remove noise, then each projection's contribution to the reconstruction is summed into each pixel in the image.

For images of a useful size, the summation operation over all pixels is the computational bottleneck. Summing each projection's contribution for a single pixel is fast; summing the effect of each projection over one million pixels for a 1024 x 1024 image – or even larger images – is intensely time-consuming.

Programmers can achieve high performance for FBP algorithms by spreading the application across an FPGA and traditional CPU combined in a heterogeneous reconfigurable system as shown in Figure 1. An FPGA contributes to a reconfigurable system's performance by allowing a programmer to explicitly and completely dedicate a device to the solution of the regular, uninterrupted streaming aspects of a program. Along with this computational intensity, an FPGA also provides parallelism which can be exploited The CPU, with its high clock rate, is an excellent workhorse for executing the irregular and conditional aspects of any program.

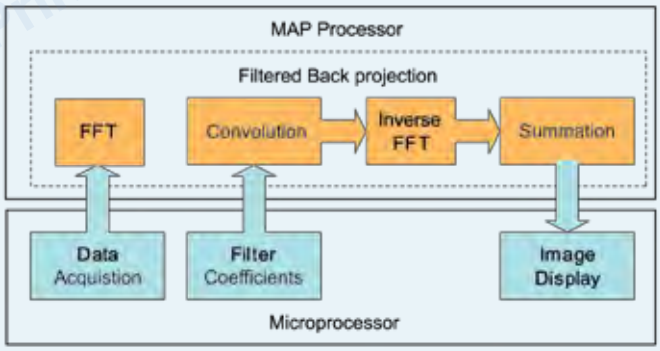

Figure 1

Equally sized sets of projection data stream into the FPGA for filtering, which is easily implemented in the frequency domain as a multiplication operation. Spatial domain projection data sets are streamed through a *fast fourier transform* (fft), multiplied by a filter coefficient matrix, and then converted back to the spatial domain using an inverse fft. These three operations are all pipelined (overlapped in time) for maximum data throughput through the FPGA.

After filtering, the backprojection (summation) step of the FBP algorithm is independent for each pixel, and so may be implemented as a set of simultaneously executing parallel summation units in the FPGA program for maximum application performance.

In a typical imaging application, the microprocessor generates appropriate filter coefficients and acquires the sensor input data. After FBP processing by the MAP processor, the final image data may be displayed or stored by the microprocessor.

Shrinking an SAR system

Given the high computational cost of the backprojection algorithm, it can be an engineering challenge to deploy such radar imaging in portable computing systems and small UAVs while keeping overall vehicle size, weight, and power (SWaP) down to acceptable levels.

During the early specification requirement studies for various UAV programs, the United States Air Force was concerned about the feasibility of deploying SAR backprojection on mid-sized UAVs. Could a reconfigurable SAR system meet both SWaP requirements and processing requirements?

To address these concerns, engineers from the Air Force Research Laboratory (AFRL) and SRC Computers jointly benchmarked the Spotlight Synthetic Aperture Radar (SAR) algorithm in 2005. The results obtained on the SRC-6 Portable MAPstation, which contains both an Intel 1.6 MHz Pentium M CPU and SRC's FPGA-based Series F MAP processor, showed a 75x performance increase using both the CPU and MAP over using just the CPU. The absolute wall clock time for the benchmark execution exceeded the Air Force requirements.

Based upon these published results, Lockheed Martin recently selected SRC to provide the Signal Data Processor (SDP) for the United States Army's Tactical Reconnaissance and Counter-Concealment Enabled Radar (TRACER) program. This system, scheduled to fly on the Warrior UAV in 2008, contains multiple MAP processors for even greater throughput.

Extending to medical imaging

Last year, discussions with customers in the medical imaging field indicated that CT scan image reconstruction used the same FBP algorithm as SAR imaging. The application engineers at SRC Computers found an open source CT scanner simulation program, CTsim, and ported the FBP portion of this medical imaging application to the SRC-6 MAPstation.

The results of this implementation, reconstructing a 1024 x 1024 image showed a 22x performance improvement with the combined CPU and Series E MAP over the MAPstation's 2 8 GHz Xeon 32-bit CPU. These results indicate that manufacturers of CT scan equipment could increase the resolution and clarity of the output of their equipment, draw significantly less power than traditional CPU-based systems, and still enjoy faster image reconstruction.

SRC-7 MAPstation: Early results

When the new SRC-7 MAPstation started shipping in 2007, application engineers at SRC began work to obtain updated application performance data on several programs, including CTsim and SAR back projection.

For CTsim, a direct recompilation of the existing Series E MAP processor code without using any of the Series H MAP enhancements yields a 29x performance improvement when compared to a 3.0 GHz Xeon 64-bit CPU. Proposed modifications to the FBP implementation in CTsim in order to take advantage of the Series H MAP (shown in Figure 2) suggest that a 60x performance improvement may be realized with minimal effort.

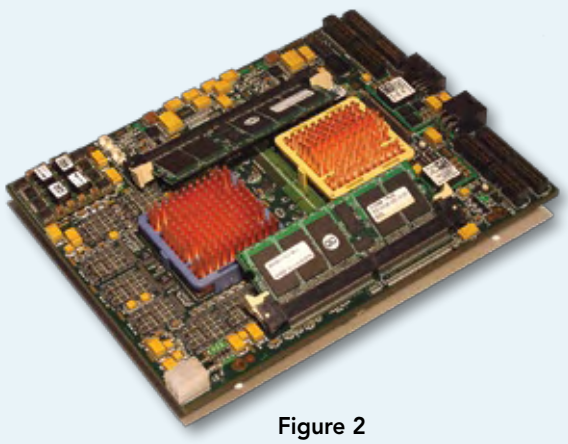

Figure 2

An FPGA contributes to a reconfigurable system's performance by allowing a programmer to explicitly and completely dedicate a device to the solution of the regular, uninterrupted streaming aspects of a program.

For SAR backprojection, optimization work is underway to port this code to the Series H MAP. Analysis indicates that a 104x performance improvement (relative to a 3.0 GHz Xeon 64-bit CPU) is a reasonable expectation when this work is complete. It is worth noting that the high-performance fftw library was used to obtain the CPU-only execution measurements. The final performance results for CTsim and SAR backprojection executing on the SRC-7 will be published as soon as they are verified.

Some of the SRC-7 system improvements that enable this performance include a 50 percent faster FPGA clock rate, a five-fold increase in system interconnect payload bandwidth (from 2.8 GBps to 14.4 GBps), and the adoption of Altera FPGAs, resulting in a more than three-fold increase in the MAP processor's floating point performance.

Reconfiguring system performance

Actual, realizable application performance on reconfigurable systems is important, but it is also good to consider other aspects of application development. FPGAs were once the domain of hardware engineers and Hardware Descriptor Languages (HDLs) like Verilog and VHDL, and programming FPGAs was notoriously difficult.

Recent advances, however, now make reconfigurable systems accessible to traditional software programmers. SRC Computers provides a unified programming environment, called Carte, which allows the programmer to use ANSI C or FORTRAN languages for combined CPU and MAP coding.

Software engineers, take note. Reconfigurable computing is spanning more industries and applications than ever before, and is no longer just a playground for hardware engineers.

PJ

David Pointer *is Director of System Applications for SRC Computers, Inc. (Colorado Springs, CO). His engineering career spans 29 years in both industry and academia, most notably at Hewlett-Packard Laboratories in Palo Alto, CA, and the University of Illinois at Urbana-Champaign. He received an M.S. in Electrical Engineering in 1982 from Northern Illinois University and an M.S. in Computer Science in 1996 from the University of Illinois at Urbana-Champaign.*

SRC Computers

719-262-0213 • www.srccomputers.com

MATHSTAR™

1 GHz Field Programmable Object Array™

FPOA Advantages:

- 1 GHz Operation
- 1 GHz Interconnect Fabric
- No Timing Closure Required
- Faster Time-to-Market

Discover our FPOA product family at:

www.mathstar.com

© 2006-2008 MathStar, Inc. all rights reserved.

SIGINT in the real world presents nuances

By Dr. Malachy Devlin

We've heard that FPGAs can do signal processing algorithms. But what about the rest of the problem: controlling analog-to-digital converters, compensating for effects of voltage and temperature, and more? Here, FPGAs get serious.

The growing use of FPGAs in SIGINT applications is enabling a significant increase in the capability of defense systems. With a wider variety of threats than ever before, the desire to monitor more of the frequency spectrum necessitates a high performance front-end processing engine. A SIGINT system that needs to monitor VHF (30-300 MHz), UHF (300-1000 MHz) communications and navigational aids, through to radar and satellite communications in the L band (1-2 GHz) frequencies, requires an analog bandwidth of 2 GHz for complete and instantaneous coverage.

Since there are a wide range of frequencies and consequently a wider range of signal modulation schemes, the processing system needs to have significant I/O bandwidth to receive the data as well as the computing performance to implement real time demodulation and analysis algorithms. Typical algorithms include Digital Down conversion, ffts, filtering, and I/Q demodulation.

DSPs short on two counts

While many of the algorithms used are familiar, it is the bandwidth of the input data that makes the implementation much more challenging. As an example, taking a simple 10 tap FIR filter and using this with a high end ADC, such as a 3 GSps part, the computing performance required is 30 GMACps. While this significantly exceeds the performance of high end DSP processors that offer around 1 GMACps, it is a fraction of the performance offered by high end FPGAs that are capable of up to 350 GMACps. This means that FPGAs have horsepower to spare to carry out additional tasks which will be required by typical applications.

Not only are the performance capabilities of DSPs under-powered for high data rate processing, but the I/O bandwidth to move these high speed data streams is unavailable. While there could be a data stream of 3 GBps from an ADC, DSPs typically provide bandwidth of 1 GBps. Clearly DSPs are not capable of handling this rate of data ingress for a single data channel. This problem is exacerbated considering antenna arrays and techniques such as beamforming are common approaches to improve the SNR of received signals.

FPGAs can tackle the problem in a much smaller space. Figure 1 shows a 6U CompactPCI card with 4 daughtercards, providing a total of 8 channels that can each capture 3 GSps at 8-bits (aggregate of 24 Gbps) into 5 FPGAs that can deliver up to 1 TeraMAC/sec of processing performance. The equivalent implementation using DSP processors would require at least a full 42" rack.

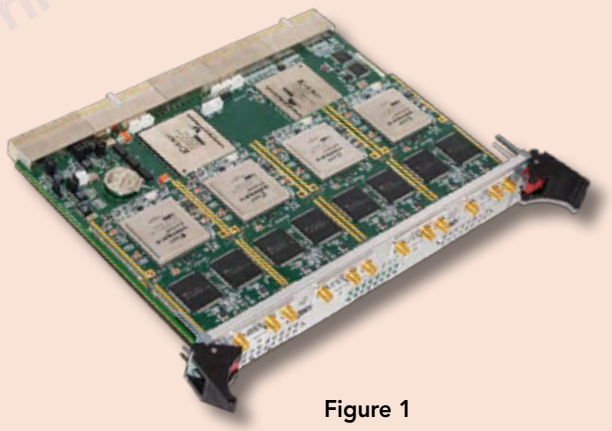

Figure 1

More control over digitization

In addition to the I/O and processing power that FPGAs provide, they bring further benefits to front-end processing applications by affording the user a finer level of control over the digitization process. FPGAs can directly control and manipulate the interface signals of ADCs and supporting peripherals such as clock synthesizers and clock path selection. Coupling high-speed data converters with FPGAs allows direct RF sampling, enabling digital manipulation of data and hence reducing complex analog RF circuitry for tuning and demodulation. Performing these functions in the digital domain rather than analog reduces sensitivity to environmental noise and improves the repeatability of processed results.

For example, incoming SIGINT data from antenna arrays must be captured with virtually no phase skew between channels. The challenge is ensuring two ADCs are accurately synchronized in order that the samples on each analog input are coincident in time. Channel synchronization is similarly critical when capturing I/Q signals, where the SNR will be degraded if the data from the two channels is misaligned resulting in reduced receiver performance and sensitivity.

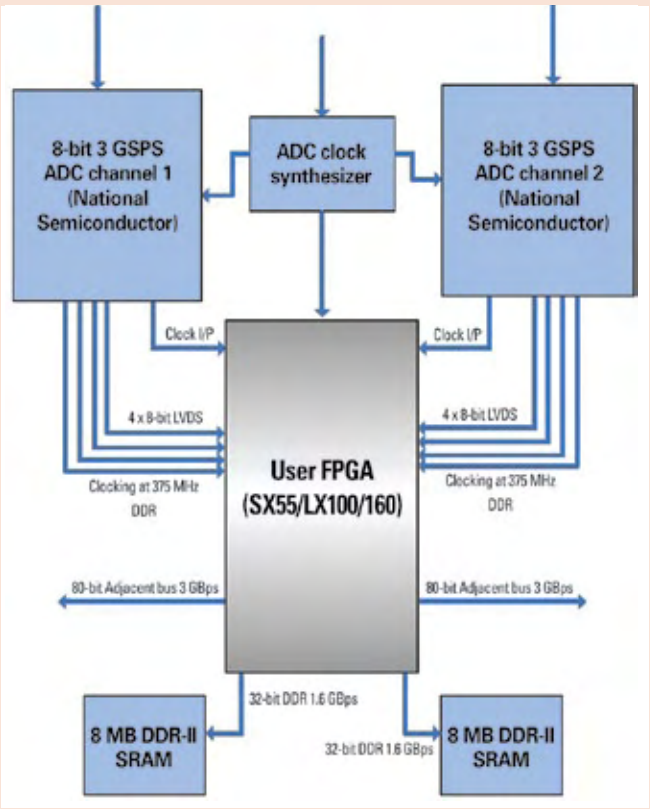

Figure 2

A lot goes into synchronicity

Figure 2 is a functional diagram of the daughtercard module (with four shown in Figure 1) with two National Semiconductor ADC083000 3 GSps ADCs connected directly to a Xilinx Virtex-4 FPGA. The key to synchronizing the ADCs is deasserting the device reset signal.

There are many parameters that need to be considered to ensure synchronicity of the channels, particularly when having to deal with the real world differences of time of arrival of the reset signals. In this specific implementation, the key factors that had to be considered in the analysis were:

1. Skew between clock inputs of the two ADCs
2. Skew between FPGA reference clock and ADC clocks
3. Reference clock delay through FPGA
4. Skew between the two ADC reset signals include FPGA clock to out skew
5. Relative jitter of reference FPGA clock to ADC clocks
6. Relative jitter between ADC clocks
7. Setup and hold of reset signal

Although the FPGA is used for handling the timing of the reset, some of these parameters are out of the control of the application developer (namely, items 1, 5, 6, and 7) and so need to be characterized in order to be managed. Fortunately, the flexibility of the FPGA means it can be used to run a dynamically configurable measurement application that measures the effect of the *uncontrollable* parameters. Using such an application, the valid window for the reset signal was measured as 306 ps, from a total period of 666 ps (assuming a 1.5 GHz clock that is used for 3 GSps data capture).

Fortunately, the flexibility of the FPGA means it can be used to run a dynamically configurable measurement application that measures the effect of the 'uncontrollable' parameters.

The critical element to maximize the reliability of channel synchronization is to deassert the reset signal in the middle of this data valid window. This shifting of the reset edge is achieved using the digital clock managers (DCMs) of the FPGA, which enable fine-tuning of the clock phase within the FPGA. The two reset signals to the ADCs are transmitted from the FPGA via flip-flops so when the phase of the FPGA clock is changed, the phase relationship between the reset signal and the ADC clocks is also changed (the ADCs are clocked directly from the clock synthesizer and are therefore independent of the FPGA). This technique enables the phase of the reset signal in relation to the ADC clocks to be changed in increments of 10.4 ps when sampling at 3 GSps.

Compensating for environmental variables

At this stage it seems a straightforward task to run the measurement design, find the center of the data valid window, and pre-configure the phase of the DCM to align the reset signal at that location. Unfortunately, environmental parameters affect skew and jitter performance.

Testing showed that the phase of the valid window for the reset signal (relative to the ADC clocks) changed with temperature. A second order non-linear relationship was found between the phase setting and the frequency and temperature values, which also varied depending on the FPGA device in use. The resultant formula for a Virtex-4 SX55 device is:

Phase(SX55) = round(mod(1.5479e-006*f*f - 0.00045165*f*t - 0.24071*f - 0.0033431*t*t + 0.27627*t + 348,256))

Where: t is the temperature in ºC
f is the clock frequency of the ADCs in MHz

This result greatly simplifies the system design as the phase setting can be computed based on two easily known variables. The alternative is to implement a calibration circuit in the FPGA and provide synchronized analog inputs on every startup, a more error prone technique.

Space saved

This application would require several hundred DSPs to implement, not to mention the need for a PLD to interface between the ADCs and processors, illustrating why FPGAs have become the first choice technology for front-end sensor processing. Not only do FPGAs provide computational prowess that reduce system SWaP by orders of magnitude, they have the I/O bandwidth and flexibility needed to implement the latest and most complex sensor interfaces.

Dr. Malachy Devlin *is CTO of Nallatech (Eldersburg, MD). He joined Nallatech in 1996, and is recognized worldwide as an industry expert on FPGA technologies. His experience includes roles with the National Engineering Laboratory (UK), Telia AB, and Hughes Microelectronics. He has a PhD from the University of Strathclyde, Glasgow, Scotland.*

Nallatech
410-552-3352 • www.nallatech.com

Kiss your FPGA good-bye

eASIC's no minimum order, Nextreme zero mask-charge ASICs are ideal for replacing your expensive and power hungry FPGAs. They are simple to design and you can have production ready devices back in only 4 weeks. With Nextreme you can also design complete systems that include embedded processors such as ARM or Tensilica cores.

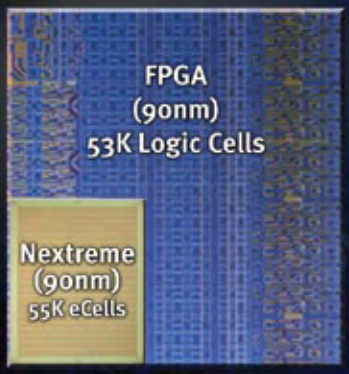

Up to 10X Lower Cost Than FPGAs

- Nextreme replaces SRAM interconnect with single via layer
  single
- Dramatically reduces die size & cost

Up to 15X Lower Power Than FPGAs

- Significantly reduce both standby and dynamic power
- Reduce power supply complexity and cost
- Eliminate expensive FPGA heat-sinks and fans

eASIC®

Learn more at

www.eASIC.com

Big images, complex processing? Think objects

By Tom Diamond

Got a 4K x 4K image to process? Equalizing, correcting, interpolating, adjusting, and filtering every pixel at 35 fps? The Field Programmable Object Array enters the fray and shrinks the processing hardware required.

Today's real-time image processing hard ware is feeling the pinch. Gigabit Ethernet links feed multi-million pixel full color space video to frame grabbers. Multi-core processors, fed by high bandwidth PCI Express ports, consume image processed data as fast as it can be delivered. It's no wonder that traditional reprogrammable devices are straining to keep pace across a range of automated inspection, security/surveillance, and professional video applications.

The Field Programmable Object Array (FPOA) is a high performance device that operates at speeds up to 1 GHz and is programmed at the object level. FPOAs are especially well suited to meet the requirements of ultra-fast high resolution image processing applications. The key to their performance is hundreds of objects that pass data and control to each other through a patented interconnect fabric.

The MathStar Arrix family of FPOAs provides 256 Arithmetic Logic Unit (ALU), 80 Register File (RF), and 64 Multiply Accumulator (MAC) objects. The objects and the interconnect fabric run on a common core clock, operating deterministically at 1 GHz in the Arrix Family's fastest device. This determinism enables a designer to select a core clock that meets the desired memory, I/O, and processing requirements.

An image processing example

In a high-end image processing implementation, 4K x 4K pixel color images are fed to an FPOA from the electro-optics of a high performance color CCD or CMOS camera as a single color channel at 35 frames per second. Use of a 900 MHz Arrix device enables external memory transfe s at the FPOA's 300 MHz top memory bandwidth.

The implementation has several stages:

- Non-linearities from process variations in the sensor array are corrected in the flat-field processing stage, equalizing pixel gain across the array and correcting for dark current flow in the sensor array.
- The Bayer demosaicing process generates RGB pixels by interpolating the color value for each pixel based on the color filter array architecture. Demosaicing increases the number of color channels to a total of three, effectively tripling the processing load.
- RGB data is then adjusted across all color-space components for the ambient lighting qualities or targeted display qualities in the color balancing processing block.
- Lastly, the image is filtered using a 3 x 3 convolution kernel before being driven out of the FPOA at the pixel rate.

Flat-field correction, Bayer demosaicing, color balancing and spatial filtering are mapped to the FPOA, as shown in Figure 1.

Flat-field correction

Flat-field correction is used to adjust image-sensor output data to ensure that constant intensity images generate constant pixel values at various levels of image intensity. This process addresses three types of pixel-based non-uniformities: gain, dark current offset, and defective pixels.

The key to this performance is hundreds of objects that pass data and control to each other through a patented interconnect fabric [running at up to 1 GHz].

For Single Personal Use Only

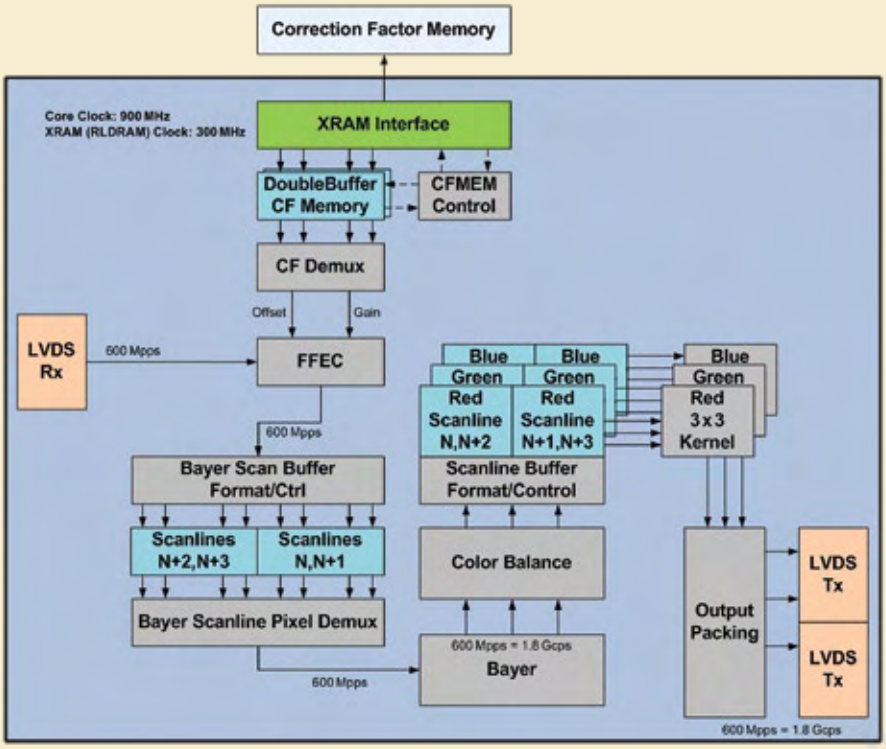

Figure 1

To rectify these non-uniformities, both a calibration and correction process must be performed. The calibration process determines the correction factors for pixel gain and offset and generates a defective pixel map. The correction process calculates an appropriate value for non-uniform pixels, and includes a dead pixel correction unit. Gain and offset correction can be performed using a single MAC object running at the core clock rate. Correction factor memory resides off-chip and can be accessed as two correction factor pairs every memory clock cycle.

Bayer demosaicing

A Bayer color filter array filter is an optical filter that routes specific wavelengths of light to a particular photo-sensor in a digital camera's sensor array. Each photosensor will detect the intensity of a certain light band. After filtering, the classical Bayer filter decomposes light into red, green, and blue components.

The digital processing associated with Bayer filters determines the missing color values for a given pixel (for example, the red and green levels for a blue sensitive sensor). Reconstructing the missing pixel colors of the CFA is known as demosaicing. This processing stage interpolates the missing color values of the color filter array (CFA) sensor pattern and generates a fully populated RGB image.

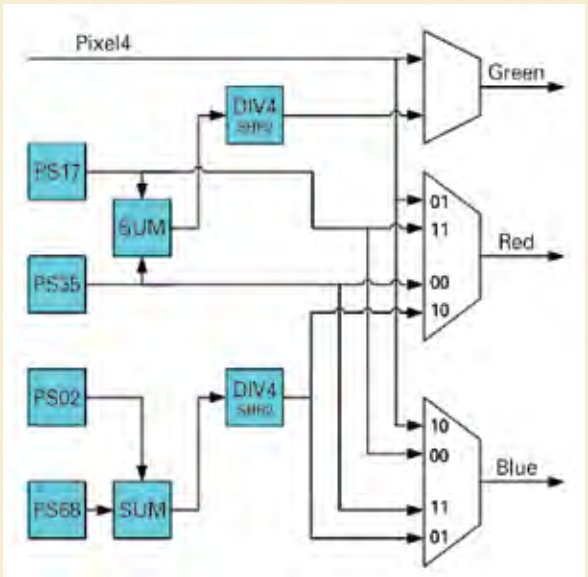

Figure 2

Pixel size at the input is a single red, green, or blue value and, at the output, consists of red, green and blue components for each pixel. Given a frame rate of 35 fps for the 4K x 4K image, this interpolation stage increases the processing throughput from 600 M to 1.8 G color components/sec. 18 ALUs support a 600 Mps input, generating three parallel outputs for each pixel (shown in Figure 2), interpolating in a 3 x 3 region about the pixel of interest.

Color balancing

The color balancing operation utilizes a nine-entry color transform matrix to correct for monitor or lighting irregularities. A fully parallel color balancer implementation at the core clock rate utilizes three MAC objects for each row of the color transform. A single RF per MAC for table distribution utilizes the local memory resources, localizes data into each MAC, and provides for support of up to 64 correction tables.

Spatial filtering

Image filtering for smoothing, gradient or edge enhancement, and sharpening is accomplished with a two-dimensional convolution using a 3 x 3 pixel mask. A fully parallel implementation of the filter is capable of calculating the result for one input every core clock. Since each color component needs to be processed separately with potentially separate masks per color component, a total of three filter processing blocks are required.

Summary

Performing flat-field correction, demosaicing, balancing, and filtering stages for a 4K x 4K, 35 frame per second, full color space frame grabber is a tall order, but is easily handled by a high performance FPOA. Further, this image processing chain utilizes just 50 percent of the objects in a single FPOA, providing room for additional image processing functions.

Tom Diamond *is marketing director for MathStar, Inc. Previously, he managed roadmap planning and product definition for Intel's LAN Access Division, and has worked in consumer electronics, test and measurement, and high volume silicon firms. Tom holds a BSME from Duke University and an MBA from the University of Washington.*

MathStar, Inc.
503-726-5500
www.mathstar.com

LatticeXP2 & MachXO
More of the Best

Live Update Programmability
Fast, secure design changes.

- ✓ Low Cost Mainstream FPGA Architecture
- ✓ Single Chip Integration
  - Instant-on through on-chip Flash
- ✓ On-chip RAM and DSP
- ✓ Flexible I/O
  - Up to 540 I/O
- ✓ Small Footprint
  - 8 x 8mm csBGA packaging

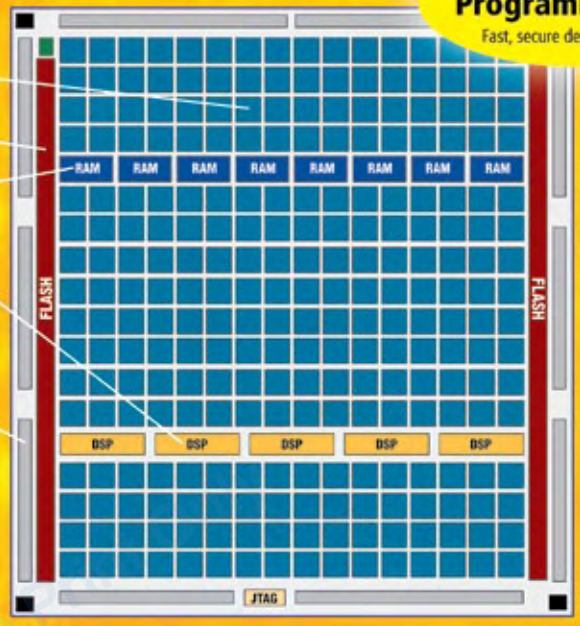

The Best Low Cost, Non-Volatile Logic

Fewer chips, smaller footprint, instant-on start-up … those are just a few of the benefits of Lattice's non-volatile programmable logic devices. Lattice offers the broadest range of non-volatile industry-standard architectures available. With on-chip block RAM, enhanced design security, and the flexibility to accelerate time-to-market, our latest LatticeXP2™ and MachXO™ devices deliver More of the Best.

LatticeXP2 devices combine an efficient FPGA fabric with non-volatile Flash cells in an architecture that was optimized from the outset with high performance and low cost in mind. LatticeXP2 FPGAs offer 5K to 40K LUTs, distributed and embedded memory, Phase Locked Loops (PLLs), pre-engineered source synchronous I/O and enhanced sysDSP™ blocks.

MachXO Evaluation Board
Complete FPGA starter kits starting at only $99!
www.latticesemi.com/99

The MachXO family combines the best attributes of CPLDs and FPGAs to serve applications that traditionally were implemented in CPLDs or low capacity FPGAs. MachXO devices include PLLs and embedded memory along with other features you wouldn't normally expect in a CPLD.

Make your next design your best with non-volatile programmable logic. Go to www.latticesemi.com/nv.

Lattice Semiconductor Corporation
More of the Best

©2008 Lattice Semiconductor Corporation. All rights reserved. Lattice Semiconductor Corporation, L (& design), Lattice (& design), LatticeXP2, MachXO, sysDSP and specific product designations are either registered trademarks or trademarks of Lattice Semiconductor Corporation or its subsidiaries, in the United States and/or other countries. Other marks are used for identification purposes only, and may be trademarks of other parties.

Programmable logic: The key to effective interface design

By Doug Morrissey

On the topic of shrinking designs using programmable logic, Doug walks us through how complex systems like media gateways are improved using FPGAs to incorporate faster, and sometimes missing, interfaces to DSPs.

Now that telecommunications OEMs are seeing sustained growth in the VoIP arena, margins are being squeezed and equipment designers are being asked to do more with less. For example, IP PBX and multiservice access nodes (MSANs) need to be price competitive at many different channel densities, and yet, the engineering cost of designing many different platforms is exorbitant.

Media gateway functionalities are best offered today using DSP technology. A few select vendors offer field-proven products that include a complete hardware and software solution. But in order to properly take advantage of these solutions, hardware designers are asked to make compromises.

The main problem is that DSPs don't always offer the right combination of DSP performance and interfaces. For example, very often it is only the largest DSP in a family that supports high-end interfaces such as SGMII Ethernet or RapidIO. When it comes time to interconnect many DSPs in larger systems, an FPGA is often used to glue devices together and perform functions such as back-pressure on packet transmission or packet classification. Hardware designers have already understood the value of using programmable logic in these scenarios.

Same DSP, but different interfaces

While different types of equipment can use the same underlying media processing technology, they each have specific requirements and constraints. Very often, the same DSP devices are targeted at all of the following applications:

- Large-scale media gateways
- IP PBX
- HMP media server off-loading
- MSAN/Access equipment

While the basic feature set may be the same, all applications require VoIP processing with a list of codecs and voice quality features; the interfaces required for each may be quite different. For example, a PCI card designed for host-based media processing (HMP) off-loading would require a PCI Express interface, whereas a MSAN may require a lightweight interface to an Ethernet PHY and to T1/E1 framers, as well as a lot of glue logic to the various components on the board. Sometimes the ideal solution is "off the beaten track." Designers may sometimes find a more cost-effective way of doing things, but find that they are limited by what the DSPs support. In the IP PBX space, certain vendors have discovered that USB is the ideal interconnect for low-cost systems, instead of Ethernet.

USB provides many advantages. First, it is hot-pluggable by design. More importantly, it offers a cost structure that is more advantageous than Ethernet. A USB hub device is less expensive and less complex than an Ethernet switch, which means that the base cost for a system is lower. Furthermore, the incremental cost is lower, because for each processing blade added to the system, very simple, low-cost RISC processors are available with USB interfaces. Most DSPs offered into the VoIP gateway market do not support USB, whereas FPGA vendors provide IP cores for such an interface. Leveraging FPGAs allows for this kind of innovation.

Combining DSPs with FPGAs

Ideally, OEMs should be able to get exactly the interfaces they need, without the cost and drawbacks of buying a fully-featured DSP. Typically, the DSPs that contain high-speed interfaces tend to have a large number of different interfaces, and be in a very large package with a very high pin count. These large packages are dense BGAs that are expensive and require a PCB with a high number of layers, which can all contribute to design delays.

By combining an FPGA with a DSP, the hardware designer can get the best of both worlds. In cases where the DSP interfaces are judged inadequate, or customization was required, DSPs are often front-ended with an FPGA. Adding an FPGA to an already fully-featured DSP is overkill. It adds unnecessary cost and board real estate, only to correct short comings in the DSP.

Ideally, OEMs should be able to get exactly the interfaces they need, without the cost and drawbacks of buying a fully-featured DSP.

A different approach can be taken when combining a DSP and FPGA. In this method, the DSP serves only as a DSP. The physical package contains a high-performance, multi-core device and uses a single memory interface to communicate with the FPGA. In this case, there is no carrying cost for the interfaces inside the DSP, keeping area and power to an absolute minimum. The DSP vendor provides turnkey FPGA designs for each major application, which can then be customized by the OEM if desired. The FPGAs used can be high-end devices or very small, low-cost devices, depending on the target application (Figure 1).

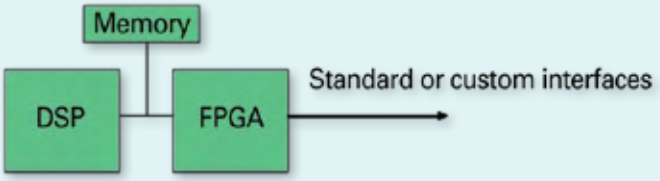

Figure 1

The main advantage to this approach is the flexibility that it offers. FPGAs can be reprogrammed in the field if necessary. Furthermore, FPGAs are always among the first devices to be released at each new process node. Therefore, they always support the latest I/O technologies with optimal power performance. As interconnect standards continue to evolve, it is possible to retarget a design very quickly to the latest fabric, whether it's PCI Express, RapidIO, or a proprietary interface.

The combined DSP and FPGA approach offers many cost savings. It is important to look at the total solution when evalua ing cost. For example, there is a plurality of Ethernet standards available: MII, RMII, GMII, RGMII, SGMII, and so on. When deciding how to interconnect many VoIP DSPs together, hardware designers can choose from a large number of Ethernet switch ICs, or consider a FPGA solution. If a certain Ethernet switch offers the ideal price point, but does not support the interfaces provided by the DSP, then a more expensive switch must be used.

More savings possible

The cost savings can be further extended beyond the electrical interface. For example, many large OEMs use proprietary headers or flags to help classify packet traffic. With an FPGA front-end, this becomes trivial and can greatly increase performance by off-loading certain tasks from the DSP. This type of header manipulation is best done in hardware by an FPGA. This offers deterministic performance, and keeps control of the design in the hands of the OEMs.

A media gateway DSP often incorporates a TDM interface that is connected to timeslot interchangers (TSIs) or T1/E1 framers. Very often, an FPGA is already present between the DSP and these TDM-based devices, or the TDM backplane. FPGAs are currently used because custom logic is required, or because the FPGAs offer more reliable I/Os that can sustain the high voltage spikes often seen on backplanes. These FPGA tasks can now be incorporated into the front-end FPGA saving more money and board area.

Another major advantage of this FPGA front-end is to make life easier for designers who again find themselves "off the beaten track." DSP manufacturers release products with feature sets and interfaces for specific market segments, usually the largest ones. It is easy to find a DSP that can serve as a generic TDM to VoIP gateway, but if the job is to design a high-performance telepresence video over IP conferencing unit, then the choice of DSPs is quite slim The amount of bandwidth required for passing compressed or uncompressed video streams between DSPs is immense. Finding a DSP with the right interface for this or any o her emerging market is quite difficult and costly.

Instead, using an FPGA front-end allows access to the latest I/O technology, as well as IP blocks for the latest standards. The designer can now get the right interfaces, without making changes to the DSP software. In organizations with large FPGA teams, the designer may choose to develop a proprietary interface that meets his or her needs perfectly and provides a competitive edge.

Media gateway designers are faced with many challenges including increased price pressure and shortened design cycles. By using a DSP that is tightly coupled to an FPGA, designers are able to get exactly the interfaces they need, without the added cost or debugging of useless I/O banks. This greatly simplifies the design and allows OEMs to add their own unique features and differentiate their system. This kind of innovation is what allows OEMs to address current and future customer requirements while maintaining a common DSP software platform.

PJ

Doug Morrissey *is Vice President and Chief Technology Officer at Octasic and has over 10 years of experience in the definition and marketing of semiconductor devices. Joining Octasic in 1999, he strategically focuses on issues within the Voice over Packet market. Previously, he worked as Marketing Manager for ATM and DSL products for Agere (formerly Lucent Technologies, Microelectronics Group), and was Senior Systems Architect at Unisys Corporation. Morrissey holds a BSc from Rochester Institute of Technology.*

Octasic Inc.

514-282-8858 • www.octasic.com

COVER PHOTOS:

Top left: Xilinx Virtex-5
Top center: Jacyl Technology XG-5000K
Top right: eASIC Nextreme
Middle left: Altera Stratix III
Middle right: Actel IGLOO
Bottom left: MathStar Arrix
Bottom middle: Lyrtech SFF SDR platform
Bottom right: Lattice Semiconductor LatticeXP2

ADVERTISER INFORMATION

Introducing the One Board with all the Right Connections

Targeting high-performance I/O and processing applications, the Model 4207 delivers an on-time solution for your real-time application using the latest high-speed serial interfaces.

On board PowerPCs, FPGA, PMC/XMC sites, and extensive memory resources are connected to the industry's fastest interfaces including VXS, Gigabit Ethernet and dual Fibre Channel ports for high-speed recording.

All board resources and interfaces are joined through a configurable zero-latency crossbar switch using protocol-transparent gigabit serial data paths.

With so many connections to choose from, you are bound to find just what you need for your next system.

- Dual PMC/XMC Sites
- Freescale MPC8641D Dual Core AltiVec PowerPC Processor to 1.5 GHz
- XC4VFX60 or XC4VFX100 FPGA
- 4 GB DDR2 SDRAM
- PCI Express, Serial RapidIO™, RocketIO™
- Dual Gigabit Ethernet Ports
- Dual 4 Gbit Fibre Channel Interface
- Fabric-transparent Gigabit Crossbar Switch
- VXS, VME64x, Gigabit Ethernet-X
- Linux and VxWorks Support
- ReadyFlow® Board Support Libraries
- GateFlow® FPGA Design Resources

Call today to have it all!
201-818-5900 or go to www.pentek.com/go/dspconnect for the 4207 brochure, technical datasheet and to request pricing.

PENTEK
Setting the Standard for Digital Signal Processing

Pentek, Inc., One Park Way, Upper Saddle River, NJ 07458
Phone: 201.818.5900 Fax: 201.818.5904 e-mail:info@pentek.com www.pentek.com
Worldwide Distribution and Support
Copyright © 2007, Pentek, Inc. Pentek, ReadyFlow and GateFlow are registered trademarks of Pentek, Inc.
Other trademarks are properties of their respective owners.

25% higher performance

29% lower power

1067-Mbps DDR3 support

FULL THROTTLE

Stratix® III

Stratix® III FPGAs, with blazing 1067-Mbps (533-MHz) DDR3 performance, are 25% faster than the competition while consuming 29% less power. With the highest density of any FPGA on the market, Stratix III devices allow you to integrate more functions in the same board space. Add Altera's Quartus® II software to the mix, and you'll complete designs in less time while meeting your power budgets. So, who wants to go faster?

ALTERA®

www.altera.com

tsmc

Copyright © 2007 Altera Corporation. All rights reserved.