Embedded COMPUTING DESIGN

www.embedded-computing.com

OCTOBER 2007 VOLUME 5 NUMBER 7

Functional Density

The pressure's on designers to boost capacity in compact devices

OpenSystems Publishing

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

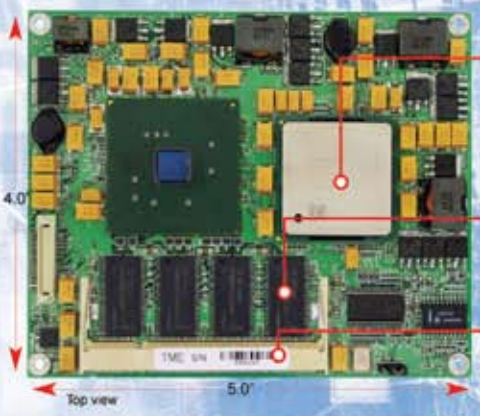

CPU, up to
2.8 GHz Pentium 4
(Low power 1.6Ghz and 2.0Ghz Pentium 4 also available)

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

SODIMM socket

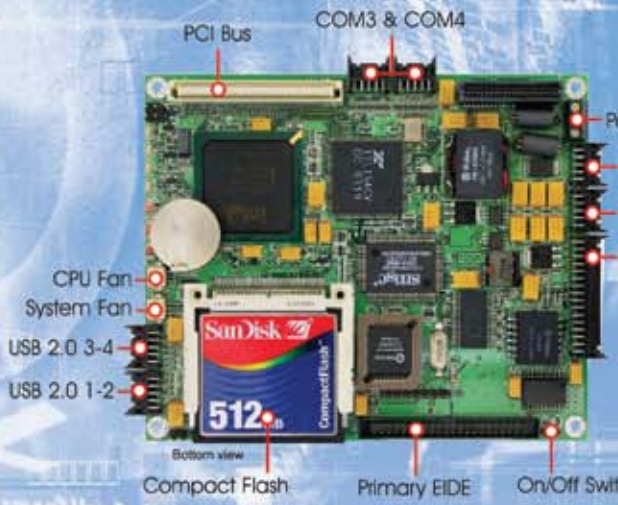

Features:

- (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

Industrial SBC Supports Wired and Wireless Communications

WinSystems' LBC-GX500 is a highly integrated, single board computer (SBC) designed for machine-to-machine connectivity with a wide variety of wired and wireless options. It provides an open and powerful platform for management of geographically distributed machinery.

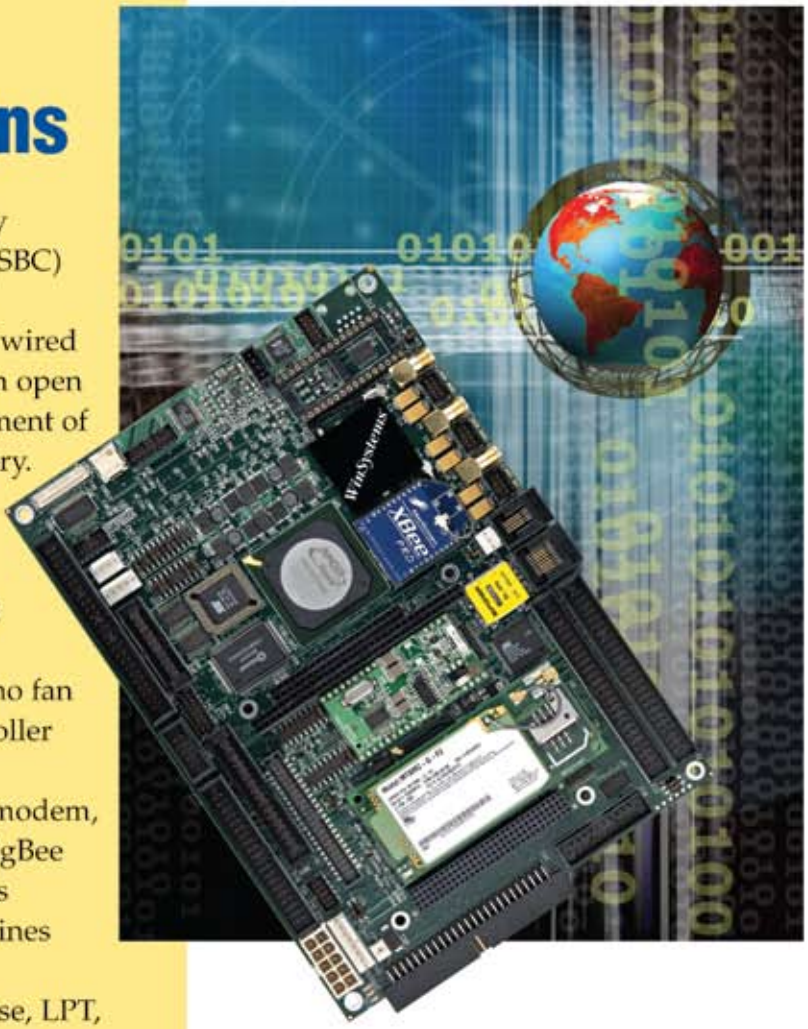

- AMD GX500@1W processor
- PC-compatible: supports Linux and Windows® XP embedded and other popular RTOS
- Operates from -40° to +85°C with no fan
- 10/100 Mbps Intel Ethernet controller
- 10 COM ports and 6 USB ports
- Socket support for 56kbps POTS modem, GPRS/CDMA cellular modem, ZigBee and 802.11a/b/g wireless modules
- 48 bi-directional TTL digital I/O lines
- Flat panel and CRT supported
- Onboard AT keyboard, PS/2 mouse, LPT, FDC, and UDMA disk controllers
- Type I and II CompactFlash cards supported
- PC/104 and PC/104-*Plus* expansion
- Optional 12-bit A/D converter, 8 SE/4 DI
- Optional Trimble GPS receiver
- EBX-size: 5.75" x 8.00" (146mm x 203mm)
- Long-term product availability
- Quick Start kits for software development

Off-the-shelf delivery, knowledgeable technical support, long-term availability and the right price makes WinSystems' LBC-GX500 the SBC of choice for your application.

Call 817-274-7553 or Visit www.winsystems.com/LBC-GX500

Ask about our 30-day product evaluation!

WinSystems®

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

Embedded COMPUTING DESIGN

www.embedded-computing.com VOLUME 5 • NUMBER 7 OCTOBER 2007

COLUMNS

DEPARTMENTS

EVENTS

ARM Developers' Conference & Design Pavilion
October 2-4 • Santa Clara Convention Center, Santa, Clara, CA
www.arm.com/developers conference

ISA Expo 2007
October 2-4 • Reliant Center, Houston, TX • www.isa.org/expo

COVER
The more capabilities embedded systems designers can squeeze into their devices, the better. Read more on the challenges of cramming functionality into space-constrained devices on page 10.

Published by: OpenSystems Publishing™

FEATURES

E-CASTS

MicroTCA
October 9, 2 p.m. EDT

RapidIO
October 10, 2 p.m. EDT

Programming the Cell Processor with IBM's new Cell Math Library and VSIPL++
October 11, 2 p.m. EDT

www.opensystems-publishing.com/ecast

E-LETTER

October: **www.embedded-computing.com/eletter**
SAS as a budding fabric
By Sam Barnett, Vitesse Semiconductor

WEB RESOURCES

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

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

Submit new products:
www.opensystems-publishing.com/vendors/submissions/np

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

Editor's Choice • Embedded Computing Design • PRODUCTS

x86 CPU redefines ultra-low power

A hummingbird requires about 1 W of power to fly. The 500 MHz VIA Eden Ultra Low Voltage processor consumes no more than that under full load. Designed to meet the low-power requirements of industrial, commercial, and mobile applications, the fanless processor achieves unprecedented speed for an x86 processor in a 1 W power envelope. Within the 21 mm x 21 mm NanoBGA2 package, the VIA Eden ULV processor can squeeze into a small, light chassis.

Based on the VIA CoolStream Architecture, the processor is manufactured using an advanced 90 nm process, enabling speeds of up to 500 MHz with 1 W peak power and as low as 100 mW idle power consumption. Integrated into the processor is the VIA StepAhead Technology Suite, which boasts performance-enhancing features including the VIA V4 bus at 400 MHz, 16 pipeline stages, sophisti-cated branch prediction, and 128 KB full-speed exclusive L2 cache, ensuring the processor's low power consumption doesn't come at the expense of performance.

VIA Technologies, Inc. • 500 MHz VIA Eden ULV processor

Generator turns waste heat into electricity

Global warming is reportedly providing us with plenty of extra heat. Imagine being able to turn that into power to run our embedded computing devices! Nextreme has developed a miniature, thin-film Thermoelectric Generator (TEG) that converts heat directly into electricity. Ideal for waste heat conversion applications, the solid-state TEG delivers power-generation densities > 3 W/cm^2 and is optimized to provide power in a form factor that can be as much as 20x thinner than bulk material alternatives. Manufactured using semiconductor fabrication techniques, TEG can be utilized in applications including automotive, military/aerospace, thermal batteries, medical implants, and wireless sensor networks.

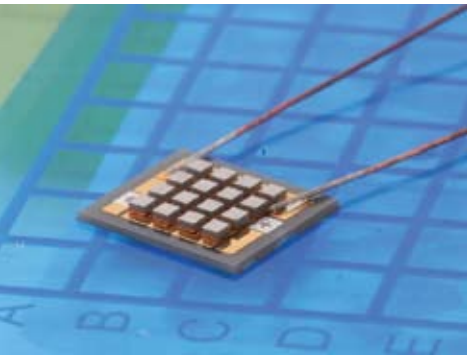

"In environments where a lot of heat is available, we have demonstrated power levels of up to 300 mW with devices that are not much bigger than a piece of confetti," said Nextreme CTO Dr. Seri Lee.

Nextreme's TEG devices generate electricity via the Seebeck Effect, where electricity is produced from a temperature differential applied across the device.

Nextreme • TEG

Embedded COMPUTING DESIGN

AN OPENSYSTEMS PUBLICATION

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
- PXI, Test & Technology
- PXI, Test & Technology E-letter
- PXI, Test & Technology Resource Guide

Editorial Director	Jerry Gipper jgipper@opensystems-publishing.com
Contributing Editor	Don Dingee
Associate Editor	Jennifer Hesse jhesse@opensystems-publishing.com
Copy Editor	Robin DiPerna
European Representative	Hermann Strass hstrass@opensystems-publishing.com
Special Projects Editor	Bob Stasonis
Senior Designer	Joann Toth
Senior Web Developer	Konrad Witte
Graphic Specialist	David Diomede
Graphic Coordinator	Sandy Dionisio
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
Assistant Managing Editor	Anne Fisher
Senior Editor (columns)	Terri Thorson
Technology Editor	Curt Schwaderer
Associate Editor	Jennifer Hesse
European Representative	Hermann Strass

Military & Aerospace Group

Group Editorial Director	Chris Ciufo
Assistant Editor	Sharon Schnakenburg
Senior Editor (columns)	Terri Thorson
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

Embedded computing's environmental impact

Jerry Gipper

Design for Energy Efficiency (DfEE) is something you're going to hear a lot about from the editors at OpenSystems Publishing in the coming months. The news is full of stories on global warming. At the same time, rising oil prices are causing us to once again take a closer look at managing existing energy sources and developing better alternate energy sources. Embedded computing is in a very interesting position because it could have an enormous impact on the environment and energy management.

Embedded computing has been directly responsible for improving the way we manage energy. Better control and monitoring of energy generation and machinery that uses energy have made these devices more effective and efficient. Automobile engines can attain greater fuel efficiency; home heating and air conditioning systems have better results with less energy usage, and power grids can better distribute power, all with the aid of embedded computing. Intelligent devices typically are more effective in managing energy.

On the other hand, embedded computing is also contributing to the environment in a negative way. Hazardous materials used in device manufacturing cause problems when using and disposing of old electronic equipment. This problem is amplified when you consider that many embedded computing devices have short life spans. Our drive to make them better makes them obsolete very quickly.

The Restriction of the use of certain Hazardous Substances (RoHS) in electrical and electronic equipment directives emerging around the globe are but a small step in the right direction. Eliminating hazardous substances can only be a good thing. Following the lead of the European Union, most developed countries have launched their own directives. In addition to RoHS, several other initiatives focus on recycling materials and equipment. Many electronics companies are starting to step up and offer recycling programs to end users. See www.electronicsrecycling.org for a lengthy list of companies providing electronic components recycling programs of one type or another. It's a good start, but more companies need to get involved.

Designers can do their part to reduce the churn of products that drives us to continuously "upgrade" our gadgets. Why can't products be designed to have longer, more useful lives so that we don't load our landfills with so much stuff? See my blog on transformers at www.embedded-computing.com for some thoughts on steps we can take to minimize waste.

Embedded computing electronics do affect the environment in a positive way. They make devices smarter and more efficient so that they consume fewer natural resources. But even here, there is tremendous room for improvement. Common technology like traffic control systems could be much smarter about traffic flow management. Just look at how much time we spend and fuel we expend waiting at stop lights or starting and stopping because a traffic control system is not smart enough to adapt to changing conditions.

In future issues, we will be discussing the effect of embedded computing on our environment. Here are just a few of the initiatives catching my interest right now.

The International Energy Agency's "1-Watt Plan" (www.iea.org) aims to reduce the standby power requirement of devices of all types to 1 W or less.

The National Electronics Action Plan search for Resource Conservation Challenge at www.epa.gov, addresses environmental concerns along the entire life cycle of electronics, including equipment design, operation, reuse, recycling, and disposal. This action plan will initially focus on computers (PCs), televisions, and cell phones.

I'll also be monitoring various legislative efforts currently in effect or being proposed that could influence designing embedded computing systems for energy efficiency.

Feel free to share your comments by e-mailing me or visiting the Embedded now blog to add your comments.

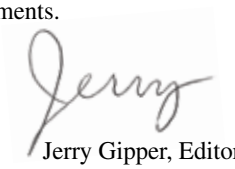

Jerry Gipper, Editorial Director

jgipper@opensystems-publishing.com

Vehicle handling under embedded control

By Hermann Strass

Europeans choosing ESC

Every year, more than 40,000 people die on heavily traveled European roads and more than 1 million Europeans are injured in road accidents. Loss of control such as skidding has been identified as the main cause, accounting for 40 percent of accidents involving serious injury or death.

It is estimated that equipping cars with Electronic Stability Control (ESC) could prevent 4,000 road deaths and 100,000 injuries per year in Europe. As a result, ESC, which is also known as dynamic stability control, electronic stabilization program, vehicle stability assist, or vehicle stability control, is becoming a promising eSafety technology. One European-wide campaign to promote the use of advanced ESC in cars, called *Choose ESC!*, was launched May 8 at the Bridgestone European Proving Ground at Aprilia near Rome. Michael Schumacher, seven-time world champion and the most successful driver in Formula 1 racing history, sponsors this campaign (information available at www.chooseesc.eu).

ESC was invented by Bosch in Germany and codeveloped by Mercedes-Benz. Figure 1 shows a first-generation ESC standard product from a 1995 series production consisting of 11 parts. Today, after 12 years of continued development, ESC models consist of only one part, as seen in Figure 2. The model on the right has an extremely powerful and fast hydraulic component.

Bosch spends about € 3,000 million (U.S. $4,000 million) each year on research. Mercedes-Benz and Bosch license ESC to other car manufacturers at no cost. Using ESC is currently voluntary in Europe, but 100 percent of all new cars are expected to be equipped with ESC by 2012.

ESC senses when a driver loses control and automatically applies braking pressure to individual wheels to help stabilize the vehicle and avoid skidding. It combines antilock brakes, traction control, and yaw control (spinning around the vertical axis) in one integrated embedded system. A powerful microprocessor system

Figure 1

Figure 2

calculates safe reactions in real time, even at speeds in excess of 240 km/h (150 mi/h) on German Autobahns. ESC is especially helpful in icy and wet conditions, such as aquaplaning or hydroplaning.

The U.S. National Highway Traffic Safety Administration (NHTSA) and the Insurance Institute for Highway Safety (IIHS) have confirmed international studies on the effectiveness of ESC, especially those involving ever-popular Sports Utility Vehicles (SUVs). The agencies report that SUVs with ESC are involved in 67 percent fewer accidents than SUVs without it. IIHS estimates that ESC could save about 7,000 lives and help drivers avoid up to 10,000 fatal crashes annually.

Events get attendees in the game

This year's Games Convention (GC) in Leipzig broke attendance records again. From August 22-26, 503 exhibitors (34 percent more than last year) from 35 countries showed their latest products to 200,000 visitors using roughly 1.2 million square feet of exhibition space. More than 2,800 journalists from 35 countries reported about this event and its 225 first-time displays of new games and equipment.

In a special exhibition at GC, Rene Meyer of Leipzig presented his collection of historic gaming equipment such as Atari 2600, Amiga 500, Super Nintendo, and other games from 1972-2007. He received a certificate and an entry into the Guinness World Records for the world's largest collection of electronic games equipment.

A record number of 900 games developers (38 percent more than last year) from 35 countries attended this year's GC Developers Conference (GCDC) August 20-22 in Leipzig. About 60 percent of the developers came from outside Germany, mostly from the United States, United Kingdom, and an increasing number from Eastern Europe. GCDC is the leading event of its kind in Europe. Information about GCDC and GC is available at www.gcdc.eu and www.gc-germany.com, respectively.

For more information, contact Hermann at hstrass@opensystems-publishing.com.

Limited space, maximum functionality

Embedded computing systems designers are often tasked with putting as much functionality into a space-challenged device as possible, whether that device is a chip, board, or system. Trying to determine which products might give the most bang for the space available in the embedded computing element can be quite difficult. Making the right choices can make or break the success of the final design. This article analyzes some of the issues and trends from components to systems that enable the best possible functional density.

By Jerry Gipper

Open up a new iPhone, and what do you see? A lot of functionality in a very small space. The iPhone designers were able to squeeze a phone, PDA, Web browsers, MP3 player, camera, game console, Web television, and many other functions into a device that fits very comfortably in the palm of your hand, and it looks good doing it.

Functional density is the term used to describe how much functionality can be placed into a given space. Today's embedded computing is all about functional density. Getting a lot of functionality into a small space reduces costs, improves reliability, enhances usability, and much more. Functional density can be driven by many factors, including the following considerations:

- **Weight:** Many applications have limitations on how much they can weigh. Electronics destined for aircraft are very sensitive to weight because it affects the cost to keep the aircraft airborne. Portable personal electronics like cell phones, laptops, GPS devices, and digital cameras are meant to be carried around. Users often will have more than one of these devices on their person. Too much weight, and devices get left at home or in the car.
- **Size:** Personal consumer devices also have a space limitation in addition to weight constraints. Too big, and again, they get left behind. Embedded electronics are often added into existing devices to improve their capabilities. Space can be at a premium if the original design had little or no room for electronics. Sometimes the space comes in a very odd shape that tends to be constrained.
- **Life-cycle costs:** This is a very different view on functional density that takes life-cycle costs into consideration. As more commercially available electronics building blocks (chipsets, boards, subsystems) become available for use in electronic systems, the life-cycle costs of a system can go up very quickly. This is because each of the components has its own life cycle that must be formulated into the total system life cycle. The system requires that work be done each time a part becomes obsolete; finding replacements, redesigning hardware and/or software, and qualification efforts all have to be re-expended. Thus, the number of COTS components directly influences the total system life-cycle costs. Reducing the number of COTS components by selecting those with the highest functional density helps reduce the life-cycle costs. In other words, fewer components, fewer failures, less cost.

Capabilities added throughout the integration chain

The types of functions that make up functional density can have different aspects that are of greater importance than others. These aspects can include computation or processing power, data storage, I/O, and the human machine interface. Some products are designed to have as much performance as possible in a given space, while others have specific types of input or output channels. Any combination in between can exist. Design goals driven by application needs determine what the proper balance of functionality should be.

Functional density can be measured at various points along the integration chain. At each point, the metrics for determining functional density and the challenges to optimizing functional density differ.

At the chip level are System-on-Chip (SoC) devices that package processing (now multiple processing cores), various types of memory, I/O interfaces, expansion buses, diagnostics, and many other functions into very small packages. Application-specific processors are often optimized even further by eliminating all unnecessary functionality. Gate density or transistor density is commonly used to indicate how capable a chip can be. Functional density implies more than gate density because it factors in the diversity, number, and type of basic building blocks used within the chip as part of the equation for determining functional density.

Board products have been the proving ground for functional density for years. Early boards had very specific, dedicated functionality on each board. One board was a processing element, a second was used for memory, a third for a particular type of I/O, and so on. A computing system would require many boards just to become a functional system. Board designers started combining multiple functions into single boards as the components for each function became smaller and more integrated. The concept of SBCs and motherboards emerged. Today, hundreds of SBC and motherboard options are available in a variety of form factors, from something smaller than a credit card to large, desktop PC-style motherboards.

The general size for a desktop or tower PC motherboard has remained fairly constant through the years, but functionality

has grown exponentially. Difficult functional density challenges have been left to laptops to squeeze the capability of a desktop into a small portable package. However, several examples of desktop technology demonstrate how the space footprint can be reduced and functionality packed in. Apple has incorporated much of the desktop PC into monitors or very small boxes that can be tucked away unobtrusively near the desktop. Apple achieves this functional density in a mostly closed environment.

New classes of micro PCs that use much smaller form factors are emerging. The Ultra Mobile PC (UMPC) is a handheld PC that has a high functional density ratio in a very small package. Others have designed functionally dense PCs that still maintain PCI slot card expansion capability. Stealth Computer Corporation has a small form factor PC with an expansion PCI card slot that is so small it can fit on the palm of your hand (see Figure 1). The chassis measures just over 2" tall with an overall footprint about the size of a hard-cover novel.

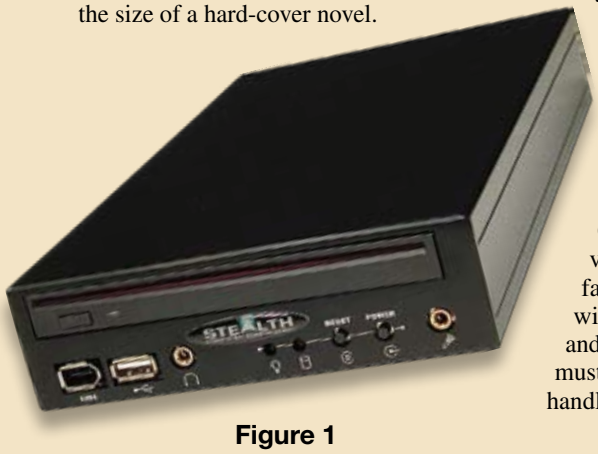

Figure 1

One of the best models of functional density can be found in the many mezzanine board products like PCI Mezzanine Cards (PMCs). The specification for PMCs dictates that they be very densely packaged. They have very limited panel space for connectors and must pack a lot of functionality into a very small space to be of any benefit as expansion cards. Figure 2 shows the DAS-429PMC/RTxD PMC card from Excalibur Systems.

Small packages pose big problems

Designing a product that has a high functional density ratio is no easy task. A by-product of functional density is thermal density. Put more electronics in a smaller or densely populated package, and the thermal density is bound to increase, often significantly. Managing the thermal density is a major challenge of a functionally dense device.

No one magic answer or solution exists when it comes to creating a functionally dense component, platform, or system.

"Thermal issues are a big concern when using dense packaging," says Phan Hoang, director of R&D at Virtium Technology. "The heat generated by memory modules is rapidly approaching the heat of a processor." Hoang also identifies other challenges. "Larger packages and modules are easier to lay out; smaller packages force different topologies. Capacitance is increased. Components need to be stacked or very closely placed, making manufacturing difficult." He notes that with very densely packaged chips and modules, manufacturing absolutely must be more automated and able to handle the closeness of components.

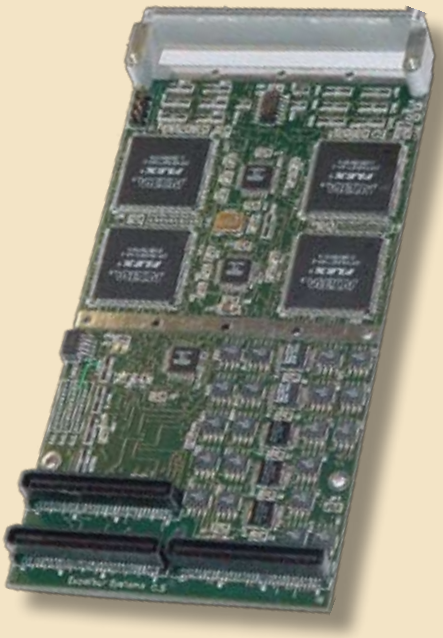

Figure 2

Designing touch screens with clear, capacitive sensors

By Andrew Hsu, strategic and technical marketing manager, Synaptics

Designers working on small systems such as mobile phones, portable digital entertainment devices, remote controls, and digital cameras continue to face major challenges in design and usability as consumers demand even smaller form factors, a better user experience, and increased functionality and storage capacity. For example, current generations of cell phones are posing User Interface (UI) and ergonomic issues because of a crowded keypad and touch-screen combination.

As a result, designers are migrating toward the extra UI space that Resistive Touch Screens (RTSs) offer. An RTS is a mechanical sensor with two layers of material typically separated by air. Pressing the top layer with your finger pushes the top layer and touches the lower layer.

The two main challenges embedded designers face with RTS are severely impaired optics of the underlying display and poor durability. Devices using RTS often fail in the field when they're dropped or the user presses the screen too hard. Once the RTS panel breaks, the device's main hardware input is eliminated, rendering the device useless. The designer's options are also limited. An RTS requires a costly top bezel to protect the edges of the screen and must be mounted on a flat surface, not beneath plastics.

A more efficient and reliable alternative is a thin, transparent, capacitive sensor touch screen embedded designers can place over any viewable surface for input and navigation. ClearPad, a patented capacitive sensor interface, offers designers a solution to overcome RTS limitations. Besides being ultra-thin, it can be integrated beneath a protective display lens or directly laminated onto

the display. The need for a heavy glass substrate is eliminated.

ClearPad's capacitive sensor module consists of a thin, transparent, finger-sensing region bonded to a flex circuit tail that contains all the sensing electronics. A finger on top of a grid of conductive traces changes the capacitance of the nearest traces. The device measures this change in trace capacitance and computes finger position. No pressure is needed to activate the capacitive sensor; a gentle stroke or glide along the surface of a capacitive pad is all that's required.

In this UI technology, capacitive sensing is combined with a transparent trace matrix. RTS and capacitive sensors use the same materials, indium tin oxide on polyethylene terephthalate. However, the capacitive sensor does not possess RTS's optical and durability issues because it is a single laminate with no air gaps to degrade optics and is solid state with no moving parts, thus contributing to its high reliability and durability. An RTS, on the other hand, is a physical switch that must flex and requires a user to press on the surface to activate it, decreasing its useful lifetime.

A touch-screen system like the ClearPad sensor is efficient because it allows the display area in a device to serve as an input and output system. Accordingly, a touch screen utilizes space more efficiently because with fewer mechanical buttons and controls taking up the space, the display can be made larger. Touch-screen controls can be easily repurposed in software to match a given application, an adaptive UI, making the device easier to use. This is especially helpful if the device is designed to have many different functions, like with a smartphone.

Because controls can appear and disappear as needed, designers can utilize the display area for many different purposes. In contrast, in devices that have mechanical buttons and controls, the real estate that contains the mechanical buttons cannot be reused even if the buttons are not used for a given application.

Chip designers have their own challenges. "No one magic answer or solution exists when it comes to creating a functionally dense component, platform, or system," states Rakesh Sethi, PhD, director of business development in the Custom SoC and Foundry Business Unit at Toshiba America Electronic Components, Inc. Functional density is affected by unit volume, which causes certain chip packages to be selected over others. High-volume parts can absorb more package development costs, making it possible for designers to consider more package options.

High channel count for analog and digital I/O drives other packaging choices. The right balance of I/O, processor, and memory subsystems and packaging options must be made as appropriate for the application. Sethi points out that using the latest semiconductor processing technologies often is not the best choice. Earlier processing technologies may offer the best return on investment. "Functional density comes indirectly," he says. "Inherently, it is very important as customers will often ask to build a whole system in a single package to save costs."

Future full of functional enhancements

The greatest improvements will continue to come at the chip level as semiconductor designers find more ways to add more capability – whether it be computational power, memory, or I/O – into smaller and smaller packages. Application-specific SoCs that can be quickly designed from extensive libraries of readily available IP blocks will evolve. These SoCs will make it relatively easy for a platform designer to quickly develop a very functionally dense device optimized to the application. Board-level products will feel the squeeze and continue to be pushed into deeper niches where application-specific SoCs may not be as practical. These will typically be low-volume applications or proof-of-concept products.

System-level solutions will continue driving larger packages using blade and motherboard technology. Cooling techniques will be stressed to dissipate the heat generated by these functionally dense systems. Networking to various locations will reduce some of the system concentration, helping minimize thermal impact.

New ways of connecting to the real world will be one of the biggest motivators for improving functional density. Our thirst for electronically assisted products will continue inspiring designers to look for better ways to stuff more and more capabilities into various packages. ECD

Annapolis Micro Systems

The FPGA Systems Performance Leader!

FOPEN **Radar Systems** **Software Defined Radio** **FLIR**
SIGINT **ELINT** **Digital Receivers** **Recording Systems**

Annapolis
Micro Systems

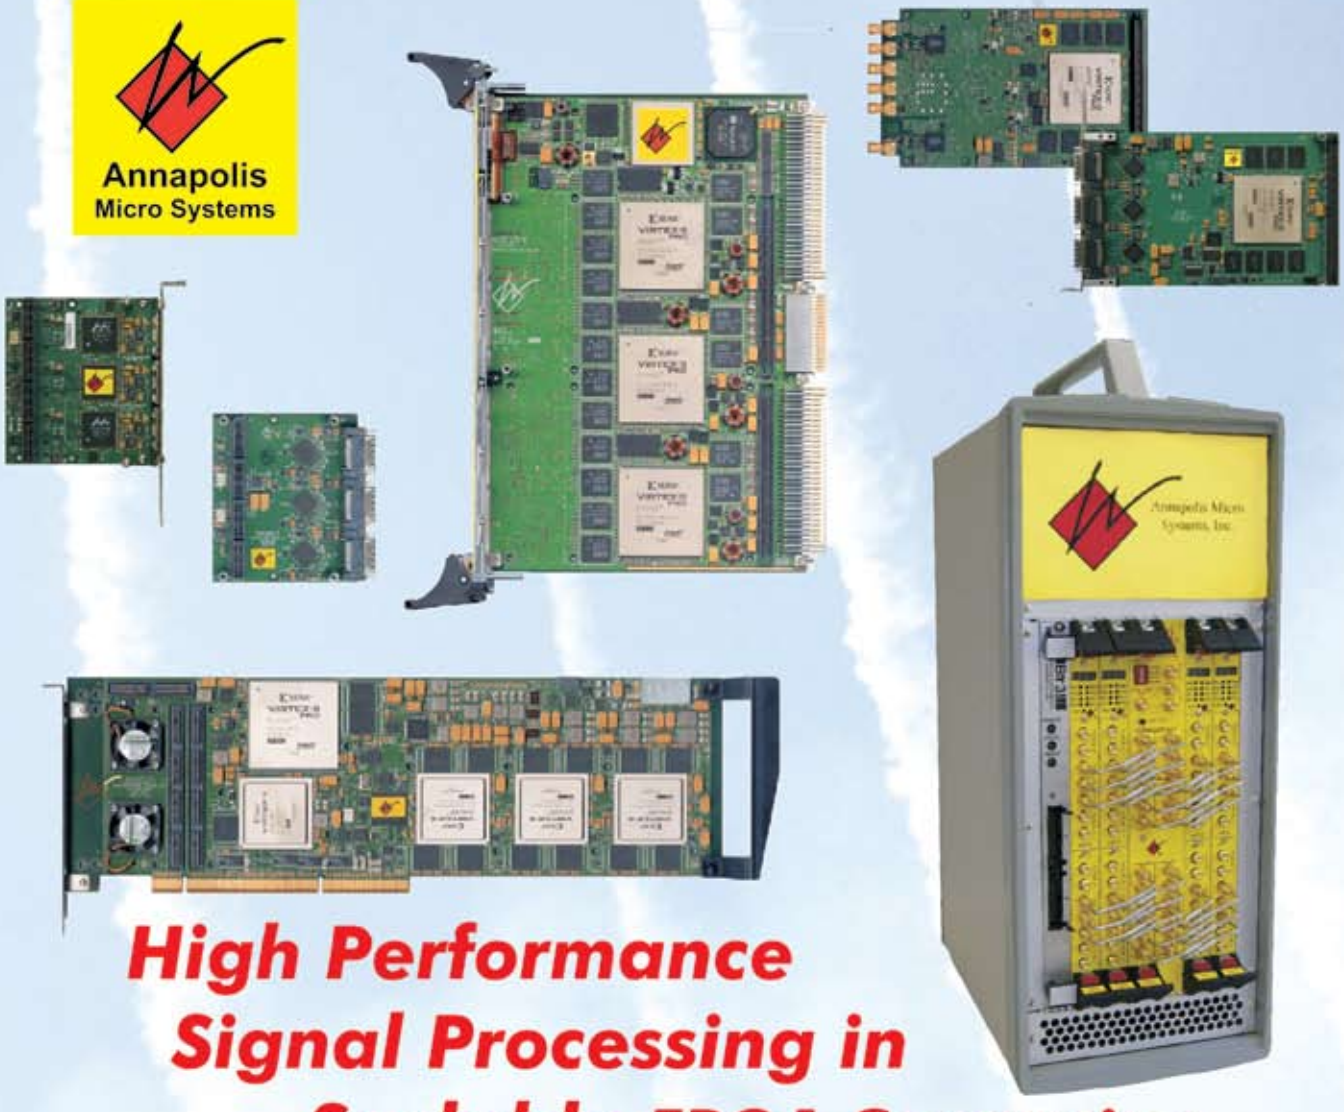

High Performance Signal Processing in Scalable FPGA Computing

Above and Beyond -----

FPGA Acceleration

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

Photo courtesy of USN

The future of embedded backplanes, system-level fabrics, and Serial RapidIO technology

By Tom Cox

Interconnects are traditionally evaluated based on technical features, compatibility, extensibility, and performance, with business advantages factored in as an afterthought. In the past two years, the embedded interconnect market has undergone a seismic shift and ultimately, two technologies – Ethernet and RapidIO – have emerged as the logical choices when examining technical leadership and performance, with PCI Express (PCIe) filling a niche spot. This market evolution has caused many OEMs to focus more intently on the business case for these three technologies. Analyzing newly published data from Crystal Cube Consulting and a Freescale Ethernet/RapidIO white paper, Tom promotes strategic discussions about taking technology decisions to the next level.

RapidIO technology is destined to become one of the dominant embedded interface technologies in a wide range of applications, most notably serving as the primary interconnect for high-speed embedded backplanes and system-level fabrics. According to a report from Crystal Cube Consulting[1], RapidIO technology will experience a five-year average Compounded Annual Growth Rate (CAGR) of 39 percent in the DSP market alone, with projected shipments of almost 78.5 million ports in 2011, as shown in Figure 1.

Several factors are responsible for the continued and rapidly progressing success of RapidIO technology in the market. These include the increased need for high-speed serial interfaces across markets and applications, the technical superiority of RapidIO links over Ethernet in embedded high-speed backplane applications and as a system-level fabric, the aggressive penetration of RapidIO interfaces integrated onto DSPs, the demise of the Advanced Switching Interconnect (ASI), and the maturity of the RapidIO ecosystem to carry the market in the long term.

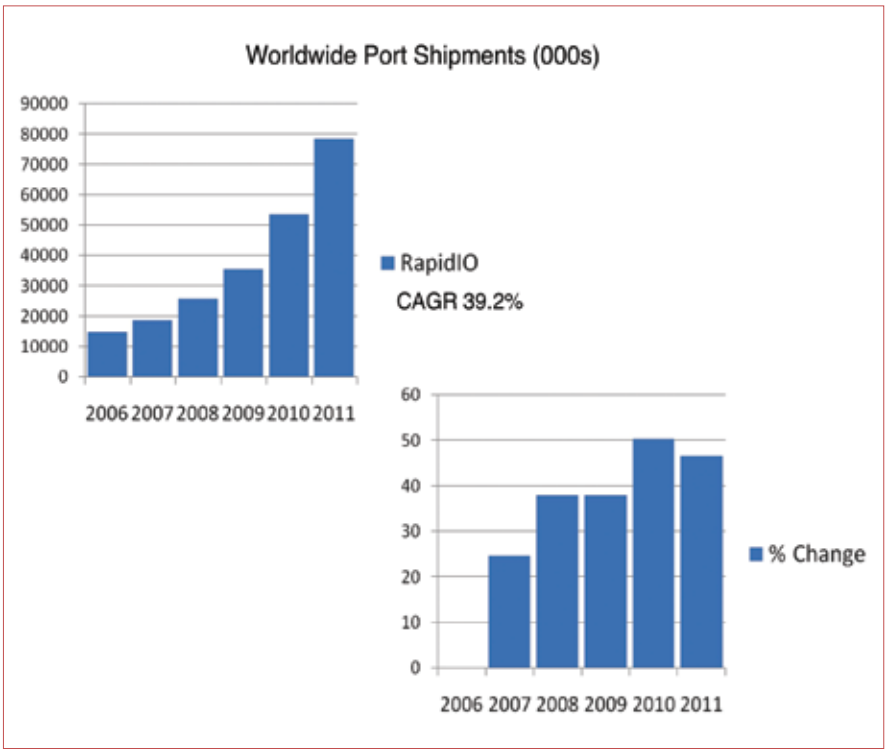

Figure 1

The need for high-speed serial interconnects

Increases in system performance necessitate innovations in interfaces that support higher bandwidth while enabling greater port density and reduced power consumption. Inherent limitations in parallel bus architectures cannot meet these needs and are the foundation of many system performance bottlenecks. For example, signal skew makes it impractical to further increase bus speeds while widening buses increases design complexity and cost because of the need for higher pin counts, more PCB layers, and a general encroachment on scarce board space.

For these reasons, the market is turning to high-speed serial interfaces such as RapidIO technology, 1 Gbps and 10 Gbps Ethernet, PCIe, and InfiniBand. Each of these technologies has particular strengths and applications in which it generally provides the best-performing and most cost-effective approach. PCIe, for example, owns the PC desktop market while InfiniBand has been repositioned as a storage area network and clustering technology.

Ethernet, while the leading interface technology in the LAN, must stretch beyond reasonable limits to extend its already broad reach into high-speed backplane and system-level fabric applications. Ethernet was designed for large-scale networks where each node could be expected to have its own powerful processing resources. As a consequence, Ethernet requires a large, flexible software stack that pushes many key functions such as guaranteed delivery and messaging to higher protocol levels. This increases overall latency, compromises reliability, and makes it inefficient (maximum utilization is about 25-35 percent in 10 Gbps links) to use in applications passing control plane traffic.

Performance can be improved with hardware-based offload engines, but these are effectively proprietary implementations, making them complex to design with and difficult to maintain while negatively affecting system interoperability. Additionally, Ethernet's often-touted economies of scale advantage doesn't apply to backplane and fabric applications that require specialized functionality. This significantly limits the number of Ethernet-based products and vendors supporting these applications.

The RapidIO specification, for its part, was specifically designed as a next-generation front-side bus that could also serve as an efficient system-level interconnect in ways that a software-based protocol like Ethernet cannot. Designed specifically for embedded in-the-box and chassis control plane applications, RapidIO provides minimal latency, limited software impact, and protocol extensibility while simplifying switch architectures and achieving data rates from 667 Mbps to 30 Gbps.

With efficient headers, hardware-based protocol processing, and key functions integrated into the base protocol – including guaranteed delivery, read/write operations, messaging, data streaming, quality of service, data plane extensions, and protocol encapsulation, all without the overhead of higher-layer protocols – RapidIO interconnects achieve 2.5 times more effective bandwidth per link than GbE at an equivalent cost. Furthermore, Serial RapidIO links can reliably consolidate both data and control planes onto a single, robust fabric, substantially simplifying system design while minimizing cost by reducing the number of interconnect technologies and ports a system must support.

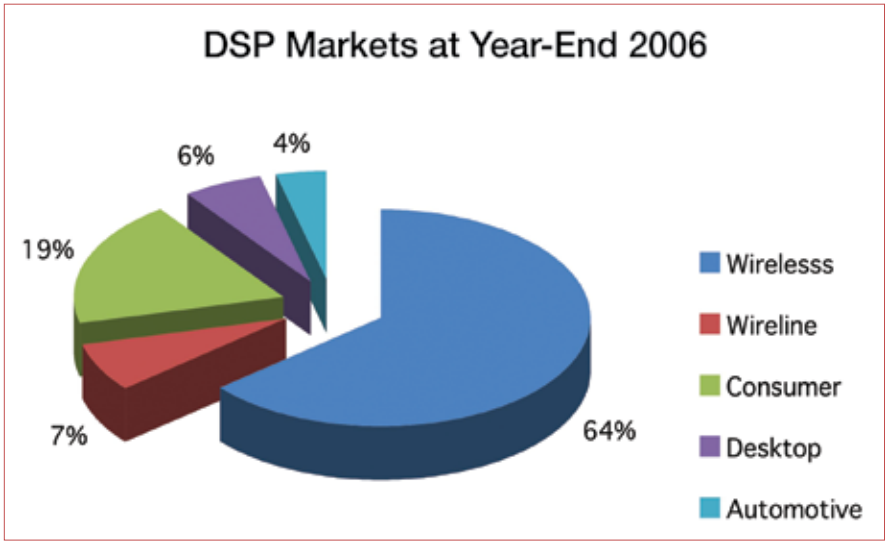

Figure 2

Aggressive market growth

One way to measure the market success of RapidIO technology in embedded backplane and system-level fabrics is to evaluate its penetration into the DSP market. DSPs are the core of many applications such as wireless communications, which represented 64 percent of the overall DSP market at the end of 2006 (see Figure 2), and therefore require a serial interconnect technology integrated directly onto the processor to achieve the greatest efficiencies when passing large amounts of data. Tracking the number of DSPs sold in particular applications makes it possible to estimate Serial RapidIO technology's overall adoption rate.

While PCIe had initial success in the DSP market, in the past two years, the DSP industry moved away from PCIe to Serial RapidIO interconnect technology. This shift arose from the recognition of PCIe's inappropriate fit for embedded backplane applications and its subsequent focus on the PC desktop market. PCIe supporters responded by developing the ASI, which was specifically designed to resolve PCIe's generally acknowledged shortcomings in fabric architectures and replace PCIe slots in the backplane market.

However, ASI did not experience the success its creators had hoped for. As Ernie Bergstrom, VP and chief analyst at Crystal Cube Consulting, states in his report, "Since the ASI architecture is essentially dead, the remaining viable architectures for high-speed backplanes are principally RapidIO technology and Ethernet."

These changes in market direction help account for the success of Serial RapidIO-based DSPs in capturing approximately 20 percent of the overall DSP market in 2006. As DSP vendors continue to introduce Serial RapidIO interfaces to more of their devices, market share is expected to increase from $1.5 billion of the overall $7.5 billion DSP market in 2006 to more than $4.71 billion of the overall $13.7 billion market in 2011, representing integration into more than 34 percent of DSPs worldwide (see Figure 3).

In addition, Serial RapidIO technology has widespread industry support dedicated to meeting the needs of developers in both the short and long terms. Serial RapidIO technology has an extensive ecosystem driven by such industry leaders as AMCC,

Freescale, and Texas Instruments, and for embedded backplane and system-level fabric applications, Serial RapidIO technology offers more product options and economies of scale than Ethernet. According to Bergstrom's report, "Technologies like the RapidIO standard have developed a sufficient ecosystem to warrant careful evaluation alongside Ethernet."

On course to greater adoption

RapidIO was designed from the start to be the most efficient interface for high-speed applications capable of meeting the throughput needs of data-intensive applications while providing the reliability required for control plane data transfers. With hardware-based protocol processing and a rich feature set defined within the base specification, Serial RapidIO stabilizes implementations, guaranteeing consistency and interoperability across the industry while leading to more choices for developers.

With strong industry support and a firm market base, RapidIO technology is well-established on its course to becoming the principal interconnect of embedded backplane and system-level fabric applications. ECD

Tom Cox *is executive director of the RapidIO Trade Association, where he is responsible for the association's future direction. Most recently, he was director of strategy at Tundra Semiconductor. Prior to that, he spent five years with IBM Microelectronics and was responsible for the PCI product line. Tom has held executive and engineering positions with ATI Technologies Inc., LSI Logic Corporation of Canada, Inc., Litton Systems Canada, and Philips Electronics Canada. He is a founding member of the PCI-SIG and has served in industry workgroups within PICMG, VITA, and NPF. Tom pursued advanced training in mathematics and electrical and computer engineering through the University of Southern California (USC), Humber College in Rexdale, and Mohawk College in Hamilton.*

RapidIO Trade Association
512-305-0070
tom.cox@rapidio.org
www.rapidio.org

References
[1] Ernie Bergstrom, Crystal Cube Consulting, "Switch Fabrics 2007: RapidIO Technology Takes the High Road to Embedded Systems," June 2007

(All figures courtesy of Crystal Cube Consulting.)

Worldwide DSP revenues vs Serial RapidIO	2006	2007	2008	2009	2010	2011	CAGR
Total DSPs ($B)	$7.5	$8.4	$9.5	$10.8	$12.5	$13.7	12.8%
% Change		12.0%	13.1%	13.7%	15.7%	9.6%	
Serial RapidIO DSPs ($B)	$1.50	$1.74	$2.20	$2.80	$3.76	$4.71	25.7%
% Change		16.0%	26.4%	27.3%	34.3%	25.3%	
% Serial RapidIO DSPs to total DSPs		20.71%	23.16%	25.93%	30.08%	34.38%	

Figure 3

DOM
DiskOnModule

Storage Leader

.Integration

.Reliability

.Performance

New!

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

pqi

Extending the reach of PCIe and other high-speed interconnects through signal conditioning

By Kenneth Curt

As data rates for multi-gigahertz protocols such as PCI Express (PCIe), Serial ATA (SATA), Serial Attached SCSI (SAS), and others continue to increase, ensuring signal integrity over long traces or cables becomes significantly more difficult. Kenneth addresses the application of active signal conditioning components to simplify system design and guarantee reliability and performance.

Modern communications standards like PCIe continue to evolve into more complex protocols with increasing data rates. Developing robust architectures around these high-speed standards introduces new challenges. For example, the PCIe Gen1 spec operating at 2.5 Gbps supports traces up to 16" over FR4 and one connector. Systems targeting 5 Gbps signaling with similar board design rules may drive just 8" of trace, which can be further decreased by approximately 1" for each via.

At 2.5 Gbps, signal integrity must be managed; at 5 Gbps, it becomes a significant concern and imperative that developers understand the limitations of standards so they can deal with signal integrity issues from the start. Many developers assume they can apply the same design techniques they did at 2.5 Gbps when moving up to 5 Gbps. Unless they approach high-performance designs realistically, they will discover the hard way that intelligent signal conditioning is essential in high-performance systems.

Signal integrity applies to a wide range of systems and standards including PCs, laptops, media centers, and large rack servers. Many standards such as SAS, SATA, 10 Gbps Attachment Unit Interface (XAUI), Fibre Channel, HyperTransport, High-Definition Multimedia Interface (HDMI), and InfiniBand are pushing the signal integrity envelope to its limits. By considering signal conditioning early in the design process, developers can reduce later engineering efforts to resolve troublesome signals, avoid expensive reworking, and meet time-to-market goals.

Extending beyond the specs

Many controllers offer some level of built-in signal conditioning to meet the minimum requirements of the applicable standard. Most of these controllers, however, are built to follow base protocol specs and support only standard signal reach limits. In addition, integrated signal conditioning tends to perform poorly compared to discrete implementations because of the difficulty mixing high-quality analog on the same die as high-speed digital circuits. For systems pushing the limits of performance, integrated signal conditioning may not be sufficient.

Moving to discrete signal conditioners with configuration flexibility enables

developers to go beyond standard specs. The resulting higher-signal quality gives developers more options, either freeing up engineering resources that would otherwise be spent managing signal integrity or allowing developers to drive the leading edge of design to a competitive advantage.

These devices have programmable characteristics enabling them to be adjusted to each signal's optimum operating conditions. Signal conditioners can be used as repeaters or redrivers to reliably extend signal reach, expanding architectural options. For instance, the 16" plus-one connector limit of 2.5 Gbps PCIe can be extended to more than 60" of FR4 and through four connectors when intelligent signal conditioning is applied. Cabled PCIe solutions of 7 m become viable with signal conditioning. For 5 Gbps Gen2 PCIe, the 8" reach limit can be extended beyond 30" with signal conditioning (see Figure 1).

In Figure 1, the top eye diagrams show a 5 Gbps PCIe Gen2 signal at its source and destination after a nearly ideal trace. With the minimal signal received, error rates will be high. The bottom eye diagram shows the signal restored to almost its original condition when a signal conditioner is placed in front of the receiver at the end of the trace.

Signal conditioning also can reduce architectural complexity for developers and their end customers. For example, blade servers must account for variable trace length over the backplane. When a configurable signal conditioner is present, blade servers can be configured automatically during the registration/learning process without slot-dependant manual tuning. This same process can be used to compensate for any signal conditioning being applied by the controller (such as preemphasis) and for aging components over time.

Active vs passive conditioning

Building a high-speed system that exhibits exceptional performance in the lab is quite a different proposition than building one that can do so in the varying and often unforgiving operating conditions of the real world.

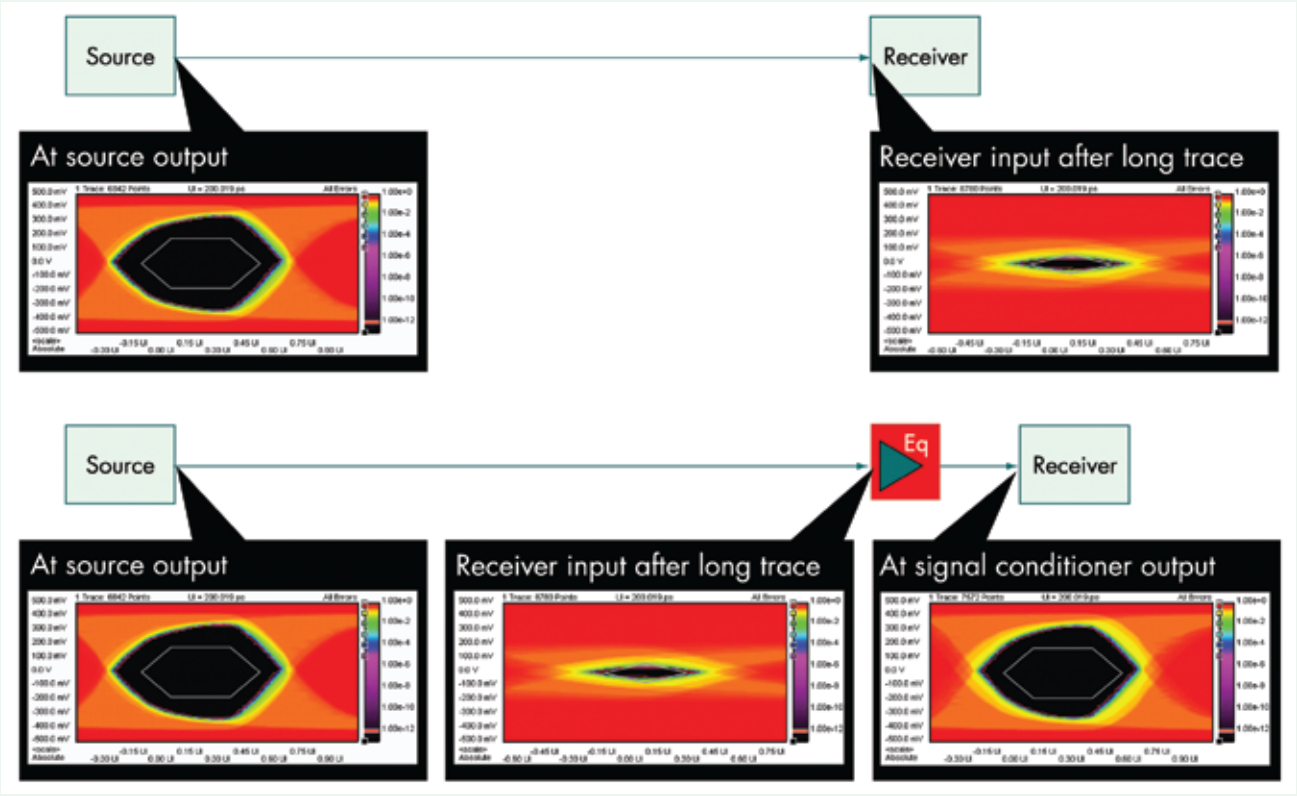

Figure 1

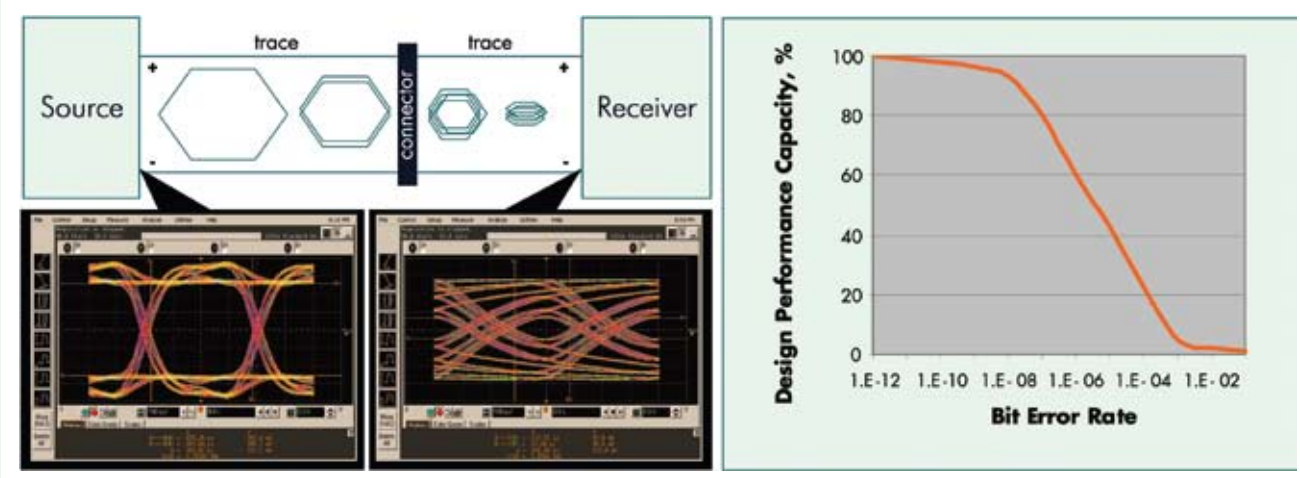

Figure 2

As signal quality directly affects Bit Error Rate (BER), it is fairly obvious when signal integrity has been compromised because the BER increases with a corresponding drop in system performance. A system will tolerate some errors, but performance typically drops fast beyond a threshold point (see Figure 2). What's not so obvious in the lab is how close to that edge a system can be and still appear to be reliable. This problem is further complicated as the number of high-speed channels increases within a system.

A system close to its limits may work fine in the lab but experience problems during manufacturing or once deployed in the field. Developers must design systems with a margin for error, which will accommodate variations in components and changes that occur as components age. While modeling can offer some insight into how much margin a system has, determining just how close a system is to the edge can be difficult to discern given the system complexity and the fact that measuring a signal changes it appreciably, skewing any margin calculations.

To maximize reliability, developers must ensure that a system is as far from the edge as is reasonable. Traditionally, this is achieved using passive conditioning techniques. For instance, adjusting trace routing or placement of other components can improve signal integrity. However, this approach often shifts the problem from one net to somewhere else on the PCB.

For marginal integrity issues, a simple LC compensation network can provide a suitable resolution, but developers must understand their design intimately to determine the components required for such a network. In the end, these networks are potentially unreliable because they introduce additional components that can vary, age, and be substituted incorrectly during manufacturing.

Other approaches such as using lower-loss PCB materials, higher-quality connectors, or replacing a difficult trace with a cable also improve signal quality but often come at an unacceptable expense. For example, it is not sound to raise the cost of an entire PCB with better materials to accommodate a few troublesome signals.

Signal conditioning allows developers to eliminate signal integrity issues early in the design cycle and avoid the inefficiencies of passing signal difficulties around the system or reworking them later. If a developer determines that a particular signal may be closer to performance limits than is comfortable, adding a discrete signal conditioner can resolve these issues and allow the developer to move on (see Figure 3).

These "bathtub curve" graphs statistically correlate signal jitter with error rate based on the eye diagrams in Figure 1. Without signal conditioning, the channel is marginal and may not reliably meet standard 10E-12 error requirements, as any shift caused by component variations during manufacturing, component aging, or changes in operating conditions could bring this channel down. Signal conditioning provides a wide margin of input tolerance, clearly stabilizing the channel to meet BER goals.

In general, it is more cost effective to directly address potentially problematic signals than disrupt the entire design through rerouting or a change of materials, especially considering the engineering cost, delayed time to market, and risk of introducing a new bug whenever a fix is made. High-performance systems tend to have higher price points, and delaying for a week of engineering to avoid adding a handful of signal conditioners and a few dollars' cost offers little advantage.

Protocol-aware signal conditioning

Signal conditioners add significant value by extending signal reach and integrity; however, to be the most robust, they must support the specifications of the protocol they are transporting to avoid introducing compatibility issues. Table 1 shows the relationship between signal conditioners and protocol compliance.

Protocol-based signals are much more than just varied electrical levels that need

Protocol compliance and signal conditioners

Electrical idle	Link error recovery
Lane discovery/link width	Hot plug and swap
Data rate negotiation and control	Channel timing and loss budget
Selectable/variable termination	Multilevel de-emphasis
Lane-to-lane de-skew within link	Round-trip delay
Beacon support	Logical idle
Power-down modes	Lane reversal
Loop back	Out-of-band signaling

Table 1

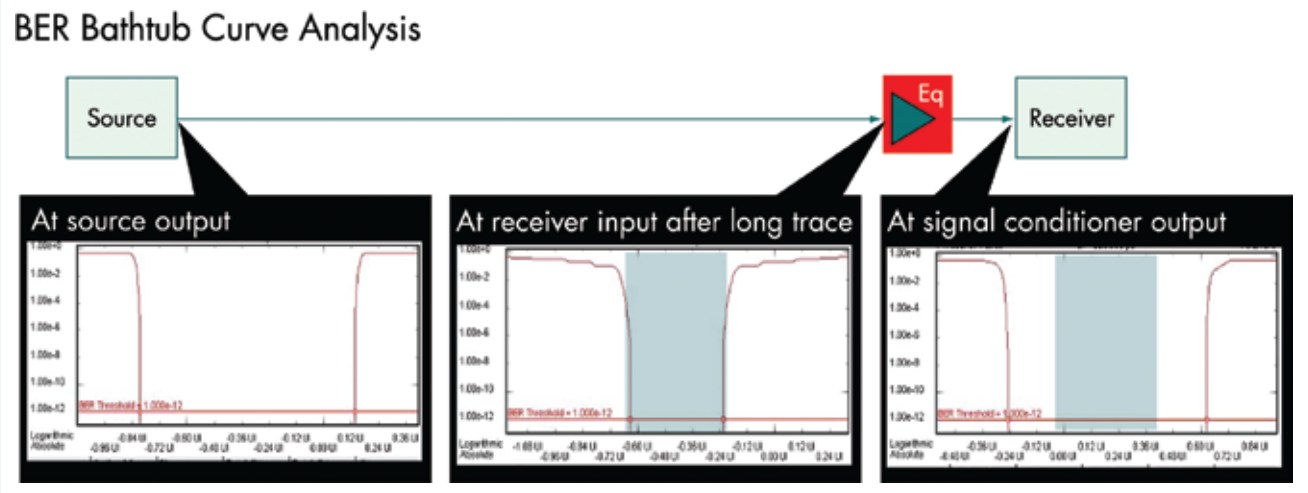

Figure 3

to be cleared of noise or amplified once they have been attenuated. At times, information in these signals is intended for the signal controller itself.

The PCIe specification, for example, does not consider the presence of signal conditioning between the host and endpoints. If a signal conditioner is not designed to a specific protocol, it may unintentionally interfere with a number of protocol mechanisms. The level of protocol compliance is an important characteristic of signal conditioners intended for use with modern communications standards. Proper signal conditioner selection can ensure compliance and minimize design issues for a particular protocol. For the most robust implementation and easiest design, a signal conditioner should be selected with protocol-specific support.

Within PCIe, several features including lane discovery can be disrupted if the signal conditioner is not protocol aware. The PCIe spec allows multiple methods for lane discovery and control. One method is to send an AC pulse from the host to an endpoint and watch the rise and fall time of the pulse to determine if a load is present on the line. If a non-protocol-aware signal conditioner is in the path, it may terminate the signal. Thus, even if there is no endpoint device attached to the signal conditioner, it will still appear to the host as if there is.

Another example of a protocol-specific issue is that of termination voltage. A generic signal conditioner may terminate to Vdd (the supply voltage), as with PCIe systems. However, in SAS systems, termination pulls to Vcm (common mode voltage).

Protocol-aware signal conditioners account for the many idiosyncrasies of standards. In the case of a PCIe-based signal conditioner, the proper termination voltage is supplied. Termination circuitry is also switched off when there is no load attached beyond the signal conditioner, thus supporting lane discovery. As a result, the signal conditioner is transparent to operation of the link while increasing signal integrity and extending signal reach.

Location, location, location

Signal conditioner placement affects much more than just signal reach and integrity. Depending on the application, it can introduce flexibility that guarantees signal integrity over varying generations of hardware, thus future-proofing designs.

Consider a consumer docking station application (see Figure 4) that may require

signal conditioning simply given the number of connectors over which PCIe or SATA signals must pass. The host system can be expected to change every year while the docking station may remain in production for three or more years. A configurable signal conditioner can allow easy optimization of critical signals to ensure reliable system interoperability as components change. Locating signal conditioning in the docking station minimizes overall expense while maintaining signal integrity and compatibility across varying generations of host systems.

The same approach can be applied to less cost-sensitive applications such as high-performance systems employing a chassis or backplane. As blades migrate to PCIe Gen2, without a way to extend signal reach, the chassis may need redesign to reduce overall signal length by bringing all components in closer to stay within the specified reach. Making an existing chassis form factor obsolete can become an expensive issue in retooling new mechanical hardware.

Rather than redesign the chassis, signal conditioners can be implemented on system boards or midplane cards, which are much less expensive to upgrade than chassis, to drive signals the required distances. Signal conditioners also can be implemented to resolve existing signal integrity issues on mezzanine or I/O controller cards.

Address at-risk signals head-on

As data rates continue to increase, intelligent signal conditioning has become essential to sustain reliability and performance in all types of high-speed applications today, from consumer electronics to high-performance communications systems. Even if systems currently do not need to exceed specified reaches or pass through extra connectors, discrete signal conditioners guarantee signal integrity to provide low error rates, high performance, and reliability.

By tackling signal integrity issues early in the design process and taking an active approach to signal conditioning, developers can avoid escalating local signal integrity issues into problems that can potentially compromise the entire system. When concerns about signal integrity arise, developers have the option to address at-risk signals directly through signal conditioners, reducing not only engineering investment and development time, but also overall cost and potential risk. **ECD**

Kenneth Curt *is product manager at San Jose, California-based Pericom Semiconductor, where he has worked since March 2006. Prior to joining Pericom, Kenneth maintained technical and management positions at Intel, Chips & Technology, Fairchild Semiconductor, and a number of other companies. He holds a BSEE from the Illinois Institute of Technology and an MBA from Santa Clara University.*

Pericom Semiconductor
408-435-0800 Ext. 294
kcurt@pericom.com
www.pericom.com

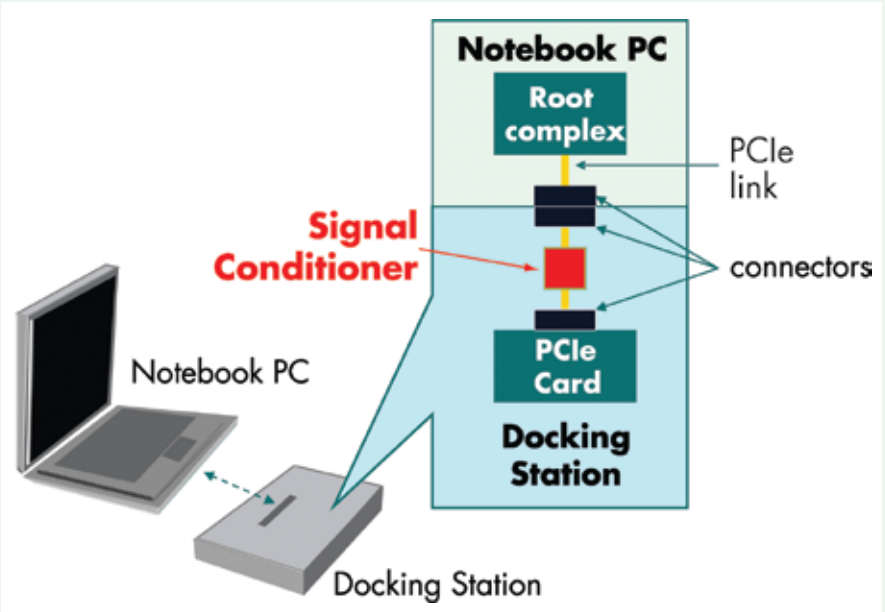

Figure 4

What Do These Embedded Applications Have In Common?

Each embedded application is unique. That's why VersaLogic offers an extensive line of embedded products with choices of size, performance, and features. That's also why **VersaLogic works closely with customers** to tailor standard products to meet their needs. **Product customization** is available in quantities as low as 100 pieces!

More importantly, embedded customers have many **common needs**. Meticulous quality standards, dedicated customer support, and reliable product delivery are just a few that VersaLogic provides. That's why customers from a wide range of industries **rely on VersaLogic** for their embedded computer needs.

Could you use **uncommonly-good support**? Call today to see if your unique application is a good match with VersaLogic's unique approach to service and support. We may have more in common than you think!

Call 800-824-3163. Visit www.VersaLogic.com.

30 YEARS 1976 - 2006 VERSALOGIC CORPORATION

VDC 2006 PLATINUM VENDOR

Rated a **"Platinum Vendor"** by Venture Development Corporation (2002-2006) for excellence in customer service.

VERSALOGIC CORPORATION

(800) 824-3163 • info@VersaLogic.com

IEEE-1588 Precision Time Protocol: Essential to next-generation test systems

By Paul Skoog

Future test environments need effective tools for enabling highly accurate synchronization over Ethernet LANs. Because IEEE-1588 can provide precise synchronization using packets, it's a logical option for synchronizing test equipment deployed on a packet network.

Next-generation test systems are converging on Ethernet as the primary interconnect for data acquisition and instrument control. The relatively recent IEEE-1588 Precision Time Protocol (PTP) standard designed to synchronize distributed and large channel count measurement systems over networks such as Ethernet is an essential part of this convergence. So what factors make PTP effective in next-generation test environments? And is PTP always the best choice in those environments, compared to sync alternatives like Inter-Range Instrumentation Group (IRIG) and Network Time Protocol (NTP)? Very often, the answer to the second question is "yes."

Master-slave synchronization

PTP is an IEEE protocol commonly deployed in Ethernet that uses packet timestamps to synchronize a slave (a PTP-enabled clock) with a master. This master (also a PTP-enabled clock) may be called a grandmaster if connected directly to a timing source such as a GPS receiver. Masters and slaves use a special packet exchange procedure to compensate for the time it takes packets to traverse the network. Devices determine packet delay by time-stamping special PTP packets' departure and arrival. They then exchange these observed times with each other in follow-up PTP packets. By comparing packets' observed departure and arrival times at each end, devices compute delays and adjust their internal clocks accordingly.

PTP also specifies special network switches called *boundary clocks* and *transparent clocks* that further mitigate packet delay effects. Boundary clocks are very precise clocks and switches that handle regular network traffic like any other switch. The clock is both a slave to a single upstream master and a master to potentially multiple downstream slaves, each of which can be a master to its own slaves. PTP packets synchronize each slave to its master; however, the PTP packets themselves do not pass through the boundary clock. Masters originate new timing packets. Without PTP packet queuing, there is no queuing delay in transferring synchronization across segment boundaries.

Transparent clocks mitigate queuing effects in a different way. They compensate for the internal queuing delays of PTP packets in the switch and adjust the appropriate PTP packets' timestamps directly. A PTP slave has no knowledge of the transparent clock and can precisely synchronize to its master with the PTP packets it receives.

PTP advantages

PTP offers key advantages over other synchronization technologies such as IRIG and NTP. IRIG, for example, employs a dedicated analog connection between the clock and each slave. PTP requires no such

point-to-point connection, potentially providing weight, installation/infrastructure, and cost benefits. PTP also can achieve submicrosecond accuracy, which IRIG cannot.

Like PTP, NTP also uses packet timestamps to sync network elements. However, time-stamping is done in software, not hardware, which causes asymmetric processing delays that reduce time transfer accuracy. Also, the lack of special switches like boundary clocks and transparent clocks to mitigate queuing effects further contributes to asymmetric delays and time transfer errors. PTP-enabled devices can autodiscover other PTP-enabled devices to form optimized timing networks, a feature NTP lacks. NTP enables NTP timeserver autodiscovery, but does not ensure that these sources or the paths to them are optimal. Table 1 compares the three technologies based on key performance metrics and features.

IRIG advantages

Even with all these benefits, however, there are still occasions when IRIG or NTP would be more suitable than PTP. For example, IRIG deployments can be simpler and less time-consuming, especially if 10 µs accuracy is sufficient and there are only a few elements to synchronize. If slaves are far apart, then it may be simpler to connect an IRIG source at each location and be done with it. PTP networks must be tested to make sure performance is within desired limits, and networks may need to be reconfigured if performance is not what is desired. It helps if PTP devices include features such as output 1 Pulse Per Second (PPS) signals to simplify measuring both packet delay and master-slave sync offsets.

NTP is generally unsuited for next-generation test instrumentation applications. Millisecond performance may be inadequate and NTP's performance can vary widely, which are primary reasons why PTP was invented for instrumentation synchronization.

Topology is key

In general, flat network topologies yield better synchronization performance than deep hierarchical networks. Using fewer cascaded network components reduces Packet Delay Variation (PDV). PDV is a critical metric of network performance because if delay did not vary, a constant offset could compensate for it. The greater the delay and the greater the delay variation, the more difficult it is to maintain synchronization between master and slave.

To maintain synchronization performance, IEEE-1588-optimized switches (that is, boundary clocks and transparent clocks) should be used where data and timing packets pass through a single egress port. Ideally, timing packets should be kept isolated from data packets until they converge at the slave.

Accurate PDV measurements can be obtained using the hardware time-stamping capability of PTP test devices. It is essential to measure the reception and transmission of packets at the slave using a reference source tightly coupled to the master. This can be done using the same external reference signal as the master (GPS, for example) or by integrating a PDV measurement capability in the master clock, such as in the integrated measurement setup shown in Figure 1.

The master-slave sync offset can be determined by comparing their respective hardware-generated PPS signals. Errors can be viewed using a frequency counter, oscilloscope, or a grandmaster equipped with an integrated time interval measurement input. Plotting a slave PPS as a histogram (Figure 2) shows the statistical

	IEEE-1588	IRIG	NTP
Peak error network type	100 ns to 100 µs	10 µs Dedicated coaxial	1-100 ms Ethernet
Typical extent	A few subnets	1 mile over coaxial	LAN/WAN
Style	Master/slave	Master/slave	Peer ensemble client-server
Protocols	UPD/IP multicast	–	UDP/IP unicast (mainly)
Latency correction	Yes	PTP networks	Yes
Network administration	Self-organizing	Configured	Configured
Hardware at time client	Required for highest accuracy	Required	No
Default update interval	~ 2 seconds	1 PPS	Varies, seconds to hours

Table 1

nature of the slave synchronization to a master, best described by its mean and standard deviation.

Figure 3 shows the effects of cascading COTS switches in a PTP timing network. This plot is a set of PPS histograms for the same slave device being synchronized separately through three individual switches and then through the cascade of the three. The cascade PPS performance closely follows the RSS (square root of the sum of the squares) of the individual switches as measured separately. This similarity shows how it is possible to characterize end-to-end network performance by characterizing its individual components, which would help in network planning or when measuring the actual cascade across its endpoints is difficult because of network size. The data also shows

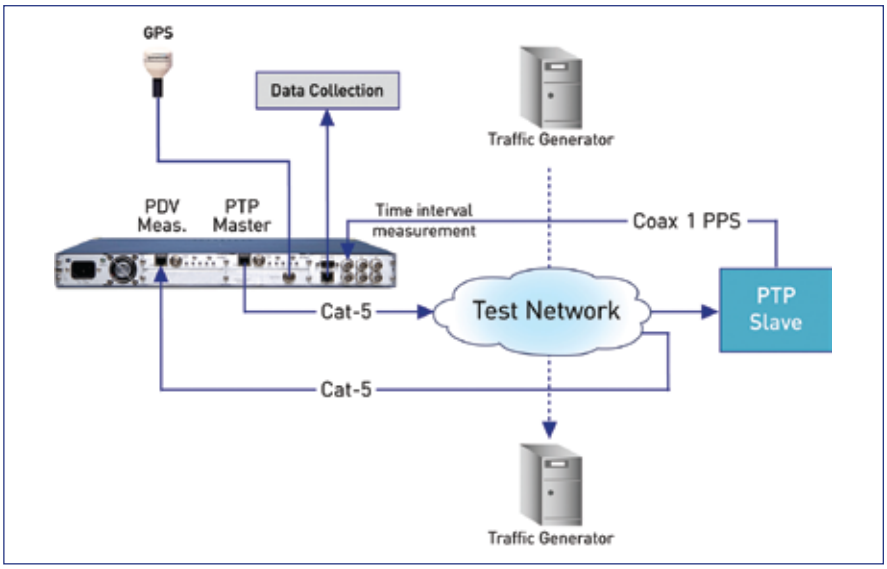

Figure 1

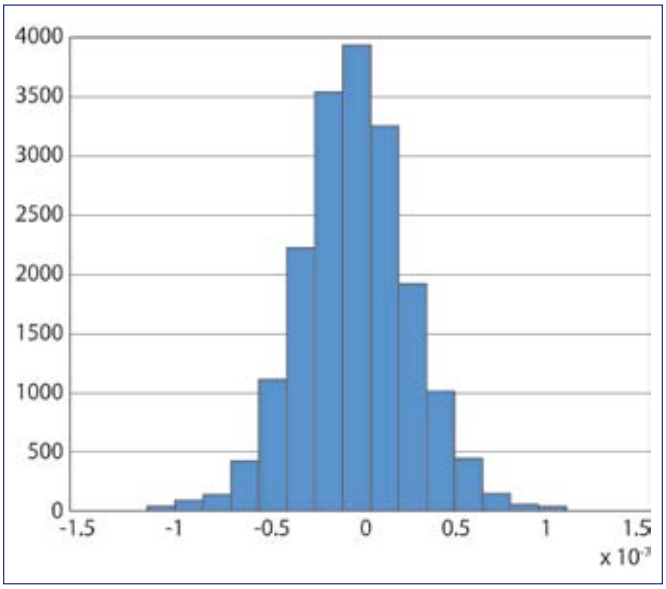

Figure 2

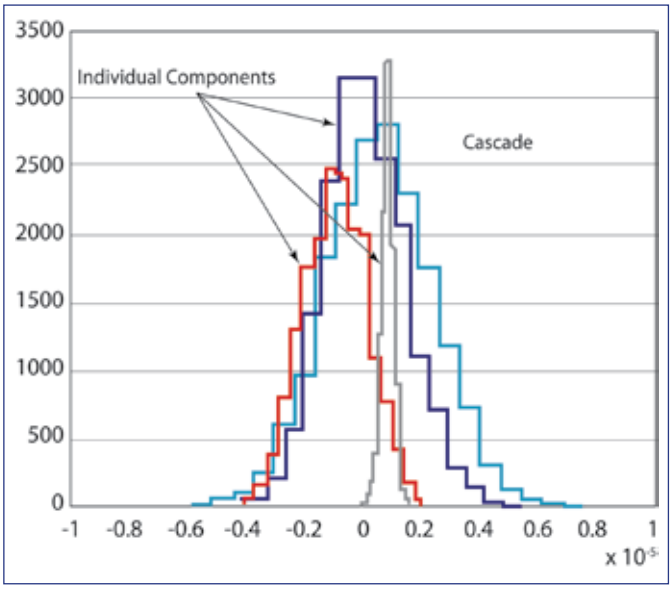

Figure 3

how a flat network benefits from tighter synchronization: Compare the tight PPS distribution for a single switch with the much wider composite distribution. Also note the varied performance of different switches.

Figure 4 shows how IEEE-1588 packets' PDV changes over the course of a workday in a production network in response to data traffic loading as measured at a single slave. In this example, packet delay varied within a range of 100-700 µs. PDV is low outside of business hours, and a "floor" of about 100 µs delay remains constant regardless of the time of day. In other words, 100 µs is attributable to nontraffic factors such as switch processing delays and distance.

Figure 5 shows the IEEE-1588 slave accuracy for the same network as Figure 4

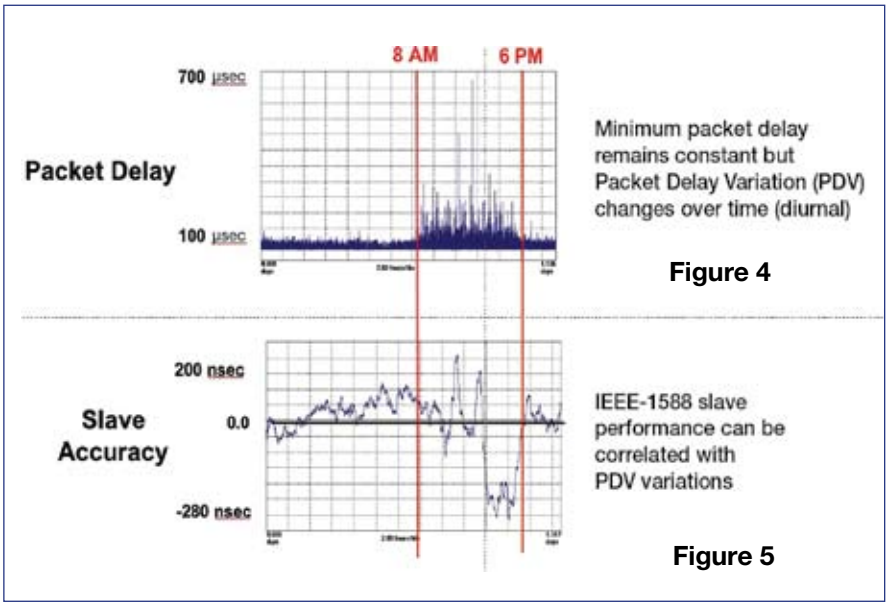

Figure 4

Figure 5

Test engineers cannot simply take precision timing for granted with PTP as they often can with IRIG.

over the same period. Clearly, PDV impacts synchronization accuracy. As PDV goes up, precision degrades, going from a range of less than 100 ns error during periods when there is almost no PDV to more than 200 ns during periods of peak PDV. If this error was not acceptable, deploying a transparent clock or boundary clock at key network choke points may improve slave synchronization accuracy.

A worthy candidate

PTP's many advantages make IEEE-1588 a worthy candidate for syncing next-generation test systems. These systems are typically deployed over Ethernet networks, which, by definition, mean much of the timing infrastructure is already in place. However, test engineers cannot simply take precision timing for granted with PTP as they often can with IRIG. While high accuracy is possible (exceeding IRIG levels, in fact), achieving and sustaining that accuracy may require some effort in terms of network configuration and testing. What will simplify that task is the steadily increasing stream of robust PTP-enabled devices now entering the market. **ECD**

Paul Skoog *manages the bus-level timing and enterprise network timeserver product lines at Symmetricom in Santa Rosa, California. Prior to joining Symmetricom in 1997, he was product manager of precision GPS positioning instruments for Trimble. Previously, he held application engineering and product management positions in the dynamic signal analysis software market. Paul holds a BSME from California Polytechnic University and an MBA from the Santa Clara University Graduate School of Business.*

Symmetricom
707- 528- 1230
pskoog@symmetricom.com
www.symmetricom.com

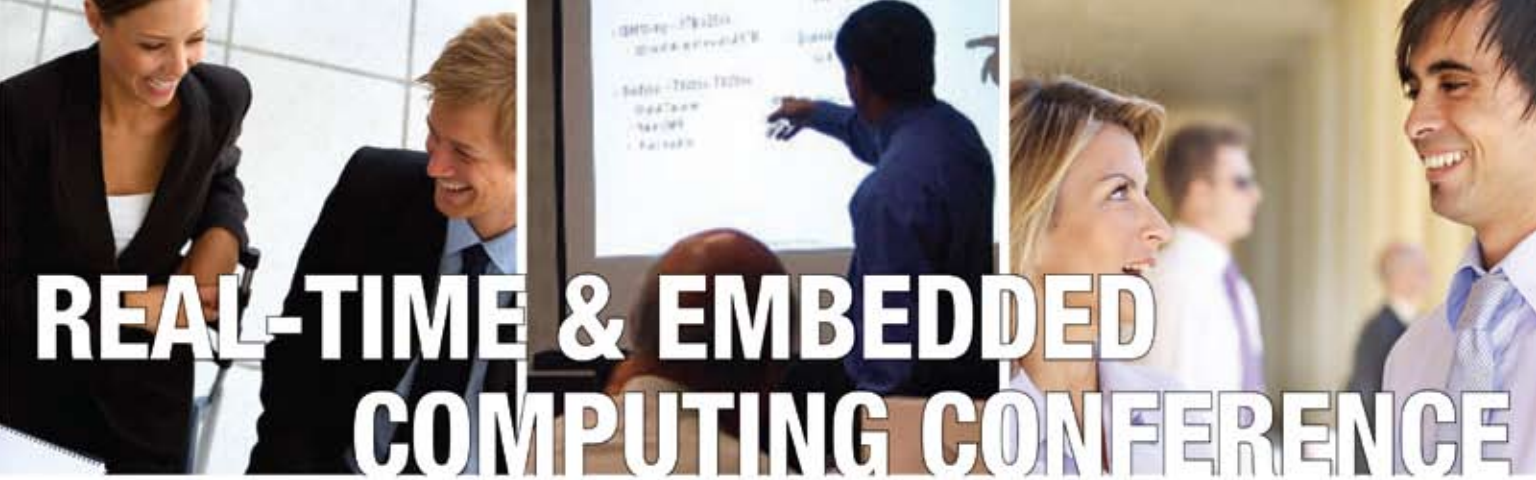

REAL-TIME & EMBEDDED COMPUTING CONFERENCE

Real-Time & Embedded computing conference

NEW TRAINING WORKSHOPS & SEMINARS

Register Today at www.rtecc.com

Enter a World of Embedded Computing Solutions

Attend open-door technical seminars and workshops especially designed for those developing computer systems and time-critical applications. Get ahead with sessions on Multi-Core, Embedded Linux, VME, PCI Express, ATCA, FPGA, Java, RTOS, SwitchFabric Interconnects, Windows, Wireless Connectivity, and much more.

Meet the Experts

Exhibits arranged in a unique setting to talk face-to-face with technical experts. Table-top exhibits make it easy to compare technologies, ask probing questions and discover insights that will make a big difference in your embedded computing world. Join us for this complimentary event!

Be sure to enter the drawing on-site for an iPod Video

AAEON · ADVANTECH · ELMA · Green Hills · intel · LYNUXWORKS · NATIONAL INSTRUMENTS · PENTEK · QNX · WIND RIVER

Join Us In...

January
Copenhagen, Denmark
Santa Clara, CA

February
Melbourne, FL

March
Huntsville, AL
Atlanta, GA
Phoenix, AZ
Albuquerque, NM

April
Boston, MA
Bethesda, MD

May
Milan, Rome Italy
Dallas, Austin, Houston, TX

June
Chicago, IL
Minneapolis, MN

July

August
Longmont, CO
Colorado Springs, CO

September
Beijing, Shanghai, China
Ottawa, ON
Montreal, PQ
Shenzhen, Xi'an, China
San Diego, Long Beach, CA

October
Tyson's Corner, VA

November
Detroit, MI
Toronto, ON
Vancouver, BC

December
Seattle, WA
Portland, OR

Waveform scanning techniques simplify embedded system designs

Recently developed measurement techniques are helping embedded system designers troubleshoot designs accurately and efficiently. These powerful techniques for characterizing and debugging designs contain microcontrollers, DSPs, FPGAs, A/D and D/A converters, and other devices that include address and data lines. Mike outlines several real-world engineering challenges and provides practical solutions illustrated by images captured live from actual embedded circuits using waveform scanning analysis. Designers can begin implementing these solutions immediately with real-time oscilloscope tools using the described techniques.

By Mike Hertz

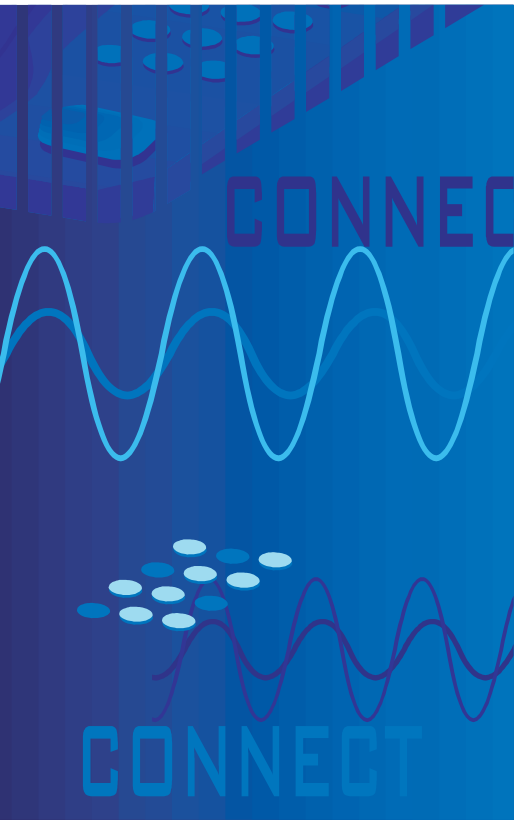

Logic failures and nonmonotonic edge detection

When digital logic errors occur in a circuit, troubleshooting the source of the logic error can be challenging. A logic analyzer cannot perform physical layer measurements such as measuring the rising edge of a single pulse. Using a traditional digital oscilloscope, digital and analog waveforms cannot be viewed simultaneously.

A Mixed-Signal Oscilloscope (MSO) is a special category of digital oscilloscope that allows for both digital and analog inputs. The newest MSOs can input 36 logic timing channels together with four standard oscilloscope channels simultaneously. However, even when using an MSO, risetime errors of a failed digital line can only be viewed if they are located.

Manually monitoring a waveform is not practical for a long period of time. In addition, only reduced waveform record lengths can be monitored manually, and analysis is limited by human error. Similarly, trigger circuits only highlight error conditions at the trigger point, and triggers cannot be used to define many types of error conditions, such as an anomalous measurement that results when cascading a chain of measurements and mathematical operators.

To solve these problems, embedded systems designers now can select a digital oscilloscope that automatically searches for sources of logic errors.

Anomaly identification using waveform scanning

A digital oscilloscope's ability to automatically search a waveform for anomalies, report the findings, and take appropriate actions based on user-defined criteria is known as *waveform scanning*. Waveform scanning is a new technology that has been deployed across a wide range of real-time digital oscilloscopes during the past year. Reliably locating the source of logic timing errors such as risetime anomalies requires an MSO equipped with waveform scanning capability.

As shown in Figure 1, the eight logic timing channels of one digital trace group are displayed in the purple Digital1 waveform. A logical false one is occurring simultaneously on all eight digital lines at the trigger point. The reason for the bit error is not readily apparent when viewing the digital trace group. With an analog scope channel probing a digital line and waveform scanning active, the scope catches each occurrence of this logic timing problem and identifies an analog runt pulse as the source of error.

Note that the waveform scanning tool has highlighted the source of error with a red rectangle on two consecutive runts with respective amplitudes of 1.3 V and 720 mV. The runt amplitude is listed on the top left corner of the display with an index number corresponding to each anomaly.

In addition, a scan overlay trace is displayed in blue. The scan overlay comprises a persistence mapping of anomalous waveforms. This is not a traditional persistence map, which typically includes all acquired waveforms. The scan overlay is a special type of persistence display that includes only those waveforms identified as anomalies.

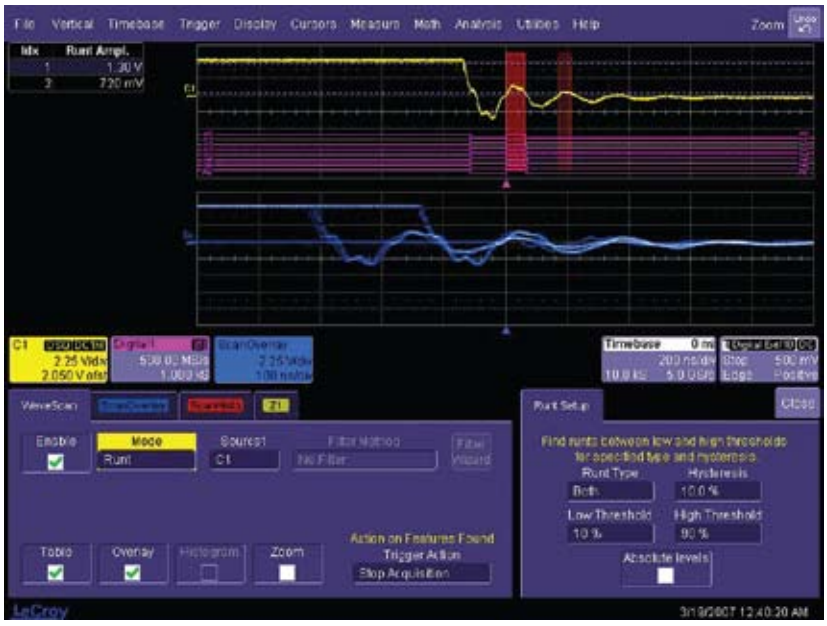

Figure 1

Monitoring physical layer characteristics automatically

Within embedded systems, other physical layer characteristics can affect digital data transmission. For example, if the data frequency (the bit rate or the inverse of the bit period) drifts over time, the clocked data read by the receiver could be misaligned and interpreted with bit errors. Waveform scanning techniques can be applied to continuously monitor a device's characteristics such as data frequency to ensure that all values are maintained within the specified range.

In Figure 2, a differential CAN signal is transmitted at a nominal data rate of 125 KHz or 125 Kbps. This serial data transmission is decoded by the digital oscilloscope to display messages transmitted from the engine, coolant system, and drive shaft nodes. In addition to this protocol layer decode, physical layer attributes must be inspected.

Waveform scanning is used to monitor the data frequency with a filter limit set to 125.01 KHz ±5 Hz. When the data frequency exceeds this tolerance band, the scope is configured to take action by stopping the waveform acquisition. In this case, five unit intervals within the acquired data block exceed the specified limits and are highlighted in red in the upper grid. The lower grid shows an automatic zoom within the coolant cyclical redundancy check that contains a data frequency of 125.012 KHz, a frequency outside of user-defined limits.

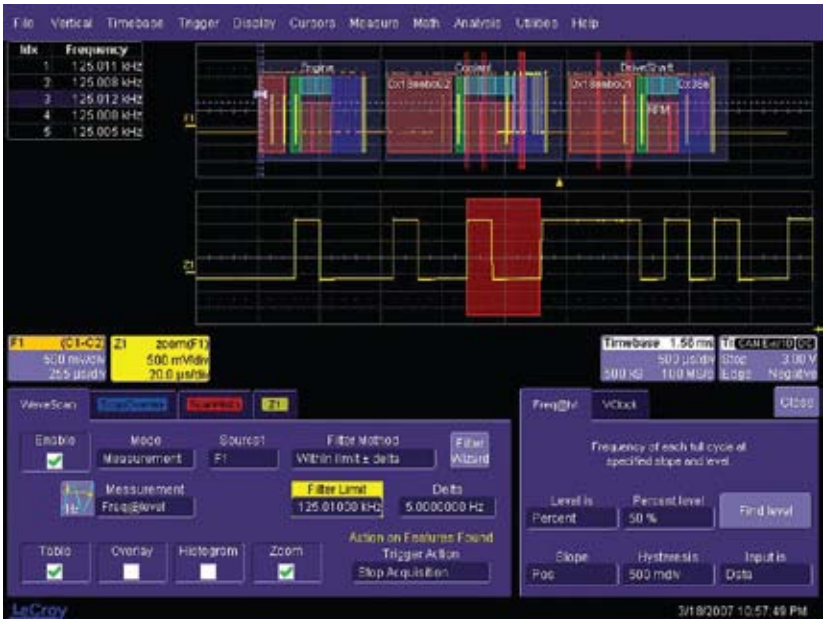

Figure 2

Fast, precise troubleshooting

Using an MSO with waveform scanning capability enables quick and accurate debugging. Automatic scans check for occurrences such as duty-cycle distortion, drifts in frequency, overly narrow widths, and intermittent events. Waveform scanning promptly finds and identifies anomalies when operating on a deep acquisition record with numerous events and constantly scans and executes actions based on user-defined criteria when operating on multiple acquisitions.

Embedded system designers facing the challenge of a timing measurement can use waveform scanning techniques to discover anomalous conditions right away and measurement filtering methods to help limit search criteria. To further analyze anomalies that the scan identified, designers can view histogrammed and overlaid measurement and data values.

These techniques also can be adapted for monitoring and troubleshooting virtually any embedded system to monitor system behavior and detect anomalies as they occur. ECD

Mike Hertz *is a field applications engineer with LeCroy Corporation in Chestnut Ridge, New York. Before joining LeCroy more than six years ago, he worked as an applications engineer with Agilent Technologies and Hewlett-Packard. He has three U.S. patents pending in the field of oscilloscope measurement design. Mike has a BSEE from Iowa State University and an MSEE from the University of Arizona.*

LeCroy Corporation
586-677-7564
mike.hertz@lecroy.com
www.lecroy.com

SATISFIABILITY

Satisfiability: A new generation of static analysis

By Ben Chelf

Source code analysis has had a mostly deserved bad reputation in software development because analyses often took too much time or produced excessive noise and a large percentage of bogus results (false positives). With low signal-to-noise ratios, most source code analysis technologies and products quickly became shelfware after a few uses. The promise of new static analysis solutions is tantalizing for developers because it offers the ability to find bugs before software is run, improving code quality and dramatically accelerating the availability of new applications. Though static analysis has historically struggled to deliver on this promise, a groundbreaking technique applied in the static analysis field may help fulfill its potential.

The software development challenges of today require companies to look for innovative ways to remove the critical flaws and vulnerabilities in software before it gets to the quality assurance lab or, worse, out into the field (see sidebar). Fortunately for developers, source code analysis is now up to the challenge. By combining breakthrough techniques in the application of Boolean satisfiability to software with the latest advances in dataflow analysis, the most sophisticated source code analysis can boast out-of-the-box false positive rates as low as 10 percent and scalability to tens of millions of lines of C/C++ or Java code.

The price of failure in software

A 2002 National Institute of Standards and Technology (NIST) study titled "The Economic Impacts of Inadequate Infrastructure for Software Testing" estimated that software errors cost the U.S. economy an estimated $59.5 billion annually. The report states that leveraging test infrastructures that allow developers to identify and fix defects earlier and more effectively could eliminate more than one-third of these costs.

For more information on this study, visit www.nist.gov/public_affairs/releases/n02-10.htm.

Software developers can use static analysis to automatically uncover errors that are typically missed by unit testing, system testing, quality assurance, and manual code reviews. By quickly finding and fixing these hard-to-find defects at the earliest stage in the software development life cycle, organizations are saving millions of dollars in associated costs.

Software DNA Map: The foundation of superior code analysis

If developers want superior analysis, they must have a perfect picture of the software because their analysis will only be as accurate as the data it is based upon.

Imagine someone gave a software development team a system of a few million lines of code – a system they'd been working on for a long time – and asked them to draw a map of the software system. To be useful, this map would need to address all the following details: how every single file was compiled for all targets, how each set of files was linked together, the different binaries that were generated, the functions within each file and the corresponding callgraph, all the different control flow graphs through each function, and on and on. If the developers attempted to do this by hand, it would be practically impossible.

Today, automating this onerous task is possible, and the process of creating such a picture, known as a *Software DNA Map*, opens the door for static analysis to significantly improve code quality and security. A Software DNA Map is an extremely accurate representation of a software system based on understanding all operations that the build system performs as well as an authentic compilation of every source file in that build system. By providing an accurate representation of an application to the function, statement, and even expression level, the Software DNA Map enables static code analysis to overcome its previous limitations of excessive false positives and deliver accurate results that developers can put to immediate use.

Static code analysis primed for advancement

Combining breakthroughs in dataflow analysis with a Software DNA Map has created substantial benefits for development organizations by enabling them to detect defects early in development. However, the performance and accuracy of these analysis solutions can still be improved. A groundbreaking new technology called *Boolean satisfiability* is poised to greatly expand static code analysis capabilities.

Boolean satisfiability

In complexity theory, the Boolean satisfiability problem (SAT) is a decision problem whose instance is a Boolean expression written using only AND, OR, NOT, variables, and parentheses. The question is: Given the expression, is there some assignment of TRUE and FALSE values to the variables that will make the entire expression true? A formula of propositional logic is said to be satisfiable if logical values can be assigned to its variables in a way that makes the formula true.

The concept of Boolean satisfiability is not a new one, but recently, efficient SAT solvers have been developed to solve very complicated formulas with millions of variables. While used heavily in the electronic design automation industry to aid in chip design verification technologies, the application of SAT solving to software analysis has remained untouched.

SAT solver

A SAT solver is a computer program that takes in a formula of variables under the operations of AND, OR, and NOT and determines if there is a mapping of each individual variable to true and false such that the entire formula evaluates to TRUE (satisfiable). If no such assignment exists, the SAT solver indicates that the formula

is unsatisfiable and provides a proof demonstrating that it is unsatisfiable.

The application of SAT to software requires that source code be represented in a form that can be plugged automatically into a SAT solver. Fortunately, with a Software DNA Map, all the necessary information from the code is available to transform it into any representation desired. Because SAT solvers deal in TRUE, FALSE, AND, OR, and NOT, the relevant portions of this program can be transformed into these constructs. Take an 8-bit variable as an example:

char a;

To represent 'a' as TRUES and FALSES, those 8 bits (1s and 0s) can be each thought of as TRUES and FALSES, so 'a' becomes an array of 8 Boolean values:

a_{-0}, a_{-1}, a_{-2}, a_{-3}, a_{-4}, a_{-5}, a_{-6}, a_{-7}

The operations that make up expressions in the code also must be translated. All expressions in code can be converted to an equivalent formula of AND, OR, and NOT. The thinking behind this is that a compiler must turn these operations into instructions in machine code and that machine code must run on a chip. The chip's circuitry is nothing more than pushing 1s and 0s (high voltage and low voltage) through a number of gates (all of which can be simplified to AND, OR, and NOT); therefore, this indicates that such a mapping exists. For example, to convert the expression:

a == 19

into a formula, the following expression would do the trick:

$!a_0 \wedge !a_1 \wedge !a_2 \wedge a_3 \wedge !a_4 \wedge !a_5 \wedge a_6 \wedge a_7$

In this example, a_0 is the high bit of 'a' and a_7 is the low bit. Plugging this into a SAT solver would render the following assignment of variables for the formula to be satisfied:

$\mathbf{a}_0$ = False (0)
$\mathbf{a}_1$ = False (0)
$\mathbf{a}_2$ = False (0)
$\mathbf{a}_3$ = True (1)
$\mathbf{a}_4$ = False (0)
$\mathbf{a}_5$ = False (0)
$\mathbf{a}_6$ = True (1)
$\mathbf{a}_7$ = True (1)

Taking that, as binary 00010011, shows that it is equivalent to 19.

Once the entire software system is represented in this format of TRUES, FALSES, NOTS, ANDS, and ORS, a wide variety of formulas can be constructed from this representation and SAT solvers can be applied to analyze the code for additional, more sophisticated quality and security problems. It is this *bit-accurate representation* of the software that enables more precise static analysis than previously was possible based solely on dataflow techniques.

Writing better code

One early application of SAT solvers for static analysis is false path pruning. When performing dataflow analysis, sometimes a defect will be reported on a path that is infeasible (or unsatisfiable) at runtime. This reported defect should be eliminated from the static analysis results since there is no possible combination of variables where it could actually happen in the software system. The methodologies for dealing with this problem within the dataflow framework pale in comparison to the abilities a bit-accurate representation and a good SAT solver can provide in this regard.

For example, with this bit-accurate representation of the source code, developers can construct a formula that is a conjunction (AND) of all the conditions in a path that lead to any given defect discovered by dataflow analysis. By solving this formula, the SAT solver

indicates if there is a set of values for the variables involved in all the conditions such that the path can actually be executed. If the SAT solver says "satisfiable," then the path is feasible at runtime. If the SAT solver says "unsatisfiable," the path can never be executed and no bug should be reported.

Consider the earlier discussion of static analysis history and the problems with excessive false positive results. By identifying false paths with SAT, developers can prune them from their static analysis results. This enables them to focus testing and analysis efforts on potential problems that have a real possibility of compromising the project at hand.

It's important to note that false path pruning is just one example of how SAT can be leveraged to provide more accurate source code analysis. Other potential applications of SAT include the ability to find problems such as string overflows, integer overflows or deadcode, and the use of assertion-based checking to identify difficult-to-find logic bugs (see sidebar). While some instances of the defects in these categories can be discovered today, SAT-based analysis allows developers to build on the success of existing dataflow analysis to reach new levels of accuracy and precision in static code analysis.

www.embedded-computing.com/search

SAT will benefit software developers and architects by fundamentally changing the way they view static analysis of their source code. Just as going from simple "grep" parsing through code to practical dataflow analysis earlier this decade was a huge eye-opener for many developers ("You mean, it can show me the whole path to get to a real bug?"), leveraging SAT solvers for static analysis can impress anyone who writes software ("How could you possibly know that about my code?"). Because of this, SAT may be the breakthrough that enables static analysis to deliver on more of its long-awaited promise.

Leveraging SAT techniques

For years, embedded software developers have looked for products that could automatically and effectively find software defects early in the development cycle. Fixing software bugs early can dramatically reduce the time it takes to bring a software product to market and potentially save millions of dollars in costly product recalls. However, tools for discovering defects automatically have had many false starts because of their failure to grasp a complete picture of the software. Older technologies incorrectly identified harmless code inconsistencies as defects and wasted developers' time.

By providing complete understanding of a given build environment and source code, a Software DNA Map can allow organizations to leverage existing SAT techniques to produce the highest-quality software possible. Leveraging SAT within static code analysis, companies can immediately find and fix software defects in source code at the beginning of the development cycle, thereby ensuring product quality and accelerating the availability of new applications to bring to market. ECD

***Ben Chelf** is CTO of San Francisco-based Coverity, Inc. Before he cofounded Coverity, Ben was a founding member of the Stanford Computer Science Laboratory team that architected and developed Coverity's underlying technology. He is one of the world's leading experts on the commercial use of static source code analysis. Ben holds BS and MS degrees in Computer Science from Stanford University.*

Coverity, Inc.
415-321-5200
bchelf@coverity.com
www.coverity.com

Potential uses of SAT in static analysis

False path pruning

A bit-accurate solution to the formula of the conjunction of all conditions on a path where a defect is about to be reported eliminates many false positives that plagued dataflow analysis solutions for years.

Integer overflow

Using a SAT solver to check key arithmetic operations (in-loop bounds, calls to memory allocation functions, as an index into an array, and so on) to verify that unexpected overflow conditions do not occur in the computation.

Symbolic checkers

Superior defect detection capabilities for the following classes of defects:

- Buffer overflow detection
- String overflows
- Deadcode

Assertion-based checking

- Software assertions can be complex and hard to check, and currently are only leveraged at runtime
- Vision: ability to write assertions using a well-known primitive library
  - No need to learn an esoteric formal language
  - Restrict assertions to checkable properties
- Result: developers will be able to find some program logic bugs
  - This is an area static checkers today simply ignore

Editor's Choice • Embedded Computing Design • PRODUCTS

www.embedded-computing.com/search

Memory modules equipped for cooling

Functional density makes cooling a challenge, and solutions that aren't cost prohibitive or don't lead to excessively engineered products are sometimes hard to find. Virtium Technology has introduced DDR2 SODIMM, SO-RDIMM, and SO-CDIMM memory modules furnished with heat spreaders to remove heat more effectively than conventional air-cooled SODIMMs. The new heat spreader is available for all SODIMM product variants using mainstream 512 Mb and 1 Gb devices.

"With these new heat spreaders, Virtium directly addresses an industry-wide problem to control excessive or less-than-optimal operating temperatures in demanding embedded environments," said Phan Hoang, Virtium Technology's director of R&D. "This approach to heat removal offers multiple benefits including better thermal performance and improved reliability in adverse thermal environments with fewer single bit errors."

Virtium's heat-sink-equipped modules are intended for use in extended temperature ranges, thus solving thermal issues critical to system performance, and are ideal as AdvancedTCA, AdvancedMC, MicroTCA, and PicoTCA memory for blade applications and PICMG SBCs.

Virtium Technology • DDR2 SODIMM memory module with thermal heat spreader

Editor's Choice • Embedded Computing Design • PRODUCTS

Body monitoring, minus the wires

The typical image of patients in hospitals with wires running from them to monitoring equipment may be a thing of the past now that wireless is gaining acceptance in health care. Toumaz Technology's Sensium Platform Integrated Development Environment Resource (SPIDER) allows customers to rapidly create and test end-to-end Sensium-based solutions. The development kit contains two Sensium hardware development boards (one for the sensor transmitter and one for the receiver base station) and a Keil 8051 compiler and JTAG debugger.

The Sensium sensor interface and transceiver platform enables intelligent, ultra-low-power wireless monitoring of multiple vital signs – for example, ECG, heart rate, body temperature, respiration, and physical activity – in real time via standard handheld devices.

Each SPIDER board contains one Sensium TZ1030 and a patchwork area for user-supplied sensor transducers. The TZ1030 includes reconfigurable sensor interfaces, a digital block with an 8051 processor and hardware MAC/network controller, and an RF transceiver. Powered by a battery or from USB/SD supplies, the board includes all the necessary RF circuitry to simplify system development.

Toumaz Technology • SPIDER

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.

ADVERTISER INFORMATION

www.embedded-computing.com/search

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

International Sales

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

Reprints and PDFs

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

www.opensystems-publishing.com

Looking for a BDM/JTAG Target Debugger?

Try LINUXSCOPE-JTD for *FREE!*

Debugger for the Eclipse™ IDE that has been optimized for use with Abatron's BDI2000 BDM/JTAG probe

FEATURES

JTAG TARGET DEBUGGER

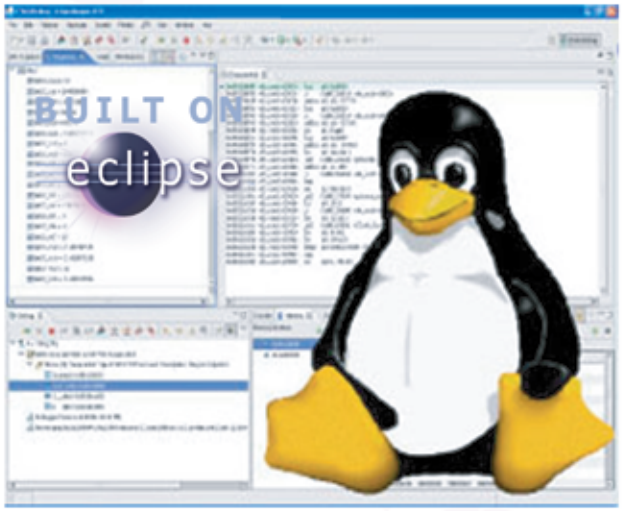

- Designed to debug BDM/JTAG targets with the Abatron BDI2000 Probe
- Flash utility simplifies flash burning process
- Allows creation of custom scripts to perform special target operations
- Designed to allow Linux kernel debugging
- Easily modify/view hardware registers
- Enhanced MI interface adds hardware breakpoint awareness into Eclipse™
- Quick setup - start debugging your target in minutes
- Eclipse™ based IDE & plug-in available allows you to standardize your tools into one platform
- Standard debug features include View/Modify memory, view disassembly, track watch points, view variable, & view stack
- Plug-in updates bringing even more features & enhancements

Visit www.ultsol.com/license-ecd.htm for your FREE evaluation copy!

For more information please contact:

USI
Ultimate Solutions, Inc.

Toll Free: 866.455.3383
info@ultsol.com
www.ultsol.com