

pulse

INNOVATIONS AND INNOVATORS

Intense LEDs take on the sun



LEDs are moving up the intensity scale for outdoor-signage uses, as the OVL series demonstrates. The devices are available in a range of colors and 3-, 4-, and 5-mm sizes.

According to Richard Saffa, vice president of visible LEDs at Optek Technology, the company's OVL series of LEDs performs in full sunlight for a variety of display and signage applications, including commercial outdoor advertising. The through-hole, oval LEDs, in 3-, 4-, and 5-mm sizes, are available in blue, green, red/orange, and red with intensities that vary depending on color and size. For example, the 5-mm, red, AlInGaP device has typical intensity of 950 mcd at 20 mA, and the green InGaN device has an intensity of 1750 mcd.

Viewing angle also depends on size; representative viewing angle is 50° at half-power and 110° maximum. Power dissipation for these 25-cent devices is approximately 130 mW.—by **Bill Schweber**

▷ **Optek Technology**, www.optekinc.com.

Book leads those on the digital dark side to the analog force

A familiar name to *EDN* readers, Bonnie Baker takes many of the subjects and themes of her popular "Baker's Best" column into more detail and depth in *A Baker's Dozen: Real Analog Solutions for Digital Designers* (Newnes/Elsevier, ISBN 0-7506-7819-4, May 2005). This clearly written, highly readable book explains both basic topics, such as A/D-converter operation, and the dilemmas that designers face, such as where, when, and how to filter signals in the signal-processing chain. Other topics include the use and validity of Spice models; digital-versus-analog-processing trade-offs and cooperation; and noise, layout, and grounding.

The \$59.95, 368-pg paperback devotes a chapter to the all-important topic of troubleshooting mixed-signal designs. Colleges don't teach this subject, but engineers usually teach themselves at the debugging bench at first using bad judgment, which leads to ex-

perience and then to good judgment. This self-teaching primer will enhance your design, decision, and debugging skills so that you can face the reality of mixed-signal systems or digital designs with their inevitable analog-world aspects.—by **Bill Schweber**

▷ **Newnes/Elsevier**, <http://books.elsevier.com>.

FEEDBACK LOOP

"Did I read that right? You can program the DSP directly in Visual Basic? Somebody actually built a VB compiler that generates native DSP code? What's next? Programming a DSP in Cobol?"

Ken Dyck, in *EDN's* Feedback Loop on www.edn.com/article/CA605784. Add your comments.


Simple transformer monitors ac currents

Is there an electrical component as versatile as the venerable transformer in its various manifestations? For measuring ac current, the 5600 series from C&D Technologies combines isolation to 500V with primary current rating to 10A. The device finds use in switched-mode power-supply feedback loops, motor-current monitoring, and other sensing applications. Bandwidth is 20 to 200 kHz, and these transformers are available in 50-, 100-, 200-, or 300-turn versions, for different sensitivity factors. Internal inductance is 5 to 335 mH, and dc resistance ranges from 117 mΩ to 8.6Ω.

The less-than-\$1 units are available as two-lead and center-tapped devices and measure 16.8×20×8.9 mm.

—by **Bill Schweber**

▷ **C&D Technologies Inc.**, www.cd4power.com.



The 5600 current-sensing transformer series offer 500V-dc isolation in 50- to 300-turn versions.


Track-and-hold device allows direct capture of signals greater than 10 GHz

The extremely wideband RTH050 dual track-and-hold device from Rockwell Scientific enables direct conversion of gigahertz-range signals and signal edges with rise times of approximately 20 psec. This unit has small-signal input bandwidth of 15 GHz and supports a sampling rate of 100 to 1000M samples/sec. The differential signal-path device incorporates cascaded track-and-hold circuits to provide a hold time that is greater than one-half the clock cycle, which reduces bandwidth requirements for the subsequent signal-handling chain. Aperture jitter is 100 fsec, and acquisition time is 200 psec. Hold-mode distortion is -42 dB for a 5-GHz signal with 0.5V p-p amplitude, rising to -30 dB for a 7.5-GHz signal; noise is less than 1 mV. Spurious-free dynamic range is 65 dB with a 1.060-GHz, 0.5V p-p signal.

According to Ron Latreille, product-line manager for mixed-signal products, the improved speed of this unit, compared with its predecessors, allows users "to digitize a broader range of signals with no increase in power con-

sumption" at 1.65W. A similar device, the RTH055, also integrates 12-dB attenuators to eliminate the need for external

components to scale down signals with excessive full-scale magnitude. The RTH050 sells for \$300 (100) and is available in a 13-lead HSD package.—by Bill Schweber
▶ Rockwell Scientific Co LLC, www.rockwellscientific.com.



Simple in concept but fast in execution, the RTH050 track-and-hold device from Rockwell Scientific offers a 15-GHz bandwidth and a 1G-sample/sec sampling rate with a dual internal structure for relaxation of interface timing.


FEEDBACK LOOP

"The Nyquist Theorem addresses avoiding aliasing of a reconstructed signal, not perfectly reconstructing the original signal. For signals approaching one-half the Nyquist frequency, severe amplitude attenuation periodically occurs as the sampled points move into and out of phase with the signal."

Konrad Kratz, in *EDN's* Feedback Loop on www.edn.com/article/CA529378. Add your comments.

LED module delivers 60 lumens

BivarOpto's LK series encompasses single-chip, 1W LEDs targeting safety, signaling, emergency, and general illumination. The 465- to 635-nm surface-mount devices, also available in white to 8000K, use a Cree XLamp die that includes an aluminum pc-board base for managing thermal issues. The lead-free encased met-




The LK Series of 60-lumen LEDs use a metal-core pc-board assembly to yield a 1W SMT design with integrated lens.

al-core module, which measures 20.3 mm sq, includes an integral lens; you can also get a two-pin header-based connector for snap-in and ease of replacement. The LED operates from -30 to $+100^{\circ}\text{C}$; maximum forward current is 400 mA. The LK series modules sell for \$4 (volume quantities).

—by Bill Schweber
▶ BivarOpto, www.bivar.com.

DILBERT By Scott Adams



Tool analyzes ICs

Giga Scale Integration claims that its new IC-economic-analysis tool can help you figure out whether a foundry's less expensive, standard 90-nm process or even its 130-nm process will meet your design requirements just as well as the same foundry's pricier, low-power, 90-nm process. The technology expands on the company's InCyte chip-estimation tool, which allows users to derive an IC specification with estimates on die size, power, performance, and leakage. The company has now taken the estimation technology a step further by creating an economic analysis add-on engine to InCyte that allows users to turn their design specification into an IC budgetary quotation and figure out the best foundry and fabric for their next design project.

Adam Traidman, Giga Scale IC's president and chief executive officer, says that the company gathered silicon-wafer-pricing and defect-density data from foundries and trade groups, package-pricing data from packages, and test- and assembly-cost data from various vendors. The InCyte economic engine's database stores that data. "We take the output of our InCyte tool, which provides technical chip estimation for die size, power consumption, and number of pins, and combine that data with the new economic data to produce this final packaged-chip cost," says Traidman.

The engine also has a life-cycle-analysis feature that allows designers to forecast and account for mask re-spins, increasing yields, and decreasing wafer and package costs.

Traidman says that Giga Scale will eventually improve the tool to allow users to analyze whether they should implement their designs in FPGAs, structured ASICs, or tradition-

al cell-based-ASIC fabrics. The economic engine is available as a \$2000 upgrade to InCyte. InCyte is available as a free download at www.chipestimate.com.

—by Michael Santarini
 ▶ **Giga Scale Integration**, www.gigaic.com.

Embedded DSP simplifies programming

Sheldon Instruments' recently released SI-C6713DSP-PC104p embedded DSP board in the PC/104-Plus form factor targets real-time test, measurement, and digital-control applications. With a typical power consumption of only 2.25W, the board features the latest 300-MHz TMS320C6713 DSP from Texas Instruments and as much as 256 Mbytes of SDRAM. An expansion site for a family of multifunction I/O modules extends the DSP power with as many as 64 analog inputs and 16 analog outputs; a timing interface with two digital synthesizers and four counters; and a digital interface that comprises 32 bits of general-purpose I/O, two quadrature encoders, two frequency counters, and two PWM outputs.

The board comes with DSP-software libraries that enable system engineers to directly program their DSP systems in Visual Basic or National Instruments' (www.ni.com) LabView. However, for those who require custom development, the conventional C/C++-based tools have source-code-distribution support. Prices for the SI-C6713DSP-PC104p board with a 300-MHz DSP and 64 Mbytes of SDRAM start at \$1985, and prices for I/O modules start at \$595.

—by Warren Webb
 ▶ **Sheldon Instruments**, www.sheldoninst.com.



Designers can easily program the SI-C6713DSP-PC104p board in LabView and Visual Basic. The board incorporates the TMS320C6713 DSP operating as fast as 300 MHz.

FEEDBACK LOOP

"Many (software engineers) are actually other types of engineers that became involved with programming in school or on the job through necessity. Often facing a time crunch, they write only enough code to make the project work, rarely giving any thought to quality code."

Joseph Travis, in *EDN's* Feedback Loop on www.edn.com/article/CA601846. Add your comments.

Physical guide routes optical guide

Even optical fibers, although relatively rugged, may still need some support and protection when designers cram them into today's tight and crowded boxes and systems. The OFFGS (optical-fiber-flexible-guide system) from Richco Inc provides the needed routing and channeling for 250- and 900-micron fibers. The flexible guide lets designers snake the fiber around the pc board and within the box and provides consistent placement and basic protection.

According to Cheryl Cummins, marketing manager at Richco, the OFFGS is available in a flat style for mounting directly to the pc board and with support posts that raise it above the pc board, so that it can go over ICs and other obstructions. A standard OFFGS piece measures 6.5 in. (16.5 cm) long and costs about \$2.

—by Bill Schweber
 ▶ **Richco Inc**, www.richo-inc.com.



Keep that optical fiber well-behaved and constrained within your pc-board-based product with a flexible guide from Richco.

Emulators speed event-based verification

Tharas Systems Inc has released two new emulator systems that use multiprocessor SOCs (systems on chips). The single-user Hammer S-Class emulator has a capacity of 16 million gates and targets hardware designers needing to speed event-based verification. The multiuser Hammer M-Class emulator has a 64 million-gate capacity and targets hardware engineers who want to verify their designs, and, through the new Virtual Connect add-on, software engineers, wishing to get a jump on software development. Richard Curtin, senior vice president of marketing and business development for Tharas, says that the two new systems improve on the company's previous offering, the Hammer 100 system, on just about all fronts.

Whereas the Hammer 100 used 128 processors, had a 32 million-gate capacity, and ran an 8 million-gate design with a testbench operating at 4.5 kHz, the new S-Class uses 1024 custom processors, has twice as many gates, and runs the same 8 million-gate design at 50 kHz. The M-Class has 4096 processors, a fourfold gate increase, and the ability to run the same 8 million-gate design at 200 kHz.

Curtin attributes the speed and capacity improvements to the company's placing 32 of its Hammer 100 custom processors on one SOC. This approach increases the capacity and exponentially speeds the emulators and allows Tharas to offer the systems in a form factor the size of a big toaster. Tharas also custom-designed the instruction set to handle Boolean operations for simula-

tion. Custom instruction-set simulation speeds emulation. The company has also worked extensively on the compiler to help users quickly load their designs onto the emulator systems. Curtin claims a 50 million- to 60 million-gate compilation time. The emulators accept Verilog- and VHDL-design files; assertions in OVL (object-verification language), OVA (Open Vera assertion), and PSL (property-specification language); and testbenches in Verilog, VHDL, or C/C++.

Tharas also offers a software add-on that allows embedded-system engineers to get a jump on software development. The company's Virtual Connect add-on has three flavors of software models targeting popular market segments: the Virtual-PC platform for graphics applications, the Virtual-Net for networking applications, and Virtual-3G for wireless applications.

Curtin says that most users are designing ASICs with 5 million to 6 million gates. "A 16 million-gate system, such as the S-Class, should cover most design groups for the next couple of design projects," he says, arguing that manufacturers will realize the value of the emulator for its increased speed versus simulation over the life of the system. Prices for the Hammer S-Class start at \$200,000 for a 4 million-gate configuration, and prices for the M-Class start at \$750,000 for a 32 million-gate configuration. Tharas also rents its emulators. Price of an S-Class rental in an 8 million-gate configuration starts at \$21,000 per month. Price of an M-Class rental in a

Module enables graphical embedded programming

National Instruments recently announced an embedded-development module for its popular LabView software that extends graphical programming to any 32-bit embedded processor. With the module, designers can describe algorithms and program their applications using LabView's graphical data-flow language. In addition, LabView's graphical front-panel controls and indicators help engineers quickly and interactively experiment and debug their embedded code.

The LabView embedded-development module includes more than 400 analysis functions for signal processing, linear algebra, curve fitting, statistics, and calculus. The module also features a structure for integrating I/O drivers and board-support packages to take advantage of specific processors. Designers can develop their applications using the LabView graphical environment and then automatically generate C code to integrate with their chosen processor tool chains.

James Truchard, PhD, chief executive officer and co-founder of National Instruments, says, "The evolution of LabView beyond the PC will open embedded-system design to

an entirely new class of developer—a more mainstream population of scientists and engineers." Prices for the module start at \$9995.—by Warren Webb

▶ National Instruments, www.ni.com/embedded.

32 million-gate configuration starts at \$52,000 per month. Prices for the Virtual Connect add-on start at \$75,000.

—by Michael Santarini
▶ Tharas Systems, www.tharas.com.

FEEDBACK LOOP

"I have this idea that I can apply engineering design to stock prices as though they were signals and make a living that way. Hey, if it works, my income won't be limited, and I can't be outsourced!"

Anonymous, in EDN's Feedback Loop on www.edn.com/article/CA529820. Add your comments.

Q&A

Todd Westerhoff

EDA industry should use what it sells

As a high-speed-design manager for Cisco Systems Inc and a former product manager for Cadence Design Systems, Todd Westerhoff feels the pain of EDA customers he used to sell to and advise. He has 25 years of experience in modeling and analyzing electronics systems.

Currently, he and his team of 10 focus on signal integrity and support other design groups at Cisco. Westerhoff earned a bachelor's degree in electrical engineering from the Stevens Institute of Technology (Hoboken, NJ). He vented his frustrations about EDA tools in an interview after a recent vendor panel sponsored by the EDA Consortium.

What do you do at Cisco?

A We use simulation tools to model the behavior of chips and the way we interconnect them at the board level. The limiting factor is that the speed of light is too slow. Circuits are switching so fast that the way a signal propagates is like a wave front. We can no longer consider circuits as lumped elements but rather as distributed structures like a wave.

Where does EDA come into the picture?

A In the past, you'd put these boards onto an oscilloscope or analyzer, but we use off-the-shelf simulation tools to increase the chance of first-pass success. The cost of mistakes in board design is prohibitive. The problems we face are so complex, you can't fix them simply by trial and error. We have to use simulation tools before we "fab" out the board.

Why do EDA tools fall short?

A The things we are trying to do, such as determine how an ASIC works

with other chips, is a fantastically complex problem. It requires multiple tools to predict the board's behavior. Even the vendors that say, "We have a full solution" don't cover the myriad details associated with modeling and manufacturing the board. The problem of building something and bringing it into manufacturing has a million angles.

Why does this problem evoke so much emotion?

A I don't know if I'd [sound off strongly in a public forum] again. But sometimes you have to shake people up to get them to do things differently. I've been saying [EDA companies should use their products in actual design] for seven or eight years, and, obviously, not a lot of people are listening. We spend a lot of time and money on getting these tools to work correctly. I don't know anybody who says their relationship with the EDA industry is a walk in the park. But we'd be a lot better off if these tools were tested in a real-world situation.

How did your transition from vendor to customer influence your thinking?

A When I got out of college, I spent 10 to 12 years as the guy marketing analog and digital simulators. Then I went out as a consultant using EDA tools, and it was a real eye opener. Instead of selling tools, I went to make a living using the tools. Many users never develop the comprehensive understanding of the tool that the vendors expect. User understanding of the tool is uneven, so everybody's perception of what needs improvement is different. The tool becomes difficult and complex to manage. They are big and complicated with many options. The user is saying "I just want to solve this problem." What grows out of that is you have tools that don't work or solve problems.

When I got into EDA around 1983, the paradigm was that, if you could get a computer to do what you wanted, you would be thankful. That [mindset] has never gone away with [EDA]. People tend to look at the problem from their own perspectives. It all changes when you turn around and try to use it, but if you don't use these things, you're not going to be able to perfect them. They're like race cars. You have to go to the track and beat them to death, or you're not going to win. It's also kind of like being in battle where everything is changing and going on around you. The last thing you need is to be in battle with tools that are unpredictable.

How do you cope?

A It's a constant problem and one we've learned to live with. The fascinating



thing is that some of the things we thought were trivial are big deals, and the things we thought were simple rough edges turn out to be big factors.

Please give an example.

A When I was a vendor, this schematic editor had all these fancy functions. Users came back and said that to put a component down and rotate it took seven mouse clicks, whereas our competitor's product required two. Vendors need to get things done repetitively right—not close, but right. You need to nail it. If it's something a user is going to do several hundred times a day, it is going to drive him crazy. And, if you're going to nail it, you've got to use it.

We can ask [a vendor] to add another bell or whistle to a simulator. [They say], "Here's this collection of capabilities. If we change it all around, maybe we can give them that feature." They might say, "We can't quite do this, but we can do something like it." That could be in 12 to 18 months, and compromises get made along the way. In the meantime, the technology moves on.—by John Dodge

For an extended version of this interview, go to www.edn.com/050707/pq&a.

 GLOBAL DESIGNER

WiCon panel opines on state of wireless standards

At the WiCon (Wireless Connectivity) World show in London in May, Mike McCamon, executive director of the UWB Forum (www.uwbforum.org), described as “an utter mess” the current situation of competing personal-wireless-connectivity standards. Partially giving rise to the description was the Bluetooth Special Interest Group’s (www.bluetooth.com) May 2005 announcement that it intends to work with developers of UWB (ultrawideband) technology to “combine the strengths of both technologies.” McCamon likens the proliferation of standards to the situation in wired networking around 1990. He looks toward a similar simplification in the same way that IP (Internet Protocol) came to dominate networking technology, anticipating that the wireless industry should be able to “clean up” the situation.

The UWB Forum stands on one side of the continuing UWB-standards wrangle, and the MBOA (Multiband OFDM Alliance) stands on the other. Nevertheless, Kursat Kimayacioglu, MBOA marketing-group member and director of wireless business development for the connectivity-busi-

ness line at Philips (www.philips.com), echoes the sentiment. He looks to a more unified wireless environment in which the Bluetooth profile becomes just one of a number using the UWB-air interface.

The show provided a snapshot of progress toward functional UWB in its native mode; Freescale (www.freescale.com) operated a streaming-video demonstration with a claimed 100-Mbps link spanning 15 to 20m. This demonstration employed the first multichip implementation of its direct-sequence UWB silicon, and the company promises further integration. Visitors could also see file transfer over a wireless-USB link using first silicon of Staccato Communications’ (www.staccato.com) single-chip all-CMOS offering. Staccato envisages the technology’s reaching the market on a time scale that would see external adapters for file transfer from PC to external hard drive early in 2006 and reaching high-end mobile-phone platforms in 2007.

—by **Graham Prophet**,
EDN Europe

► **Wireless Connectivity World**, www.wiconworld.com.



BITS & PIECES

► **Yokogawa Electric Corp** (www.yokogawa.com) of Japan has released the DL9000 series digital-oscilloscope platform for the 1-GHz band. For more, see www.edn.com/050707p1.

► **Wipro Technologies** (www.wipro.com), the software services arm of India’s Wipro Ltd, has created an SOC (system-on-chip)-based MP3 player that can operate either as a stand-alone chip or as part of a mobile phone. For more, see www.edn.com/050707p2.

FPGAs implement high-end image-processing applications

New FPGAs with more DSP resources and embedded-processing capabilities have made the global image-processing market more competitive. According to Rahul V Shah, ASIC manager at elfochips, an ASIC-design and -verification-services provider, designers can offload software-implemented algorithms, such as DCT, static Huffman, AES (Advanced Encryption Standard), color-space conversion, and gamma correction, to FPGA hardware. He claims that this approach improves system performance. Shah also notes that the conventional manner of implementing algorithms by software limits its performance due to serial data processing. Increasing frequency beyond certain limits causes system issues. FPGAs have flexible architectures and dedicated DSP blocks. Coupling these benefits with parallel processing strikes a proper balance in system performance and cost. FPGAs extend the flexibility to reprogrammability, resulting in a quick turnaround time.

According to Shah, parallel processing in hardware is impossible because hardware processing can execute instructions only one at a time. “If you want to run a DCT along with static Huffman with the same processor, one process at a time will execute,” he says. “However, in hardware, because everything runs in parallel, you can have DCT and SHF [super-high-frequency] running in parallel at the same time without any performance hit.” He adds that FPGAs provide the flexibility to upgrade to new standards and reprogram devices. For example, you can modify any system-level application to DDR and then move to DDR-2.

Designers can also reshuffle images to create applications, such as video cell phones, set-top boxes, LCD projectors, keyboards, videos, mice, and digital cameras and camcorders, among others. “Implementing DSP algorithms for image-processing blocks, such as DCT, AES, and static Huffman, requires a huge amount of memory, multiplier, and accumulator blocks,” says Shah. For example, a DCT at 133 MHz can take as many as 64 multiplication and addition operations. Designers can map these multipliers onto the hardware to perform multiplication and addition operations, rather than do it sequentially in software. Static Huffman and AES cores have high memory requirements for storing coefficient values and performing mathematical operations on incoming data. Handling these operations in software slows down system performance and overloads the CPU. Offloading these tasks to the hardware means that the memory stores the image, and the CPU performs other control operations.

Implementing a dynamic Huffman algorithm in hardware would be a bad idea, because the algorithm requires dynamic calculations, which need a large amount of hardware, so proper partitioning between hardware and software is necessary. This approach is more cost-effective and provides better performance than running the algorithm in software.

—by **Pradeep Chakraborty**, *EDN Asia*

► **elfInfoChips**, www.einfochips.com.

07.07.05

Hybrid DSP/ μ Ps meet low-cost BoM targets

Freescale has extended the scope of its digital signal controller devices—microcontrollers that have a combined control and signal processing function set—with a range of very low-cost, flash-based parts. Designers can use them in motor control, domestic appliance, power supply and consumer product applications. In the 56800 family, these chips have part numbers in the MC56F8000 format. They use the 56800E core (which includes single-cycle multiply-accumulate) and are housed in 32-pin packages giving up to 26 I/Os. They offer the 6-output PWM controller used elsewhere in the DSC range, running at 96 MHz (three times the basic clock of 32 MHz) to give fine resolution on the PWM output for precise motor control.

There is a dual-channel analogue-to-digital converter with conversion time of 1.125 μ sec, that can be synchronised with the PWM stage. External parts count can be minimised by use of an internal oscillator that is stable to

$\pm 2.5\%$ over the industrial temperature range, with parts that operate over the full automotive range to be added at a later stage. Other peripherals include four timers and a range of interfaces. Any peripheral that is not in use can have its external pins re-allocated as I/O.

Development support is provided via the Metrowerks CodeWarrior suite, with a low-cost development kit (\$100) and a free version of the software available by download. The free version is limited to 16k of code space, but the smallest parts in the family have a flash memory block of that size and so can be programmed with the free download. Support includes free downloadable reference designs in a number of application areas. The software also offers a library of pre-compiled software routines for functions that include motor control, communications, DSP operations and mathematics. The chips will be priced at \$2.99 (10,000).—by Graham Prophet

► **Freescale**, www.freescale.com

Multi-standard software-defined-radio adds DRM format

DRM (Digital Radio Mondiale, www.drm.org) is a proposed standard for digital radio that enables broadcasters to reuse frequencies below 30 MHz as the conventional services there (AM, LW, and "Short Wave") are listened to by progressively smaller audiences. DRM uses compressed audio and QAM digital modulation to achieve robust transmission and reception with added service features. Broadcasters can achieve extended coverage—in some cases continent-wide—with the same power that previously covered a more limited area, or can cover the same geography with reduced power. They can also package multiple channels in the same spectrum allocation. A further attraction for operators is that broadcasters often only need a modest upgrade to their transmitters installations. Depending on the data rate and codec used, quality of "near-FM" standard is achieved. Many broadcasters are already experimenting with transmissions.

Now, RadioScape—which has for some time produced modules to receive DAB digital radio—has introduced the RS500 module that adds DRM to DAB (and FM/RDS, LW, MW and SW). It includes the same features as the earlier RS300 module, forming the core of a complete digital radio with pause, rewind and record to MMC memory card, plus full access to electronic programme guide data. DRM permits features such as frequency shifts, where the broadcaster moves the transmission between his allocated frequencies in the course of a day to cope with the variable nature of propagation below 30 MHz, with receivers enabled to track the channel shifts. RadioScape recently changed its profile to be a manufacturer of complete modules rather than a supplier of silicon and provider of licensed module designs.

RadioScape builds the module around the same Texas Instruments TMS320 DSP chip used in the earlier design, now coded TMS320DRM350 (which will also be available direct from TI).

RadioScape suggests a full-featured digital radio with DRM can now be marketed for around € 250. Such a set will manage all of the sources of a particular station, offering the listener a selection of all possible sources of



the programme material. Digital radio, according to RadioScape, will come to mean (in the territories where they are in use) the sum of DAB and DRM services. DRM is a permitted format in many regions of the world including the US, but so far is little used in the US, the main momentum in that market for digital radio having gone into satellite-based subscription services. RadioScape suggests that one area where it may gain a foothold is in the provision of local—city-wide, or even more local—services to special-interest or specific-language audiences. With DRM, these can be served with very low radiated power, and a high level of frequency reuse in geographical terms.—by Graham Prophet

► **RadioScape**, +44 20 7317 1969, www.radioscape.com

► **Texas Instruments**, www.ti.com

Standard FPGA hardware will power Teraflop computer

A consortium based in Edinburgh, Scotland is to build and operate the largest computer so far planned, based on reconfigurable computing

techniques and on FPGA hardware. The machine will be a 64-node configuration and will be in the “supercomputer” class, aiming for operating speeds in excess of one Teraflop. It will be built with COTS (commercial-off-the-shelf) hardware. The project is planned to take two years and will be based at the Edinburgh Parallel Computing Centre (EPCC) at Edinburgh University, which already hosts Europe’s largest “conventional” supercomputer. Partners in the alliance include FPGA IP specialists Algotronix; reconfigurable systems suppliers Alpha Data and Nallatech; Xilinx; and

the Institute for System Level Integration, a joint research initiative between several universities in Scotland.

A computing system will achieve the fastest execution of any algorithm when the algorithm is directly implemented in hardware. An FPGA-based computing system offers a way to approach this ideal, using reconfigurable hardware, while retaining a general-purpose computer’s ability to tackle any computation problem. The principle has been known for a long time, and commercial companies (Nallatech and Alpha Data among them) have

used it to build successful product ranges. This is believed to be the first time that the technique will have been applied on such a large scale.

One impediment has been the difficulty of achieving the same flexibility as a conventional computer in handling any arbitrary problem. For a conventional computer, code is written and compiled. For a reconfigurable hardware machine, the user has to partition the problem to determine which parts are best suited to running in hardware and software; then generate the configuration code for the hardware components, and compile the software code. The alliance will aim to demonstrate the progress that has been made in that process by selecting and porting three existing super-

computer applications to the new machine. From this work, the partners will create tool kits and other IP and make them available for license, and it is planned that this activity will become the centre of a new “cluster” of FPGA-based computing activity. Once running, researchers with computation-intensive academic projects will be able to reserve time on the computer.

Beginning on page 38 of this issue of *EDN Europe*, Dr Malachy Devlin of Nallatech explains why he thinks that programmable logic technology has now reached a level of maturity that makes a configurable computing project on this scale realistic.—by Graham Prophet

► **EPCC**, www.epcc.ed.ac.uk; **Algotronix**, www.algotronix.com; **Alpha Data**, www.alphadata.co.com; **ISLI**; www.sli-institute.ac.uk; **Nallatech**, www.nallatech.com; **Xilinx**, www.xilinx.com

Program DSPs directly in LabView

National Instruments has extended the capabilities of the LabView graphical development environment by adding a module that directly generates executable code for certain Texas Instruments DSP chips and development boards. Previously, designers might have developed signal processing algorithms in LabView but would have had to transfer them to an environment such as TI’s own Code Composer Studio to compile the code. The DSP Module is aimed at both DSP developers and engineering students (as a learning environment), providing a graphical, system-level context in which to directly programme their hardware. It will program TI TMS320C6711 and C6713 evaluation boards, in-

cluding NI’s own Speedy-33 board with signal processing functions, and elements such as spectral analysis and filtering, without a separate DSP compiler.

Programming is by the established LabView drag-and-drop graphical user interface. NI says that the module also cuts the time needed to make incremental changes to a project. It integrates with the LabView Digital Filter Design Toolkit, that includes tools for modelling and creating software-based filters, as well as FPGA code-generation capabilities to offer a further implementation alternative.

—by Graham Prophet
► **National Instruments**, www.ni.com, +44 1635 523545



Wireless protocols adopt UWB—but UWB wrangle drags on

An audience at the WiCon (Wireless Connectivity) show in London in late May heard the current situation of competing personal wireless

connectivity standards described as “an utter mess”. The speaker was Mike McCamon, Executive director of the UWB Forum; the occasion was a panel discussion in the WiCon Conference. Stimulating the debate was the announcement a few days earlier from the Bluetooth SIG (Special Interest Group) that it intends to work with developers of UWB (ultra-wideband) technology to “combine the strengths of both technologies”. McCamon likens the present proliferation of standards to the situation in wired networking around 1990 and looks towards a similar simplification in the same way that IP came to dominate networking technology, anticipating that the wireless industry should be able “to clean up” the situation.

McComan is not a disinterested party, as the UWB Forum stands on one side of the continuing UWB standards wran-

gle, with the Multiband OFDM Alliance ranged opposite. Nevertheless, a similar sentiment was echoed by Kursat Kimayacioglu (of MBOA and Philips), who looked to a more unified wireless environment in which the Bluetooth profile becomes just one of a number using the UWB air interface: “They can be converged; I expect a demonstration of Bluetooth on a UWB platform in 2006.” In this vision, shared by several vendors participating in the show, wireless versions of multiple familiar formats will be married to UWB by having the upper layers of their protocols talk to UWB as a common underlying physical layer. USB (the most talked-about at WiCon), plus 1394 and IP would have equivalent status to Bluetooth in this scenario, and each would have UWB as a short-range, high data-rate transport option.

For the Bluetooth SIG, Executive Director Mike Foley sees the exercise from the reverse direction; enabling a UWB radio to work in a Bluetooth environment. Citing an example of how such a hybrid might be used in a product, Foley imagines an MP3 player. The player would use the regular Bluetooth (the latest version 2.0 Extended Data Rate, for full stereo reproduction) to route music to the headphones; but when it’s time to download Gbytes of content from PC to player, it would switch to a “turbo mode” to get very fast transfer of bulk data via UWB.

The show provided a snapshot of progress towards functional UWB in its native mode; Freescale operated a streaming-video demonstration with a claimed 100 Mbit/sec link spanning some 15-20m of the exhibition floor. This employed the first (multi-chip) implementation of its direct-sequence UWB silicon, with further integration promised. From the MOFDM camp, visitors could see file transfer over a wireless-USB link using

first silicon of Staccato Communications’ single-chip all-CMOS offering. Staccato envisages the technology reaching the market on a time scale that would see external adaptors (for, say, file transfer from PC to external hard drive) arrive early in 2006; reaching high-end mobile phone platforms in 2007.

Elsewhere, the WiCon event was notable for the resurgence of interest in Bluetooth, spurred by the release of version 2.0 EDR. Test equipment business always being a good indicator of true design activity, Agilent cites testing to that standard as its most in-demand product in the sector. The company was also demonstrating the use of paired and phase-coherent signal sources, analysing a co-channel MIMO data link, a capability that will be required for the (also much-discussed) embryonic 802.11n variant—another technology currently locked in a technology-proposal stand-off.

—by Graham Prophet