Reed Electronics Group

EDN

VOICE OF THE ENGINEER

JUNE 21

Issue 13/2007

www.edn.com

EXPLORING MEMORY ARCHITECTURES: PILLARS OF PROCESSING PERFORMANCE

Page 44

DESIGN FOR DEBUGGING: THE UNSPOKEN IMPERATIVE IN CHIP DESIGN

Page 35

The evolution of SERDES designs for cell phones

Page 57

Digital Power for a Changing World

Adaptable, High-Density, Multi-Functional POL Controller

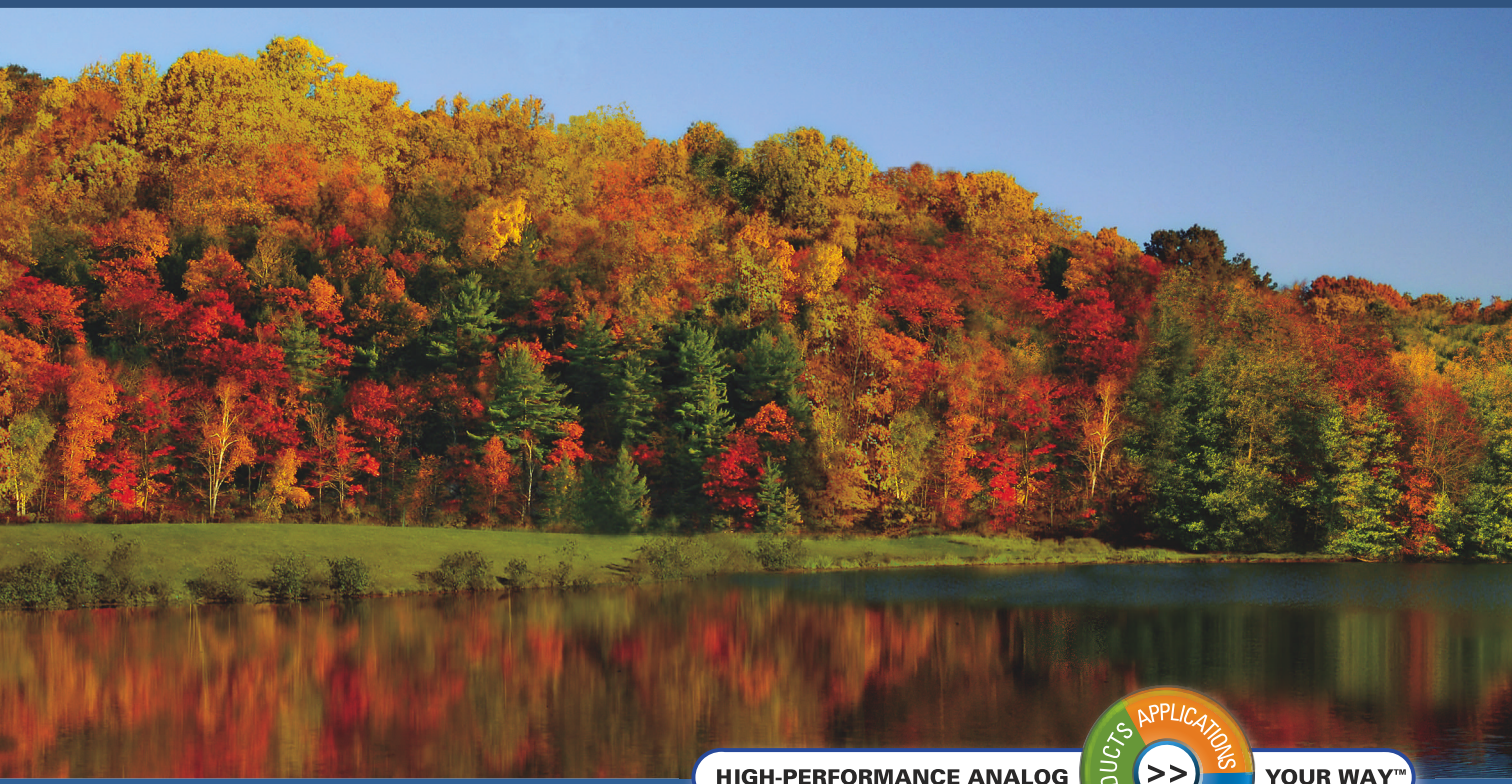

HIGH-PERFORMANCE ANALOG >> YOUR WAY™

PRODUCTS APPLICATIONS SUPPORT

The new UCD9240 Fusion Digital Power™ point-of-load (POL) controller from Texas Instruments gives designers faster time to market without sacrificing features or performance. This flexible and adaptable multi-rail controller provides high power density, dynamic power supply, load-optimized phase management and configurable sequencing.

Phase Management Efficiency Benefits

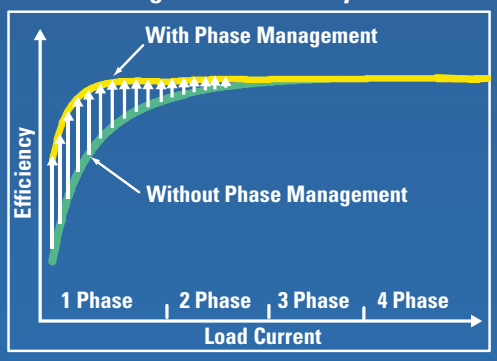

High-Performance Analog >>Your Way

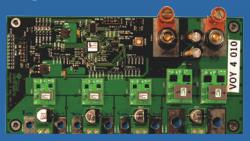

For evaluation modules and data sheets, visit >>
www.ti.com/digitalpower
1-800-477-8924 ext 9240

Technology for Innovators™

Texas Instruments

In Stock and at Your Door

CLICK OR CALL TODAY!

Access to more than a million products now available online!*

Same-day shipment, next-day delivery on orders entered by 8:00 PM CST.

Quality Electronic Components, Superior Service

www.digikey.com

1.800.344.4539

701 Brooks Ave. South • Thief River Falls, MN 56701 • USA

*New product added daily.

90% of the time spent measuring

is in setting up the measurement.

Agilent N6705A DC power analyzer

- *View critical multiple-bias-supply turn-on/turn-off sequences*
- *Display voltage and current versus time*
- *Generate DC bias supply transients*
- *Log data from seconds to days*

Setting up and programming multiple test instruments is holding you back. Free yourself and get right to your DC bias testing with the new Agilent N6705A DC power analyzer.

The N6705A integrates the capabilities of up to four power supplies, a DMM, an arbitrary waveform generator, a scope, and a data-logger into a single instrument.

Every function is controlled from the front panel, and a single graphical display reveals all your measurements. Even the controls are designed to save you time. The N6705A includes an easy and intuitive interface, with color-coded controls and convenient button placement to give you quick access to the functions you need.

Don't be held back. Time savings are within reach.

LXI

u.s. 1-800-829-4444
canada 1-877-894-4414

www.agilent.com/find/dcpoweranalyzer

Agilent Technologies

Comprehensive Analog Portfolio

High-Performance Analog >> Your Way™

AMPLIFIERS

- Audio Amps
- Buffer Amps
- Difference Amps
- High-Speed Amps
- Instrumentation Amps
- Isolation Amps
- Log Amps
- Low-Voltage Op Amps
- Power Op Amps
- Precision Op Amps
- Programmable Gain Amps
- Video Amps

DATA CONVERTERS

- Audio ADCs
- Delta-Sigma ADCs
- Pipeline ADCs
- SAR ADCs
- Audio DACs
- Current-Steering DACs
- Delta-Sigma DACs
- Precision DACs
- General-Purpose DACs
- Audio SRCs
- Data Acquisition Systems
- Modulators/Filters

POWER MANAGEMENT

- Battery Management
- Charge Pumps
- DC/DC Switching Controllers
- DC/DC Switching Converters
- Digital Power
- Display Drivers
- Hot Swap
- LED Drivers
- Linear Regulators
- MOSFET Drivers
- Plug-In Modules
- Power Factor Correction
- Power Supply Control

INTERFACE

- 1394
- CAN
- Circuit Protection
- Digital Isolators
- Display Interface
- LVDS/MLVDS
- PCIe/PCI
- RS-485, 232 & 222
- SCSI
- Serializers
- Deserializers
- Transceivers
- UARTs
- USB

RF & ANALOG COMPONENTS

- 4-20 mA Transmitters
- Analog ASSPs
- Analog MUXs
- Clocks & Timers
- Comparators
- Current Shunt Monitors
- Digital Up Converters
- Digital Down Converters
- Fan Controllers
- ISM Band
- References
- Switches
- Temp Sensors
- ZigBee™

HIGH-PERFORMANCE ANALOG >> YOUR WAY

PRODUCTS APPLICATIONS SUPPORT

Texas Instruments offers a comprehensive high-performance analog portfolio from amplifiers and data converters to power management and interface. Whether you need faster speed, higher precision, lower power or smaller size, TI has analog ICs to differentiate your design. That's High-Performance Analog >> Your Way.

OPA211	**Low-noise, low-power, 36V bipolar op amp** (1nV/√Hz; 3.5mA) in small MSOP-8 package
OPA333	**Lowest power, zero-drift op amp:** 2µA offset, 17µA current, SC70 and SOT23 packages
ADS8422	**Fastest 16-bit SAR ADC (4MSPS)** 2LSB INL, 92.5dB SNR
ADS5546	**Fastest 14-bit ADC (190MSPS)** 73.2 SNR at 70MHz IF; 84dBc SFDR at 70MHz IF; 1.1W power
ISO721	**High-speed digital isolator:** 4000V (peak) isolation; signaling rate of 0-150Mbps; high electromagnetic immunity; low input current requirement; failsafe output
TAS5261	**PurePath Digital™ amplifier** capable of driving more than 300W
TPS63000	**96% efficient buck-boost converter** with up to 28% greater run-time; 3 x 3 mm² QFN package
TPS40140	**Stackable, scalable, flexible PWM buck controller** generates 10A to 320A of output current and enables maximum efficiency up to 16 phases

www.ti.com/analog

Technology for Innovators™

Texas Instruments

Interface

High-Performance Analog >> Your Way™

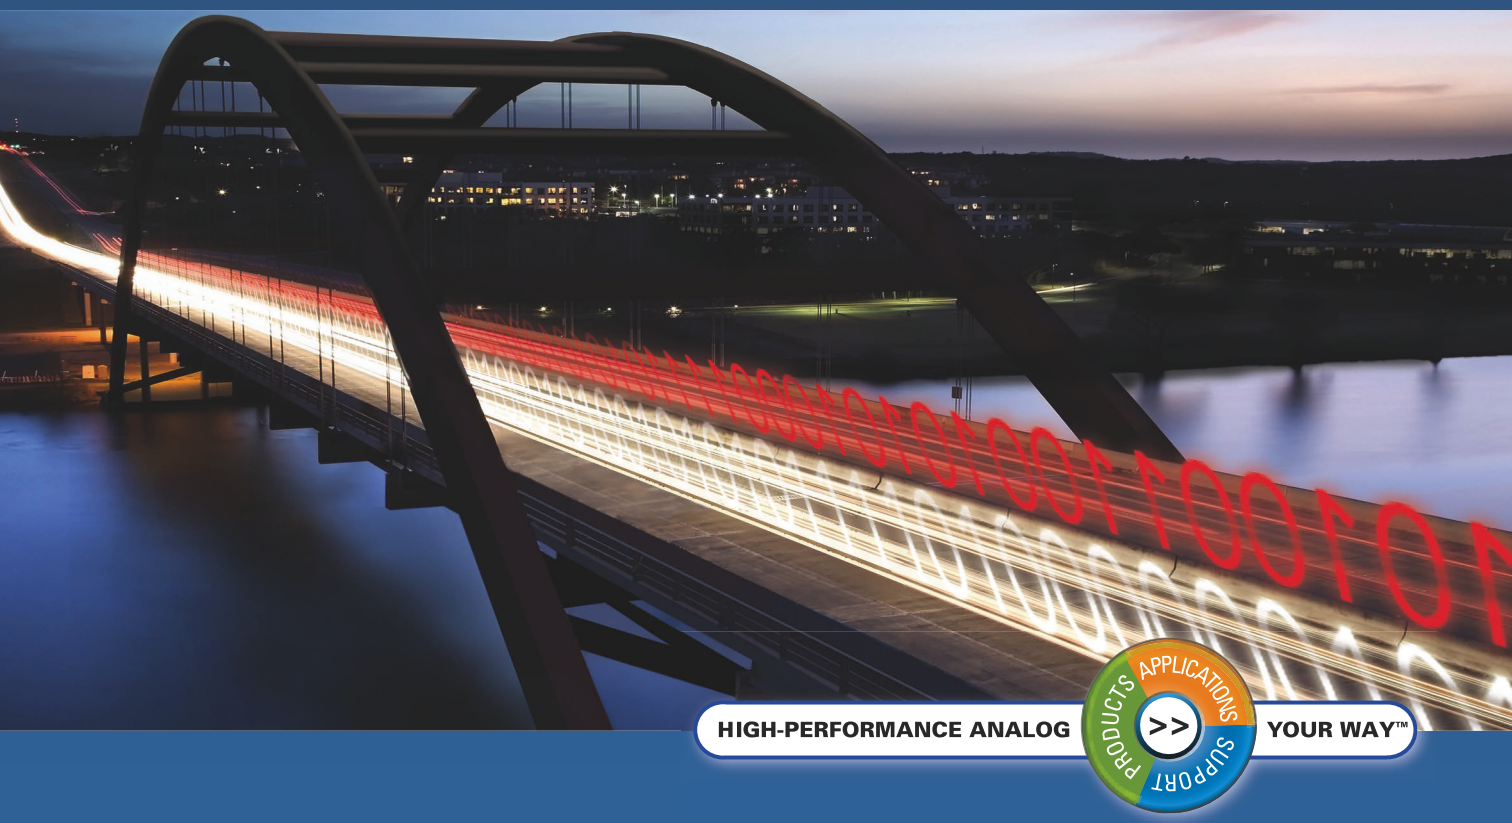

Interoperability, speed, distance, low jitter, magnetic immunity – regardless of your design considerations, TI has an interface product that meets industry standards, satisfies your application needs and helps you win in the marketplace. That's High-Performance Analog>>Your Way.

SN65LVCP404	**Fastest** 4x4 crosspoint with VML output, 4.25Gbps, 30-ps jitter
TMDS351	**Smallest** 3:1 HDMI/DVI switch with selectable input equalization
SN65MLVD3	**Industry's first/only** single-channel M-LVDS-to-LVTTL translator, 250Mbps
TLK1221	**Smallest** Gigabit Ethernet SerDes in 6x6mm QFN, 0.6 to 1.3Gbps
ISO7220/1	**Longest life** – 28 years, dual 4kV isolation, 150Mbps, lowest skew 0.2ns
SN65HVD235	**First in 3-V CAN bus**, 16kV ESD, auto-baud loopback and low-power standby
SN65LBC184	**Highest ESD-rated** RS485 with TVS, 30kV IEC, 400-W protection
SN65HVD23	**Only** RS485 with receiver equalization 25Mbps over 200 meter cable lengths
XIO2000A	**Highest level** of integration delivers seamless PCI to PCI Express® migration
TUSB3410	**Most flexible** USB-to-serial bridge featuring programmable 8052 solution
TL16C255x	**Lowest supply-voltage** 1.8V UART enables interface to latest processors
TPD3E001	**Lowest leakage** (1nA max), 15kV ESD Protection, 1.5-pF

Ease the selection process with samples, evaluation modules and the Interface Selection Guide >>

www.ti.com/interfaceproducts

1.800.477.8924 ext.1105

Technology for Innovators™

Texas Instruments

Precision Amplifiers

High-Performance Analog >> Your Way™

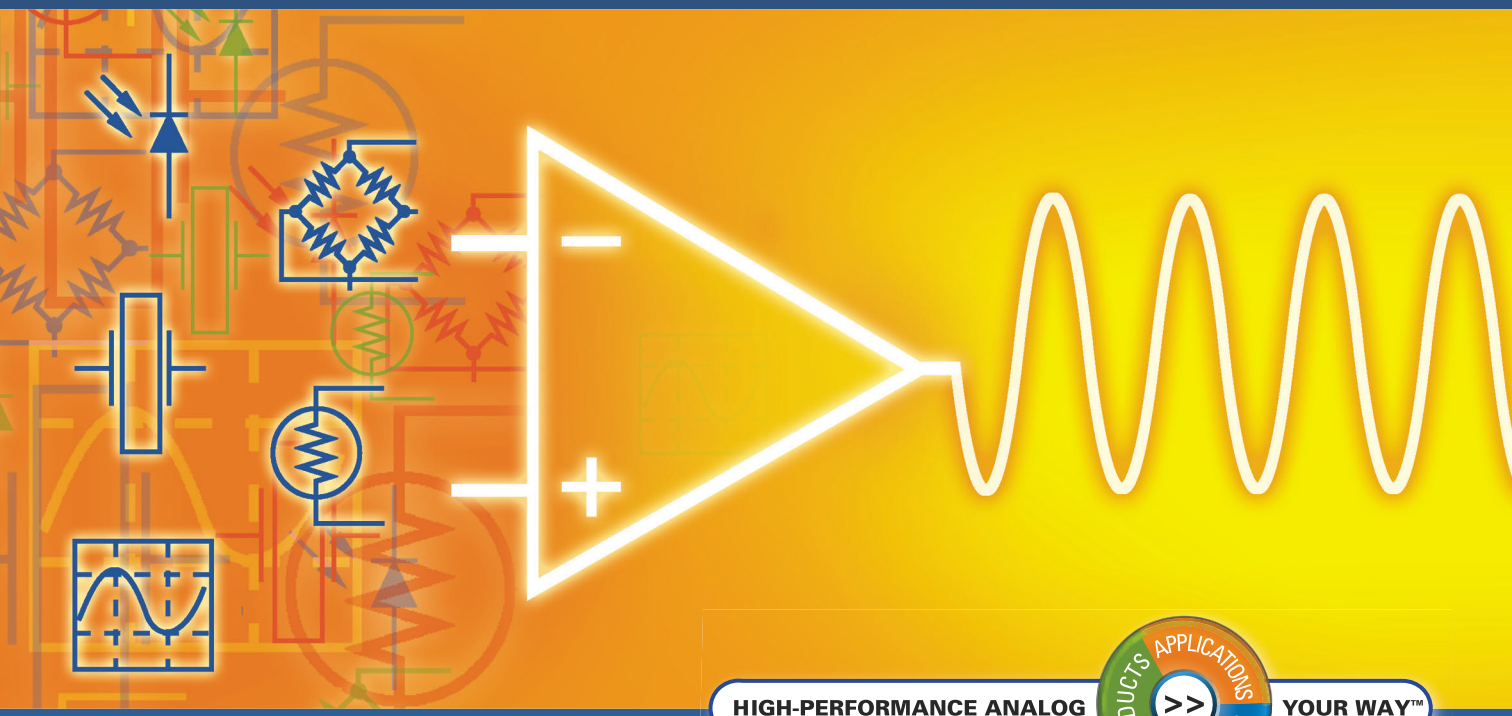

HIGH-PERFORMANCE ANALOG >> YOUR WAY™

PRODUCTS APPLICATIONS SUPPORT

TI's comprehensive portfolio of high-precision op amps features low power, low noise, high speed, low voltage and small packages. Whatever your op amp challenges, TI has the products to satisfy a broad range of applications. That's High-Performance Analog >> Your Way.

OPA211	**Lowest Power, 1.1nV/√Hz, 36V** Op Amp: 3.6mA/ch, bipolar, 80MHz, 1µV/°C drift
OPA827	**Lowest Noise, JFET, 36V** Op Amp: 18MHz, 4.5nV/√Hz at 1kHz, 250µV V_{OS}, 1µV/°C drift
OPA365	**Zero-Crossover, RRIO** Op Amp: 50MHz, 4.5nV/√Hz, 0.0004% THD+N, 100dB CMRR
OPA333	**Lowest Power, Zero-Drift** Op Amp: 17µA I_Q, 2µV V_{OS} 0.02µV/°C drift
OPA727	**eTrim™ CMOS** Op Amp: 20MHz, 15µV V_{OS}, 0.6µV/°C drift, +4V to +12V
OPA355	**Fastest CMOS** Op Amp: 200MHz, 360V/µs slew rate, RRO within 100mV
OPA569	**Best Swing-to-Rail Power** Op Amp: 2.4A I_O, 150mV output to rail, +2.7V to +5.5V
OPA656	**High-Speed** Op Amp: 230MHz GBW, 100µV V_{OS}, 2µV/°C drift, 2pA I_B
INA326	**Precision, Single-Supply Instrumentation** Amp: 20µV V_{OS}, 0.1µV/°C drift

Ease the selection process with samples, online simulation and filter design tools >>

www.ti.com/precisionamps

1.800.477.8924 ext.1102

Technology for Innovators™

Texas Instruments

Power Digital ICs with High-Efficiency LDOs

Low-Input LDOs Deliver 80% Efficiency with Less Than 100 mV Load Transient Response

Typical Application Circuit

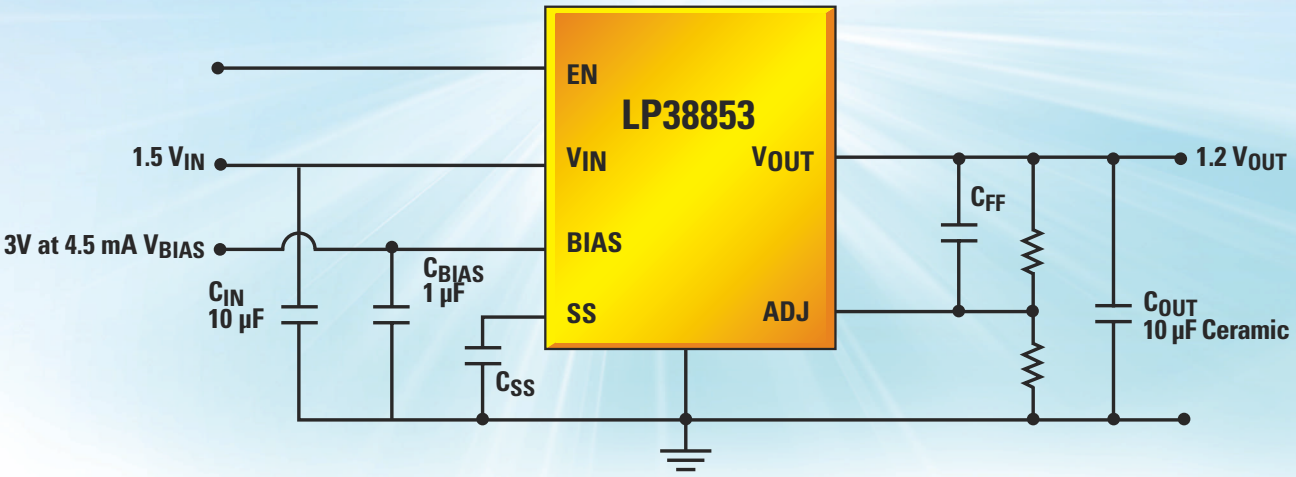

LDOs for Powering Digital ICs

Product ID	Load Current (mA)	V_{OUT} Range (V)	V_{IN} Range (V)	Features
LP38853/56/59	3000	0.8 to 1.2	1.3 to 5.5	Enable, soft start, adj. output
LP38852/55/58	1500	0.8 to 1.2	1.0 to 5.5	Enable, soft start, adj. output
LP38851	800	0.8 to 1.2	1.0 to 5.5	Enable, soft start, adj. output
LP5952	350	0.5 to 2	0.7 to 4.5	Tiny micro SMD-5 packages
LP5951	150	1.3 to 3.3	1.8 to 5.5	Small SOT23-6 and SC70-5 package
LP3991	300	0.8 to 2.8	1.65 to 3.3	1% output accuracy

Ideal for powering DSPs, FPGAs, digital cores, graphics processors, and post regulation

For FREE samples, datasheets, and more information on LDOs contact us today at:

ldo.national.com

Or call: 1-800-272-9959

National Semiconductor

The Sight & Sound of Information

EDN

contents

6.21.07

The evolution of SERDES designs for cell phones

Exploring memory architectures: pillars of processing performance

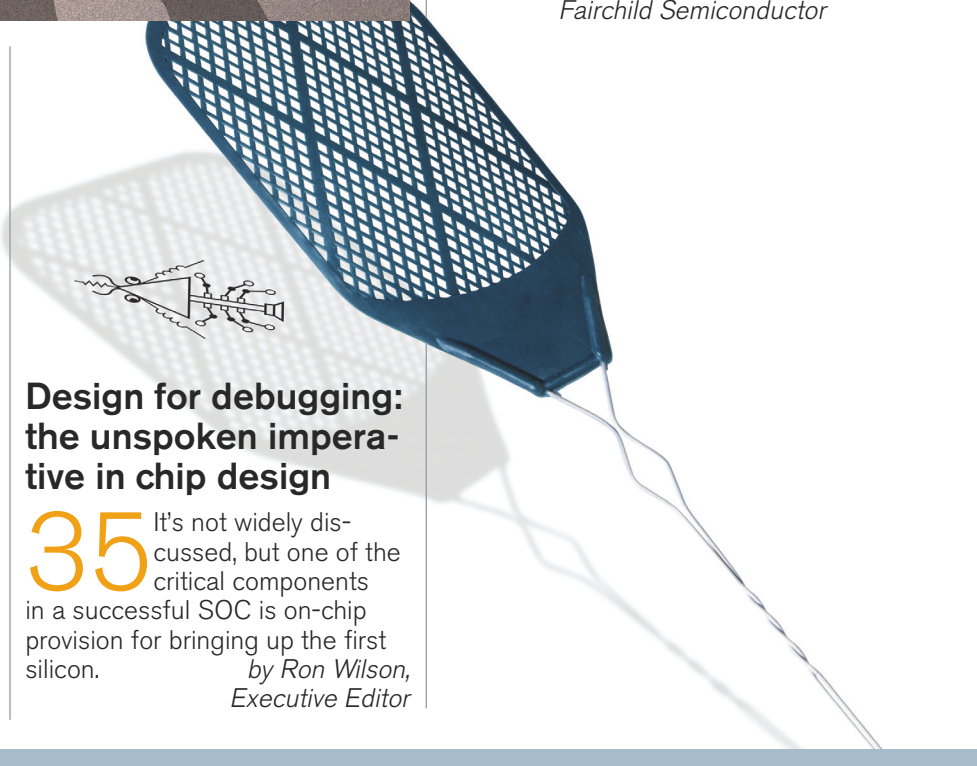

Design for debugging: the unspoken imperative in chip design

DESIGNIDEAS

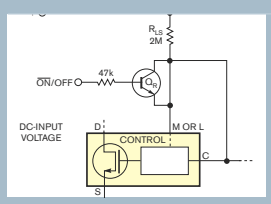

Industry-Leading Voltage-Level Translation ICs

Auto-Direction-Sensing Translators

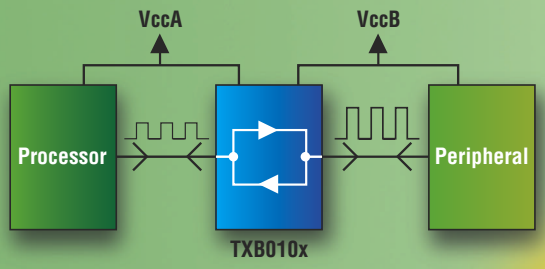

Application-Specific Translators

(Memory Card Interface)

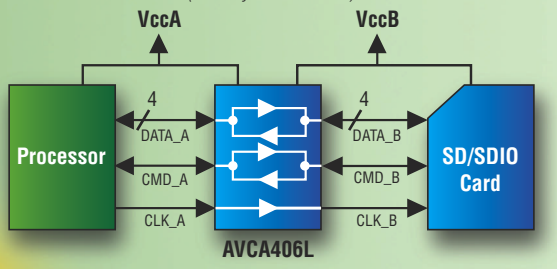

Auto-Direction-Sensing Translators for Open-Drain Applications

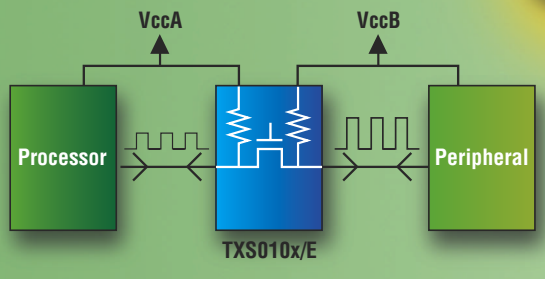

Dual-Supply-Configurable Translators

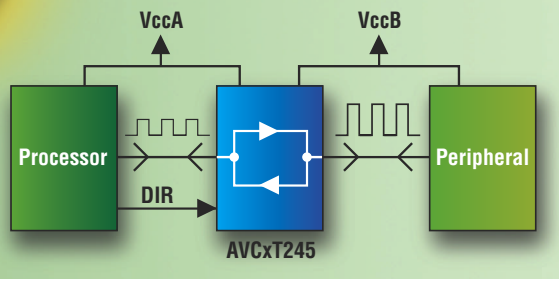

TI Has the Translation Device for You.

Texas Instruments offers a wide-ranging selection of voltage-level translation devices from auto-direction-sensing and dual-supply-configurable devices to application-specific memory card interface translators. When mixed-voltage issues become a concern in your design, TI has the translation device that fits your requirements.

Device	Bit Width	VccA (V)	VccB (V)	ESD Protection	Smallest Package
Auto-Direction Sensing Translators*					
TXS0102	2	1.65 to 3.6	2.3 to 5.5	8-kV HBM	8-bump WCSP
TXS0104E	4	1.65 to 3.6	2.3 to 5.5	10-kV Air Gap Discharge (IEC)	12-ball BGA
TXB0104	4	1.2 to 3.6	1.65 to 5.5	15-kV HBM	12-ball BGA
TXB0108	8	1.2 to 3.6	1.65 to 5.5	15-kV HBM	20-ball BGA
Application-Specific Translators (Memory Card Interface)					
AVCA406		1.4 to 3.6	1.4 to 3.6	15-kV Air Gap Discharge (IEC)	48-ball VFBGA
AVCA406L		1.2 to 3.6	1.2 to 3.6	6-kV HBM	20-ball VFBGA
CF4320H		1.65 to 3.3	3.3 to 5	15-kV HBM	114-ball LFBGA
Dual-Supply Configurable Translators*†					
SN74AVC1T45	1	1.2 to 3.6	1.2 to 3.6	2-kV HBM	6-bump WCSP
SN74AVC2T45	2	1.2 to 3.6	1.2 to 3.6	8-kV HBM	8-bump WCSP
SN74AVC8T245	8	1.2 to 3.6	1.2 to 3.6	8-kV HBM	24-pin QFN
SN74AVC16T245	16	1.2 to 3.6	1.2 to 3.6	8-kV HBM	56-ball VFBGA
SN74AVC32T245	32	1.2 to 3.6	1.2 to 3.6	8-kV HBM	96-ball LFBGA

** Other bit-width options available. † Bus Hold feature available.*

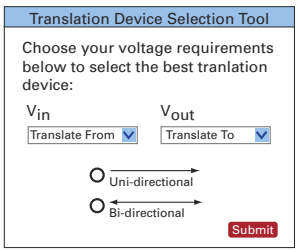

Finding the right translation solution is easy with **TI's Translation Selection Tool**.

www.ti.com/trans1

800.477.8924, ext 365

Technology for Innovators™

Texas Instruments

contents 6.21.07

pulse

Dilbert 22

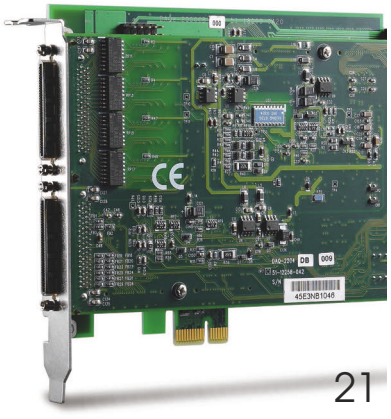

21

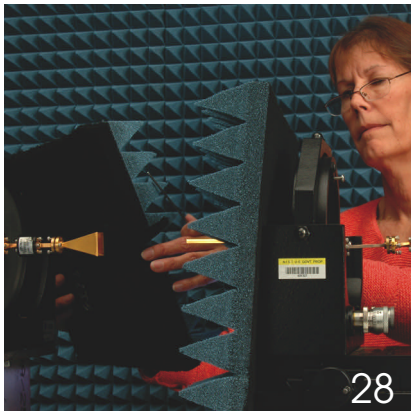

28

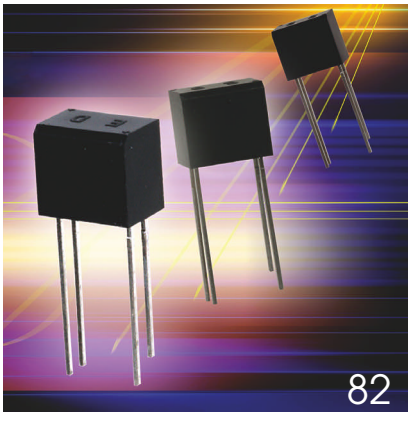

82

DEPARTMENTS & COLUMNS

PRODUCT ROUNDUP

EDN® (ISSN#0012-7515), (GST#123397457) is published biweekly, 26 times per year, by Reed Business Information, 8878 Barrons Blvd, Highlands Ranch, CO 80129-2345. Reed Business Information, a division of Reed Elsevier Inc, is located at 360 Park Avenue South, New York, NY 10010. Tad Smith, Chief Executive Officer. Periodicals postage paid at Littleton, CO 80126 and additional mailing offices. Circulation records are maintained at Reed Business Information, 8878 S Barrons Blvd, Highlands Ranch, CO 80129-2345. Telephone (303) 470-4445. POSTMASTER: Send address changes to EDN®, PO Box 7500, Highlands Ranch, CO 80163-7500.

Rates for nonqualified subscriptions, including all issues: US, $165 one year; Canada, $226 one year, (includes 7% GST, GST#123397457); Mexico, $215 one year; air expedited, $398 one year. Except for special issues where price changes are indicated, single copies are available for $10 US and $15 foreign. Publications Agreement No. 40685520. Return undeliverable Canadian addresses to: RCS International, Box 697 STN A, Windsor Ontario N9A 6N4. E-mail: Subsmail@ReedBusiness.com. Please address all subscription mail to EDN®, 8878 S Barrons Blvd, Highlands Ranch, CO 80129-2345. EDN® is a registered trademark of Reed Elsevier Properties Inc, used under license. A Reed Business Information Publication/Volume 52, Number 13 (Printed in USA).

LeCroy invites you to unusual events.

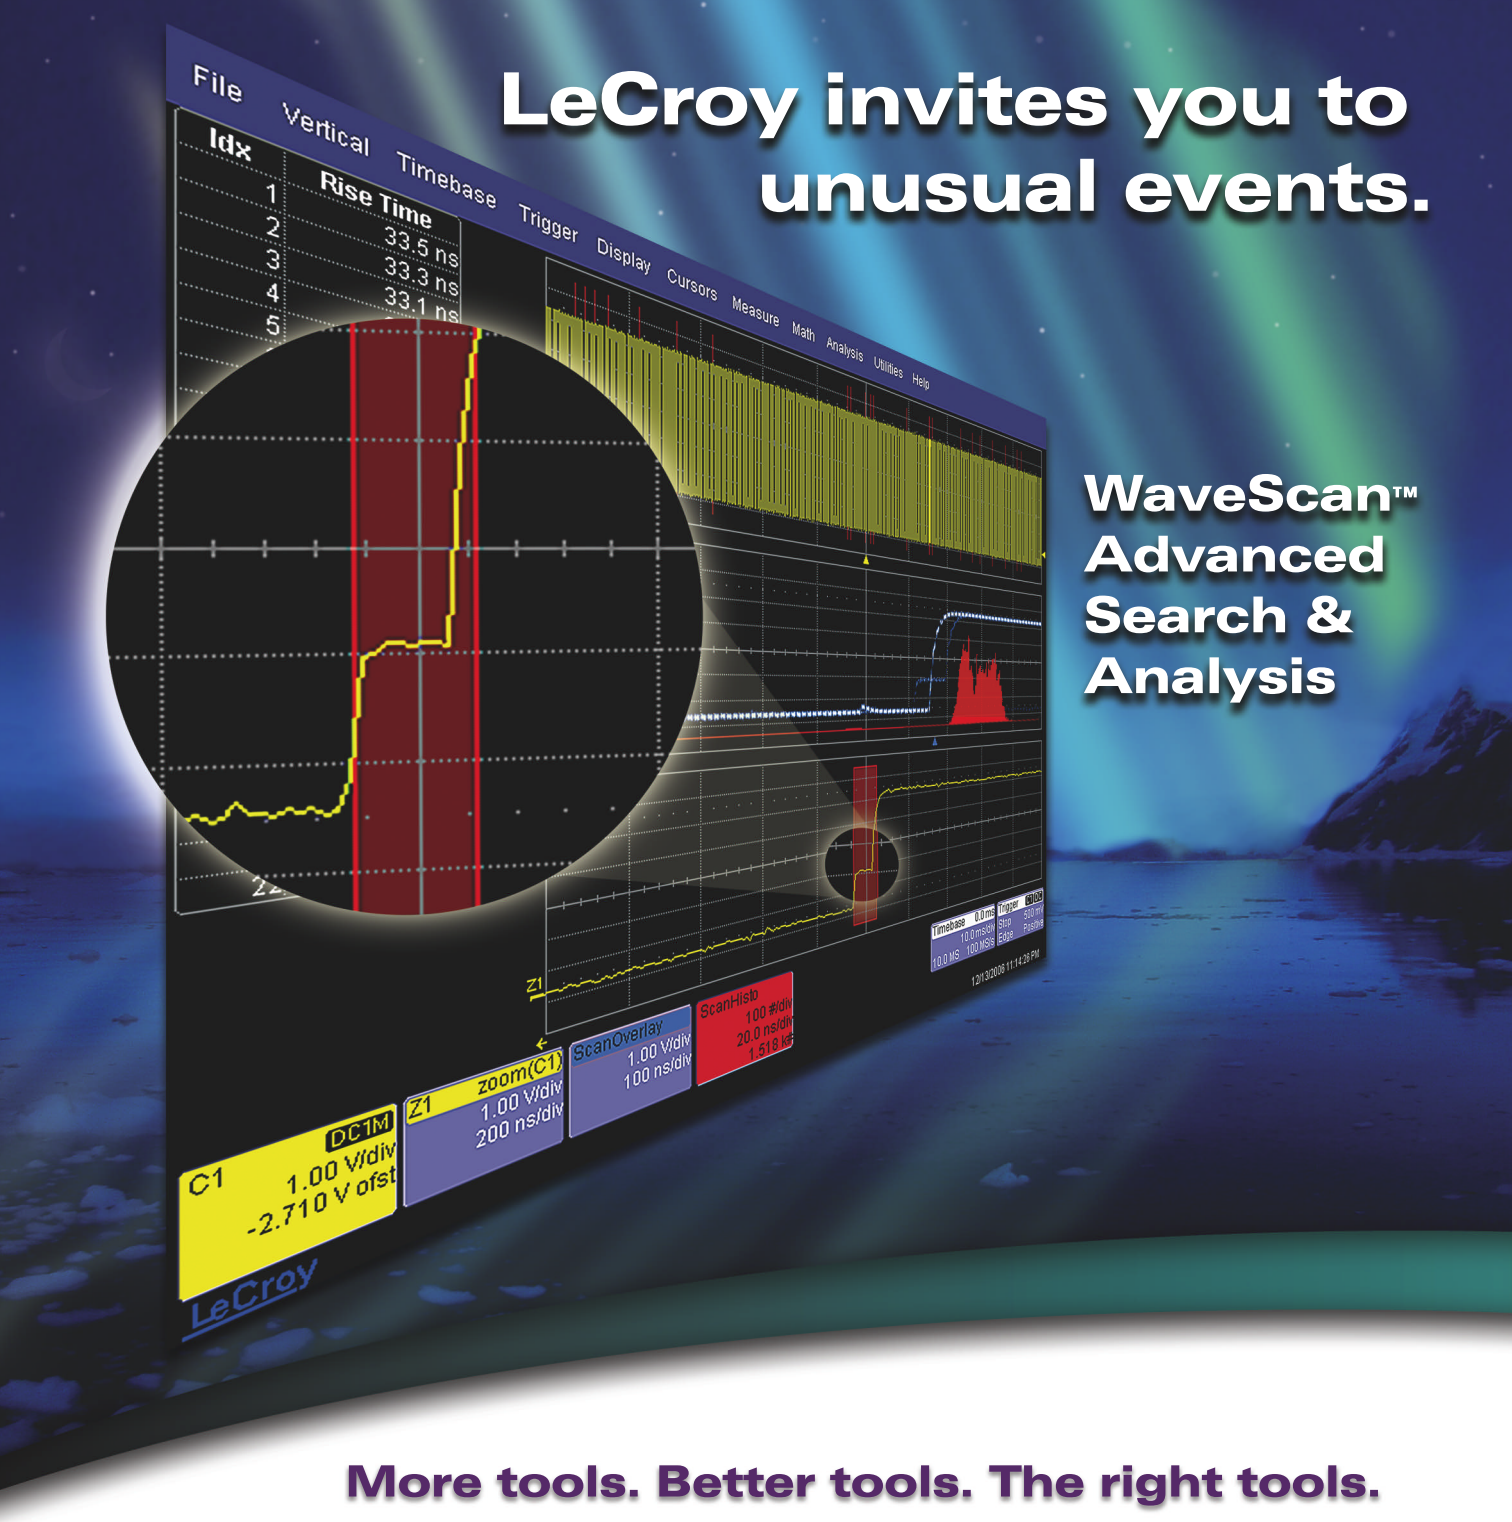

WaveScan™ Advanced Search & Analysis

More tools. Better tools. The right tools.

WaveScan, a powerful new tool only for LeCroy oscilloscopes, continuously monitors live acquisitions for unusual events – even if you're not there. Select from more than 20 search criteria, including frequency, rise time, runt, and duty cycle. WaveScan also locates rare occurrences in a single capture, or even a saved waveform and marks events for quick identification. Debug faster. It's that simple.

EDN innovation 2007 FINALIST

EDN HOT 100 products WINNER

EDN Magazine has included LeCroy's WaveRunner® Xi and WaveSurfer® Xs with WaveScan in it's 'Hot 100 Products' list. WaveScan is also an EDN 2007 Innovation Award Finalist.

www.lecroy.com
Keyword: unusualevents
1-800-5-LeCroy

LeCroy

EDN online contents

www.edn.com

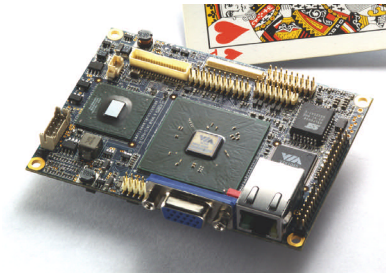

ONLINE ONLY

Check out these Web-exclusive articles:

VIA defines Pico-ITX form factor
→www.edn.com/070621toc7

Interconnect intellectual property moves in chip-architecture circles
→www.edn.com/070621toc3

See an excellent video of a bendy organic display
→www.edn.com/070621toc4

Brion's Tachyon box offers virtual tapeout
→www.edn.com/article/CA6447806

Lattice offers 90-nm, flash-based FPGA
→www.edn.com/article/CA6445681

Demos on Demand posts tool demos for low-power design
→www.edn.com/070621toc1

Regularity intensified: Put some fiber in your network diet
→www.edn.com/070621toc2

Rave review for free PCB-layout manual
→www.edn.com/070621toc5

Sierra releases IC place-and-route tools for 45-nm era
→www.edn.com/article/CA6446622

Will DFM become part of the semiconductor-manufacturing-equipment ecosystem?
→www.edn.com/070621toc6

READERS' CHOICE

A selection of recent articles receiving high traffic on www.edn.com.

Taking a bite out of power: techniques for low-power-ASIC design
→www.edn.com/article/CA6442429

Five cool things you can do with an accelerometer
→www.edn.com/070621toc9

Accurate USB 2.0 temperature sensor needs only a handful of parts
→www.edn.com/article/CA6442433

In-between spaces: Understanding inductance in high-speed digital circuits
→www.edn.com/article/CA6442436

Comparator detects position of peaks and valleys in a waveform
→www.edn.com/article/CA6442431

Intel, AMD diverge on multicore strategies
→www.edn.com/article/CA6445861

Lithium-ion batteries power next generation of electric vehicles
→www.edn.com/article/CA6445393

Intel, ST form independent flash company
→www.edn.com/article/CA6445428

Micron, Blackstone, China and the DRAM world
→www.edn.com/070621toc8

Audiocast: Intel pans IBM's air-gap technology
→www.edn.com/article/CA6445842

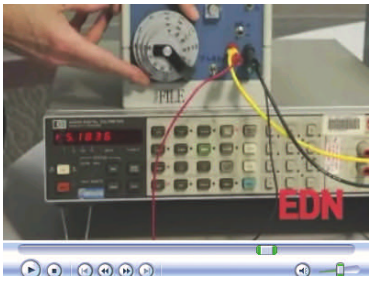

EDN TECH CLIPS

Video Design Idea: Build your own laboratory precision voltage reference
Mark Thoren, mixed-signal-application engineering manager with Linear Technology, demonstrates an amplifier-based circuit design for a relatively inexpensive precision-voltage source.

VITA 46: A new direction in board-level design
Learn about this emerging VMEbus-replacement technology in a video produced for *EDN* by a core member of the VITA 46 specification-development team.

Video Design Idea: Measure nanoamps to ensure accurate computer clocks
Jim Williams, staff scientist with Linear Technology, explains why PC clocks are invariably wrong and how engineers can surmount the extreme measurement challenge involved in solving the problem.

→**www.edn.com/techclips**

PART SEARCH

EDN's Part Search, powered by SupplyFrame, provides the latest production prices, products specs, availability information, and lead times. Use this new tool for cross-referencing, getting instant quotes, and connecting quickly with authorized distributors.
→www.edn.com/partsearch

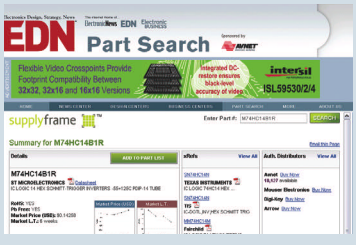

Flexible, easy-to-use true-differential VNAs eliminate guesswork

By Jochen Simon

As increasing numbers of RF devices depend on integrated-circuit technology, RF designs' use of differential topologies is becoming more and more commonplace. Moreover, once a project team decides to base a product on a custom RFIC (radio-frequency integrated circuit), it often becomes a "no-brainer" to have the IC perform circuit functions that, in an earlier era, would have been handled by separate passive components. In general, IC implementations of these functions use active, rather than passive, circuit elements. The result of this RFIC revolution is thus an explosion in the use of active as well as differential (also known as balanced) circuits at radio frequencies.

This situation presents a challenge for engineers who must accurately characterize RF devices. The mathematics that underlies an understanding of device characteristics assumes that the DUT (device under test) is linear. But, whereas passive devices sometimes exhibit nonlinearity, active devices are much more likely to be nonlinear—at least to a degree that calls into question the validity of parametric values derived through assumptions of linearity. In addition, characterizing nonlinear devices often requires larger numbers of more difficult and time-consuming measurements than does characterizing linear devices.

The assumption of linearity also forms the basis for the mathematics by which engineers have traditionally extended inherently single-ended measurements to differential devices [1]. You might say, then, that the RFIC revolution and the resulting proliferation of nonlinear, differential circuit structures has mounted a multi-pronged attack on the classical methods of RF-component characterization and on the architecture of VNAs (vector network analyzers), the instruments at the heart of most highly accurate, parametric RF measurements. Until now, nearly all such instruments have, as a practical matter, been limited to fundamentally single-ended measurements; differential results were derived virtually, which means mathematically by a linear superposition of single-ended measurements.

Although true differential VNA measurements were possible through the use of baluns and hybrids [2], or by adding two signal sources to a conventional VNA [3], these techniques were limited by the lack of components with sufficiently ideal characteristics over a wide bandwidth. Furthermore, ready-to-use systems with true differential-measurement capability were not commercially available.

Figure 1: The Rohde & Schwarz ZVA-series of 300-kHz to 40-GHz RF vector network analyzers offer a true-differential source and measurement option that takes the guesswork out of characterizing active differential components, including structures common in modern RFICs.

R&S ZVA to the rescue

With its R&S ZVA series of 300-kHz to 40-GHz RF VNAs (Figure 1), which offer an optional TruDi (true-differential) measurement mode, leading VNA supplier Rohde & Schwarz is coming to the rescue of RF design engineers around the world who need new tools to deal with the RFIC revolution's new measurement challenges.

TruDi operation of the R&S ZVA is based on a second signal source that is installed in instrument versions with four test ports. This source can produce a signal whose amplitude is equal to and whose phase is opposite that of the signal from the first source. The R&S ZVA can also produce signal pairs that have adjustable relative amplitude and phase, thus enabling the simulation of real-world signals that are not perfectly balanced—that is, signals that contain common-mode as well as normal-mode (differential mode) components or even purely common-mode components—with no differential components. It is also possible to sweep the phase difference and magnitude imbalance to quickly reveal how a device responds to signal imperfections. No such measurements were even theoretically possible through the use of baluns and hybrids.

Going deeper into the details of the new R&S ZVA TruDi option, it is noteworthy that the user-defined amplitude and phase relationship of the signals refers to the VNA's outgoing single-ended waves, which serve as incoming (a-) waves for the DUT. This relationship is measured with full-system error correction applied at each point of each sweep. According to this measurement, signal amplitude and phase are re-adjusted before the desired measured quantity is acquired, thus taking time-varying DUT reflection into account. The error terms that are used during wave correction are derived from a conventional full n-port calibration, which also defines the reference plane where the user-defined

amplitude and phase relationship is valid. DUT topology can be defined with a high degree of flexibility, which means that any two single-ended ports can be combined to form a balanced port. Because one internal R&S ZVA source serves two single-ended test ports (ports 1 and 2 belong to one R&S ZVA source; ports 3 and 4 to the other), the only limitation for a balanced port is the requirement that the constituent single-ended ports have different sources.

An R&S ZVA with the TruDi option supports three operating modes, which are listed below in order of increasing complexity:

1. Source of fully error-corrected differential and common-mode outgoing waves of amplitude and phase defined with respect to a reference plane that is established by conventional n-port system-error calibration. A user can shift the reference plane by defining and activating a transmission-line offset, which may be lossy. You can also sweep the amplitude and phase relationships between the single-ended signals and calibrate the signal power.
2. Source function as in 1, plus fully error-corrected measurement of all single-ended or balanced-wave quantities defined in the DUT topology. After a conventional (single-ended) power calibration, you can even determine the amplitudes of differential and common-mode wave quantities.
3. Source function as in 1, plus fully error-corrected measurement of all mixed-mode S-parameters defined for the DUT topology. For this measurement, you stimulate each balanced port in two modes (differential and common-mode) and each single-ended port in one mode.

A mouse-click switches modes

Switching between true-differential and virtual-differential (that is, mathematical superposition) measurement modes requires no more than a mouse-click. After you set up two measurement channels, you can even measure the same DUT in both modes at the same time. A comparison of the true and virtual differential modes appears in Figure 2, which shows the transmission of a SAW filter, a representative passive linear device. As you might expect, there is no noticeable difference in measurement results between modes.

As an additional measurement example, Figure 3 shows the gain of a differential low-noise amplifier (LNA) for WCDMA applications in a power sweep. At small source-signal levels, hardly any difference exists between the two modes. With increasing source power, however, the amplifier reaches 1-dB compression at an input power of about -10.5 dBm compared with about -6.5 dBm in the virtual mode. For other differential devices, you can also observe the opposite behavior—that is, earlier compression in virtual mode.

In Figure 3, the stimulus axis of the true-differential trace has been shifted by 3 dB with respect to the axis of the virtual mode. This shift is necessary for a comparison of the two modes because the differential-stimulus amplitudes must be the same. In virtual-differential mode, the R&S ZVA power setting refers to a single-ended source signal, whereas in TruDi mode, the setting relates to a differential-source signal. Because half of the power of the single-ended source signal is contained in its unwanted common-mode component, the power must be increased by 3 dB.

The new R&S ZVA TruDi option provides system designers with a valuable tool that allows them to measure—rather than guess at—key characteristics of differential devices. Above all, designers can investigate the large-signal behavior of these devices under real operating conditions. Only now have such measurements become possible.

References:

[1] D. E. Bockelman, W. R. Eisenstadt: Combined Differential and Common-Mode Scattering Parameters: Theory and Simulation. IEEE Trans. MTT, Vol. 43, No. 7, Jul. 1995, pp. 1530-1539

[2] D. E. Bockelman, W. R. Eisenstadt: Pure-Mode Network Analyzer for On-Wafer Measurements of Mixed-Mode S-Parameters of Differential Circuits. IEEE Trans. MTT, Vol. 45, No. 7, Jul. 1997, pp. 1071-1077

[3] J. Dunsmore: New Methods & Non-Linear Measurements for Active Differential Devices, IEEE MTT-S Digest 2003, Vol. 3, pp. 1655-1658

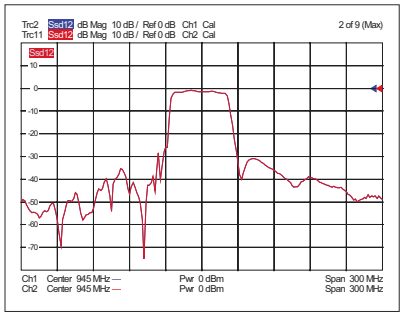

Figure 2: You can't see the transmission of a SAW filter measured in virtual-differential mode (blue trace) because the true-differential-mode measurement (red trace) completely coincides with it.

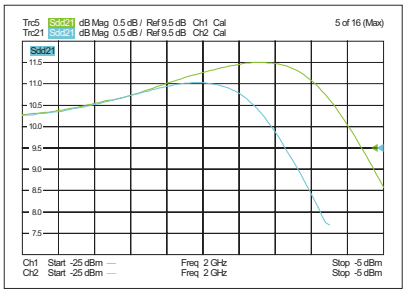

Figure 3: In this power sweep of differential-amplifier gain, hardly any difference exists between the virtual and TruDi mode results at low power, but at higher source power, the amplifier reaches 1-dB of compression at about -10.5 dBm in TruDi mode versus about -6.5 dBm in virtual mode.

Author´s Biography

Jochen Simon received his Dipl.-Ing. and Dr.-Ing. degrees from the Technical University of Munich. Since 1993, he has been involved in hardware and software development of vector network analyzers at Rohde & Schwarz, Munich, Germany.

For more information visit: http://www.rohde-schwarz.com/product/ZVA.html

BY MAURY WRIGHT, EDITORIAL DIRECTOR

DAC stumbles looking to find firm future

The June 2007 rendition of DAC (Design Automation Conference) in San Diego again revealed an uncertain direction for the chip-design community. Moreover, organizers appear to be looking for some way to stop the trend of a shrinking exhibition and attendee list. This year, the organizers in part tried to promote an automotive theme, although the effort has come off flat at best. For sure, the conference continues to offer valuable education to chip designers. But it's the chip designers—especially in leading-edge processes—that are shrinking in number and certainly working for different types of companies from those they worked for just a few years ago.

EDN and TSMC hosted the EDAC (EDA Consortium) executive reception and panel on Sunday to kick off the show. That panel, which *EDN* Executive Editor Ron Wilson and *Electronic News/Electronic Business* Editor in Chief Ed Sperling co-moderated, pitted a group of senior IC-design managers against a group of marketing executives from the fab and EDA market. Wilson summarized some of the key points from the panel in his blog (www.edn.com/070621ed2). I'd like to add a few of my own thoughts on that panel and relate them to what I see as another lackluster DAC.

First and foremost, I think you can draw some conclusions from the chip-designer side of the panel. Major fabless-IC vendors Broadcom and Qualcomm were represented, but companies such as Cisco, Apple, and Nokia were not on the panel. ASIC-design house eSilicon was represented. But there were none of the ASIC designers working for end-product companies that would have been front and center on such a panel just a few years ago.

Now, I've been quick to claim in the past that the shrinking number of ASIC starts was not as bad as advertised. I still think the analysts that track those statistics fail to account for the ASSPs that are increasingly designed by the likes of Qualcomm. I'm sure that Cisco is still doing a fair amount of ASICs. But in what technology is Cisco working, and how does that align with the presumed hot issues at DAC?

The major issue at DAC this year continues to be how to verify SOCs (systems on chips) built in the latest process technologies. The panel clearly indicated that Broadcom and Qualcomm care about that. But who else cares? The list is short. Intel cares but uses internal tools. Memory vendors need the latest process technology, but they don't need to fully verify memory blocks with ECC and other technologies that correct errors.

Ironically, I'm not sure that the EDA vendors really care. Mentor has certainly done well with verification tools. But reading between the lines at the Sunday panel, tool vendors may not believe that there's sufficient interest—read "customers"—to innovate the system-level approach that's really needed in state-of-the-art chip design. Now, TSMC and other fabs clearly need the fabless chip vendors to fill their factories, and, certainly TSMC, among other fabs, has deepened its level of support for leading-edge chip designs.

Meanwhile, I suspect that the likes of Cisco are churning right along at less ambitious process nodes. The big communication-equipment vendors can't fit everything into one chip anyway, and they are relying on off-the-shelf ASSPs in some cases. Apple likely has no interest in designing an SOC for, say, the iPod. It wants to churn the technology twice a year anyway to capture the upgrade market as well as new customers. The company never sells enough of a single model to justify a custom SOC.

In short, this DAC has been more of the same. As I wrote last year in "DAC disappoints: Maturing segment needs a jumpstart" (www.edn.com/article/CA6356786), there have been no real significant innovations in EDA. The verification story is getting stale. In his blog, Wilson suggests that preverified IP (intellectual property) could perhaps solve the problem. Alas: The news isn't necessarily good there, either. At the end of a Monday panel entitled "Just who is providing the IP?" longtime EDA-industry follower and moderator Peggy Aycinena pessimistically noted that the panel on IP was almost exactly like one held five years ago, and that observation clearly calls into question industry progress on IP reuse. EDN

Contact me at mgwright@edn.com.

MORE AT EDN.COM

Go to www.edn.com/070621ed1 and click on Feedback Loop to post a comment on this column.

FAIRCHILD
SEMICONDUCTOR®

Greater efficiency, less space. Bright ideas in lighting ICs.

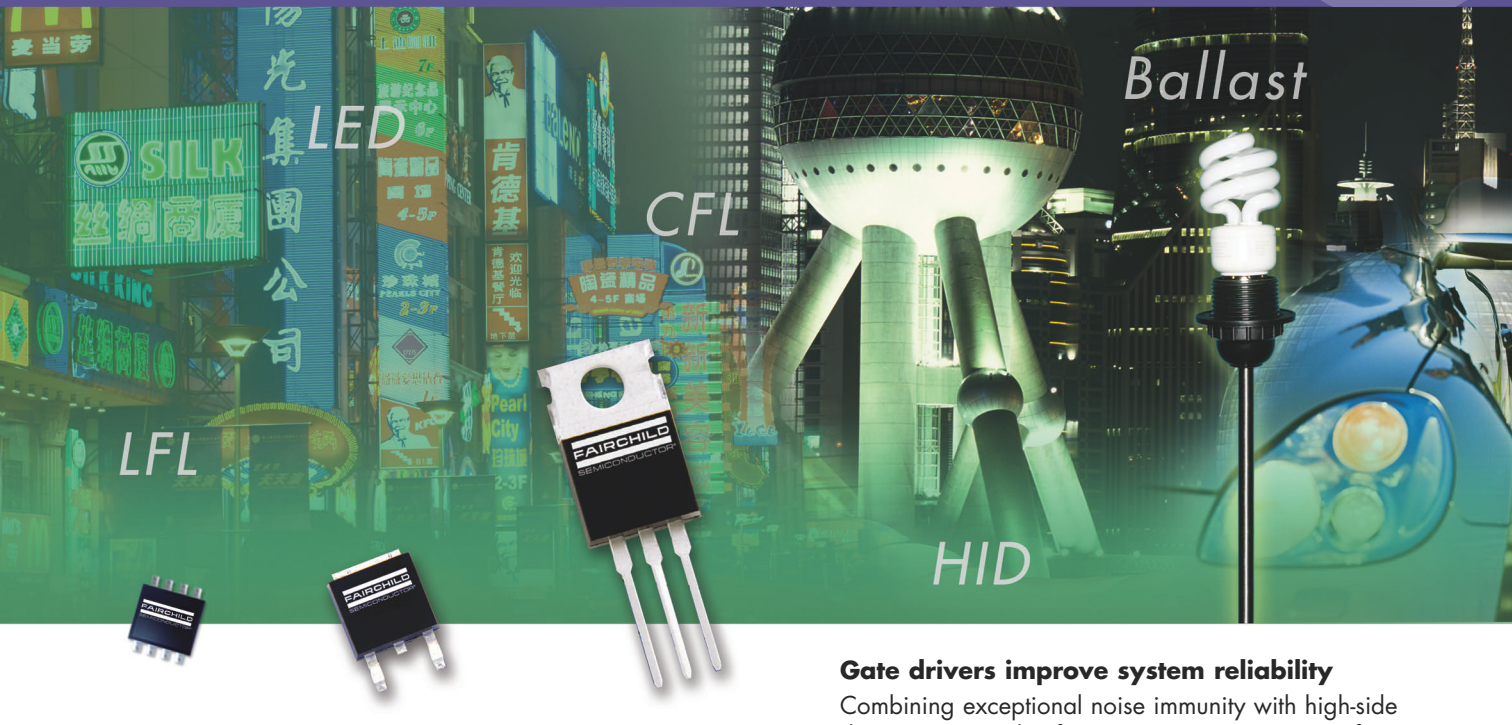

Integrated ballast control ICs simplify designs

Now you can eliminate components and complexity in ballast designs. For example, our new FAN7711 features all the optimum power functions as well as circuitry protection required for CFL and LFL applications. Fairchild gives you more functionality, in less space, with improved reliability—combined.

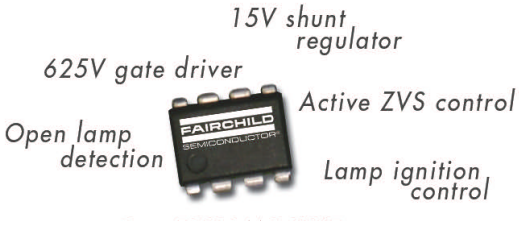

FAN7711 Electronic Ballast Control IC

Gate drivers improve system reliability

Combining exceptional noise immunity with high-side driver operation that features negative V_S swings of –9.8V (V_{BS} = 15V) to protect against negative noise, Fairchild's HVICs enhance system reliability and minimize power consumption.

High-performance SuperFET™ MOSFETs

SuperFET MOSFETs minimize conduction and switching losses to help improve overall system efficiency. Their combined low on-resistance, DPAK packaging and ability to withstand high speed switching transients make them the ideal MOSFETs for ultra-slim ballast designs.

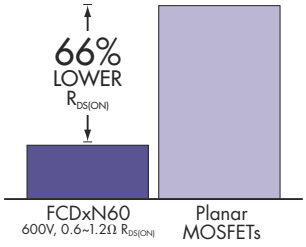

To learn more about our complete lighting solutions—including PFC ICs, IGBTs and transistors—please visit www.fairchildsemi.com/lighting.

FAIRCHILD SEMICONDUCTOR CELEBRATING 50 10 YEARS

www.fairchildsemi.com

the power franchise™

Buyers & Sellers:
Meet Your Match

BuyerZone is the leading online marketplace for business purchasing. Millions of buyers use BuyerZone's free Request for Quotes service to quickly get multiple bids from top vendors who can meet their needs. Sellers get connected to prospective new customers with near-term buying plans. It's a win-win situation – resulting in billions of dollars in purchasing transactions each year.

Try BuyerZone now. It's the smart move for your business!

Visit www.BuyerZone.com
or call 1.800.601.5224

BuyerZone's Request for Quotes service is available for more than 125 purchase categories, including:

- Office Equipment
- Marketing Services
- HR & Benefits
- Telecommunications
- Office Furniture
- Financing Solutions
- Information Technology
- Retail & Restaurant
- Industrial
- Construction

…and more

BuyerZone
Where Smart Businesses Buy and Sell

A division of
Reed Business Information.

www.BuyerZone.com

ANALOG edge℠

Powering and Dimming High-Brightness LEDs with the LM3405 Constant-Current Buck Regulator

Application Note AN-1644

Vinit Jayaraj, Principal Design Engineer

The LM3405 is a 1A, constant-current buck regulator designed to provide a simple, highly-efficient solution for powering LEDs with the highest power density. The low-feedback reference voltage of 205 mV (typical) allows the use of LEDs with large forward voltages. The LM3405 LED driver comes with an integrated 1A high-side power switch, internal current limit, over-voltage protection and thermal shutdown. The LM3405 driver utilizes current-mode control with internal compensation offering ease-of-use and predictable, high-performance regulation over a wide range of operating conditions. This article describes a few of the many application solutions for the LM3405 LED driver. The LM3405 device is offered in a 6-pin Thin-SOT (TSOT) package and operates at an internal switching frequency of 1.6 MHz, allowing the use of small-value inductors, thereby saving board space.

Typical Application Circuit

Figure 1 shows the typical application circuit using the LM3405 regulator driving a single LED. The boost voltage can be derived from V_{IN} or V_{OUT}. The BOOST to SW voltage (VBOOST-VSW) should never exceed the operating limit of 5.5V, and must be greater than 2.5V for optimum efficiency.

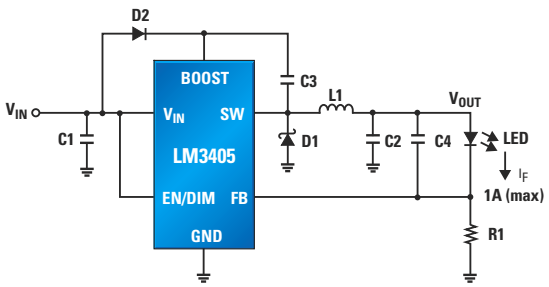

Figure 1. Typical Application Circuit

In the above circuit, the maximum V_{IN} is 5V and therefore the boost voltage is derived from V_{IN}. The LM3405 data sheet provides a detailed description of various circuits for boost-voltage generation. The LM3405 driver tightly regulates the FB-pin voltage to 205 mV (typical) and this allows the LED current (I_F) to be set by the following equation:

$$I_F = \frac{205\ mV}{R1}$$

R1 (min) is 200 mΩ such that I_F (max) is limited to 1A. In case of an over current, the internal current limit will trigger and turn off the internal power switch on a cycle-by-cycle basis. In case of an over voltage (sensed at the FB pin), the internal power switch is turned off. The LM3405 regulator also has a built-in thermal shutdown (with hysteresis) which also turns off the internal power switch if the die temperature exceeds a typical value of 165°C.

PWM Dimming Using the EN/DIM Pin

The LEDs can be dimmed by applying a PWM-logic signal to the EN/DIM pin of the LM3405 driver as shown in *Figure 2*.

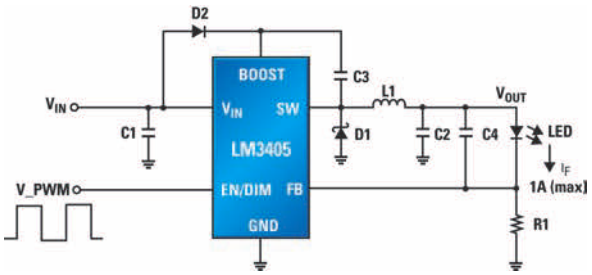

Figure 2. PWM Dimming of LEDs by Applying a PWM Signal to the EN/DIM Pin

A logic high at V_PWM enables the IC, and a logic low disables the IC. In this manner, the LED current is turned on and off. In order to eliminate flicker, the lowest PWM-dimming frequency is normally chosen to be above 100 Hz. The upper end of the PWM frequency is determined by the turn-on delay of the LM3405. If the V_PWM signal has an ON-time of T_{ON} and a period of T_{PER}, the average I_F is given approximately by the equation below. Note that the delay in enabling and disabling the IC is not included in the following equation. Typically the time delay from the instant where EN/DIM=1 to the instant where I_F is fully established is approximately 100 µs (see *Figure 4*).

$$I_F\ (average) = \frac{205\ mV}{R1} \times \left[\frac{T_{ON}}{T_{PER}}\right]$$

National Semiconductor
The Sight & Sound of Information

The average LED current is therefore controlled by T_{ON}, T_{PER}, or both and is perceived by the eye as a brightness change. ***Figure 3*** shows the measured average LED current for varying duty cycles and frequencies.

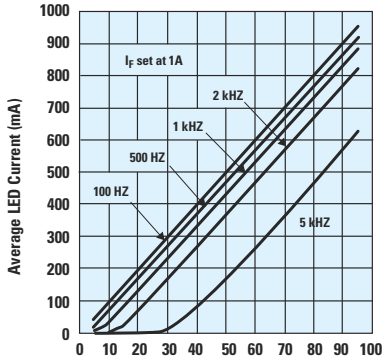

Figure 3. Average LED Current versus Duty Cycle and Frequency of V_PWM at the EN/DIM Pin

The startup delay of the LM3405 device is shown in ***Figure 4*** for an LED-current setting of 1A. This is the delay from the instant when EN/DIM=1 to the instant when I_F settles to 900 mA (90% of the set LED current).

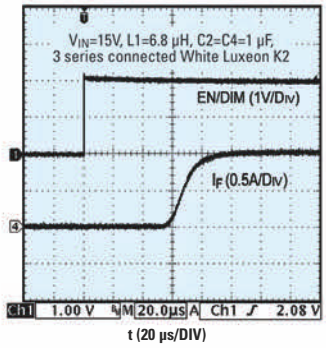

Figure 4. Startup Response of the LM3405 to an EN/DIM Signal with I_F Set at 1A

Deriving a Self-Biased Boost Voltage When Driving Two or More LEDs in Series

In a typical application where two or more LEDs are driven in a single string, the boost voltage is derived from a separate, external low-voltage source so as to meet the (VBOOST-VSW) voltage requirement. This is because when powering two or more LEDs in series, V_{IN} will be greater than 5V and therefore V_{IN} cannot be used to provide VBOOST. In this case, V_{OUT} will also be high and cannot be used to derive VBOOST. ***Figure 5*** shows another

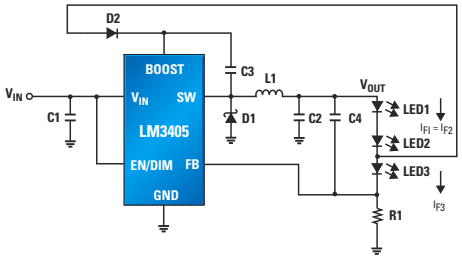

Figure 5. Deriving a Self-Biased Boost Voltage from the LED String When Driving Two or More LEDs in Series

approach for deriving a self-biased boost voltage from the LED string itself without the need for generating a separate low-voltage supply.

The anode of the boost diode D2 is connected to the anode of LED3. The voltage at the anode of LED3 is approximately equal to 205 mV plus the forward voltage VF of LED3. This voltage is approximately in the range of 3V to 4V (depending on the I_{F3} current setting) and meets the BOOST to SW voltage requirement. It must be noted that the current through LED1 and LED2 will be slightly larger than the current in LED3 due to the fact that the average charging current for the boost capacitor C3 will now be provided through LED1 and LED2. Therefore, the LED currents are:

$$I_{F3} = \left[\frac{205\ mV}{R1}\right]$$

$$I_{F1} = I_{F2} = I_{F3} + \overline{I_{C3}}$$

where I_{F1}, I_{F2}, and I_{F3} are the currents in LED1, LED2, and LED3, respectively and $\overline{I_{C3}}$ is the average current charging C3 over one switching cycle. This approach can be used if exact current matching (and hence brightness) is not required between I_{F1} and I_{F3} or between I_{F2} and I_{F3}.

For information on driving multiple LED strings and parallel LEDs with the LM3405 device, please go to:
www.national.com/ae3

Conclusion

This article describes a few application solutions using the LM3405 buck regulator. These solutions provide the user with highly-compact, driving and dimming solutions for high-brightness LEDs. The LM3405 data sheet may be referred to for a detailed description of operation and component design guidelines. ■

National Semiconductor
2900 Semiconductor Drive
Santa Clara, CA 95051
1 800 272 9959

Mailing Address:
PO Box 58090
Santa Clara, CA 95052

National Semiconductor
The Sight & Sound of Information

EDN

PUBLISHER, *EDN* WORLDWIDE
Alan Robinson
1-408-345-4450; fax: 1-408-345-4400;
aarobinson@reedbusiness.com

EDITORIAL DIRECTOR, *EDN* WORLDWIDE
Maury Wright, 1-858-748-6785;
mgwright@edn.com

EXECUTIVE EDITOR
Ron Wilson, 1-408-345-4427;
ronald.wilson@reedbusiness.com

MANAGING EDITOR
Kasey Clark
1-781-734-8436; fax: 1-303-265-3053;
kase@reedbusiness.com

EDITOR IN CHIEF, EDN.COM
Matthew Miller
1-781-734-8446; fax: 1-303-265-3017;
mdmiller@reedbusiness.com

SENIOR ART DIRECTOR
Mike O'Leary
1-781-734-8307; fax: 1-303-265-3021;
moleary@reedbusiness.com

EMBEDDED SYSTEMS
Warren Webb, Technical Editor
1-858-513-3713; fax: 1-858-486-3646;
wwebb@edn.com

ANALOG
Paul Rako, Technical Editor
1-408-745-1994;
paul.rako@reedbusiness.com

EDA, MEMORY, PROGRAMMABLE LOGIC
Michael Santarini, Senior Editor
1-408-345-4424;
michael.santarini@reedbusiness.com

MICROPROCESSORS, DSPs, TOOLS
Robert Cravotta, Technical Editor
1-661-296-5096; fax: 1-303-265-3116;
rcravotta@edn.com

MASS STORAGE, MULTIMEDIA, PCs AND PERIPHERALS
Brian Dipert, Senior Technical Editor
1-916-760-0159; fax: 1-303-265-3187;
bdipert@edn.com

POWER SOURCES, ONLINE INITIATIVES
Margery Conner, Technical Editor
1-805-461-8242; fax: 1-805-461-9640;
mconner@reedbusiness.com

DESIGN IDEAS EDITOR
Charles H Small
edndesignideas@reedbusiness.com

SENIOR ASSOCIATE EDITOR
Frances T Granville, 1-781-734-8439;
fax: 1-303-265-3131;
f.granville@reedbusiness.com

ASSOCIATE EDITOR
Maura Hadro Butler, 1-908-347-9605;
mbutler@reedbusiness.com

EDITORIAL/WEB PRODUCTION MANAGER
Diane Malone, Manager
1-781-734-8445; fax: 1-303-265-3024
Steve Mahoney, Production/Editorial Coordinator
1-781-734-8442; fax: 1-303-265-3198
Melissa Annand, Newsletter/Editorial Coordinator
Contact for contributed technical articles
1-781-734-8443; fax: 1-303-265-3279
Adam Odoardi, Prepress Manager
1-781-734-8325; fax: 1-303-265-3042

CONTRIBUTING TECHNICAL EDITORS
Dan Strassberg, strassbergedn@att.net
Nicholas Cravotta, editor@nicholascravotta.com

COLUMNISTS
Howard Johnson, PhD;
Bonnie Baker; Joshua Israelsohn;
Pallab Chatterjee

PRODUCTION
Dorothy Buchholz, Group Production Director
1-781-734-8329
Kelly Jones, Production Manager
1-781-734-8328; fax: 1-303-265-3164
Linda Lepordo, Production Manager
1-781-734-8332; fax: 1-303-265-3015

EDN EUROPE
Graham Prophet, Editor, Reed Publishing
The Quadrant, Sutton, Surrey SM2 5AS
+44 118 935 1650; fax: +44 118 935 1670;
gprophet@reedbusiness.com

EDN ASIA
Raymond Wong, Managing Director/
Publishing Director
raymond.wong@rbi-asia.com
Kirtimaya Varma, Editor in Chief
kirti.varma@rbi-asia.com

EDN CHINA
William Zhang, Publisher and Editorial Director
wmzhang@idg-rbi.com.cn
John Mu, Executive Editor
johnmu@idg-rbi.com.cn

EDN JAPAN
Katsuya Watanabe, Publisher
k.watanabe@reedbusiness.jp
Kenji Tsuda, Editorial Director
and Editor in Chief
tsuda@reedbusiness.jp
Takatsuna Mamoto, Deputy Editor in Chief
t.mamoto@reedbusiness.jp

The EDN *Editorial Advisory Board serves as an industry touchstone for the editors of* EDN *worldwide, helping to identify key trends and voicing the concerns of the engineering community.*

EDN. 225 Wyman St, Waltham, MA 02451. **www.edn.com**. Phone 1-781-734-8000. **Address changes or subscription inquiries:** phone 1-800-446-6551; fax 1-303-470-4280; subsmail@reedbusiness.com. For a free subscription, go to **www.getfreemag.com/edn**. Reed Business Information, 8878 S Barrons Blvd, Highlands Ranch, CO 80129-2345. Include your mailing label.

Reed Electronics Group

The Capacitance Company

KEMET

CHARGED.™

You can get capacitors from a lot of suppliers.

But there's only one place you can get the most capacitor technologies across the most capacitor types. Along with full design collaboration services and fast custom design and prototyping. Along with the industry's best service and on-time delivery.

Along with consistent, high part quality and reliability. Along with global availability.

Along with hundreds of new parts every year – some of which push the envelope in ways even you haven't thought of yet. From the most ambitious R&D program in the business.

The only place you can get all of that? Here at KEMET, The Capacitance Company.

If it's anything to do with capacitance, we can do it.

www.kemet.com

ALUMINUM CERAMIC ELECTROLYTIC FILM PAPER TANTALUM

Power and Control in One Chip

QUAD Output, 5A and 10A PWMs Combine Integration and Flexibility

Design Your Entire Power System With an Easy-to-Use Graphical Tool (GUI)

Create "Custom" Device Configurations with Non-Volatile Programming I²C Interface for Dynamic "On-the-Fly" Output On/Off/Voltage Control

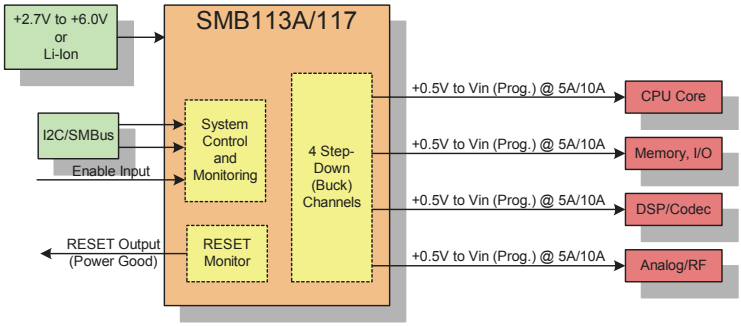

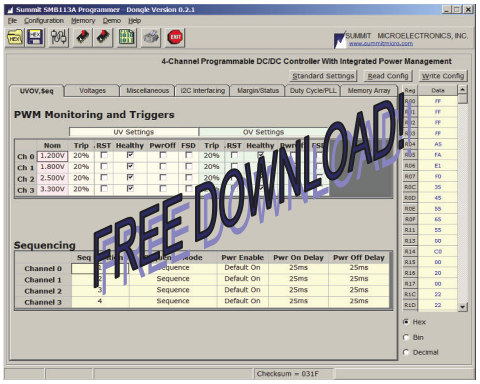

Develop application with convenient programming tool and upload HEX file
Summit ships pre-programmed "custom" units in high volume

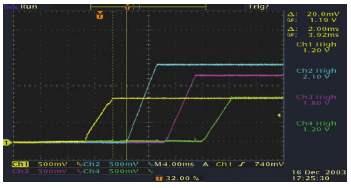

Sequencing

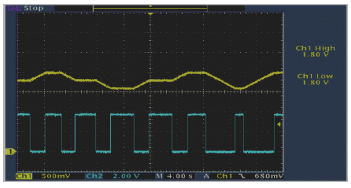

Dynamic Output Voltage

Features

- +2.7V to +6.0V or Li-Ion Battery Input *
- ± 1.5% Output Voltage Accuracy
- >90% Efficiency
- I²C Interface + Programmable Configuration
- Dynamic Output Voltage Control
- Pushbutton Enable Input
- Programmable FET Driver "Deadtime"
- Programmable Enable-Sequencing-Output Slew Rate
- UV/OV Monitors with RESET Cor
- +0°C to +70°C or –40°C to +85°C

Applications

- Cable/Satellite Set-Top Boxes
- PVR/DVR Boxes
- Digital HDTV
- Slimline Notebook PC's
- SOHO Networking Equipment
- Automotive/Marine Navigation

Features/Parameters	SMB113A	SMB117
Step-Down Outputs	4	4
Control Interface	I²C	I²C
Input Voltage Range (V)*	2.7-6.0	2.7-6.0
Output Voltage Range (V)	0.5-VIN	0.5-VIN
Switching Frequency (kHz)	800	400
Source/Sink "DDR" Mode	Y	Y
Max Output Current (A)	5	10
Operating Temp Range (°C)	-40 to +85	-40 to +85
Package Size (mm)	5x5	5x5
Package Type	QFN-32	QFN-32

For more information see:
www.summitmicro.com

SUMMIT MICROELECTRONICS, Inc.
"Programmable Power for a Digital World"

*For higher input voltages – contact factory

EDITED BY FRAN GRANVILLE

pulse

INNOVATIONS & INNOVATORS

PCI Express speeds data acquisition

With high-speed, high-volume industrial-automation, laboratory-measurement, actuator-control, and acoustic applications in mind, Adlink Technology recently unveiled a line of PCIe (PCI Express) data-acquisition cards. With a choice of input ranges and sampling speeds plus a differential mode to maximize noise suppression, the DAQe-2200 series samples as many as 64 analog channels with a maximum data-transfer rate of 250 Mbytes/sec over the one-lane PCIe bus. The DAQe-2200 series comprises the 12-bit DAQe-2204 ADC card with a 3M-sample/sec sampling rate; the 16-bit DAQe-2205 ADC card with a 500k-sample/sec sampling rate; and the 16-bit DAQe-2206 ADC card with a 250k-sample/sec sampling rate.

The cards include analog and digital triggering, two 12-bit analog outputs, 24 programmable digital-I/O lines, and two 16-bit timer/counters. The DAQe-2200 series also features a system-interface bus to allow the synchronization of as many as four cards, autocalibration that adjusts gain and offset, and high-speed transfer of data directly to and from memory through the DMA controller. The DAQe-2200 series of cards has prices starting at $895, with volume discounts available.

—by Warren Webb

▷**Adlink Technology Inc**, www.adlinktech.com.

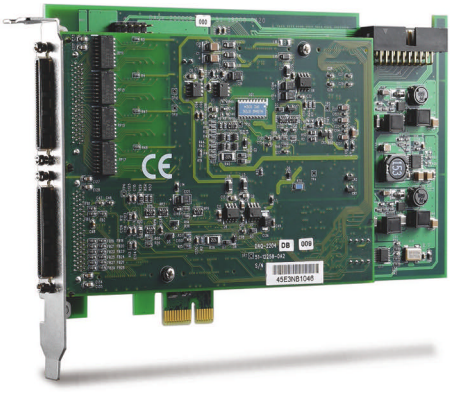

The DAQe-2200 series of 64-channel PCIe data-acquisition cards offers designers a choice of input ranges and sampling speeds.

Tiny cube packs three gyros, three accelerometers, and more

The ADIS16355 inertial-measurement unit provides 6DOF motion measurement, comes factory-calibrated, and measures less than 0.75 in.^{3}

IMUs (inertial-measurement units) rely on a combination of accelerometers and gyros to measure acceleration and rotational change and to calculate position and motion. Typically finding use in aircraft and missiles, IMUs usually measure several inches along a side and have prices starting in the thousands of dollars—and that's before you add in the complex calibration procedure of their internal sensors. In contrast, the ADIS16355 IMU, a new member of Analog Devices' iSensor family, measures less than 0.75 in.^{3} and costs $359 (1000). It incorporates three MEMS (microelectromechanical-system)-based gyros and three MEMS-based accelerometers to provide three axes of angular-rate sensing, three axes of acceleration sensing, and thus 6DOF (6°-of-freedom) motion measurement.

An important system-integration cost for IMUs is their calibration for temperature and motion, often requiring expensive and complicated testing on 3-D, rotating calibration platforms. The ADIS16355 undergoes calibration at the factory, in which operators program the calibration coefficients into each unit through its SPI (serial-peripheral interface). This step allows designers to integrate the IMUs into their system with no additional temperature or motion calibration. The ADIS16355 includes full-temperature-range calibration, with a bias-temperature stability of 0.005°/sec/°C; the ADIS16350 version provides a room-temperature calibration with a bias-temperature stability of 0.03°/sec/°C and costs $275 (1000).**—by Margery Conner**

▷**Analog Devices**, www.analog.com.

EDA start-up GateRocket brings hardware-based accelerator to FPGAs

EDA start-up GateRocket Inc wants to bring hardware-assisted verification to designers programming the biggest FPGAs from Xilinx (www.xilinx.com) and Altera (www.altera.com) and claims that its mixed-hardware-and-software approach will provide a 10- to 1000-times speedup in the simulation of these FPGAs. EDA veteran Dave Orecchio, GateRocket's chief executive officer, says that Altera's Stratix and Xilinx's Virtex devices are so large that most designers must simulate their FPGA designs before they can program them into the FPGA. Traditionally, FPGAs haven't required this step. Designers could simply program their designs into the targeted FPGA and then test that their design runs as they expect. If the design doesn't run as they expect, they simply change the design and reprogram the device, repeating the procedure until it works correctly.

"These big FPGAs are too big for the 'blow-and-go' methodology where users design the chip and try it," says Orecchio. "The methodology is divergent. Some people doing small chips still use that methodology, but we're focused on the larger chips where that methodology doesn't work. And most designers using blow and go will tell you that, once they program the chip and it's in their end system, they can't figure out what's wrong with it." As a result, many designers working with the largest FPGAs test their designs with simulation before they program the FPGAs. But Orecchio claims that conventional simulation is too slow and, in some cases, inaccurate. "It's pretty common that, if you are working with third-party IP [intellectual-property] blocks, the simulation model you get from the IP vendor is very different from the physical version of the IP," says Orecchio. "That can cause a lot of problems."

The company's 5.5-in., disk-drive-sized RocketDrive simulation accelerator contains a version of the targeted FPGA. The "device-native" product contains the same FPGA you are targeting for your end design, meaning that, when users program their designs into it, they get a true representation of the functions of the designs running on the targeted FPGAs, says Orecchio. So far, the company fields versions of the RocketDrive containing either an Altera Stratix II device or a Virtex-4 device, depending on the targeted FPGA.

> "These big FPGAs are too big for the 'blow-and-go' methodology where users design the chip and try it. ... That methodology doesn't work.

The company plans to offer new versions of the RocketDrive supporting Virtex-5 and Stratix III.

Chris Schalick, GateRocket's vice president of engineering and chief technology officer, developed the RocketDrive, which users can insert into their workstations or PCs. Alternatively, they can buy a rack of RocketDrives in a PC-tower configuration, connecting the drives to their workstation through a PCIe (PCI Express)-interface card. A user can place an entire FPGA or a part of one into the RocketDrive system and then run tests on it, with the testbench or the remainder of the design running on the user's choice of commercial-logic simulator, such as Incisive, ModelSim, or VCS. Designers can also use the system for software development.

The package also includes the RocketVision software, which facilitates the simulator-to-RocketDrive connection and RocketDrive programming. RocketDrive is compatible with FPGA-vendor programming tools and third-party FPGA-synthesis tools.

GateRocket doesn't yet have any customers, but Orecchio says the company is just now launching the product, and the prospective market is big and growing. Traditionally, FPGA vendors sold most of their largest FPGAs to ASIC designers, who use the FPGAs to prototype ASIC designs. That scenario is now changing, according to Orecchio. Most of the mammoth FPGA-design starts are targeting the use of high-end FPGAs in end products, such as medical and networking equipment, automotive applications, and test equipment. He also points out that the number of FPGA starts is increasing, whereas ASIC starts have been decreasing in number. "From my analysis in looking through Altera and Xilinx annual reports, roughly 24% of their customers are using these large FPGAs," says Orecchio. The portion of that 24% who are planning to use FPGAs in end products is GateRocket's target audience. Designers could use the technology for prototyping smaller ASICs or for either FPGA or ASIC in-circuit emulation. But, for now, the company focuses on supporting its use as a simulation accelerator for big FPGAs. Prices for GateRocket's RocketDrive start at $25,000.

—by Michael Santarini

▷**GateRocket Inc**, www.gaterocket.com.

DILBERT By Scott Adams

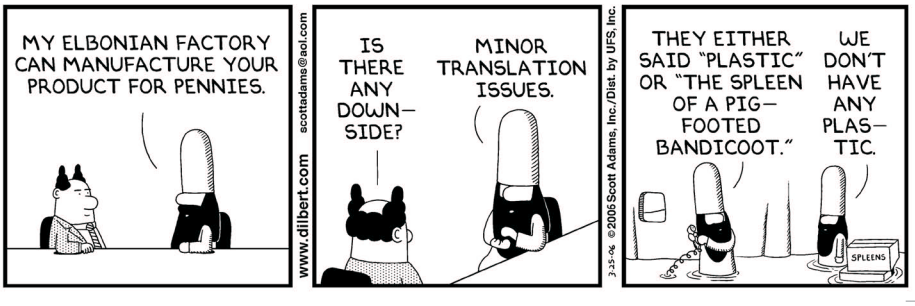

PC/104 Embedded Machine to Machine Connectivity

Keep in touch with your industrial application by using WinSystems' wired and wireless modular PC/104 modules. Small, rugged, and operational over extended temperature ranges, they are perfect for transportation, pipeline, instrumentation, MIL/COTS, and homeland security applications.

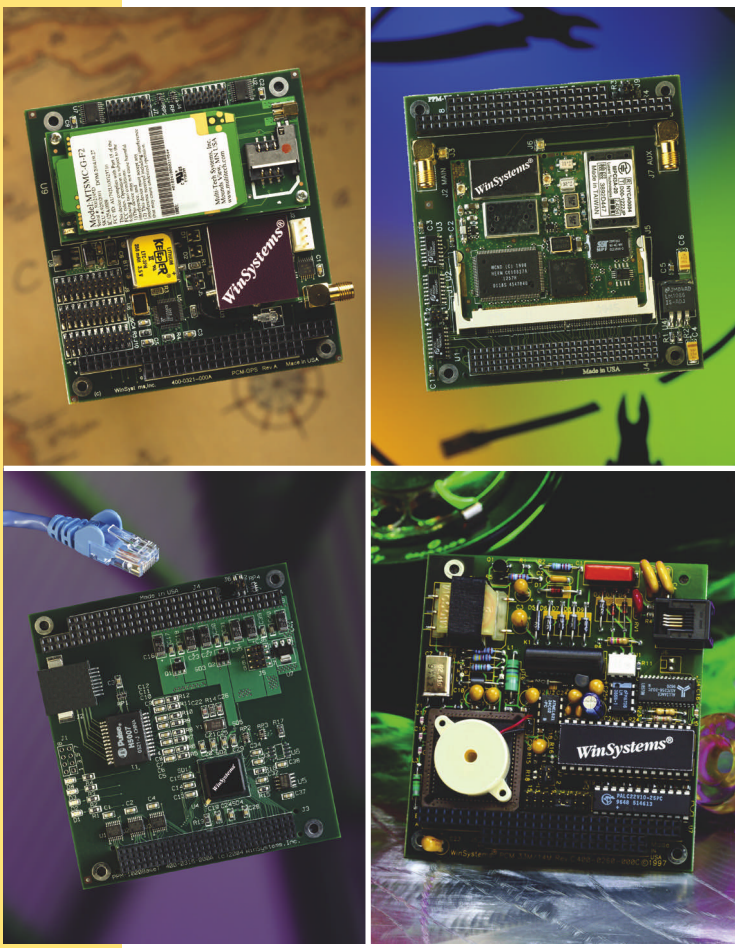

- Wireless
  - GSM/CDMA Cellular
  - 802.11 a/b/g Networking
- Ethernet
  - 10/100 Mbps Ethernet
  - Gigabit Ethernet
- Modem (POTS)
- USB 2.0
- RS-232/422/485 Serial I/O
- Other I/O Expansion
  - A/D, D/A
  - Digital
  - GPS
- Embedded Platforms Supported
  - PC/104 – 3.6 x 3.8 inches
  - EPIC – 4.5 x 6.5 inches
  - STD Bus – 4.5 x 6.5 inches
  - EBX – 5.75 x 8.00 inches
- O/S Support
  - Windows® XPe and CE
  - Linux 2.6
  - x86-compatible RTOS
- Extended Temperature Operation

Off-the-shelf delivery, knowledgeable technical support, and long-term availability makes these products the right choice for your application.

epic

Microsoft® Windows® Embedded Partner

VDC Platinum Vendor

Extended Temperature

Call 817-274-7553 or
Visit www.winsystems.com

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

EZ-Color controller family simplifies HB-LED-lighting-system designs

All HB (high-brightness) LEDs are not created equal: Even within the same part number, their performance, including wavelength, brightness, and power, can vary. HB-LED manufacturers use a "binning" process to test each LED and sort it into a category, which becomes part of the specification. HB LEDs from different bins have visibly different colors, and designers must allow for these colors in the system's RGB-color control. Having the capacity to allow for and adapt to different bins can save as much as 15% of the component cost, as opposed to locking your design into one bin and then finding that its price is at a premium or that it is no longer available. In addition, an HB LED degrades with temperature, and different colors degrade at different rates. Developing these color-creation and -control algorithms is a complex, time-consuming part of the design process.

To ensure consistent color and brightness, the EZ-Color controller uses the PSoC Express visual-design tool to allow for HB-LED-bin characteristics, as well as temperature feedback from the display to account for LEDs' age degradation.

Addressing that challenge, Cypress has combined the new EZ-Color family of HB-LED controllers with the PSoC Express embedded visual-design tool to allow designers to input a supported vendor's HB-LED part number and bin. The PSoC Express software then applies the manufacturer's bin specifications and temperature-feedback algorithms to program the EZ-Color controller to compensate for both HB LEDs' binning characteristics and their temperature response. The family comprises the CY8CLED04/08/16, which drive as many as four, eight, or 16 HB-LED channels.

The EZ-Color controller uses Cypress' proprietary Prism spread-spectrum power-control technology, rather than a fixed-frequency PWM signal to control LED brightness. The advantage of Prism over a PWM approach is a reduction in low-frequency flicker and EMI. The CY8CLED04 sells for less than $4 (high volumes).

—by Margery Conner

▷**Cypress**, www.cypress.com.

06.21.07

SIGRITY I/O PLANNER SPANS IC, PACKAGE, AND PCB

Sigrity Inc last month introduced OrbitIO Planner, a tool for dynamic I/O planning across IC, package, and PCB (printed-circuit-board) design. To date, the company has offered tools for IC, package, and PCB power and signal-integrity analysis and packaging-design tools it acquired from Synopsys (www.synopsys.com). Now, the company is offering a tool to help I/O designers get an up-front estimate of how their IC layouts will affect pin placement in the package and, ultimately, the targeted PCB.

According to Jiayuan Fang, president of Sigrity, designers changing the I/O in an IC layout traditionally have also had to implement the IC in the package and PCB layouts. But these three disciplines have become separate; thus, propagating and weighing the impact of changes across all three have been difficult and time-consuming. OrbitIO Planner aims to solve this problem. The tool employs a single data model that connects IC-, package-, and PCB-layout environments and allows users to make real-time evaluations and do what-if explorations of I/O assignment and connectivity options.

Users feed the tool LEF (layout-exchange-format) and DEF (design-exchange-format) files from IC-layout environments, input data from Sigrity or Cadence (www.cadence.com) package-layout tools, and input data from PCB-layout tools. "The tool tries to find the best location for I/O cells and the best net assignments for the pads and balls on the package to minimize the number of interconnection layers and crossover of the interconnects," says Fang. Once users have derived an I/O plan, they can use the tool to feed the I/O data back to the PCB and IC tools in their native formats.

OrbitIO Planner includes features to help designers monitor I/O hierarchy and relationships. The three disciplines often use different syntax and naming conventions to describe nets, so with OrbitIO, Sigrity has included a feature to help users manage syntactical differences in net names among domains, including scenarios of one-to-many or many-to-one net mapping.

OrbitIO Planner supports multiple package types, including SIP (system in package) and POP (package on package). It also supports wire-bond, flip-chip, surface-mount, and mixed-die-attachment methods.

Beta customers are currently using OrbitIO Planner, and Sigrity expects to have the product ready for the mass market in the third quarter of 2007 at a starting price of $58,000 for a one-year subscription.

—by Michael Santarini

▶**Sigrity Inc, www.sigrity.com.**

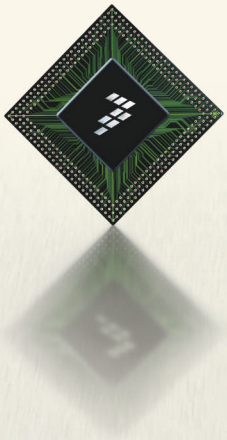

It lets you move between 8-bit and 32-bit as fast as the market demands.

The Controller Continuum. Only from Freescale.

Your market won't tolerate design holdups, so why should you? Accelerate your time to market with the unprecedented design flexibility found in the 8-bit and 32-bit solutions of the Freescale Controller Continuum. Pin-for-pin compatibility, shared peripherals and a common set of design tools let you move between the two architectures with the greatest of ease. It's the shortest distance between your design ideas and success.

Accelerate your design process at freescale.com/continuum

freescale™ semiconductor

Microprocessor Forum recap: balancing power and performance

Balancing processing performance with energy efficiency continued to be the theme at this year's Microprocessor Forum, which took place last month in San Jose, CA. The forum is an opportunity for companies not only to present new processors, but also to provide a first technical look at upcoming processor architectures.

Intel (www.intel.com) started off the forum with several presentations that provided a benefits summary of the upcoming devices using the company's 45-nm process; a technical overview of the 45-nm Penryn core microarchitecture; and explorations of the technical implications of emerging demands for tera-scale computing, including more scalable memory architectures, improving management of caches, and QOS (quality-of-service) memory-resource distribution.

A presentation from Mears Technologies (www.mearstechnologies.com) described the MST (Mears silicon technology), which uses a channel-replacement technology to derive an enhancement in drive current, along with a 60% reduction in gate leakage, compared with an 85-nm PNO (plasma-nitrided-oxide)-gate CMOS.

AMD (www.amd.com) introduced the Griffon architecture and the Puma platform targeting mobile computing. Both include a new integrated memory controller featuring a DRAM prefetcher, fixed analog- and I/O-voltage planes, and separate voltage planes with independent frequency and voltage scaling for each CPU core. The architecture also features an on-die north bridge; HyperTransport 3 I/O links with support for dynamic link-width scaling; autonomous hardware power management; and multipoint thermal control, which uses multiple on-die thermal sensors to automatically reduce the system's power state when the temperature exceeds defined limits.

A presentation by Denso (www.denso.co.jp/en) focused on future requirements for microprocessors in automotive embedded systems, including body electronics and navigation. Freescale (www.freescale.com) introduced the MPC5121e, which integrates a power-architecture-based processor core with a multimedia-acceleration core, a 2- and 3-D-graphics-acceleration core, a three-plane blend-display-interface unit, and a multiclient DRAM controller. Parimics (www.parimics.com) presented its hardware-and-software approach for image analysis, which relies on layers of arrays of PEs (processing elements); there are 76,800 PEs per die.

Renesas (www.renesas.com) presented information on the SH-Navi2V, which integrates a 600-MHz SH4A processor core with a 2-D-graphics accelerator and an image-recognition engine; the image-recognition engine supports 200 APIs (application-programming interfaces).

NTT DoCoMo Communications Laboratories (www.nttdocomo.com) covered cell-phone-technology requirements for Super 3G LTE (long-term evolution) and research for 4G wireless systems. IP-Flex (www.ipflex.com) presented its dual-core, dynamically reconfigurable processor, which it based on the DAPDNA (digital-application processor/distributed-network architecture)-IMX. The presentation included details on the 955-PE matrix and the software-development flow to target this type of architecture.

Nvidia (www.nvidia.com) discussed the SM multithreaded multiprocessor with eight SP-thread processors, which can handle 768 threads in hardware. The company also touched on the CUDA (compute-unified-device-architecture) programming model. Stream Processors (www.streamprocessors.com) detailed its Storm-1 SP16HP architecture, which aligns with computationally intensive streaming applications on parallel data with limited locality of reference. Stretch (www.stretchinc.com) presented its S6 architecture, which adds video surveillance to the list of target applications with its second-generation ISEF (instruction-set-extension-fabric) products.

AMCC (www.amcc.com) presented its dual-core Titan, a power-architecture-based processor. The company implemented the processor with Intrinsity's (www.intrinsity.com) four-phase-clock Fast14 technology, which can deliver 2-GHz operation on each core with a power consumption of 2.5W per core at 1V on the TSMC (www.tsmc.com) 90-nm GT.

ARM (www.arm.com) presented its first processor core, the Cortex-M1, which will reside in an FPGA device; the core will be available for Actel (www.actel.com) FPGA devices this July. ARM also presented information about multicore support for the ARM Version 7 architecture, which touched on cache coherence, a coherence-accelerator port, an 8-bit extension to the ARM Version 6 exclusives for SMP (symmetrical-multiprocessing) support, and support for multiple paravirtualized operating systems. MIPS (www.mips.com) introduced its 74K core, which builds on the MIPS32 24KE.

Qualcomm (www.qualcomm.com) presented information about its Scorpion processor, an instantiation of the ARM Version 7 microarchitecture in a 65-nm LP technology with a VeNum (vector-numerics) multimedia engine, and the Snapdragon mobile platform. Tensilica (www.tensilica.com) presented its software-programmable Diamond 388VDO dual-core video decoder/encoder that uses a stream and a pixel processor to improve the performance of CABAC (context-adaptive-binary-arithmetic-coding) compression.

—by Robert Cravotta

▷**Microprocessor Forum**, www.instat.com/mpf/07.

FEEDBACK LOOP

"People end up losing their jobs when the guys with the golden parachutes decide to make another nickel by outsourcing everyone but themselves."

—Robert Weise, in *EDN*'s Feedback Loop, at www.edn.com/article/CA6437123. Add your comments.

Bidirectional current monitors that make sense.
In measurement and control, analog is everywhere.

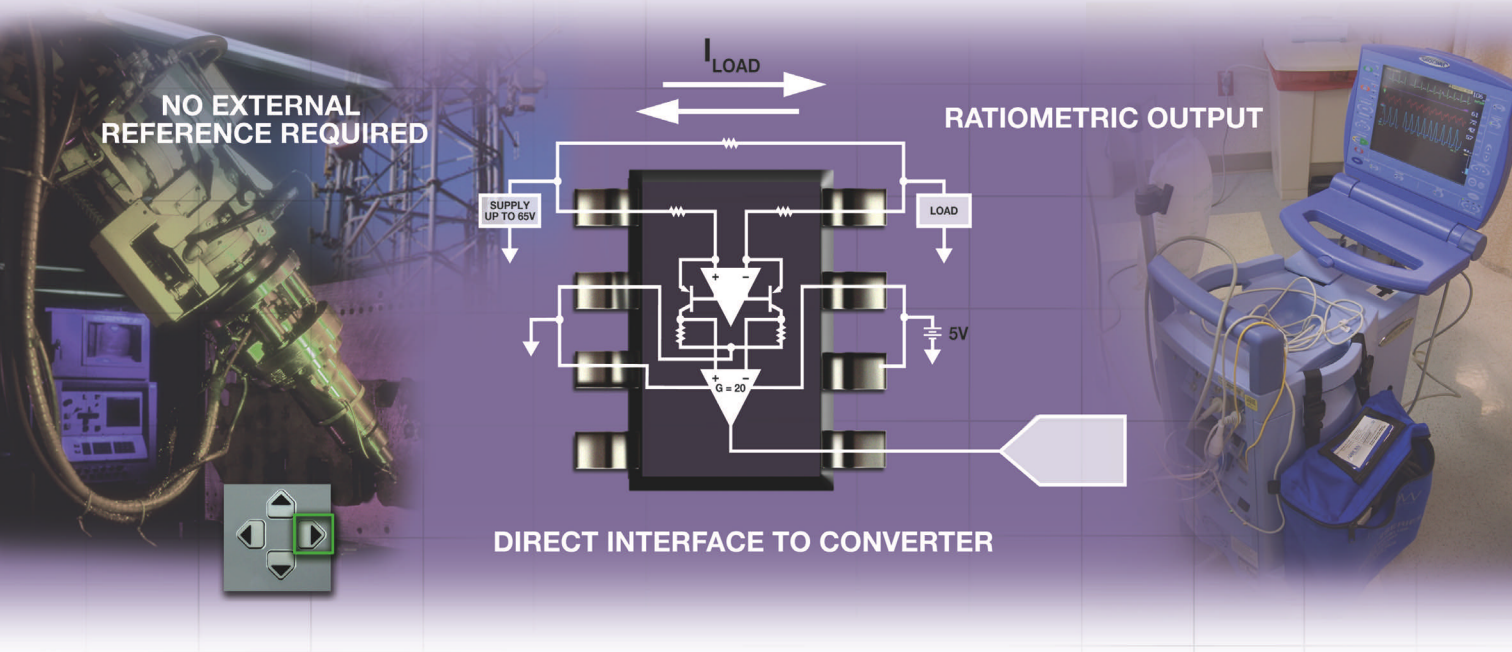

AD8210
Bidirectional Current Sense Amplifier

- –2 V to +65 V operating common-mode voltage range
- –40°C to +125°C operating temperature range
- 120 dB typical CMRR at dc
- 10 ppm/°C typical gain drift
- Gain of 20
- 500 kHz bandwidth
- Price: $1.79/1k

Typical ≤ 1 µV/°C Offset Drift

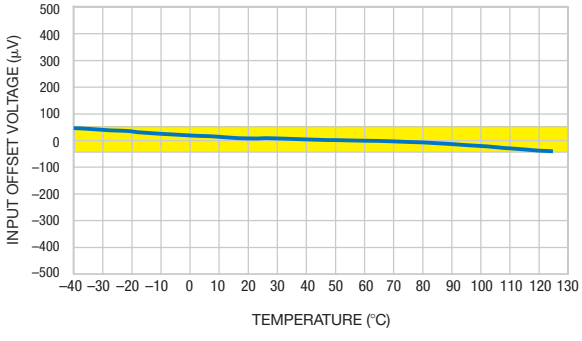

Precision over temperature—guaranteed

With our new AD8210 bidirectional current sense amplifier, you get the industry's most integrated, flexible, and accurate current sense monitor. With features like ±8 µV/°C maximum offset drift and 20 ppm/°C maximum gain drift guaranteed across the entire –40°C to +125°C temperature range, it represents a milestone in precision current sense monitoring. The AD8210 improves your ability to meet your error budget and to control more precisely a wide range of loads in communications, industrial, and medical applications. Two other bidirectional monitors—the AD8205 (gain of 50 V/V) and the AD8206 (gain of 20 V/V)—offer 50 kHz operation and excellent performance over temperature.

All three amplifiers feature an innovative ratiometric output offset architecture that inherently improves the accuracy of your ADC and your system. With a typical 5 V single-supply, each device can be configured for both bidirectional and unidirectional current sensing. Excellent output accuracy is maintained throughout the input voltage range through the use of a proprietary thin film precision network.

For more information on ADI's current sense amplifiers and monitors, please visit *www.analog.com/currentsense* or call 1-800-AnalogD.

analog is everywhere.™

www.analog.com/currentsense

ANALOG DEVICES

RESEARCH UPDATE

BY MATTHEW MILLER

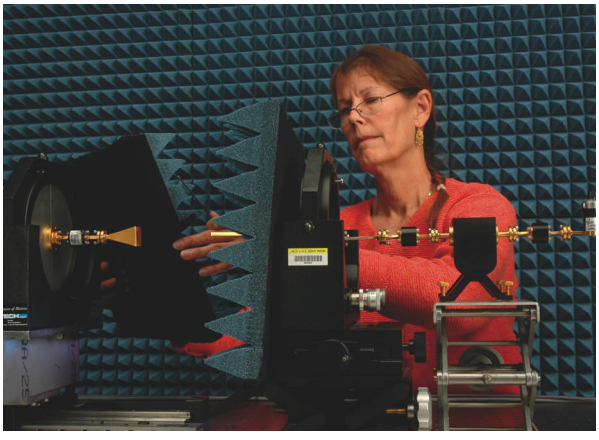

A NIST engineer prepares a table-top system to characterize a 94-GHz horn antenna (courtesy Geoffrey Wheeler, NIST).

Table-top setup characterizes antennas in 60- to 110-GHz range

Researchers at NIST (National Institute of Science and Technology) have developed a table-top apparatus that can accurately characterize the gain and polarization of antennas operating in the 60- to 110-GHz frequency range. The researchers state that their innovation will contribute to the proliferation of higher frequency antennas in both military and civilian applications. For example, automotive collision-avoidance radar operates at 94 GHz and requires an antenna small enough to integrate into a car's bumper, the researchers say.

The researchers achieved the antenna's size using NIST's extrapolated-gain technique, which makes it possible to compute far-field characteristics based on measurements of near-field behavior. The approach employs oversampling and averaging techniques to minimize the impact of scattering and range imperfections, according to NIST, which has used the approach in lower frequency characterization for decades.

The new system uses two antennas—one that is fixed onto the table and another that moves along a rail—to make its measurements. A laser tracker controls alignment and positioning between the two with uncertainty of less than 20 microns at 1000 points/sec.

NIST scientists assert that the system will play a critical role in the "measurement-traceability chain," because researchers and engineers from a variety of organizations and companies use NIST's facilities to calibrate the systems that the engineers then use to calibrate other systems in the field.

▷**National Institute of Science and Technology**, www.nist.gov.

MILITARY-FUNDED-RESEARCH PROGRAM SEEKS TO BEAT THE HEAT

Do you think your system has heat problems? Then, consider the challenge the military faces with space-constrained, electronics-packed, high-powered vehicles that must function without failures in extreme temperatures. A new program under the auspices of the DOD (Department of Defense) MURI (Multidisciplinary University Research Initiative) program will seek to attack the heat issue at the chip level.

The $7.5 million project, which funds research at the University of Virginia, Arizona State University, the University of California—Berkeley, the University of Illinois, and Rensselaer Polytechnic Institute, aims to explore microfabrication, nanotechnology, and system-control approaches to reducing chip temperatures in large-scale, high-temperature electronic systems. For example, researchers at Rensselaer Polytechnic will look into the mechanisms that control heat transfer in various liquids flowing through microdevices, compare microchannel configurations, and explore the impact of nanoparticles suspended within liquid coolants. The school will also develop a large-scale system simulator that the military will use to virtually prototype cooling systems before implementation.

▷**United States DOD**, www.defenselink.mil.

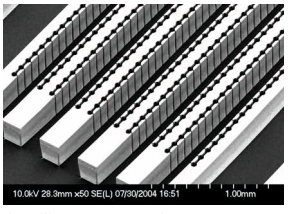

A military-research program will examine the use of microchannels to control heat in distributed, high-power systems.

06.21.07

MEMS DESIGN CAN REACH 300 MICRONS, USE POWER TOOLS

Students at the University of Oklahoma developed the winning MEMS (microelectromechanical-system) design in a contest that Sandia National Laboratories runs. The design can raise an arm 90° and 300 microns above its substrate. The design, *parvissimus bracchius*, Latin for "smallest arm," provides power and real estate at the end of the arm for microtools, such as a measuring device, a grabber, or a cutter. The design relies on thermal actuators and a set of ratchet teeth to jack up the arm and features a "shark-jaw" grabber tool that can further extend to seize its prey using an accordionlike telescoping structure.

▶**Sandia National Laboratories, www.mems.sandia.gov/ua/contest.html.**

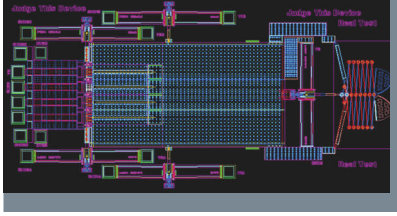

This MEMS design by students at the University of Oklahoma can raise its arm to employ microtools, such as a "shark-jaw" grabber (right).

Precision ADCs with size and power at a record low.
In data conversion, analog is everywhere.

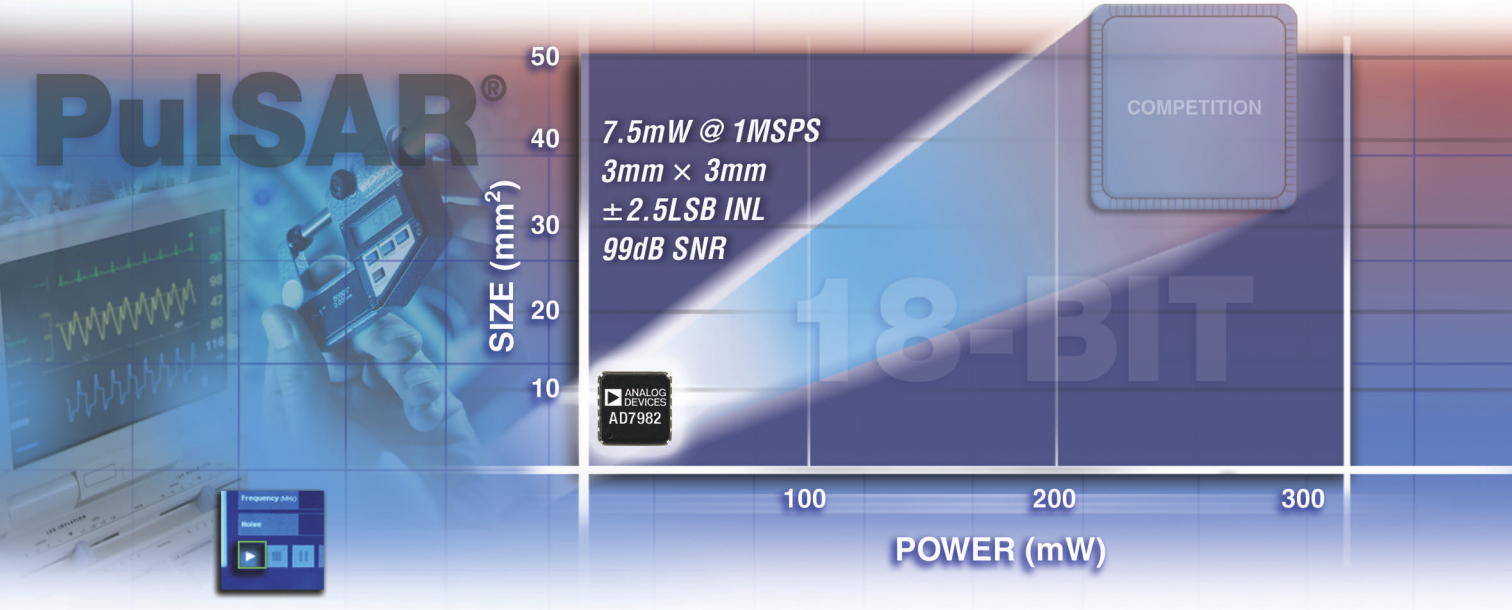

AD7982

- 18-bit, 1 MSPS, ±10 ppm INL max, 99 dB SNR
- Low power scaling with throughput: 7.5 mW @ 1 MSPS, 75 µW @ 10 kSPS
- Tiny packages: 10-lead MSOP or LFCSP
- Serial SPI interface with daisy chain
- True differential input ±5 Vpp or ±2.5 Vpp

AD7980

- 16-bit, 1 MSPS, 7.5 mW, ±30 ppm INL max
- Same pinout as AD7982
- 5 V or 2.5 V input with ground sense

Recommended ADC Drivers

- ADA4841-x family of Op Amps
- Low noise and distortion at 1.1 mA
- Rail-to-rail outputs and dc precision

18-bits, 1 MSPS, 7.5 mW, 9 mm²: breakthrough performance for medical, industrial designs

In medical, industrial, instrumentation, and sensor equipment, low-power consumption and portability are critical to performance optimization. That's why our new PulSAR® successive approximation analog-to-digital converters represent milestones in precision data conversion technology. Compared to competitive offerings, these ADCs use less than 5% of the power, and are just 20% the size, while maintaining uncompromised performance in other areas such as ac specifications and dc linearity. This family of converters is ideal for challenging applications that require low heat dissipation, extended battery life, and high channel integration.

For more information about Analog Devices' PulSAR technology, please visit *www.analog.com/pulsar-small* or call 1-800-AnalogD.

Part Number	Resolution (Bits)	Sample Rate	Max Integral Linearity (LSB/ppm)	SNR, noise rms (dB/ppm of FSR)	Power @ 100 kSPS	Price ($U.S.) @ 1k
AD7982	18	1 MSPS	±2.5 LSB, ±10 ppm	99 dB, 4 ppm	750 µW	23.00
AD7690	18	400 kSPS	±1.5 LSB, ±6 ppm	102 dB, 2.8 ppm	4.4 mW	19.50
AD7691	18	250 kSPS	±1.5 LSB, ±6 ppm	102 dB, 2.8 ppm	4.4 mW	14.50
AD7980	16	1 MSPS	±2 LSB, ±30 ppm	91.5 dB, 9.4 ppm	750 µW	19.50
AD7693	16	500 kSPS	±0.5 LSB, ±8 ppm	96.5 dB, 5.3 ppm	3.6 mW	18.00
AD7685	16	250 kSPS	±2.5 LSB, ±38 ppm	93.5 dB, 7.5 ppm	1.35 mW	6.50
AD7942	14	250 kSPS	±1 LSB, ±61 ppm	85 dB, 20 ppm	1.25 mW	4.75

All products are pin-compatible in 10-lead MSOP or 10-lead 3 mm × 3 mm LFCSP.

analog is everywhere.™

www.analog.com/pulsar-small

ANALOG DEVICES

BY HOWARD JOHNSON, PhD

Finger the culprit

Last week, at a class in Rochester, NY, one of my students asked, "What is the most difficult kind of problem to debug?"

My answer came quickly: "Something that fails every two weeks." If a device fails less often, you can pretend it isn't happening and ship the product anyway. If it fails more often, you stand a much better chance of tracking down the source. Every two weeks is about the worst it can be.

When debugging a rare mode of failure, never attempt a direct fix. The test cycles associated with each attempted improvement will kill your development schedule. Your first order of business is to *make the problem worse*. Discover what triggers the failure event, and increase the rate of failure to something more reasonable. After that, you can attempt solutions.

You can always make a system fail using a hammer, but that scenario is not what I'm suggesting. Find some control that makes the system fail in the same way, with the same symptoms—just more often. Then you have a good handle on the problem. Finding two or three mechanisms that make the system fail would be ideal.

Digital products often fail due to inadequate timing margins or coincidences of timing, so start your search there.

Suppose your system comprises several large ICs, A through E, all fed by a central clock-repeater chip. Consider a bus carrying data from A to B. If you retard the clock for A, you stress the setup time at B. Retard the clock at B, and you stress the bus timing in the opposite direction. If the bus incorporates a robust timing margin, small adjustments in the clock timing should produce no errors. On the other hand, if your bus timing is marginal, then this technique pinpoints the culprit.

For a timing-adjustment approach to work, you must arrange an error counter. When an error occurs, your test setup must record it but keep moving. If the system stops every time it hits one error, it becomes almost impossible to debug. A bell or gong sound at each failure works conveniently. (Use earbuds to avoid annoying your lab mates.)

Your first order of business is to *make the problem worse.*

Clock-timing adjustments can pinpoint problems with crosstalk as well as with bus timing. Clock timing affects crosstalk because it slightly changes the relative time of arrival of aggressive voltage spikes. If you can move the noise spike out of the clock window, then the spike no longer matters.

So, how do you change clock timing? Sometimes, just putting your finger on a clock trace adds enough parasitic capacitance to retard the clock edges. A little experimentation quickly teaches you how to calibrate your finger.

Microwave engineers perform such tests in a somewhat more controlled way. They like to glue a ¼-in.-square bit of copper onto the end of a wooden stick or pencil and touch that to the trace. The capacitance to ground of that bit of metal produces a small phase adjustment in the circuit. If you need to advance the timing, use a negative-delay circuit (**Reference 1**).

What if your clock traces aren't on the surface where you can touch them? Oops! That's an important point about board layout: Each clock trace must be accessible, somehow, somewhere, for the purpose of adjusting the clock timing.

Systems with two or more clock domains complicate the testing process. As two clocks precess in phase, problems may occur at only one phase relationship. To test for this scenario, rig up an external phase-locked dual-clock source with a knob that intentionally adjusts the phase relationship of the two clocks. Connect this device to your system and use it to dial around the phase circle, looking for a phase relationship that causes more errors than normal. For instance, adjust the two clocks straight on top of each other, or offset slightly, trying to stimulate various modes of ground bounce, board crosstalk, or metastability that you believe might influence your system.

If you find a phase relationship that greatly increases the error count, lock it down and then go find that bug! EDN

REFERENCE

1 Johnson, Howard, "Negative delay," *EDN*, Aug 30, 2001, pg 24, www.edn.com/article/CA152889.

Howard Johnson, PhD, of Signal Consulting, frequently conducts technical workshops for digital engineers at Oxford University and other sites worldwide. Visit his Web site at www.sigcon.com or e-mail him at howie03@sigcon.com.

Our 16-bit MCUs give you tons of options in every area but one.

They only come in black.

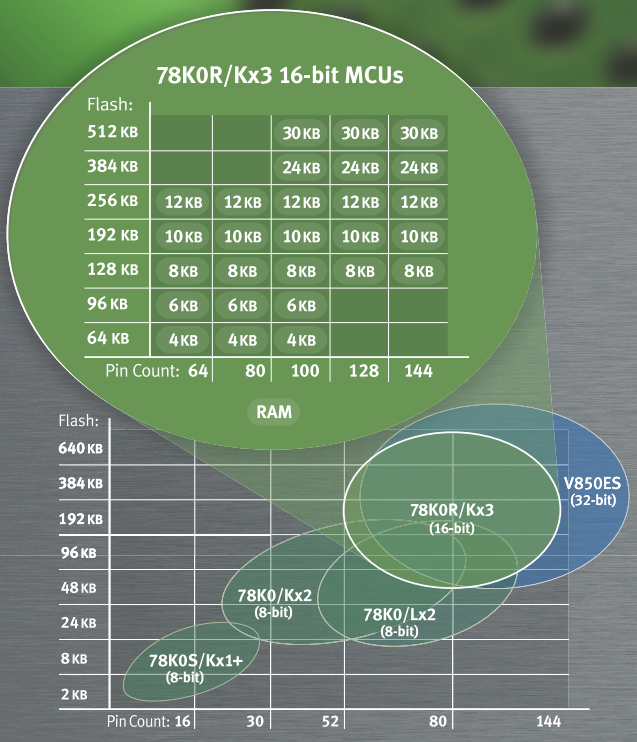

More choices in on-board flash.

NEC Electronics America makes it easy to get exactly what you want. Our microcontrollers give you strength in memory, power and performance. The 78K0R product lineup includes 30 devices available in 64 to 144 pins and 64 KB to 512 KB of memory—the widest range of on-board flash memory available. Compared to the leading 16-bit solutions available, these MCUs will save you 30% or more on power consumption. They deliver 13 Dhrystone MIPS at 20 MHz. Plus, migration from our 8- to 32-bit microcontrollers is simple — even the colors are compatible.

Learn more at www.am.necelcom/16bitmcus and order a FREE starter kit while supplies last. Restrictions apply.

Empowered by Innovation

NEC

BY JEROME JOHNSTON • CIRRUS LOGIC

The eyes have it

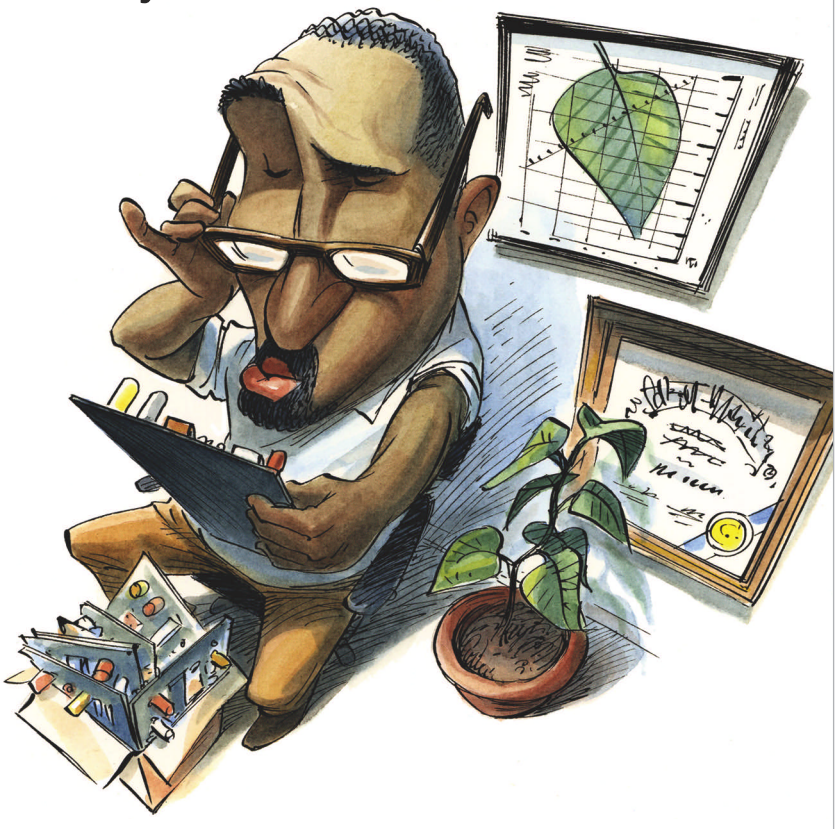

DANIEL VASCONCELLOS

In the mid-1970s, I joined a company that designed and manufactured scientific instruments for studying plants. LI-COR (www.licor.com) provided instruments such as the world's only portable leaf-area meter, used to measure the area of a leaf without detaching it from the stem of a plant. My first assignment was to develop a new light-measuring instrument—a quantum radiometer photometer. The core was to be a chopper-stabilized amplifier used in a transimpedance configuration. Back then, there were no monolithic ICs to perform this task. The design involved using MOSFET and JFET transistors as the chopping switches along with some monolithic amplifiers surrounded by resistors and capacitors. The design was challenging. An early problem was that the offset of one of the amplifiers drifted for the first 100 hours of operation, causing the instrument to lack long-term stability. The cause was an ion-contamination problem at the IC manufacturer. After solving this problem, the final amplifier achieved about 50 pA of input current, offset drift of less than 50 nV/°C, and noise of approximately 30 nV/$\sqrt{\text{Hz}}$.

LI-COR manufactured various light sensors using different-colored filter glass in front of a photodiode to shape the spectral response. We input the photodiode output to the chopper-stabilized amplifier configured as a transimpedance amplifier. The amplifier had several gain options that a rotary-switch selection of various resistor settings determined, ranging from approximately 1 kΩ to more than 50 MΩ. The output of the amplifier drove a precision analog meter as an indicator.

Later, a production technician asked whether I would help him troubleshoot a bunch of the amplifier boards that had failed to pass their performance test. He had been unable to determine their cause of failure and had accumulated about 40 boards in a box.

The first amplifier had a peculiar behavior. When the gain was set with a feedback resistor greater than 10 kΩ, the amplifier output would saturate to its supply rail. With a feedback resistor of less than 10 kΩ, the amplifier would behave within specification. Failures on the production line are always more challenging to troubleshoot, because the circuit has never worked normally. This one ended up being a microscopic solder short between the positive and negative inputs of the amplifier.

However, most of the failures were boards that seemed to fully work but had excess noise. In examining signals on all the nodes throughout the circuit, everything looked normal, but the meter readout was unstable. Although its mean reading was correct, it erratically jumped around. I examined the circuit operation, then looked at the defective board alongside one that was working properly. Then I saw it: a resistor installed backwards.

To minimize board space, the assemblers mounted the resistors, most being axial-leaded ⅛W resistors, on the board standing up. The first amplifier stage had two resistors to set its gain. The resistor going from the negative input of the amplifier to ground had been assembled with the lead going to the amplifier input facing up in the air. In this configuration, the lead became an antenna that allowed the injection of interference directly into the input of the first amplifier stage—the source of the excess noise.

We inspected the rest of the amplifier boards; most had the resistor in backward. We added a note to the board's test procedure to carefully check the position of this resistor before beginning the electrical tests. EDN

Jerome Johnston is a principal engineer at Cirrus Logic.

PRECISE, RELIABLE, HIGH BRIGHTNESS LED DRIVER ICs

For Constant LED Current Regulation

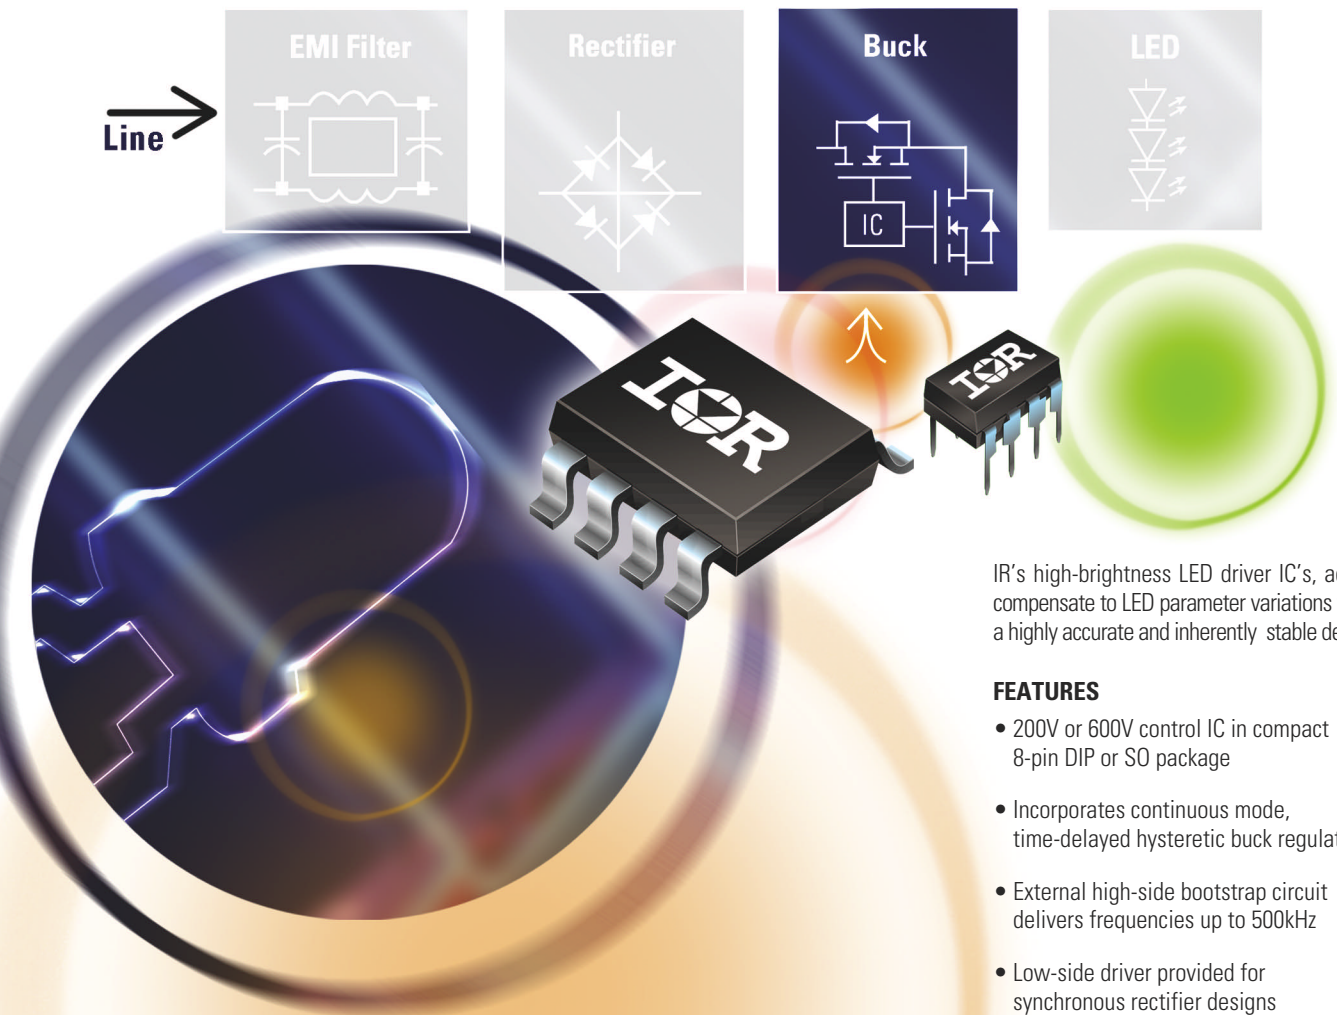

IR's high-brightness LED driver IC's, adapt and compensate to LED parameter variations to enable a highly accurate and inherently stable design.

FEATURES

- 200V or 600V control IC in compact 8-pin DIP or SO package
- Incorporates continuous mode, time-delayed hysteretic buck regulator
- External high-side bootstrap circuit delivers frequencies up to 500kHz
- Low-side driver provided for synchronous rectifier designs
- Micro-power startup of less than 500µA
- 140ns deadtime for continuous current regulation
- Auto restart, non-latched shutdown and PWM dimmable capability

Part No.	Package	Voltage	Load Current Regulation	Micro-power Start-up	Deadtime	Frequency
IRS2540PbF	DIP8, S08	200V	+/-5%	<500µA	140ns	<500kHz
IRS2541PbF	DIP8, S08	600V	+/-5%	<500µA	140ns	<500kHz

for more information call 1.800.919.7898 or visit us at
www.irf.com/lighting

International
IOR Rectifier
THE POWER MANAGEMENT LEADER

The 1nside Edge
that shapes the Future

Technology ahead of everyone - Samsung

Samsung Electro-Mechanics manufactures the world's fastest, brightest and thinnest products through high tech development. Imagine your world and Samsung will make it real. **Samsung Electro-Mechanics** will be your most reliable partner.

☑ **LED/LED BLU** : Specialized High Power LED (1,3W Cool/Warm White) and Middle Power LED (RGB, R/G/B, Cool/Warm White) for signage, electric sign, scene and mood lighting
▶**Martin Cho** (martinch@samsung.com)

☑ **Camera Modules** : Excellent image quality with high-resolution CMOS technology and advanced opto-mechanics, like auto-focus (Piezo and VCA), 5MP 3X Zoom and DIS
▶**Yoon-Chul, Park** (y.c.park@samsung.com)

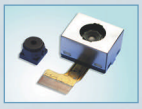

☑ **HDI/FCBGAs** : High density interconnected PCB for mobile phone and FC BGA for CPU, chipset and memory, combining circuit miniaturization and build-up technology
▶**Man-Hun, Hur** (martin.hur@samsung.com)

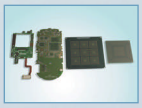

☑ **Chip Components** : Leading edge chip components including world-first ultra small MLCC based on nano materials and creative design simulation ▶**Sang-Koo, Kang** (semksk@samsung.com)

☑ **Precision Motors** : Precision-geared solution and vibration motors based on integrated mechanical technology and hall technologies ▶**Brian Yoo** (byoo@samsung.com)

☑ **Digital Tuners** : Total solution for multimedia device, including TV, set-top box and portable device; multiple functions coupled in single package or very small form factor possible
▶**Young Lee** (young20@samsung.com)

☑ **Antennas** : CBA (Ceramic Block Antenna at 6*2*1.5t) simplifies mobile tuning with easy connectivity. IMA (In Mold Antenna) provides enhanced features like space, cost, and performance using IML technology

Network Modules : Wireless connectivity solution; Bluetooth, wireless LAN, GPS, Wimax, ZigBee, UWB, NFC, and various combo solutions for current and LTE.
▶**Hyun-Seok, Jeong** (jamejung@samsung.com)

☑ **Power Supplies** : High efficiency, high power density, and high quality in Flat Panel Display (LCD-TV, PDP-TV) IT , computer and renewable energy devices
▶**Su-Jung, Jung** (sj282.jung@samsung.com)

www.sem.samsung.com

SAMSUNG
ELECTRO-MECHANICS

SAMSUNG

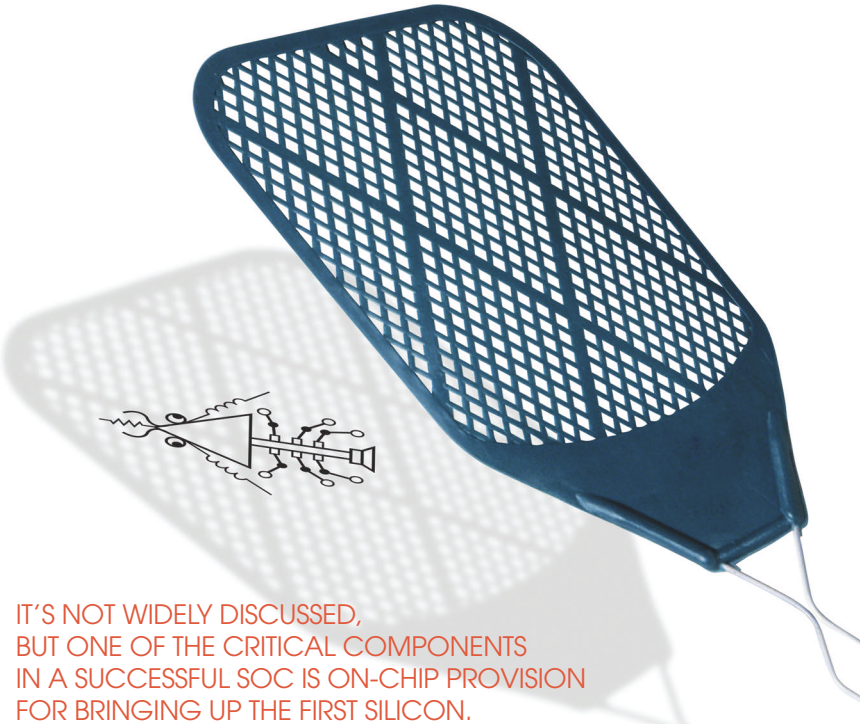

IT'S NOT WIDELY DISCUSSED, BUT ONE OF THE CRITICAL COMPONENTS IN A SUCCESSFUL SOC IS ON-CHIP PROVISION FOR BRINGING UP THE FIRST SILICON.

Design for debugging: the unspoken imperative in chip design

BY RON WILSON • EXECUTIVE EDITOR

Testing and debugging present different problems. In testing, the goal is to determine as quickly as possible whether the chip is working correctly, with high, but not absolute, certainty. Chip-design teams now universally recognize that doing so requires the addition of DFT (design-for-test) circuitry on the chip, and third-party-tool and IP (intellectual-property) companies can aid in this purpose.

Debugging is quite another story. The goal of debugging is not simply to determine that the chip is not working, but to find out why it is not working. This inquiry is not confined to a few seconds on a test floor, but may last weeks. It is not automatic, but requires the participation of the chip-design team. And it occurs at discrete points in the design cycle: during first silicon bring-up, during reliability studies, and during field failure analysis.

Given this profile, you might think that a good DFT strategy would be sufficient to meet the needs of silicon debugging—and, in fact, it often was. But with the growing complexity of SOC (system-on-chip) designs, leading design teams report that they are dedicating more and more planning, implementation, and silicon area to circuitry that supports debugging rather than test.

"Ten years ago, when we were designing with three layers of metal, this was not a big issue," observes Bay Microsystems' Senior Vice President of Engineering Tony Chiang. "If there was a problem with the chip, you could probe the metal directly to watch the circuitry, and, with a focused-ion-beam system, you could even rewire it. Now, with nine metal layers and 0.2-micron metal pitch, that's simply not possible. We have to make the circuitry observable and controllable from outside the chip, without exceeding our cost goals or overrunning our schedule."

That situation, in a nutshell, describes the world of designing for debugging.

A PANORAMA OF TECHNIQUES

Debugging is not entirely divorced from DFT. Broadcom, for example, has a corporate-level team of about 70 engineers who work with all the chip-design teams in the company on both debugging and testing, according to that company's senior director of test-development engineering, Kris Hublitz. And Hublitz repeatedly cites DFT vendor LogicVision as a key partner in Broadcom's chip-debugging strategy.

Others agree. "Design for debug is not too distant from manufacturing test," says David McCall, a vice president at CSR (Cambridge Silicon Radio). "Both start at about the same point."

That point, numerous design managers emphasize, is the quest for controllability and observability. In debugging, as in manufacturing test, the fundamental problem is to place the circuitry into a known state, start it running, and observe its behavior. In the days of medium-scale integration, boundary-scan techniques could adequately accom-

plish this task. Because chips had little internal state, you could thoroughly test them by sending the inputs through a known series of states, clocking the circuit, and observing the outputs.

With the advent of the microprocessor, things became more complicated. Microprocessors have lots of internal state, so simply forcing the inputs to a known vector and watching the outputs is not particularly informative. Early on, the industry tried a variety of techniques to make microprocessors debuggable—from providing scan for each cloud of logic between registers to relying on the same sort of trace, breakpoint, and single-step functions that minicomputers use for software debugging. A combination of those two did the trick.

Designers use the same suite of tools on the digital portions of SOCs today. A separate collection of techniques serves analog and mixed-signal circuits. But no one approach can encompass an entire complex SOC. So, the process of designing for debugging comprises partitioning the system into independently debuggable blocks, implementing a debugging strategy for each block, and integrating these strategies into a plan for the full chip that keeps the user interfaces for individual blocks similar and minimizes the silicon resources that the circuitry requires. As a final step, designers must double-check that, using these debugging resources, the fully integrated operation of the chip is both controllable and observable, because you cannot observe some bugs by looking at functional blocks in isolation.

AT A GLANCE

- For complex ICs, silicon debugging requires on-chip hardware.
- The need for on-chip hardware goes beyond the needs of design for testing.
- Digital and analog blocks demand different strategies.
- Teams should have a full plan for bringing up silicon during architectural design.

DIGITAL SOCs

The most basic form of SOC is a CPU core surrounded by simple, often nonprogrammable, peripheral blocks and memories. In most cases, the CPU core is third-party IP, and it comes at least with the option of an internal debugging kernel, the inclusion of which the software-development team usually insists upon. This kernel combines with the normal DFT circuitry that the design team implements for the peripherals to give sufficient observability and controllability to isolate problem nets. You can use the debugging kernel in the CPU core to stimulate all but the asynchronous portions of the core and to capture results. The kernel can also, by making the CPU read and write peripheral registers, stimulate and observe the peripherals, usually allowing designers to pinpoint a malfunction down to the level at which scan chains can take over.

But such simple SOCs are not that common today (**Figure 1**). More often, the chip will have several CPUs, or a host CPU core and several other formidable processor cores of various kinds. Even some of the peripheral controllers may be sufficiently complex that simply stimulating them via the CPU and observing the results won't be sufficient to diagnose them. And there will be multiple clock domains, often not synchronized with each other. Such chips call for sterner measures.

In this case, several tactics are available. One, suggests Broadcom's Hublitz, is simply to make the inputs and outputs of all the major functional blocks accessible to pins on the chip. This approach can mean a great deal of multiplexing. In designs that incorporate a large number of I/O and memory interfaces, the die may be pin-limited before the inclusion of any additional access for debugging purposes, so designers find that they have to reuse pins for debugging access. And simply bringing out the inputs and outputs of very complex blocks may be no more useful than exercising them with the host CPU core; designers may need to bring out internal signals, as well.

All of this multiplexing and routing adds up and may simply be impractical. Further, the resulting additional interconnect may mean that even if all the blocks are physically accessible from the pins, they may not be accessible at speed. And this problem is a serious one. "We find that we must test circuits—and especially the interconnect between blocks—at-speed," Hublitz says. "This is particularly true at 65 nm. Otherwise, you'd be lucky to find half the faults in the chip."

Hublitz emphasizes that a good DFT strategy supported by ATE (automatic test equipment) can enormously help the debugging process. "We do our first debugging pass on the ATE systems," he says. "After we are sure the chip won't melt, we give it to the design guys on the bench, and, from there, it's a joint effort." Hublitz says that the chip may repeatedly go back to Broadcom's test floor so that the ATE sys-

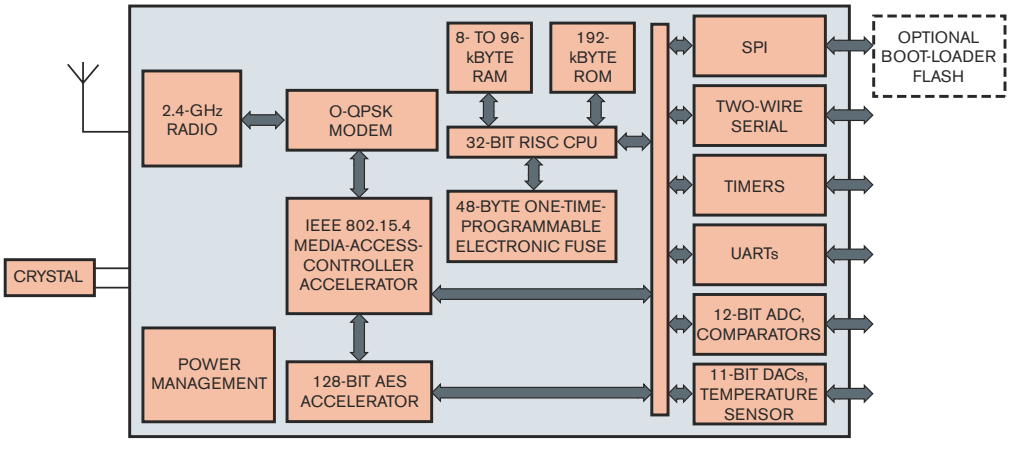

Figure 1 Single-chip network devices, such as this Jennic 513x device, present debugging engineers with a plethora of digital, analog, and RF challenges.

Micrel's Two New RF Platforms Deliver Best "One-Two" Punch In Remote Actuation

Distance = MICRF112YMM + MICRF211AYQS

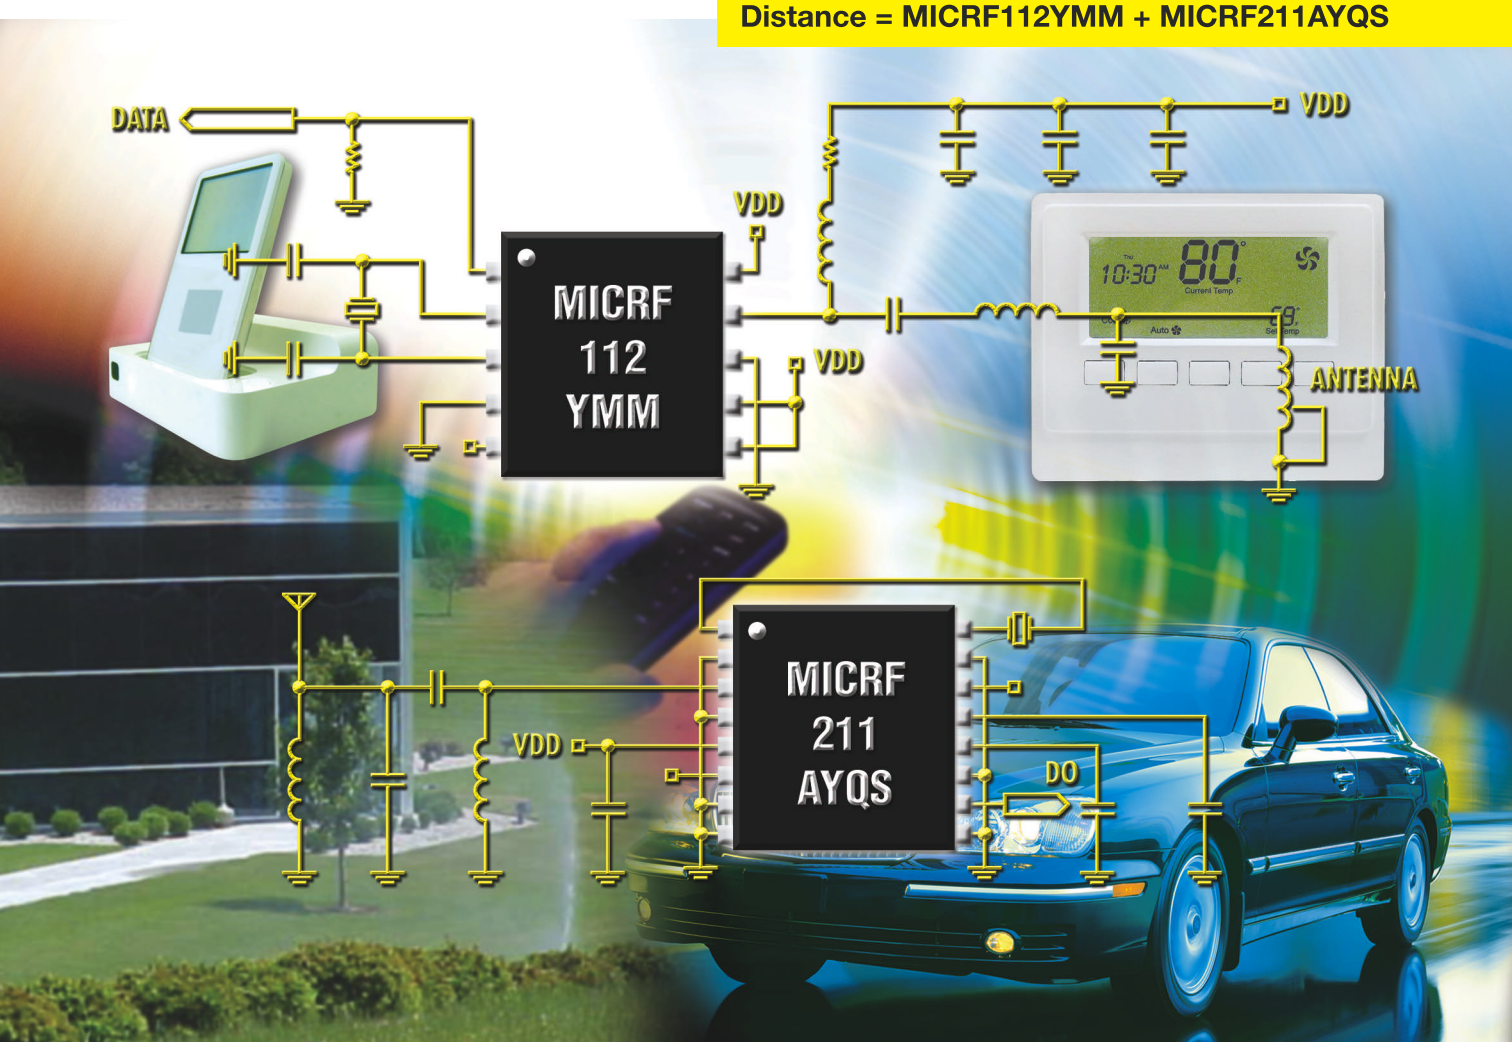

While most integrated transmitters operate above 2.0 volts, the MICRF112YMM operates down to 1.8 volts, squeezing more juice out of batteries than the competition. The MICRF112YMM offers best-in-class integration, performance and ruggedness, bringing value to the remote actuation market in terms of quality, reliability and repeatability.

The MICRF211AYQS is the latest addition to Micrel's QwikRadio® receiver family. Offering best-in-class integration, packaging and performance, the MIC211AYQS receiver solution provides performance, reliability, repeatability and low cost of ownership to the remote actuation market.

TX: MICRF112YMM

Frequency Range	300 ~ 450MHz
Data Rate	100kbps OOK NRZ
TX Power	+10dBm
Operating Voltage	1.8V ~ 3.6V
TX Current OOK	7mA Typical
Package Type	MSOP10

RX: MICRF211AYQS

Frequency Range	380 ~ 450MHz
Data Rate	10kbps
Sensitivity	-110dBm
Operating Voltage	3.0V ~ 3.6V
RX Current	6mA
Package Type	QSOP16

Applications:

MICRF112YMM
- Automobile alarms
- Consumer/Entertainment remotes
- Environmental monitoring
- Entertainment remotes
- Lighting and power switches
- Remote Keyless Entry (RKE)
- Tire Pressure Monitoring System (TPMS)

MICRF211AYQS
- Automobile alarms
- Environmental monitoring
- Entertainment remotes
- Lighting and power switches
- Remote Keyless Entry (RKE)

For more information, contact your local Micrel sales representative or visit us at www.micrel.com/ad/micrf112_211.
Literature: 1 (408) 435-2452 Information: 1 (408) 944-0800

MICREL®
Innovation Through Technology™

www.micrel.com

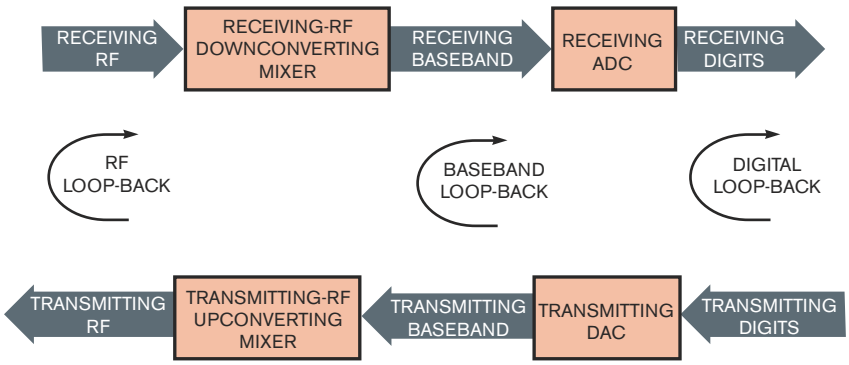

Figure 2 A series of loop-back connections renders this CSR transceiver more visible to the debugging engineer.

tems can collect large amounts of data or perform at-speed checks for the bring-up team. "It really helps to have your own ATE capability in-house," he reports. "We have 28 systems, and we add a new one about every quarter. They are all primarily for debug, and new silicon has priority on the equipment."

Even with access to ATE systems, however, some signals and states are inaccessible to a probe card. They require another strategy: internal stimulation and logic analysis. Sometimes, the only effective way to stimulate a block at speed and to capture its behavior is with circuitry built into the block itself. Bay, which organizes its network-processing chips as strings of independent processors, widely uses this technique, according to Chiang. Important blocks may have their own debugging kernels, including single-step and breakpoint capability and trace buffers to capture internal state in real time. This approach permits what Jun-wen Tsong, Bay's director of logic design, describes as a multistage-verification process.

"First, we exercise the chip at the module level. In this mode, each module is isolated: We can inject enough state to start it running and then observe its behavior stand-alone." These tests must be conducted at full clock speed for them to be accurate. In this way, the designers wring out each stage in the string of processors. At this time, designers also isolate the I/O ring from the internal blocks so that inputs go directly to the output FIFOs. Once Bay's designers have independently verified the I/O ring and the internal blocks, they reunite the two and test the chip as a whole.

Obtaining meaningful data from full-chip, at-speed operation takes planning, however. The debugging kernels in the individual processors must be able to recognize not only local instructions and data words, but also big-picture data types important to the operation of the chip: packets and data grams. Additionally, a 36-bit bus runs through the entire chip, bringing out critical signals from any block to the package pins in real time, so that debugging engineers can watch the operation of blocks while the chip is processing packets at full speed. In addition, hardware monitors specific assertions, such as FIFO full/empty assertions, in real time. Broadcom has a similar approach. Hublitz says that the company's wireless-LAN chips have enough internal debugging hardware that engineers can track vector magnitude all the way through the chip, from input to baseband to output.

Once they've isolated a problem to a function within a block, debugging engineers can turn to a lower level of diagnostic tools, based on familiar DFT strategies. "We have clock control for triggering and single-stepping within the blocks, and scan for what we consider to be the important signals," explains Bay's distinguished engineer and silicon architect, Barry Lee. "Ideally, we can see exactly how a particular pipeline is executing down to the pin and register level."

THE ANALOG CHALLENGE

When analog circuits are involved, everything changes. "We partition the analog apart from the digital circuits for debug," Lee explains. "The debug tech-

Compact. Versatile. Connected.

STELLARIS® LM3S6965 Ethernet Evaluation Kit

features the Stellaris® LM3S6965
ARM® Cortex™- M3-based microcontroller
with fully integrated 10/100 Ethernet (MAC+PHY).
The kit includes evaluation versions of popular software tools, spanning the design spectrum from evaluation...to prototyping... to application-specific design. Perfect for a variety of applications including connected motion, embedded web servers, network-based sensing, and building automation.

...just $79!

LUMINARY MICRO™

www.LuminaryMicro.com

niques for the two are different. In the analog world, you want to peel open the loop-back paths. And you may have to bring everything out to the package pins." Because the primitives of activity in analog circuits are not synchronized to a clock, there is no way to capture them.

Analog, like digital, designers have seen their ability to probe and experiment with their designs vaporize with shrinking geometries, observes Analog Devices' fellow, Paul Ferguson. "We used to simply have a laser cutter on the probe station, and, if we wanted to modify a circuit, we did. Later, as geometries shrank, we moved to focused ion-beam systems. But they are only really useful at pitches of about 250 nm or greater. That means, practically speaking, that if you are working in a 65-nm process, you can make changes only to the top two metal layers."

This situation has led to an interesting change in analog-design style, according to Ferguson. "We were recently doing a PLL for a 90-nm design and found that we would have to complete the VCO [voltage-controlled oscillator] before we had really proper models. So, we brought some lines for adjusting the gain and some other parameters up to the top metal layers where we could get at them. It really helped in the debug process."

Matt Ball, mixed-signal project engineer at single-chip-radio vendor Jennic, also emphasizes the importance of bringing up critical analog signals where you can get to them. "We put in as much programmability and digital trimming as possible," he says. "Some things have to be metal-trimmed, though, and we bring all those locations up to a single mask level for accessibility."

Beyond bringing up live signals to the top metal layers or to the package pins, today's analog designers have other weapons at their disposal to set and observe the state of their circuits. The foremost is the reality that at fine geometries, there is intense cooperation between the analog circuits and the digital circuits that calibrate and monitor them.

CSR's McCall says that in its designs, ADCs monitor many points in the analog circuitry for digital supervisory circuits. These points naturally give debugging engineers access to the behavior of the analog section simply by bringing

MORE AT EDN.COM

Go to www.edn.com/070621df and click on Feedback Loop to post a comment on this article.

the outputs of the converters to the outside of the package. "Often, important analog signals are going to get digitized at some point anyway," says Ball. "So, why not pick up the samples, filter them with your on-chip DSP, and output the result so we can see it?"

Designing a filter or an amplifier so that digital circuits can trim all its important electrical characteristics might seem like massive overkill. But it can make the difference between first-time-working silicon and silicon that has to have two new metal masks before debugging can even start on the digital portion of the design. And, given the increasing variability with which designers must contend at processes smaller than 90 nm, this much digital trimming may be necessary anyway, just to yield an adequate number of working chips.

But how do you see to trim? For signals of reasonable precision and frequency—the IF (intermediate-frequency) signals in a radio chip, for instance—you can simply use careful routing and analog multiplexers to bring the signals out of the package in test mode. "At IF, the buffers can be quite good," Ball says. "You can get signals from important nodes out to pins and see what you need to see." Analog Devices' Ferguson agrees. "For debug purposes, you often don't need more fidelity than an analog multiplexer can provide; you can see oscillations or a 20% gain error quite well."

If you can't bring the signals out of the package, you can sometimes route them to on-chip data converters. "We will typically have an auxiliary ADC on the chip to monitor die temperature, battery voltage, and the like," Ferguson explains. "We put a huge [multiplexer] on the front of it and use it to examine other nodes in the analog section during debug. But be wary: The extra measurement circuits you put in can corrupt something else. For instance, turning on the multiplexer to observe a node may add just enough capacitance to stabilize a circuit that was oscillating. And, if you inadvertently cross power domains with your debug signals, you can introduce

DON'T WASTE YOUR TIME, MONEY, AND EFFORT ON YESTERDAY'S RF SOLUTIONS

RF SIGNAL GENERATOR ■ RF SIGNAL ANALYZER ■ RF POWER METER

The Only Available Lab Quality, Handheld **RF Power Meter**

Introducing the industry's fastest, most flexible, and most cost-effective RF solutions.

Faster test execution **speed**.
Quickly change test requirements with **SDR architecture**.
Unmatched ease-of-use with **touch-screen GUI**.

Visit www.keithley.com/at/295 to request a *FREE* copy of *Speaking RF: Wireless Communication Terminology.*

KEITHLEY

A GREATER MEASURE OF CONFIDENCE

publisher@keithley.com • 1-800-588-9238 • Fax 440-248-6168

sneak current paths that you hadn't counted on."

Ball echoes the warning that you have to be selective with this approach. "The 10 or 20 fF you incur by buffering an analog signal can change the behavior of a node," he concurs. Jennic tends to build its debugging provisions around only those areas that have presented problems before, such as bandgap cells. "We tend to put in bypass circuits, just in case," Ball adds. This conservatism can minimize the chances of disrupting functional circuits.

With planning, good luck, and a bit of elegance, it is possible to reuse functional blocks for debugging purposes. Many analog signals terminate in a data converter and so are at least partially observable through it. Ferguson points out that you can easily switch sigma-delta converters to operate as filters, providing visibility into the incoming analog signal. Alternatively, you can carefully route their bit streams out to pins, giving observability of both sides of the converter. Once you digitize the data, you can use CPU or DSP blocks to condition and compress it or to test assertions against it.

It is also possible to build debug intelligence—the equivalent of a simple network analyzer, for example—into a block. A loop-back path can use a transmitter and a receiver to check each other (**Figure 2**), and a little more circuitry can extract the resulting analog waveforms. "On our Gigabit PHY [physical-layer] designs, we are capturing some analog data in the PHY block," Broadcom's Hublitz reports.

OUTLOOK FOR THE FUTURE

It is not hard to imagine a scenario in which, during early system design, each functional block receives enough self-test capability to diagnose itself during at-speed operation, down to the level at which the DFT-scan chains can take over. This approach would usually require an input buffer or a signal generator to stimulate the block, an output capture register or ADC to observe it, and enough internal breakpoint and trace capability to reveal the block's internal workings. Some SOC design teams now do this planning. Actual implementation then becomes a compromise between the level of debugging support the architects feel necessary and the amount of overhead the design can tolerate.

Taking the concept a step further, a designer of an elegant system can come up with ways to repurpose some functional blocks to serve as signal sources or capture devices for other blocks. The auxiliary ADC is an excellent example, but more such opportunities exist. The addition of a fast data converter, for instance, might turn a signal-processing block into the equivalent of a network analyzer or a digital oscilloscope. A few additions to the control logic might convert a buffer-SRAM array into a trace buffer.

In this manner of thinking, the functional blocks on the chip become a pool of debugging resources, available with only the resetting of a few multiplexers and mode switches. But this process requires forethought. Such an organization impacts floorplanning and global routing, so it must occur at the outset of the design rather than during late implementation.

And it is a process that could stand some tool support, as well, Ferguson argues. Elaborate tools exist to automatically install such structures as scan chains, scan controllers, and vector generators. And, DFT hardware can be indispensable in register-level diagnosis of problems. But no such tool support exists to create debugging structures. Ferguson, for one, would like at least to see an inspection tool that would rate a mixed-signal block for observability and controllability, as well as to scan it for simple mistakes. Ideally, a tool could traverse a design and propose a debugging architecture and process. But that matter is for the future. **EDN**

FOR MORE INFORMATION

Analog Devices
www.analog.com

Bay Microsystems
www.baymicrosystems.com

Broadcom
www.broadcom.com

Cambridge Silicon Radio
www.csr.com

Jennic
www.jennic.com

LogicVision
www.logicvision.com

You can reach Executive Editor Ron Wilson at 1-408-345-4427, ronald.wilson@reedbusiness.com.

RAQ's

Rarely Asked Questions

Strange but true stories from the call logs of Analog Devices

Contributing Writer John Ardizzoni is an Application Engineer at Analog Devices in the High Speed Amplifier Group. John has been with Analog Devices for 4 years, he received his BSEE from Merrimack College in 1988 and has over 27 years experience in the electronics industry.

Should you butcher your circuit or your test gear? Capable engineers should be prepared to question, and modify, anything, not just their own designs.

Q. While probing a new high-speed amplifier board, I noticed the output had huge overshoot and was ringing like a bell. I had paid close attention to the layout, optimizing supply bypassing, grounding, parasitic reactances, and trace terminations. Where did I go wrong?

A. You've identified the common causes of ringing and overshoot (decoupling, grounding, parasitics, and terminations). Assuming careful design, there may be nothing wrong with your circuit (really!). The problem may be with testing it, and the culprit may be your oscilloscope probe.

Many high-speed amplifiers have difficulty when driving capacitive loads (which introduce a pole in the feedback response, lowering phase margin and causing instability—but that's a RAQ for another day). Probes can add about 10 pF (for a typical 10× passive probe) to the measurement node. This additional capacitance may well cause ringing and overshoot. See if a lower capacitance probe helps. Active probes typically have less capacitance than passive probes, so try using one. Or you could try a passive probe with a higher attenuation factor (100x); these also have lower capacitance.

It's more likely that the inductance of the wire to the ground clip of the probe is responsible, though. The parasitic inductance of the wire and the capacitance of the probe form a resonant (tank) circuit. Tank circuits are commonly found in… that's right... oscillators. Fast edges may have enough energy to excite the tank circuit and make it ring.

So, be brutal and amputate the wire. To do so you will need to dismantle the probe. Remove (unscrew) the plastic sleeve that surrounds the probe tip, exposing the outer metal jacket of the probe—which is its ground connection—and then remove the ground clip lead. This leaves a stripped down probe with an exposed ground shield. Such a probe is very suitable for high-speed measurements.

To use it, simply place the tip on your test point while touching the outer metal to ground. If you can't conveniently make a direct ground connection, wrap a short length of bare wire around the probe's outer metal jacket a few times. The free end of the wire, which should be as short as possible, is touched to the closest ground point.

You will be impressed with the improvement in the measurements you can make with this modified probe. It may not be pretty, but it works like a charm.

To learn more about high-speed amplifiers,
Go to: http://rbi.ims.ca/5393-101

Have a question involving a perplexing or unusual analog problem? Submit your question to:

raq@reedbusiness.com

For Analog Devices' Technical Support, call 800-AnalogD

SPONSORED BY

ANALOG DEVICES

BY ROBERT CRAVOTTA • TECHNICAL EDITOR

EXPLORING MEMORY ARCHITECTURES: PILLARS OF PROCESSING PERFORMANCE

PROCESSOR-BASED SYSTEMS RELY ON MULTIPLE, HETEROGENEOUS MEMORY SUBSYSTEMS TO DELIVER BETTER SYSTEM PERFORMANCE, POWER, AND COST EFFICIENCIES.

A key performance characteristic of processor architectures is how much application-specific work they can perform per unit of time. The EEMBC (Embedded Microprocessor Benchmark Consortium) benchmark, unlike Dhrystone MIPS (millions-of-instructions-per-second) scores, describes the performance of processors executing tasks in embedded-system applications. Version 1.0 of the EEMBC benchmarks does not capture the system-level influences, such as the memory subsystems, of processing performance because the benchmarks can often run from within the processor's L1 cache. However, EEMBC's second-generation, system-level benchmarks for networking and digital entertainment more realistically stress even those processors with large cache memories.

It is increasingly important to consider the system-level impact of the memory subsystems of a processor because the types and sizes of the memories and access methods in the system define the upper limit of a processor core's performance. According to Gerard Williams III, a fellow at the processor division of ARM, a processor with an ideal memory system never misses in the cache and has ideal access to the bus. Chip designers must first understand the processor's IPC (instructions-per-cycle) capability and then try to implement a memory architecture that minimizes the performance loss. This performance loss can result from caching or memory-access effects, such as miss rate due to capacity misses, cache size, or conflict misses.

A well-matched memory subsystem can in the best case merely preserve the processor's maximum IPC rate, whereas a mismatched memory subsystem can drastically reduce the processor's performance by starv-

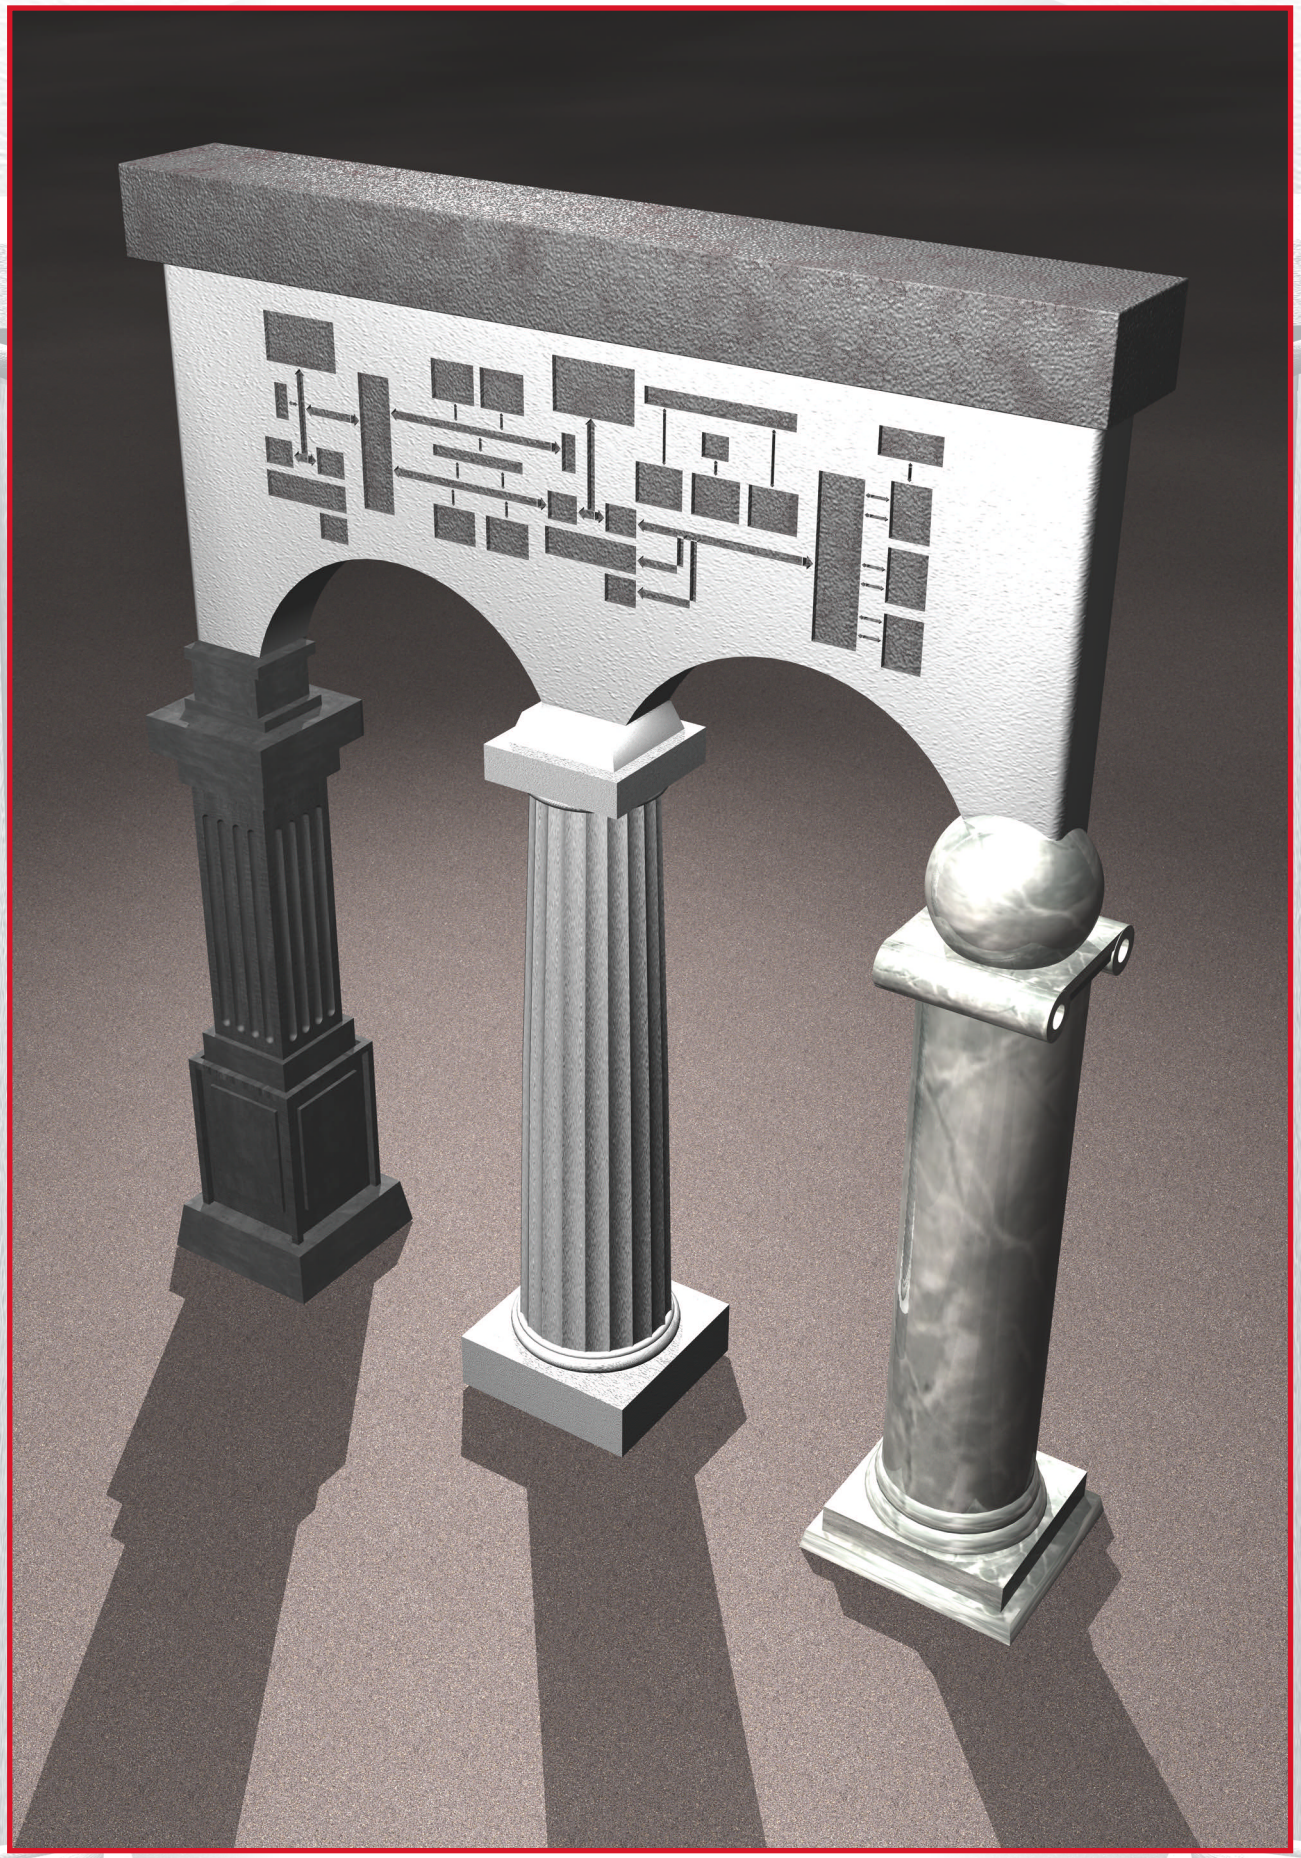

ing and idling the core's execution units. Building and implementing memory subsystems that do not adversely affect the processor core's performance continue to be challenging because the performance gap between processor logic and the memories is widening with each process-technology reduction. In essence, the improvement in memory-access latency, the time it takes to receive the first bit of a memory request at each process-technology step, gets smaller than the commensurate clock-rate improvement of the processor's core logic.

Likewise, the best performance impact software developers can accomplish with insightful placement of program instructions and data within the memory subsystem is to preserve the processor's maximum IPC rate. However, mismatching the placement of program instructions and data in the memory subsystem with the application's usage scenario can significantly degrade the processor's performance. Freescale's application note on preventing M1 memory contention provides an example that exhibits a worst-case 54% processor-performance degradation due to memory contention that the developer can avoid with better placement of the data buffers (**Reference 1**).

In general, compilers and profiling tools can provide limited assistance with global optimizations for placing instruction and data in memory. Green Hills' optimizing compilers support the reordering of functions in memory to optimize cache hits. Texas Instruments' CodeSizeTune profile-based-compilation tool assists a developer's ability to explore configurations by automating the building and profiling of different versions of the software with different compiler settings that affect code size and execution speed (**Figure 1**). In general, though, for many high-efficiency and real-time-constrained systems, the burden falls on the software developer to understand the memory subsystem to avoid incurring unnecessary BOM (bill-of-materials) costs because of the system's inefficient use of processing and memory resources.

TOLERANCE OF LATENCY

A primary concern when implementing memory architectures is making the processor tolerant of the access latencies of the memories the system uses. A properly designed memory subsystem can mask much of the system's memory-access latency and provide a sufficient read/write throughput rate—that is, the memory-access time for subsequent data in the same block of data—to support continuous access. This scenario avoids starving the processor's execution units of instructions and data. Memory designers must also balance masking the memory's access latency against the silicon area of the memory, the total power the memory consumes, and the ease of use of the memory by software developers and tools (see **sidebar** "Ease of use").

Direct drivers of memory-access latency are the time it takes to perform address decoding, activating the appropriate word line, sensing the bit line, and driving the output from the sense amps. The address-decoding latency is the time it takes to latch the address and decide which word line requires activation; this process takes O(n log n) time as a function of the size of the memory's row and column addresses. So, as the memory structure gets bigger, so too does the time to decode the addressing. The word-line-activation latency is the time it takes to raise the word line; it is primarily an RC delay related to the length of the line, with longer lines driving longer delays. The bit-line-sensing latency is the time it takes for the sense amplifiers to detect the cell contents. The bit-line architecture, the RC of the sense-drive line, the cell-to-bit capacitance ratio, and the sense-amplifier topology all affect bit-line-sensing latency. The output-driving latency, an RC delay, drives the time it takes to propagate the data from the sense amps to the output.

Memories and the logic to manage them dominate the silicon area of many processor-based devices. As a result, memories can be the largest components of the device's silicon cost and the largest consumer of both dynamic and static power in the system. The many types of volatile and nonvolatile memory available involve numerous trade-offs, and the system designer must balance and manage the key parameters to deliver good enough memory performance at lower cost and power consumption.

To balance masking memory-access latency, silicon cost, and power consumption, processor-based devices usually rely on a hierarchical memory struc-

AT A GLANCE

- Memory subsystems and software can best impact a processor by merely preserving its theoretical maximum performance.
- The processor's architecture is a first-order driver of the options available to design the memory subsystem.
- A processor's tolerance of latency is a balance of the speed, cost, and power consumption of implementing a hierarchy of fast and expensive or slower and cheaper memories.
- Preserving the mechanisms that provide tolerance of memory-access latency is still a mostly manual exercise for developers.

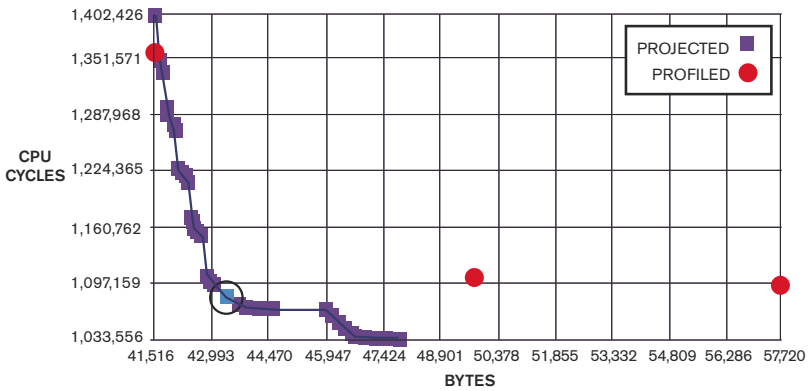

Figure 1 Software developers can graphically explore and compare the effect of configurations and compiler settings with CodeSizeTune to best select the balance between code size and processing cycles (courtesy Texas Instruments).

Broaden Your Oscillo Scope

with NI PXI Digitizers

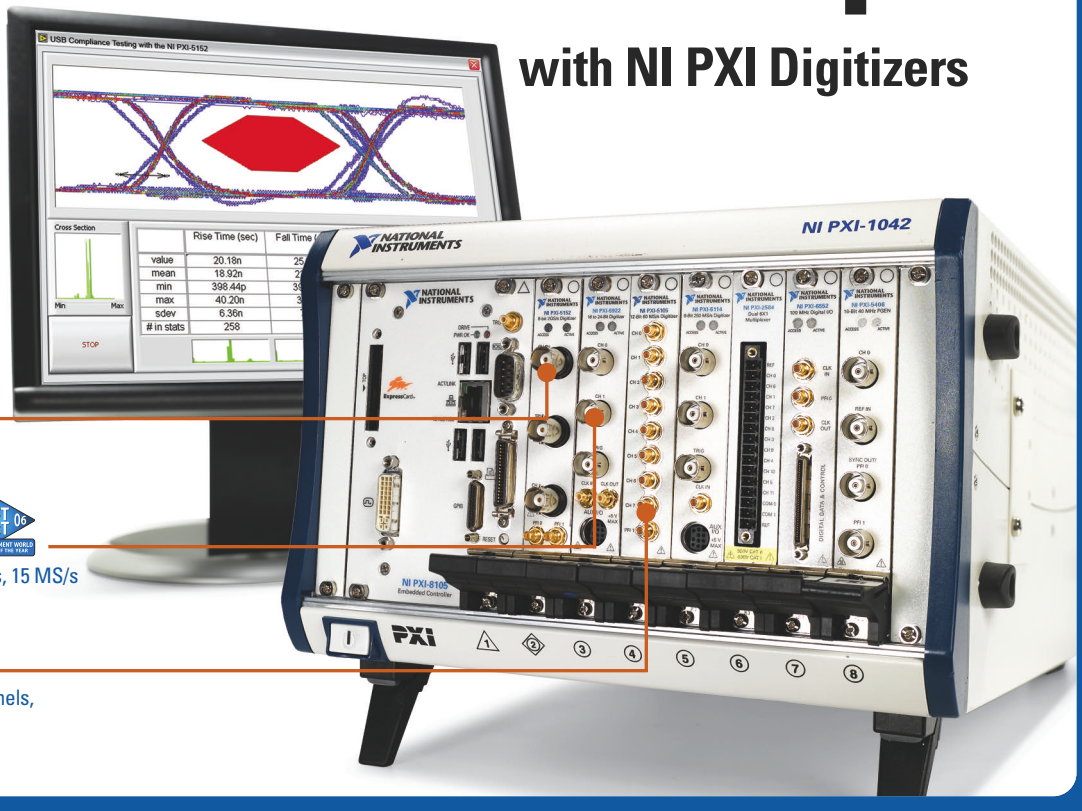

Higher Speeds
2 GS/s, 300 MHz
NI PXI-5152 *(NEW)*

Higher Resolution
From 24 bits, 500 kS/s to 16 bits, 15 MS/s
NI PXI-5922

More Channels
60 MS/s, 8 simultaneous channels,
Lowest price-per-channel
NI PXI-5105 *(NEW)*

The most complete range of modular PXI and PCI-based digitizers/oscilloscopes.

Integrated Software

- Interactive benchtop measurements
- Single NI-SCOPE driver for all NI oscilloscopes
- Support for National Instruments LabVIEW and LabWindows/CVI, Microsoft Visual Studio, and NI TestStand

Tight Synchronization

- Create a 136-channel oscilloscope for high-channel-count applications
- Add digital lines to your oscilloscope for mixed-signal applications

From DC to RF, NI modular instruments power measurements from prototype to production.

Define your test system with NI modular instrumentation.

Multimeters	7½ digits, 1000 V
Audio Analyzers	24 bits, 500 kS/s
Digitizers/Oscilloscopes	**8 to 24 bits, up to 2 GS/s**
Signal Generators	Up to 16 bits, 200 MS/s
High-Speed Digital I/O	Up to 400 Mb/s
RF	Up to 6.6 GHz, 20 MHz RTB
Switching	Multiplexers, matrices, digital I/O, counters
Multifunction I/O	Analog I/O, digital I/O, counters

To read the "Top 10 Things to Consider when Selecting a Oscilloscope/Digitizer," visit **ni.com/oscilloscopes**.

(800) 891 8841

NATIONAL INSTRUMENTS™

ture that places smaller amounts of faster yet more expensive memories closer to the processor core and larger amounts of slower and less expensive memories farther from the processor core (**Figure 2**). After the processor registers, which are the fastest and scarcest memory resources in the system, memory hierarchies may use local memories or TCMs (tightly coupled memories), multiple layers of caches, and volatile and nonvolatile on- and off-chip memories.

Modern optimizing compilers are competent at managing the use of the processor registers, but they are weaker at managing and optimizing the other memories. This situation is due partly to the fact that optimizing the use of the registers works well as a tactical exercise with a local view of the program code. To optimize the use of the other memory structures, such as the TCMs, in a processor-based system requires a more global view of the system, and this capability is still emerging in most compilers.

Local memories or TCMs connect

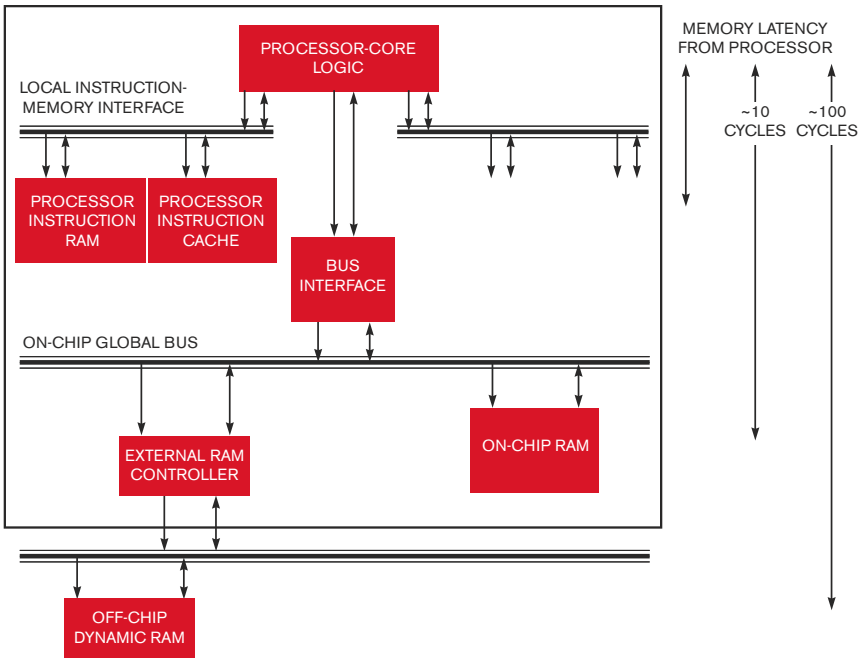

Figure 2 The memory-access latency may increase tenfold for each level of the memory hierarchy from the processor (courtesy Tensilica).

EASE OF USE

Ease of programming is a feature that is important to software developers. A flat address space that hides the memory hierarchy makes it easier for the developers to program. Brian Boles, digital-signal-controller-division technical fellow at Microchip Technology, shares that, "Generally, it is easier for a compiler to target an application to a generalized memory structure." It is more difficult for compilers to optimally allocate the code and data to application-specific memory structures without visibility to the global and dynamic characteristics of the application code.

For sophisticated applications that require operating systems, such as Linux, the memory architecture may need to support virtual addressing. However, a consideration for developers using heavy operating systems to meet time-to-market schedule pressures is the potential loss of insight into how to partition the software to take advantage of on-chip resources to save power and cost. Part of the conflict is balancing and determining how much of the on-chip memory the operating system requires to operate mainly out of on-chip memory and how much of the memory this approach leaves for the application. "To date, general-purpose operating systems do not have hooks to specify the complete physical-to-memory-system mapping so as to facilitate the most optimum use of the underlying memory system," says Phil Ames, segment-marketing manager for the Embedded and Communications Group at Intel. "However, it is common in embedded designs to hand-tune the software to make best use of the memory system."

Managing each different class of memory may require specialized software. For example, small-block NAND flash (528 bytes per page) usually requires different flash-management software from large-block NAND flash (2112 bytes per page). One approach to manage this situation is to modularize the software into layers so that the software developer has to rewrite as little as possible when changes are necessary. According to Doug Wong, member of the technical staff at Toshiba's memory-products group, "NAND flash appears to be the first commodity memory to add significant intelligence to the memory device itself in order to make it easier to use." Toshiba's LBA-NAND and eMMC-compliant embedded NAND both contain built-in controllers that perform NAND-management functions, such as block management, wear-leveling, logical-to-physical-block translation, and automatic error correction. This approach significantly reduces the burden on the system architect or software engineer in managing the NAND-flash device for an FFS (flash file system) or for FTL (flash translation layer).

RENESAS M16C

M16C- The broadest platform with true code and pin compatibility

42 to 144pin, 24K to 1M Byte Flash & 1K to 48K Byte RAM

Renesas Technology

No.1* supplier of microcontrollers in the world

presents a wide range of M16C microcontrollers. M16C is the only fully code-compatible platform in the industry that addresses the entire 8-bit through 32-bit price/ performance application space.

Your application can range between 24K Bytes and 1M Bytes of code size, with between 42pins and 144pins of package size, while keeping the same code base and development tools.

The consistency and compatibility of the M16C Platform enables you to reduce your development time while still allowing the flexibility to adapt to changing system requirements.

M16C Product Roadmap

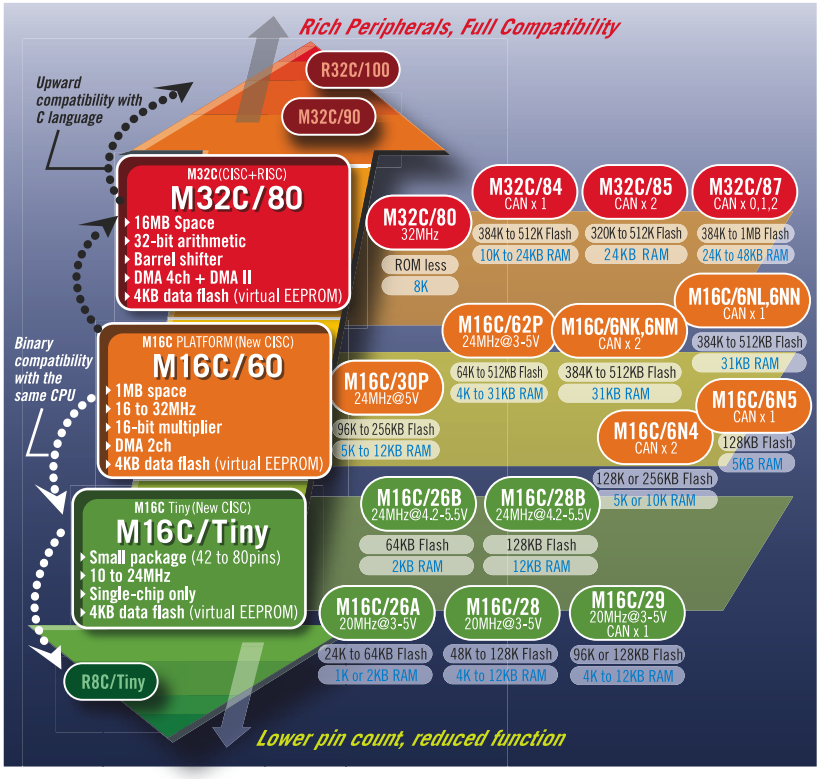

HOT Products | **M16C/29 Group**

M16C/29 Block Diagram

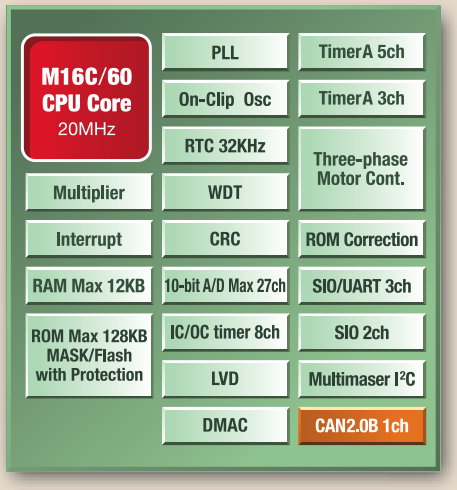

Top Reasons To Select M16C

- **Broad Platform** - Wide range of selection; 24KB to 1MB on-board Flash, 42pin to 144pin
- **Compatibility** - Pin and Code compatibility across the platform enables you to upgrade/downgrade MCU in same package without changing the board design and peripherals.
- **Powerful** - High speed interrupts (M32C: 0.61u sec in 32MHz) , optimized instructions for 1cycle operations, Hi-speed hardware multiplier.
- **Versatile** - Specialized on-chip peripherals; CAN, LIN, 3ph-PWM, DMA, etc
- **Efficient** - Multiple clock sources and multiple power saving modes
- **Quiet** - Built-in noise cancellation circuits provides excellent EMI/EMS characteristics that outperform IEEE standard spec
- **Easy** - Same Tool Chain for evaluation and development across the board
- **Reliable and secured** - Trusted flash and built-in fail safe features such as oscillation-stop detection circuit, protect registers, enhanced WDT, etc.

*Source:Gartner Dataquest (April 2006) "2005 Worldwide Microcontroller Vendor Revenue" GJ06333

#1MCU REACH FURTHER

Get Started Today -

Go online and register to be eligible for a FREE Starter Kit

www.america.renesas.com/ReachM16C/a

Renesas Technology Corp.

Everywhere you imagine. RENESAS

to the processor core through local- or dedicated-memory buses for access performance similar to that of cache memories. Memory-access determinism is a key difference between TCMs and caches. Cache-line locking manually and temporarily enables a cache at the line level to act as a TCM. Program-instruction and code access through a TCM is deterministic, but, with a cache, the designer must consider the worst-case scenario of cache misses. "A typical rule of thumb for the penalty of a cache miss is an order of magnitude longer access latency than the previous level," says David Fisch, director of architecture at Innovative Silicon. "An L2 memory access has 10 times longer latency than an L1 cache access, and it also has a 10 times shorter latency than an L3 memory access." However, using TCMs puts the onus on the software developer to manually manage that memory space, usually with a DMA controller, so that the necessary code and data are in the TCM when the processor needs them.

Cache comprises less of a faster memory to mask the latency of a larger amount of slower memory. Slower memory is denser and, hence, cheaper. Caches rely on the premise of temporal and spatial locality to mask the memory-access latency of the slower memory. "Temporal locality" describes the premise that, if the processor requests some data, then the processor will soon need that same data again. By keeping a copy of the data in its storage, the cache can avoid going to the slower memory. "Spatial locality" describes the premise that, if the processor requests code at a memory location, then it is highly likely that the next processor request will be the code at the next memory location or close to that location. By prefetching some amount of the data near the currently requested data at the same time as the original fetch, the cache can have the next few data locations in its store without incurring the latency of another fetch from the slower memory.

Larger caches usually mean fewer cache misses at the cost of more silicon area. Increasing the cache-set associativity, which refers to the number of locations where a given memory can reside in the cache, almost always reduces cache misses. The cache line's length

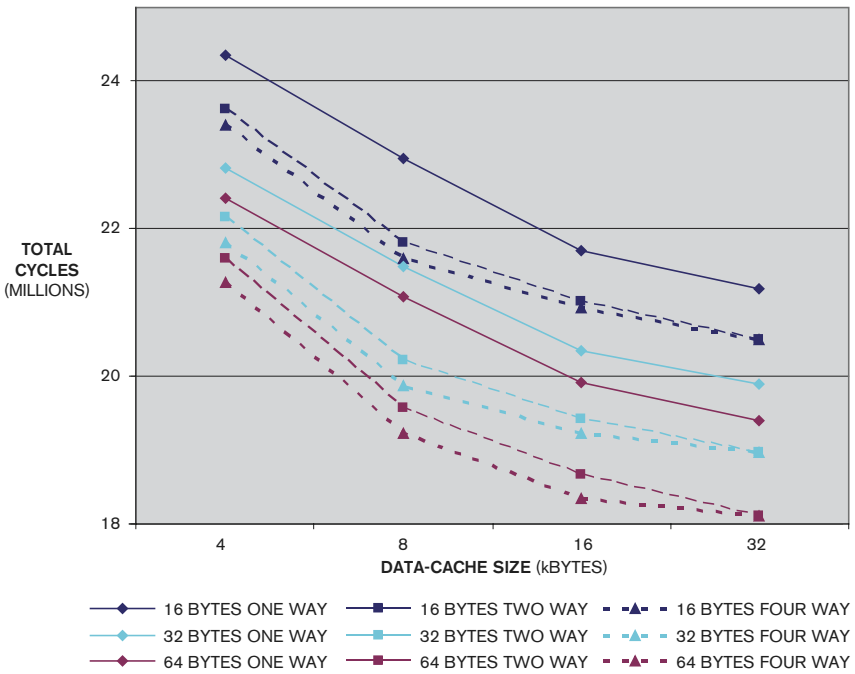

Figure 3 A data-memory-stall graph based on different cache sizes, line lengths, and associativity levels for a JPEG-encoding application shows cache configurations from a 4-kbyte, direct-mapped, 16-byte-line configuration with a load-miss rate of 13.4% to a 32-kbyte, four-way-set-associative, 64-byte-line configuration with a load-miss rate of 1.9% (courtesy Tensilica).

can vary positively or negatively based on an application's behavior. According to Bill Huffman, chief architect at Tensilica, "Configuring caches is an iterative task that is highly dependent on the application set that will execute on the processor."

Balancing the various cache parameters can be a complex process that involves trade-offs between silicon area and miss rates (**Figure 3**). In the **figure**, the explored cache configurations span from a 4-kbyte, direct-mapped, 16-byte-line configuration with a load-miss rate of 13.4% to a 32-kbyte, four-way-set-associative, 64-byte-line configuration with a load-miss rate of 1.9% for a JPEG-encoding application (**Reference 2**). Even though the larger cache is better, there is a diminishing return of benefits for the 32-kbyte cache. There is a larger performance benefit from increasing the cache-line size than from doubling the size of the cache; the longer cache lines reduce silicon cost. Also, although higher cache-set associativity is better, in this example, going from two-way to four-way-set associativity yields fewer benefits. In short, no clear rule of thumb exists for configuring caches.

DECISION DRIVERS

The processor-core architecture is the first-order driver of the memory-architecture options that a designer has. The reason is that the designer builds the core with assumptions about how the memory components interface with and complement the core. Von Neumann and Harvard architectures are two common processor architectures that model and implement different ways to view and access memory. Processors based on a von Neumann architecture model the system memory as a single storage structure that hosts both the program instructions and data; a single bus interface services all program and data accesses. Processors based on a Harvard architecture model the system memory for program instructions and data as physically and logically separate storage structures with separate bus interfaces—one for instructions and the other for data. The Har-

Intersil Video Drivers

High Performance Analog

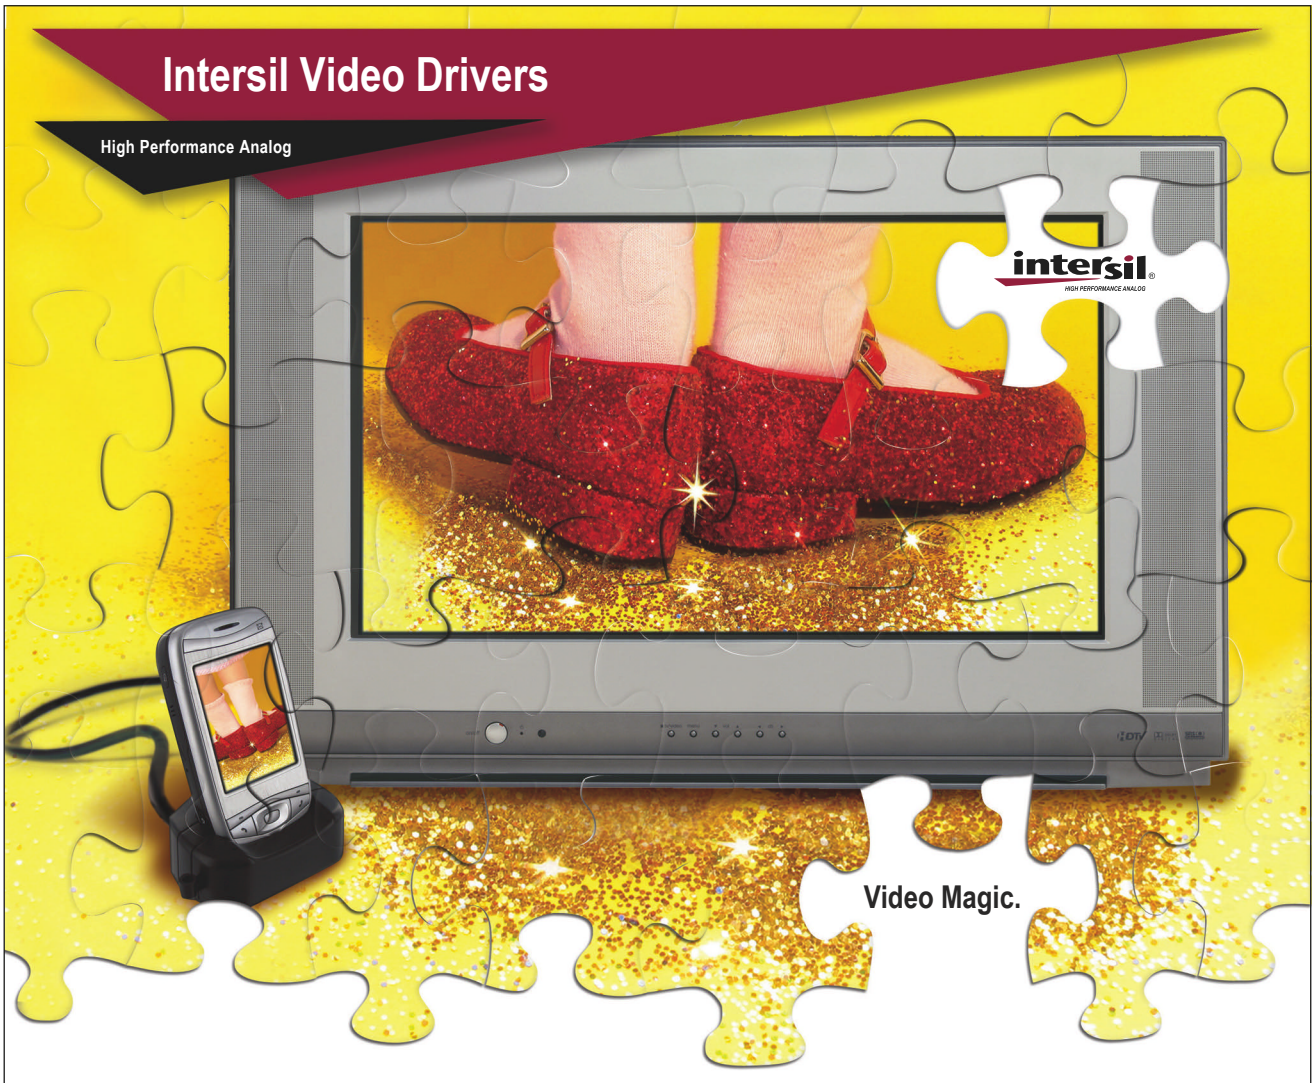

Intersil's **Video Drivers in space-saving µTQFN and chip scale packages** deliver quality video from tiny little devices.

Intersil's ISL59114, ISL59115, ISL59116 and ISL59117 reconstruct S-video to composite video using a fraction of the power consumed by competitive devices.

9MHz reconstruction filter removes aliasing noise from video signals. Perfect for composite and S-video signals.

Summer amplifier (ISL59114/ISL59116) creates a composite video signal.

Integrated gain of 2 buffer gives great output current drive and compensates for double-terminated video loads.

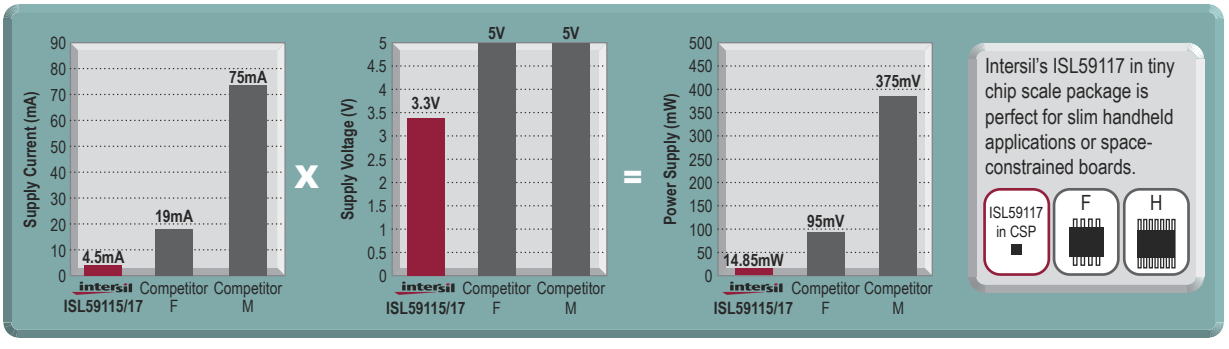

Intersil's ISL59117 in tiny chip scale package is perfect for slim handheld applications or space-constrained boards.

Go to **www.intersil.com** for samples, datasheets and support

Intersil – Amplify your performance with advanced signal processing.

intersil®
HIGH PERFORMANCE ANALOG

vard architecture supports simultaneous access for program instructions and data, whereas the von Neumann architecture does not.

To choose an optimized memory design, a designer must also understand the application's behavior and requirements. Considerations for the memory design are: How will application data enter and exit the system, and will the processor directly load the data or will an external agent, such as a DMA controller, load the data into the processor's local RAM? You must ask similar questions about outputs: Will the processor directly drive the output ports, or will an external agent, such as a DMA controller, transfer the data from the processor's local RAM to an I/O interface. Other questions include: What is the application's start-up scenario, can the system make efficient use of special memory interfaces, and can the on-chip-memory resources accommodate all or even just the performance-sensitive code and data of the application?

DESIGNERS CAN TAILOR THE PROCESSOR TO THE KNOWN CONSTRAINTS OF THE APPLICATION THEY ARE TARGETING.

The application start-up requirements affect where you can store the initialization code and through what interfaces the system can retrieve it. On-chip OTP (one-time-programmable) ROM is useful for storing boot code because it is small with high silicon density. It supports fast start-up because it needs no wait time after start-up to begin executing. The initialization code could reside in and execute in place from flash memory; it could also reside in off-chip memory and be shadowed into on-chip instruction RAM, which can result in longer system start-up. If the application code and data can reside in on-chip memories, it might be unnecessary to support off-chip-memory interfaces. If the performance-sensitive program code can fit into local memories, the designer may not need to implement caches.

MORE AT EDN.COM

- For an article by Senior Technical Editor Brian Dipert about memory, go to www.edn.com/article/CA181882.
- For a related article about allocating memory and dealing with fragmentation, visit www.edn.com/article/CA421501.
- For an article on handset-memory considerations, visit www.edn.com/article/CA498769.
- Go to www.edn.com/070621cs and click on Feedback Loop to post a comment on this article.

Designers can tailor the processor to the known constraints of the application they are targeting to include only the amount of random and nonvolatile memory the application requires. The sizing and parameters of TCMs, caches, or special memories target the application. Processors targeting a wider set of applications typically implement a generalized memory architecture that includes the maximum resource requirements of the set of applications with variants of the device offering fewer resources to meet lower costs. For systems with similar processor-core architectures, the memory subsystem becomes a higher order driver for differentiating the system's deliverable processing performance, power consumption, and price (see **sidebar** "Multiple options" at the Web version of this article at www.edn.com/070621cs).

Memory controllers abstract the implementation of the memory block they service so that it appears as a data pipe to the processor system. They contain the logic necessary to read the memory block and, as appropriate to the type of memory they service, write, refresh, test, and correct errors. For on-chip memories, the memory controller can manifest the company's proprietary innovations, which differentiate its processor device from a similar device from its competitor. As a result, most processor vendors are unwilling to discuss the specifics of their memory controllers. They hint at techniques for use in memory controllers, including using wide data buses, multiplexed or staggered access of banks, buffering, pipelining, and transaction reordering, as well as specialized

Intersil High Speed Analog

High Performance Analog

Intersil offers a wide portfolio of High Speed Analog, including the industry's first Triple 1GHz Current Feedback Op Amp.

High Speed Analog:

- Op Amps
- Comparators
- Differential Drivers
- Receivers
- Pin Drivers
- A/D Converters
- D/A Converters
- High Speed Serial Comm
- Clock Generation/Synthesis
- RF Products

Go to **www.intersil.com** for samples, datasheets and support

EL5367

World's Fastest and Only Triple 1GHz Current Feedback Amplifier

Handles ultra-high resolution video with room to spare. 1GHz gain of 1 bandwidth and 800MHz gain of 2 bandwidth into a 150Ω load. 6000V/µs typical slew rate.

EL5167

World's Fastest Amplifier in Tiny SC-70 Package

Get blazing speed and reduce board size. We've packed 1.4MHz performance in an SC-70 package. This is the smallest and fastest amp with just 9mA power consumption.

EL5104

Get Current Feedback Performance with Voltage Feedback Control

Intersil has eliminated the nasty tradeoff between ease of use, DC accuracy, and pure speed. Get up to and above 700MHz with virtually unlimited slew rate, almost zero overshoot, and low power consumption.

Intersil – Amplify your performance with advanced signal processing.

intersil®
HIGH PERFORMANCE ANALOG

and speculative access patterns.

In addition to the characteristics of the implemented memory, system-level factors that affect the design and efficiency of memory controllers include how physical addresses map to the internal representation of the memory system; the type of addressing patterns, such as burst, random, and concurrent-access patterns; the mix of reads and writes; and how unused memory enters low-power modes. Its primary usage model normally dictates the architecture of a memory controller, such that a graphics or a multimedia memory controller might optimize for sequential accesses whereas a memory controller for an embedded-communication-system application might optimize random accesses over a large memory span. For those embedded memories with system-level reliability requirements, the memory controller, for additional complexity, can provide ECC (error-correction-code) protection.

The traffic pattern at the memory controller differs significantly between single- and multiple-processor-core systems. A memory controller for a single-core system may use a stream, but the memory controller for a shared memory in a multicore system might need the ability to handle multiple streams and random traffic. For multicore designs, the memory architecture must enable fast and efficient message passing and data sharing between processors. Although different approaches exist for accomplishing these goals, no single configuration is efficient for all types of communications. Fast, point-to-point channels and queues are essential to exchange short and critical messages, whereas shared memory is better for sharing larger data structures. When using shared memories, users need programming support for synchronization and memory management.

As more embedded systems incorporate multiple cores, especially heterogeneous cores, as part of the design, development tools will most likely evolve to better assist developers with the spatial and temporal placement of code and data to sustain better latency tolerance and squeeze out better performance in increasingly complex designs. The development tools must assist developers in better understanding the global behavior of the system and matching that behavior with memory subsystems available in the system. Otherwise, memory and chip designers must continue to incorporate ever-more-complex control algorithms in their memory controllers to invisibly compensate for software signers' and development tools' lack of visibility into the behavior of the memory system. **EDN**

REFERENCES

1 Schuchmann, David, "Tuning an Application to Prevent M1 Memory Contention," Application Note AN3076, Freescale Semiconductor Inc, May 2006, www.freescale.com/files/32bit/doc/app_note/AN3076.pdf?fsrch=1.
2 "How to Optimize SOC Performance Through Memory Tuning," White Paper, Tensilica, www.tensilica.com/products/WP_optimize_memory-tuning.htm.

FOR MORE INFORMATION

Altera
www.altera.com
ARC International
www.arc.com
ARM
www.arm.com
Atmel
www.atmel.com
Cambridge Consultants
www.cambridgeconsultants.com/asic
CriticalBlue
www.criticalblue.com
EEMBC
www.eembc.org
Freescale
www.freescale.com
Green Hills Software
www.ghs.com
Hi-Tech Software
www.htsoft.com
Innovative Silicon
www.innovativesilicon.com
Intel
www.intel.com
Intrinsity
www.intrinsity.com
Kilopass Technology
www.kilopass.com
Microchip
www.microchip.com
MIPS
www.mips.com
NXP Semiconductors
www.nxp.com
Qimonda
www.qimonda.com
Renesas
www.renesas.com
Samsung
www.samsung.com
STMicroelectronics
www.st.com
Stream Processors
www.streamprocessors.com
Tensilica
www.tensilica.com
Texas Instruments
www.ti.com
Toshiba America Electronic Components
www.toshiba.com
Virtium Technology
www.virtium.com
Xilinx
www.xilinx.com

Intersil Display Products

High Performance Analog

Improve your display performance with Intersil's high-performance mixed signal and power ICs.

Whether you're powering a small, medium, or large panels, we've got what you need. Everything from the industry's best programmable buffers and your DC/DC sub-system to RGB backlight drivers and award-winning ambient light sensors.

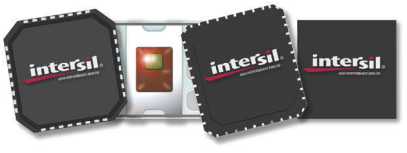

Display Products:

- Ambient Light Sensors
- Programmable Gamma Buffers
- V_{COM}-Adjusting DCPs
- DC/DC Sub-systems
  - Boost
  - Charge Pumps
  - System Power
- Backlight LED Drivers

Go to **www.intersil.com** for samples, datasheets and support

Intersil – An industry leader in Switching Regulators and Amplifiers.

intersil®
HIGH PERFORMANCE ANALOG

Fluid synchronization of design, IP and manufacturing for your custom SoC development.

Leading companies trust Toshiba to efficiently communicate with their internal teams, 3rd-party IP vendors, tool providers, and manufacturing to synchronize the complex flow required for timely and accurate, first-time designs of custom SoCs/ASICs. From calculators to mobile handsets to TVs, Toshiba's seamless end-to-end integration, wide variety of IP and Six Sigma quality processes ensure that our deep commitment to your success enables you to deliver market-leading products on time and on budget. Together we can change the world.

Test your SoC IQ and win premium Noise Cancellation Headphones online. Qualified engineers can win one of four prizes (while supplies last) when you take the advanced technologies SoC IQ Test 2 by S.R. SoCrateez, now at SoCworld.toshiba.com. Everyone wins something!

TOSHIBA
Leading Innovation >>>

BY MICHAEL FOWLER • FAIRCHILD SEMICONDUCTOR

The evolution of SERDES designs for cell phones

APPLYING A SYSTEMS APPROACH TO SILICON DESIGN LEADS TO SUCCESS.

Slimmer cell phones with more functions and features have created a number of system-design challenges. The problems are well-known: managing energy consumption and minimizing noise while moving data around in a handset. The situation provides an opportunity for silicon designers to tackle a system-level-design problem. By applying experience designing SERDES (serializers/deserializers) for devices from laptop computers to cell phones, Fairchild's design team learned valuable lessons about the systems approach.

A product initially targeting cell phones required an in-depth study of the systems and in-depth discussions with several customers. What Fairchild now calls its μSERDES (micro-SERDES) family had design requirements that differed greatly from versions developed previously for computers. Differences between the previous SERDES and the cell-phone application included issues that the Fairchild team faced for the first time, such as serialization of a microcontroller interface, power as a critical design factor, the need for bidirectionality, and the fact that a radio would be an integral part of the application using the SERDES.

The needs and perceptions in this market differed so greatly from those in other applications that the semiconductor supplier and systems manufacturers even used different terms. The systems manufacturers worried about ESD (electrostatic-discharge) events that would temporarily cause the LCD to stop working, and the semiconductor manufacturers worried about ESD events that would permanently destroy the device. The difference in views over this one small question of ESD was sufficiently interesting to warrant the summary that **Table**

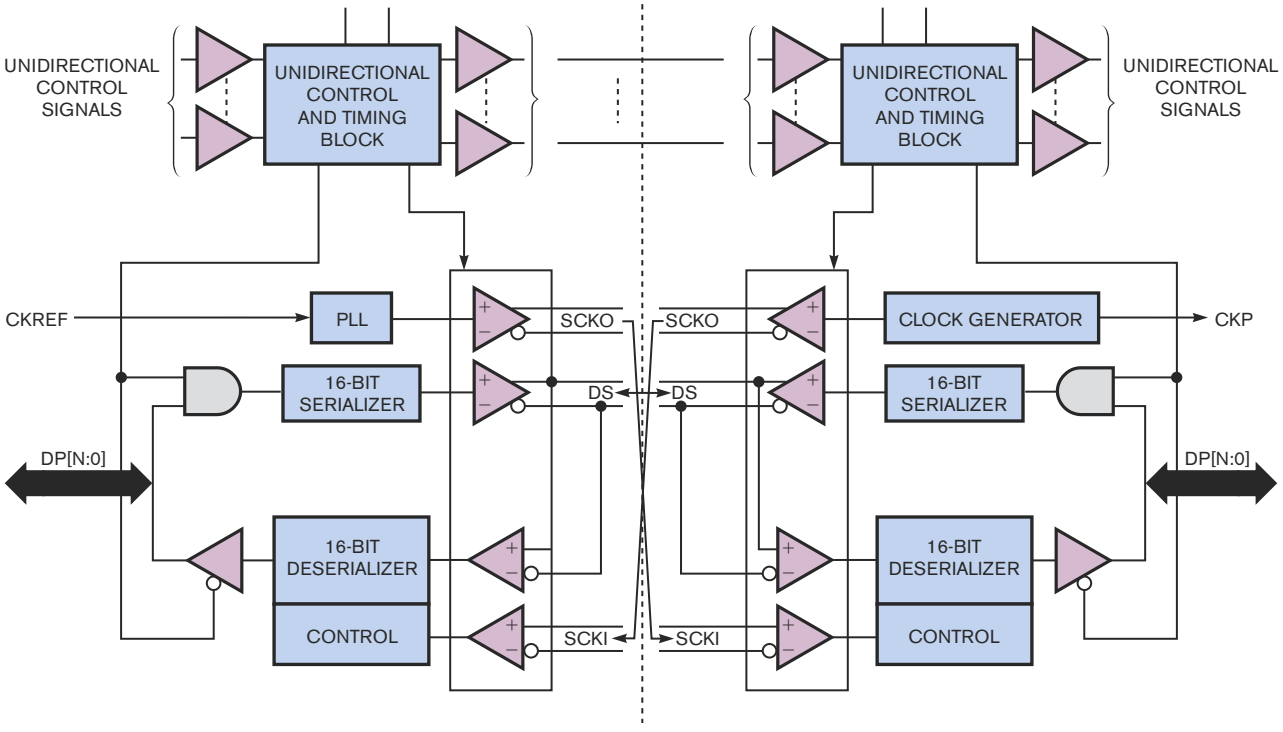

Figure 1 The basic building blocks of a bidirectional, 16-bit SERDES include a PLL and a clock generator.

1 provides. Needless to say, all the first-time technical issues and supplier-user nomenclature and perspective differences made for several surprises along the way.

IDENTIFYING THE OPPORTUNITY

It all began at a customer trade show in 2001, when a product engineer from a cell-phone manufacturer and one of Fairchild's engineers began discussing the flip, or clamshell, phone design. The application intrigued the Fairchild engineer, who had been involved with the novel implementation of SERDES for laptop computers. In many ways, the situations were the same: The displays were in different places from the processors. In the computer application, SERDES technology allowed the application to pass government EMI (electromagnetic-interference) testing, such as FCC (Federal Communications Commission) requirements. The increasingly higher frequency interfaces were using LVCMOS (low-voltage-CMOS) signaling that created a very-high-EMI environment that could no longer meet the EMI limits. With this previous experience, the Fairchild engineer concluded that a SERDES approach could be as important for the cell-phone applications as it had been for laptop computers. As it turned out, the EMI issue was a much larger concern than for the laptop from a systems perspective. However, it was also a much smaller concern, when selecting devices, than two other key issues: susceptibility and the number of wires.

USER INPUT

As the Fairchild design team became more focused on providing a SERDES option for cell phones, its members realized that an active engagement with the key cell-phone customers' design engineers would be crucial to success. So, the team traveled the world to discuss the application requirements with every major cell-phone manufacturer. It soon became obvious that this application was not standard for a SERDES. The cell-phone SERDES would require some radical changes to be acceptable.

After the initial discussions with customers, the team found that the system may not always be synchronous and that there may not be a defined free-running clock to load and transfer the data. An important design consideration was a bidirectional datapath, which was required to serialize a microcontroller datapath; a clock, synchronous with the data, would *not* be available. In addition, write as well as read cycles would be beneficial.

These initial requirements indicated a radically different implementation of a SERDES and turned out to be just the initial design surprises. The team would need to decouple the parallel-input interface from the serial interface as far as timing was concerned. Data would not be coming at convenient,

TABLE 1 COMPARISON OF ESD REQUIREMENTS

Parameter	Previous requirement	Cell-phone requirement
ESD	Destructive test	Phone glitch
EMI	−70 dBm	−120 dBm
Susceptibility	NA	No bit errors

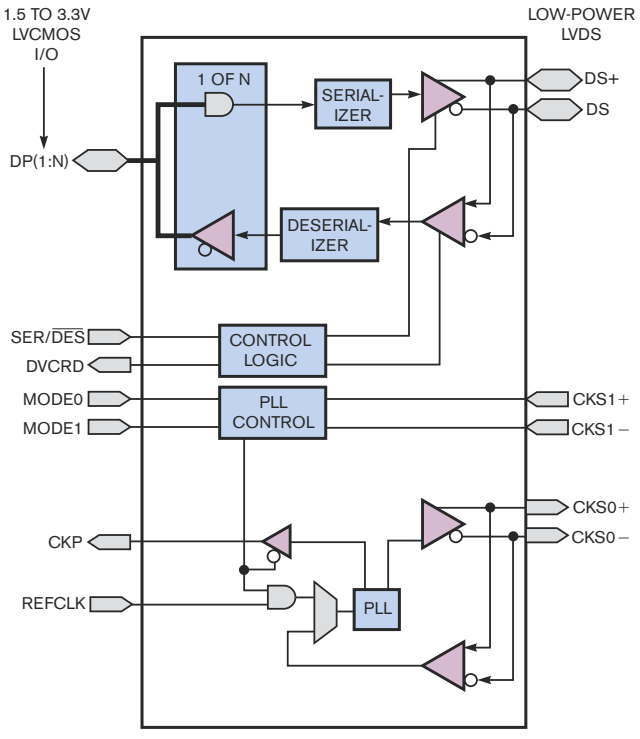

Figure 2 The first-generation µSERDES required significant trade-offs in design.

regular clock intervals. Data would arrive whenever the microcontroller "felt like" sending it. For a read operation, while data was going in one direction, which happened to be the wrong direction, control signals still needed to go in the opposite direction. Both of these requirements were well beyond the capability of a standard run-of-the-mill SERDES.

A high-end SERDES, such as an Ethernet SERDES, had some of these qualities but was far too expensive for this type of application and did not match all of the requirements. Packetizing the data, including header information, error detection, and error correction for transferring information to a display approximately 10 cm away was clearly overkill. The design needed to be a basic SERDES that could transfer data in both directions and use an asynchronous signal to latch the incoming data.

Another key insight—another design surprise—occurred at this time. Although previous SERDES-application designers had to worry about FCC regulations, this one had much worse issues to deal with concerning EMI. Previous applications had to be below radiation levels of −60 dBm to make sure the system did not interfere with other systems. The cell phone could not tolerate EMI at such levels, due to the radio. The µSERDES design could not interfere with radio transmission and reception from an antenna that may be only a quarter of an inch away. So, −60 dBm was no longer the target; −120 dBm was the new target.

The next specification that was out of line was power con-

Embedded Ethernet
Any Application, Any Location

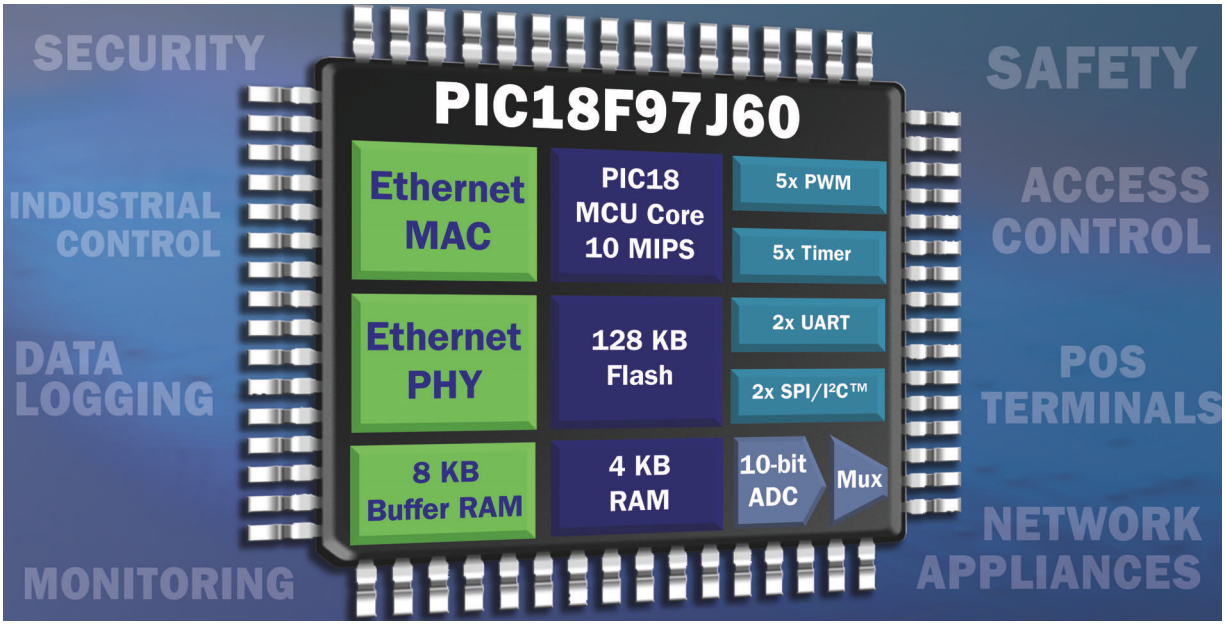

Monitor, control or re-program your application remotely with ease using the integrated Ethernet PIC18F97J60 family and FREE TCP/IP software.

GET STARTED IN 3 EASY STEPS
at www.microchip.com/ethernet.

1. **Learn about our Ethernet devices in 20 minutes**
   Take advantage of our Ethernet web seminars.
2. **Download our FREE TCP/IP software**
   Our TCP/IP stack is available in source code for flexible and optimized code size.
3. **Check out our low-cost Ethernet tools**

Evaluate the PIC18F97J60 family with the ***PICDEM.net™ 2 Demo Board*** – available for only $129.98 when you use coupon code **ETHTOOL** during checkout at www.microchipDIRECT.com.

Purchase and program your Ethernet PIC® microcontrollers and related development tools at.... microchip DIRECT www.microchipdirect.com

Device	Pins	Flash (KB)	Features
PIC18F97J60	100	128	10-BaseT Ethernet 12 KB RAM (8 KB dedicated Ethernet) 5x 16-bit timers 10-bit ADC, 16 channels 2 analog comparators 2 UART with LIN protocol 2 SPI, 2 I²C™ Industrial Temperature -40° to +85°C
PIC18F87J60	80	128	
PIC18F67J60	64	128	
PIC18F96J65	100	96	
PIC18F86J65	80	96	
PIC18F66J65	64	96	
PIC18F96J60	100	64	
PIC18F86J60	80	64	
PIC18F66J60	64	64	
ENC28J60	28	8K RAM	MAC, PHY, SPI Interface

Visit www.microchip.com/ethernet for additional devices supporting connectivity solutions and related development tools, application notes, reference designs, web seminars and more!

Or you may consider...

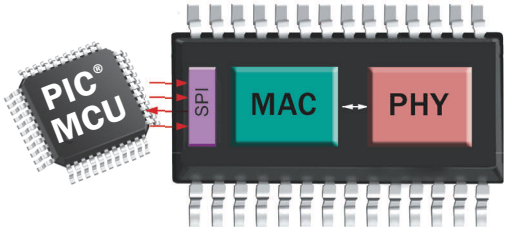

Adding Ethernet to any application with Microchip's ***ENC28J60*** *stand-alone Ethernet controller with full software support for PIC18, PIC24 and dsPIC® DSCs.*

MICROCHIP
www.microchip.com/ethernet

Microcontrollers • Digital Signal Controllers • Analog • Serial EEPROMs

sumption. The power consumption of the typical SERDES was off by orders of magnitude from the requirement. Battery life is one of the most important factors in cell-phone design. Every milliamp of operational current is a concern for talk time, and every microamp of current is a concern for standby time. As the team dissected SERDES implementations, it became clear that there were three main components of power consumption in the SERDES pair. Both the digital logic in the device and the serial link would consume significant current, as would the PLLs (phase-locked loops). The team had determined that it would use differential signaling for this application due to the minimum radiation it provided, as well as its high bandwidth. At this point, the engineers had not realized the importance of this decision. For the most significant reduction in power, the PLLs needed further investigation. Both the operational current and the standby current were unacceptable. Standby current was off by three orders of magnitude.

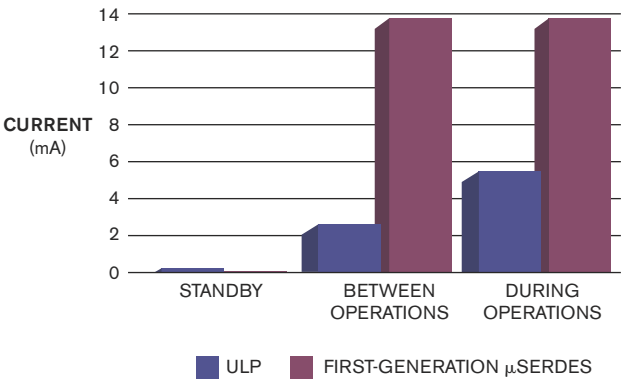

Figure 3 The 24-bit FIN324, a ULP device, consumes roughly 6 mA per pair of devices for a 5.44-MHz pixel interface.

A PRELIMINARY SPEC

All of this information was beginning to take form in an initial product specification. The engineers needed a bidirectional device with serialization for at least a camera and an LCD, current consumption of less than 5 µA at standby and less than 7 mA in operation per device, and less than −120-dBm EMI. Another key item for this market was that most of the engineers who would use the device had never dealt with serialization, so it was also critical that they make the device as easy to use as possible.

A rough sketch of a device showed how it would operate and meet most of the criteria. However, a few questions still remained concerning how to provide the key features. **Figure 1** shows the approach the engineers began presenting to customers.

The **figure** depicts a design close to the device that went into production for the initial generation. It is bidirectional and can work with a microcontroller interface or a pixel interface. Of course, some questions still remained about how to hit some of the parameters, such as power, EMI, and microcontroller-read operation, as well as the separation of parallel and serial clocking. For example, EMI needed to be much lower than previous SERDES with a target of −120 dBm versus a typical target of −70 dBm. It was also necessary to provide an approach for a standard pixel interface with pixel clock, vertical synchronization, horizontal synchronization, and pixel data. But that exercise was rather straightforward, and the team expected it to be similar to the previous SERDES.

Now, the engineers had some important insights into the application and the customers' needs. It was clear that every customer required both pixel and microcontroller interfaces that varied according to platform needs. In some applications, the baseband processor required a microcontroller interface to the LCD, whereas other baseband processors required the

TABLE 2 DESIGN TRADE-OFFS AND DECISIONS

Parameter	First option	Second option	Decision
No. of serial signals	Send the bit clock separately, which involves using more wires but less power; embed the word clock	Embed all of the clocking and use a single signal in each direction	Because power is so critical to cell-phone design, the team decided against embedding all of the clocks, which would have required a PLL's consuming a lot of current. The engineers embedded the word clock without consuming additional power. They considered a number of alternatives that they deemed unworkable because of radio-susceptibility concerns.
Parallel interface	Pixel interface	Microcontroller interface	The team decided to use both. One device handles all of the LCD interfaces at the customer level, enhancing marketability.
No. of bits for the LCD-interface device	27 bits to support 24 bits of RGB data	22 bits to support 18 bits of RGB data	The team selected the main device for the LCD, the FIN24AC, for 22 bits. Minimizing the size of the device seemed more important than supporting 24 bits of RGB. Display vendors were not yet offering 24-bit RGB.
Read interface for microcontroller	Provide a single device that could support read, write, and pixel interfaces	Separate the device that could do the read	The team separated the device that could do the read. For most customers, the write-only choice was satisfactory.
Packaging	MLP	BGA	The team initially selected MLP because it did not require a substrate, which incurs additional cost. The package seemed small. Customers quickly informed us that the package was too big. The design ended up requiring both packages, because the MLP is excellent for mounting on a flex cable.

Simple Control for Complex Applications

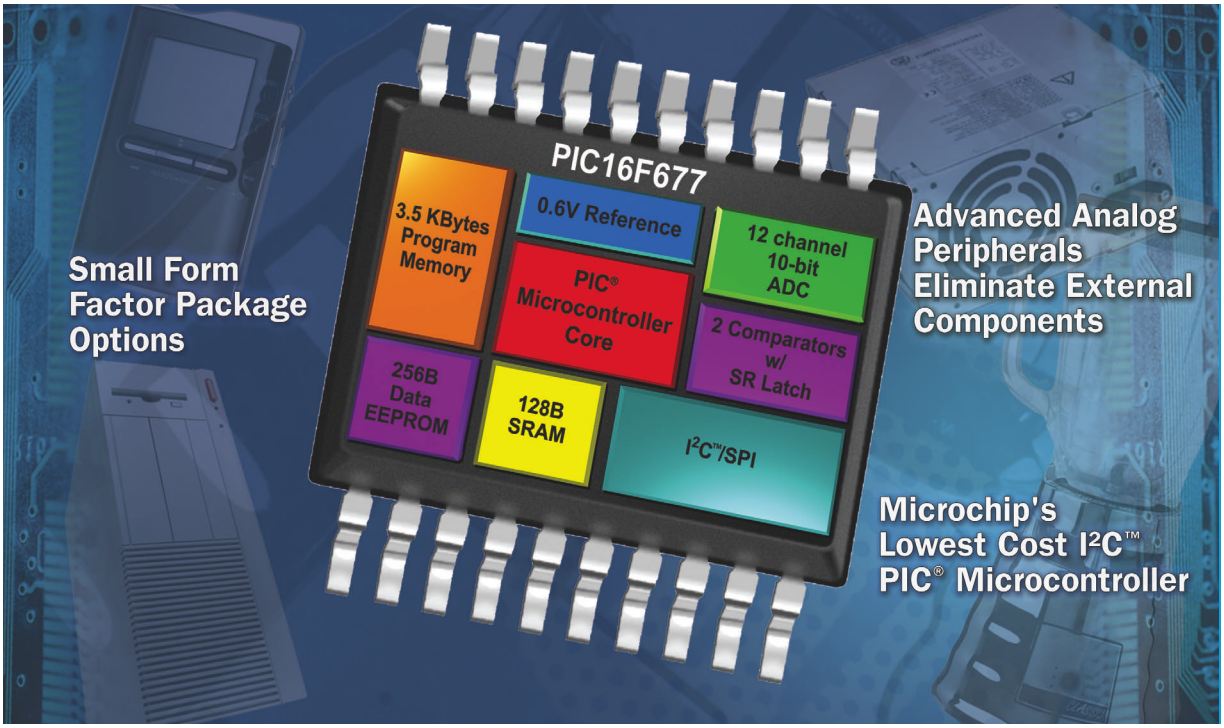

Microchip's 20-pin, general-purpose PIC® microcontrollers are flexible and easy to use – saving you design time and development costs.

Whether you're designing an MP3 player or a flight control system, Microchip's low-pin count devices give you the features you need to get the job done, while making configuration easy. A simple architecture, a cost-saving integrated peripheral set, and easy-to-use development tools allow you to spend your time perfecting your application – instead of debugging it.

20-PIN DEVICE OPTIONS

	Device	Program Memory Flash (words)	Data Memory SRAM (bytes)	Data Memory EEPROM (bytes)	10-bit A/D (ch)	Timers 8/16-bit	Other Feature20-s	20-pin Package Options
NEW!	PIC16F631	1024	64	128	0	1/1		PDIP, SSOP, SOIC, 4mm x 4mm QFN
NEW!	PIC16F677	2048	128	256	12	1/1	SPI/I²C™	PDIP, SSOP, SOIC, 4mm x 4mm QFN
	PIC16F685	4096	256	256	12	2/1	ECCP+	PDIP, SSOP, SOIC, 4mm x 4mm QFN
	PIC16F687	2048	128	256	12	1/1	EUSART, SPI/I²C™	PDIP, SSOP, SOIC, 4mm x 4mm QFN
	PIC16F689	4096	256	256	12	1/1	EUSART, SPI/I²C™	PDIP, SSOP, SOIC, 4mm x 4mm QFN
	PIC16F690	4096	256	256	12	2/1	ECCP+, EUSART, SPI/I²C™	PDIP, SSOP, SOIC, 4mm x 4mm QFN
	PIC16F785	2048	128	256	12	2/1	Op Amps, CCP, Hardware PWM	PDIP, SSOP, SOIC, 4mm x 4mm QFN

microchip DIRECT
www.microchipdirect.com

Get started today with the **PICkit™ 2 Starter Development Kit** for only $49.99!

MICROCHIP
www.microchip.com/easy

Microcontrollers • Digital Signal Controllers • Analog • Serial EEPROMs

more standard pixel interface. Because almost every system manufacturer used processors from multiple suppliers, a single manufacturer almost always required both interfaces.

So, it would be advantageous to provide a single device that could provide an approach for both interfaces. In addition, the same device needed to provide both the parallel-to-serial and the serial-to-parallel transition. This approach would be a huge advantage to customers, because they could purchase one product rather than four or six devices. It would reduce inventory requirements and provide a volume-pricing advantage. Design decisions specified that the device would be bidirectional, with each device containing both a serializer and a deserializer for microcontroller operation, and that the serial I/O would be differential to reduce EMI.

One of the critical breakthroughs for the design team involved the differential I/O. All differential I/Os, although called current mode, used voltage-level receivers. The receivers looked for a voltage difference between the two signals to

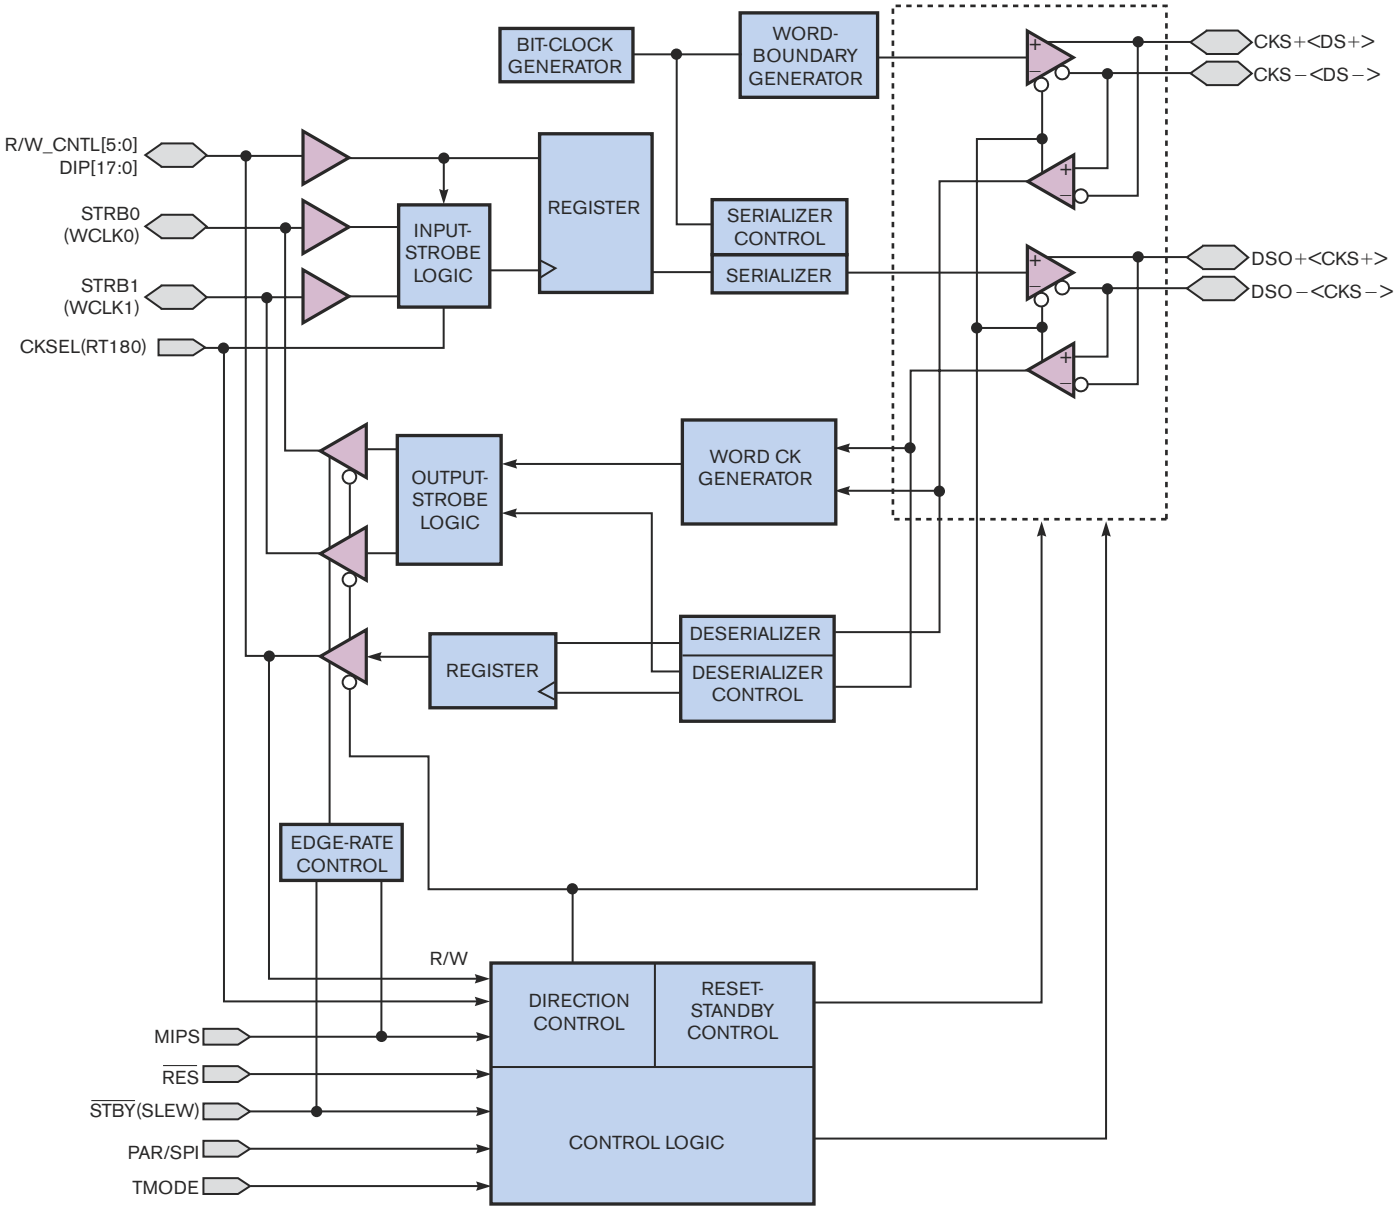

Figure 4 Because the 42-ball BGA can still accommodate dual interfacing, distinct control signals, and the elimination of the external clock, the package size of the ULP remains unchanged.

AC TO IC

Any Voltage • Any Current
Any Power Architecture

RoHS Compliant

ZM7300 Digital Power Managers

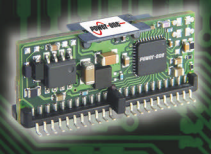

Z-Series Digital POLs

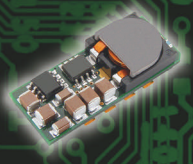

Y-Series Analog POLs

DPA & IBA Bricks

Applications	Solutions
Digital Point-Of-Load Power Management and Conversion	I^2C compatible Z-7000 Series reduces power system components, traces, and development time by 90%.
	Pin-strapable Z-1000 POLs provide digital power management capabilities without software or external controllers.
Analog Point-Of-Load Conversion	Y-Series POL converters are available in industry-standard and high-performance Power-One footprints.
Board-Level DPA and IBA	The industry's broadest range of bricks, from 1/16 to full, includes bus converters, multi-output, and PoE solutions.
System-Level DPA and IBA	AC-DC front ends in hot-swap power shelf and chassis-mount configurations.
Centralized Architectures	Single and multiple output models include flexible modular solutions.
Industrial and Transportation	Rugged cassette style and DIN-Rail mount AC-DC and DC-DC products.
CompactPCI	AC-DC and DC-DC in 3U and 6U form factors.

AC-DC Front Ends

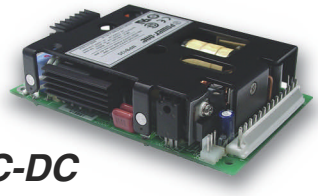

AC-DC Embedded

DIN Rail

Cassette

CompactPCI

power-one®
Changing the Shape of Power

www.power-one.com

determine a one or a zero. After careful consideration, the design team decided to look instead for current direction. Like a lot of good ideas, the advantages seem obvious in retrospect. Current direction is more stable than voltage levels, so the amount of drive necessary to distinguish between a one and a zero should be significantly less. In fact, it was possible to go as low as a 50-mV swing rather than the minimum of approximately 250 mV of voltage differential. That value represents five times less drive across the same 100Ω termination. The team termed this new I/O technology CTL (current-transfer logic).

CONFIRMATION/FEEDBACK

The new technique required another round of customer visits to explain this concept and seek further insight into the application requirements. In discussions with customers on this round, the team gathered a few more key details. The camera and the LCD were in a similar location, and the datapaths on both ends were in a similar location. The engineers had a bidirectional µSERDES, so why not try to use a single pair to perform both sets of operations? One of the most critical pieces of information, essentially another design surprise, came to light at this time.

A knowledgeable engineer at one of the key customers to which we were presenting CTL and the advantages of the low swing for EMI and power consumption mentioned that the low swing could be a problem due to susceptibility. "Susceptibility" is the electrical signal's vulnerability to radio noise that could disturb the signal enough to cause data corruption. This insight was absolutely key. Now, there were barriers on both sides of the signaling equation. On one side, the team had a radiation issue, so it needed to make the signal as small as possible. On the other side, the team had radio noise that could cause significant disturbance to the signal if it were too small.

The engineers were getting close to a reasonable solution, and the customers were intrigued with the idea of using a single pair of µSERDES to even further reduce the number of wires on the flex without adding another pair of devices. An interesting point here is that it was becoming clear that the customers had little fear of EMI. They had always had to deal with it, so EMI was less of an issue than the number of wires. Reducing the number of wires was the struggle for system designers due to the advent of the flip phone and the need to send all of these signals through a tiny hinge between the flip and the base of the phone. They could have gone with LPLVDS (low-power low-voltage differential signaling) as the easier path, but they had great confidence in CTL. There would be considerable advantages of supplying a signaling structure that would emanate less than −120 dBm, because it would reduce the need for shielding on the flex for the differential signals. More important, the shielding could affect the characteristic impedance of the flex. In addition, CTL could provide power savings and is also robust for susceptibility due to the receiver structure.

THE CUSTOMERS WERE INTRIGUED WITH THE IDEA OF USING A SINGLE PAIR OF µSERDES TO EVEN FURTHER REDUCE THE NUMBER OF WIRES ON THE FLEX WITHOUT ADDING ANOTHER PAIR OF DEVICES.

FINAL DESIGN DECISIONS

The approach for the microcontroller-interface clocking was to completely separate the timing of the serial and the parallel clock. The team still needed to generate the serial clock through PLL-based multiplication and the main control signal—for example, write enable—to strobe in the parallel data. The extrapolation of this insight would lead to the development of a second-generation device, ULP (ultralow-power) µSERDES.

On the power front, CTL would significantly reduce power over LPLVDS, but those savings were not enough. The team needed to eliminate the use of at least one of the PLLs, which would provide a significant power savings. Between eliminating one of the PLLs and reducing current with the use of CTL, the team could almost halve the current. The problem was that the team needed to send three sets of information between the serializer and the deserializer: the data, the bit clock to clock the individual bits as they arrived at the deserializer, and the word clock or word boundary that identified the first bit of a word for the deserializer.

Unless the engineers were willing to send three signals over six wires, they would need a PLL at the deserializer to provide either the bit clock or the word clock. Otherwise, they would have to come up with a seriously innovative approach. They concluded that it was critical to discard one of the PLLs, and the only place to eliminate one without placing an unreasonable burden on the customers was at the deserializer. The simplest approach was to send the serial data along with a bit clock to clock the data into the deserializer. Now, the challenge was how to send the word clock without using a PLL or adding another differential signal.

One of the engineers determined that the team could solve the problem if it could do something clever with the two signals to encode the word boundary. Adding two bits to the serial stream produces a unique pattern to define a word boundary. At the point in time of these two bits, if the clock did not toggle but the data did, a singular sequence would occur that you could easily identify as the word boundary. Even better, you could identify it exclusively for every word. Most schemes for encoding a clock take a PLL and thousands of words to find the word boundary through a repetitive recognition algorithm. This ingenious scheme allowed immediate recognition of the word boundary. If the system lost the word boundary due to a huge noise event, it would lose only one word.

DESIGN TRADE-OFFS

The team had now knocked down just about every known barrier to provide an initial SERDES approach for cell-phone designers. The units satisfied the needs of the customers and

Half the Size, Half the Loss
All the Power

***Direct* 48 V: 1.x V conversion in 91 A/in² with 91% efficiency**

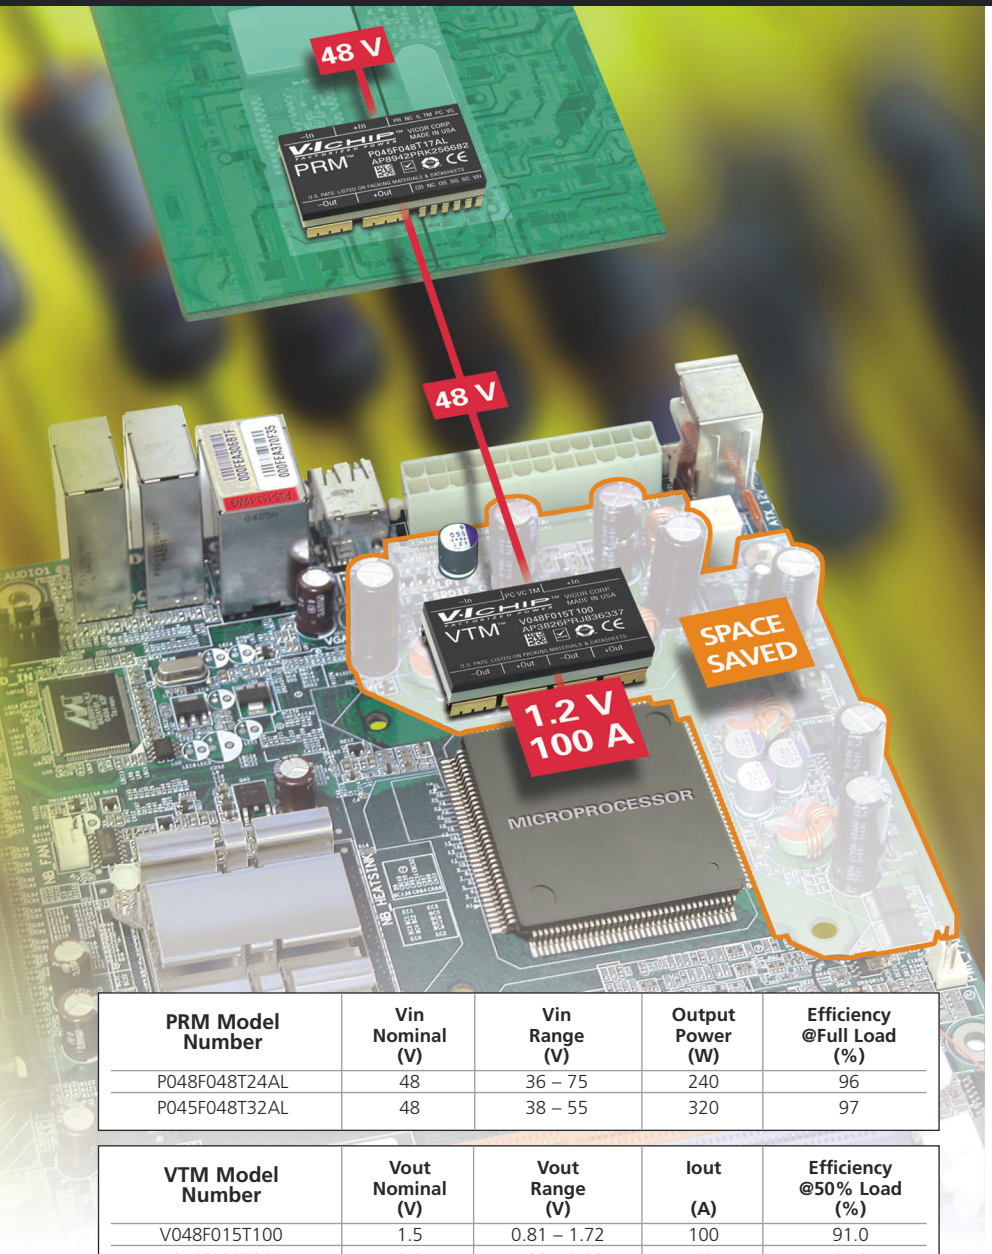

32.5 x 22.0 x 6.6 mm
1.28 x 0.87 x 0.26 in

RoHS

VICHIP FACTORIZED POWER

- **Direct 48 V to load conversion**
- **High efficiency for energy savings upstream**
- **Small size, less parts**
- **Fast transient response**
- **Fast time to market**

The new PRM+VTM system provides the smallest, highest efficiency, highest density complete DC-DC conversion for high performance 48 V power applications such as servers, telecom and ATE.

A PRM regulator +VTM current multiplier chipset delivers up to 100 A or 300 W for 48 V direct-to-load density of >500 W/in³.

Call for design support.
800-735-6200
Order evaluation boards
now at: vicorpower.com/vichip.

PRM Model Number	Vin Nominal (V)	Vin Range (V)	Output Power (W)	Efficiency @Full Load (%)
P048F048T24AL	48	36 – 75	240	96
P045F048T32AL	48	38 – 55	320	97

VTM Model Number	Vout Nominal (V)	Vout Range (V)	Iout (A)	Efficiency @50% Load (%)
V048F015T100	1.5	0.81 – 1.72	100	91.0
V048F020T080	2.0	1.08 – 2.29	80	94.2
V048F040T050	4.0	2.17 – 4.58	50	94.8
V048F120T025	12.0	6.50 – 13.75	25	95.1

For full range of products, visit vicorpower.com/vichip

Products in mass production and available from stock.

800-735-6200 **vicorpower.com/fpaedn** VICOR

provided state-of-the-art power, EMI, and susceptibility, and they were usable with a microcontroller or a pixel interface.

The engineers needed to make a number of trade-offs on this first-generation product (**Figure 2**). **Table 2** lists some of the most important ones. The main parameters involved power, size, and cost—all of which were equally important for the cell-phone design.

The major disappointments on this project were that the half-duplex approach that the team originally envisioned for using a single pair of µSERDES was just too cumbersome to implement at a system level, and the read operation did not work as well as the team had expected.

Because of the systems approach and close interaction with customers in the design phase, the company started shipping products as soon as it received approval for production. The team was happy with the success of this project. There were three variations to accommodate the camera and LCD, as well as a product for special applications. A 12-bit device targeted the camera, which is almost always an 8-bit-pixel interface. This device has also seen use in 8-bit-LCD-pixel interfaces. A second version targeted the typical 16- or 18-bit-LCD-pixel interface. A third part targeted microcontroller usage.

THE APPLICATIONS GROUP ACQUIRED MULTIPLE CELL PHONES FROM CUSTOMERS, TOOK THEM APART, INSERTED THE µSERDES PAIR INTO THE DESIGNS, AND DEMONSTRATED THAT THEY WORKED.

Because the SERDES concept was so new to cell-phone designers, the engineers developed detailed application notes on the modes of operation and provided evaluation cards to ease testing. To further simplify the implementation of technology that was a significant departure from previous cell-phone designs and because schedules were so critical, they provided cell-phone hookups showing the use of the µSERDES in the applications. The applications group acquired multiple cell phones from customers, took them apart, inserted the µSERDES pair into the designs, and demonstrated that they worked.

THE NEXT STEP

In parallel with the initial project, the team had begun identifying the next-generation product. This next product improved the initial design concept but required development time that was outside the scope of the initial project. Recently introduced ULP devices use the separation of the parallel and the serial clocking for even greater reduction in power consumption.

Because there was no need for synchronizing the parallel

Less Is More

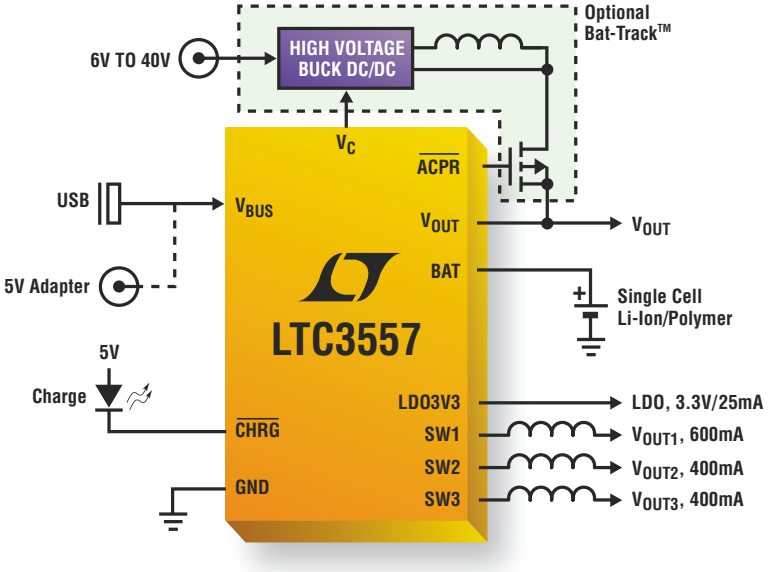

No Hassle Portable Power

Our LTC®3557 Power Management IC includes all the "hard-to-do" power functions, including PowerPath™ control, battery charging, power selection and multiple output regulators. With its linear PowerPath front end, the LTC3557 provides seamless transition between multiple power sources, as well as high efficiency charging and instant-on operation. The regulators, control and charging are all packed in a low profile 4mm x 4mm QFN package and needs no microprocessor for operation.

Features

- Seamless Transition Between Input Power Sources: Li-Ion Battery, USB, 5V Wall Adapter or High Voltage Buck Regulator with Battery Tracking to Minimize Heat
- Internal Ideal Diode Provides Low Loss PowerPath
- Triple Adjustable High Efficiency Step-Down Switching Regulators (600mA, 400mA, 400mA I_{OUT})
- Pin-Selectable Burst Mode® Operation for Long Run Time
- 1.5A Maximum Charge Current with Thermal Limiting
- Low-Profile (0.75mm) 4mm x 4mm 28-Pin QFN Package

LTC3557 Demo Board

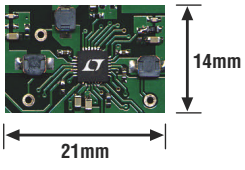

Actual Size

No µController Needed for Charging or PowerPath Control

Info & Free Samples

www.linear.com/3557

1-800-4-LINEAR

Free Portable Solutions Brochure

www.linear.com/portsolutions

LINEAR TECHNOLOGY

and serial data streams, it was possible to eliminate the second PLL. Eliminating all PLLs had been a goal since the beginning of the first-generation product. The team realized that it could use just an internally generated clock to provide the clocking for the serial stream. It significantly lowers power and provides a microcontroller interface that does not require a system clock.

Eliminating the second PLL was a breakthrough on its own. It turned out to be even better when the team realized that the clock could always run at the maximum frequency of the specifications. In this case, the serial stream could burst across the interface, and the pair of devices could go into a pseudo-power-down state, burst standby, for most of each transfer cycle instead of staying powered up continuously.

Engineers often regard dynamic power as dependent on the frequency of operation. This assumption is inaccurate. The reality is that dynamic power depends on the number of edges per unit of time. Because the team was transferring a specific number of bits, or edges, per unit of time, the dynamic power did not suffer due to the higher speed. Instead, the burst standby after the transfer considerably reduced the static power. Furthermore, for a microcontroller that is not in continuous operation, the average power would drop drastically due to the burst standby between operations.

MORE AT EDN.COM

Go to www.edn.com/ms4229 and click on Feedback Loop to post a comment on this article.

As a comparison, the first-generation 24-bit FIN24xx had state-of-the-art power consumption, using roughly 14 mA per pair of devices for a 5.44-MHz pixel interface with a constant clock. The 24-bit FIN324, a ULP device, consumes roughly 6 mA per pair of devices for a 5.44-MHz pixel interface (**Figure 3**). A microcontroller interface requires an assumption of percentage of time of operation. A reasonable estimate is that the microcontroller operates roughly 30% of the time. This estimate gives a power number of approximately 3 mA if the microcontroller cycle frequency is the same 5.44 MHz.

In the end, the engineers' experience showed them that you can apply expertise in SERDES design from one market—the computer industry—to a different market with different requirements. But making this transfer required intensive research, dialogue with customers, and a willingness to shift the focus of design from the chip to the end system (**Figure 4**). It also meant accepting that design communities have not only different concerns, but also different language for discussing them. EDN

AUTHOR'S BIOGRAPHY

Michael Fowler is a member of the technical staff and a member of the µSERDES design team of Fairchild Semiconductor. You can reach him at michael.fowler@fairchildsemi.com.

LINEAR TECHNOLOGY

DESIGN NOTES

High Linearity Components Simplify Direct Conversion Receiver Designs – Design Note 418

Cheng-Wei Pei

Introduction

A direct conversion radio receiver takes a high frequency input signal, often in the 800MHz to 3GHz frequency range, and utilizes one mixer/demodulator stage to convert the signal to baseband without going through an intermediate frequency (IF) stage. The resulting low frequency (baseband) signal spectrum has useful information at frequencies from DC to typically a few tens of MHz. Designing these receivers requires the use of very high performance analog ICs. High performance direct conversion radio receiver signal chains for applications such as cellular infrastructure and RFID readers require high linearity, low noise figure (NF), and good matching between the in-phase and quadrature (I and Q) channels.

The Right Components for the Job

Linear Technology's LT®5575 direct conversion demodulator has a combination of excellent linearity and noise performance. The most important linearity specification for direct conversion mixers is the 2nd order intercept point (IIP2) due to the 2nd order distortion product falling within the baseband output spectrum, and the LT5575 boasts 54.1dBm at 900MHz (60dBm at 1900MHz). The LT5575 also has high 3rd order linearity and a low noise figure of 12.8dB.

The LTC®6406 is a fully differential amplifier with low noise (1.6nV/$\sqrt{Hz}$ at the input) and high linearity (+44dBm OIP3 at 20MHz) in a small 3mm × 3mm QFN package. External resistors set the gain, giving the user maximum design flexibility. The low power consumption (59mW with a 3.3V supply) means that using two amplifiers for I and Q has minimal effect on the system power budget. The LTC6406 maintains high linearity up to 50MHz, which is perfect for WCDMA receivers and other wideband applications.

A Basic Receiver Design

One common design challenge when using active demodulators is level-shifting the outputs, which can have a DC level close to V_{CC}, to a usable DC level within the input range of the analog-to-digital converter (ADC). Fortunately, the LTC6406's rail-to-rail inputs make interfacing with the outputs of the LT5575 simple and direct. The LTC6406 also includes an extra feedback loop (controlled

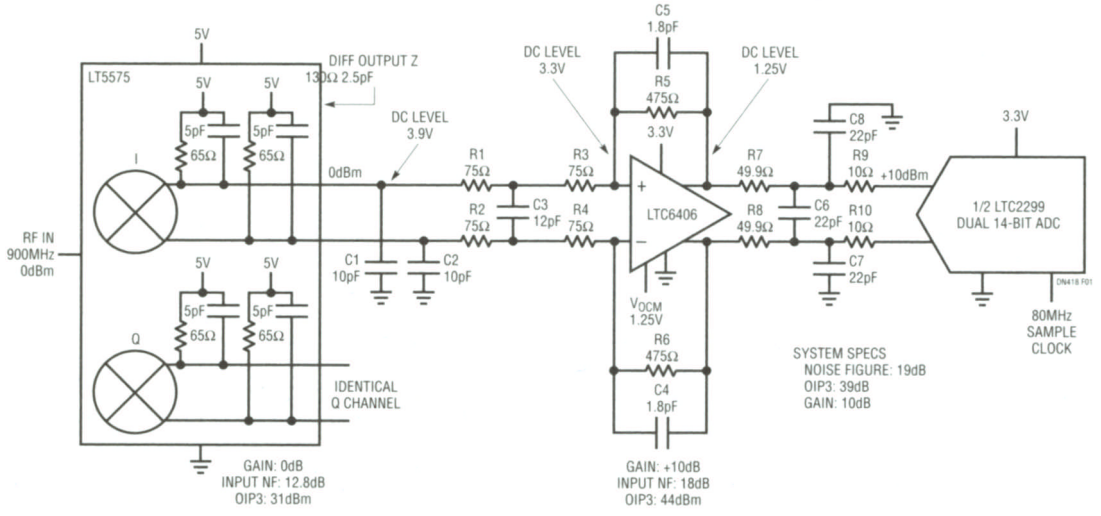

Figure 1. The LT5575 Demodulator and LTC6406 Amplifier Driving an LTC2299 14-Bit ADC. System Bandwidth is Approximately 40MHz. Overall System OIP3 was Measured to be 39dBm

by an external V_{OCM} voltage) that independently sets the output common mode DC level, regardless of the input voltage.

Figure 1 shows a basic receiver circuit with the LT5575 demodulator and the LTC6406 followed by an LTC2299 14-bit ADC. An RC lowpass filter at the output of the demodulator filters undesired out-of-band signals, and another RC lowpass filter before the ADC antialiases and limits noise bandwidth. The DC voltage at the LTC6406 inputs is 3.3V, the same as the supply voltage.

Adding Free Gain to the System

For signal chains that require more gain, the LTC6401-8 differential amplifier/ADC driver is a good complement to the LT5575 and LTC6406. The LTC6401-8 has higher linearity (50dBm OIP3 at 20MHz) and 2.5nV/$\sqrt{Hz}$ of input noise in a 3mm × 3mm QFN package. It contributes gain and linearity without significantly impacting the noise figure. Figure 2 adds the LTC6401-8 (also available in 14dB, 20dB and 26dB flavors) to the signal chain to drive the LTC2299. The higher linearity of the LTC6401-8 increases the combined system OIP3 to 45dBm. In addition, 8dB of gain is added with no significant degradation to the noise figure. The 400Ω input impedance of the LTC6401-8 is not a heavy load for the LTC6406, which enables direct coupling of the two amplifiers with minimal signal loss (from series resistors, etc.).

A More Selective Filter

There are three places where a filter can be implemented in the circuit of Figure 2: after the mixer, in between the two amplifier stages, and prior to the ADC. Each has its trade-offs, but the simplest design places the filter after the mixer. This topology attenuates unwanted signals earlier in the signal chain, which preserves the IP3 of the following stages and allows for more gain through the system. An LC filter at the demodulator output minimally affects the distortion and noise figure of the system, whereas LC lowpass filters can present a heavy load impedance to a feedback amplifier output near their resonant frequencies. For reasons outside the scope of this article, it is tricky to design LC networks at the input of a high speed sampling ADC.

A concern when designing the LC network is the need to preserve the I and Q gain/phase matching of the LT5575 (0.04dB/0.4° mismatch), which necessitates using low tolerance LC components (±2% inductors and ±5% capacitors). The frequency response and group delay of the system are almost entirely determined by the LC filter.

Conclusion

Signal chain devices that offer high linearity and excellent noise specifications can greatly simplify the design of high frequency receivers—speeding up entire design cycles.

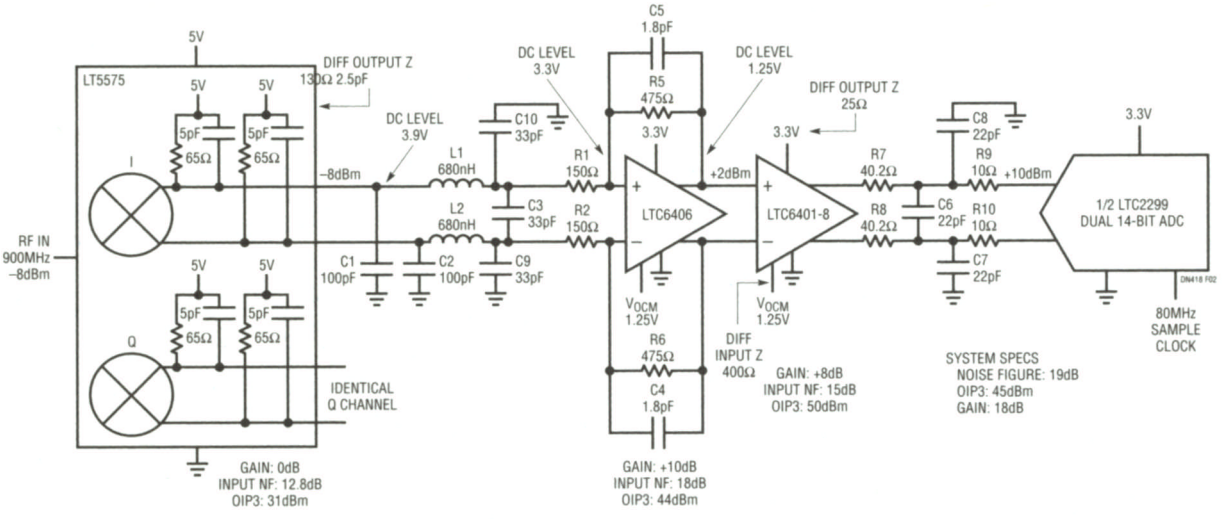

Figure 2. The LT5575 Demodulator with a 20MHz Lowpass Filter followed by LTC6406 and LTC6401-8. System OIP3 is Measured to be 45dBm at 920MHz RF with a 900MHz Local Oscillator. System Noise Figure (NF) Adds up to Approximately 19dB

Data Sheet Download

www.linear.com

For applications help,
call (408) 432-1900, Ext. 2525

Linear Technology Corporation
1630 McCarthy Blvd., Milpitas, CA 95035-7417
(408) 432-1900 • FAX: (408) 434-0507 • www.linear.com

dn418f LT/TP 0607 305K • PRINTED IN THE USA

Simplify PowerPath ON/OFF Control

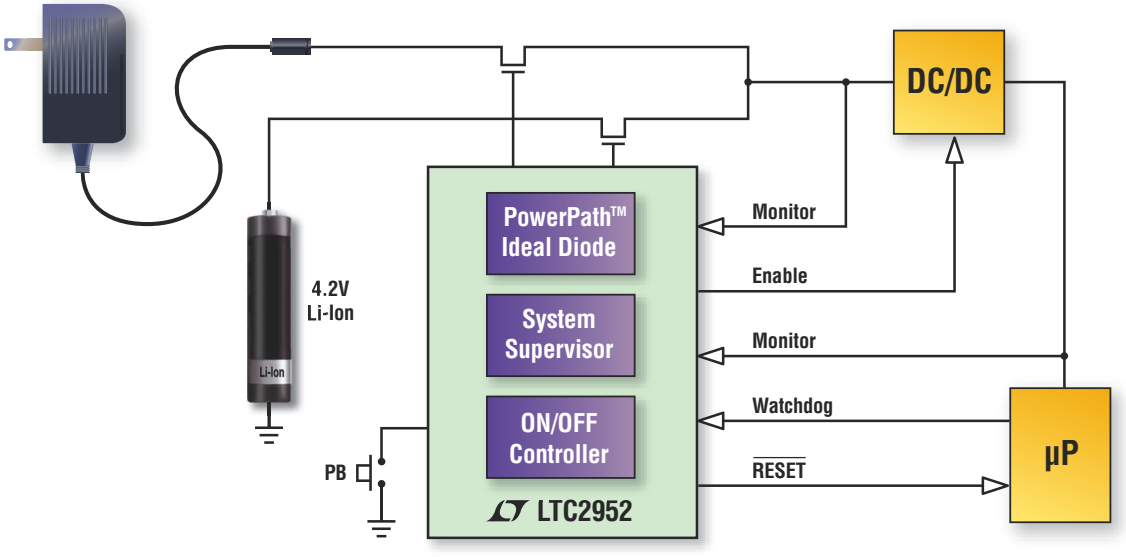

3 in 1: Integrated Ideal Diode Controllers, Pushbutton & Supervisory Functions

The LTC®2952 is a versatile power management IC with push button ON/OFF control of system power and system monitoring features. The power management function includes two power path controllers for applications requiring an ideal diode-OR function. The push button controller provides enable control for DC/DC converters, processor interrupt logic and adjustable ON and OFF debounce times. Robust supervisory functions make the LTC2952 an ideal fit in high reliability applications. Power-fail and watchdog reset functions provide further enhancements to ensure system integrity.

Features

Part No.	Supply Voltage (V)	Supply Current	ON Timer	OFF Timer	Kill Timer	Comments	Package
LTC2950	2.7 to 26	6µA	Adj	Adj	1024ms	Active high enable output (LTC2950-1), active low enable output (LTC2950-2)	TSOT-8, DFN-8
LTC2951	2.7 to 26	6µA	128ms	Adj	Adj	Active high enable output (LTC2951-1), active low enable output (LTC2951-2)	TSOT-8, DFN-8
LTC2952	2.7 to 28	25µA	Adj	Adj	Extendable	Push button power path controller with system monitoring	TSSOP-20, QFN-20
LTC2954	2.7 to 26	6µA	Adj	Adj		Interrupt logic for menu driven applications. Active high enable output (LTC2954-1), active low enable output (LTC2954-2)	TSOT-8, DFN-8

Info & Free Samples

www.linear.com/2952
Literature: 1-800-4-LINEAR
Support: 408-432-1900

Free Industrial Signal Chain Brochure

www.linear.com/indsolutions

LINEAR TECHNOLOGY

Single Port PoE

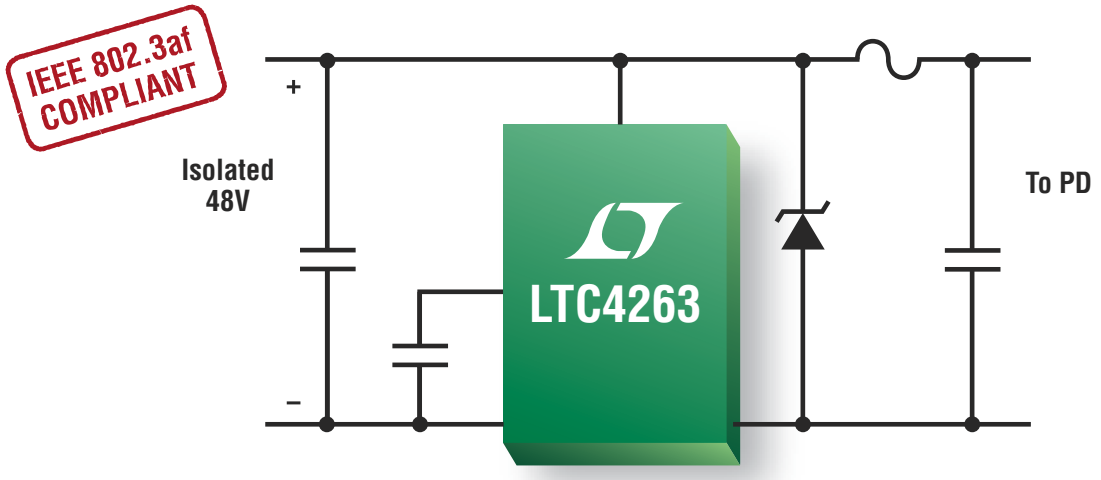

Fully Autonomous Controller with Integrated FET Simplifies PSE Design

The LTC®4263 is a complete, single port Power over Ethernet controller for Power Sourcing Equipment. It integrates a high-voltage isolation FET, AC-/DC-disconnect circuitry, a fully compliant detection and classification engine, and a novel multiport power management scheme, requiring only a few external components. The fully autonomous operating mode requires no software and port status is optionally displayed using a flashing LED. The LTC4263 combines integrated features and simplicity with the ruggedness and reliability of Linear's PoE product family.

Features

- IEEE 802.3af Compliant
- Fully Autonomous Operation without Microcontroller
- Thermally Protected Internal MOSFET
- Power Management Works Across Multiple Ports with Simple RC Network
- Precision Inrush Control with Internal Sense Resistor
- Powered Device (PD) Detection and Classification
- LTC4263: $2.95 each in 1k qty

Small Footprint, Simple Design

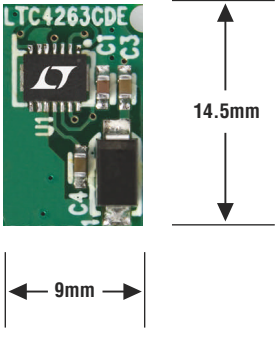

Info & Free Samples

www.linear.com/4263
Literature: 1-800-4-LINEAR
Support: 408-432-1900

LEAD-FREE RoHS COMPLIANT

LINEAR TECHNOLOGY

EDITED BY CHARLES H SMALL
AND FRAN GRANVILLE

designideas

READERS SOLVE DESIGN PROBLEMS

Add a grounded-switch feature for Topswitch on/off control

Robert N Buono, Aeolian Audio LLC, Bloomfield, NJ

The Power Integrations Topswitch family (www.powerint.com/topgxproduct.htm) of integrated flyback-regulator ICs provides exceptional performance in small, low-pin-count packages. For the lowest-pin-count packages, the multifunction, or M, pin serves multiple purposes, including on/off control and undervoltage- and overvoltage-input detection. Other package types include an L pin, which also provides this function. The application notes and data sheets show how to implement the various features available at these pins. For example, to allow remote on/off control and still preserve undervoltage and overvoltage functions, the application drawings show an NPN transistor, Q_R, which connects between the M or L pin and the Control pin (**Figure 1**). To turn off the regulator, Q_R must be biased on. To achieve this goal requires a base voltage of 2.6V dc or greater.

The circuit in **Figure 2** provides a new feature that allows you to switch the regulator on or off using a grounded switch that is sometimes more convenient to implement than a switch that references to the Control pin. In the case of a mechanical switch, this circuit would require no external power to implement this function. This feature is important in applications in which the Topswitch power supply is the only source of power. This circuit does not disturb the functioning of the undervoltage and overvoltage functions of the M or L pin. To understand the functioning of the circuit in **Figure 2** requires an explanation of the internal workings of the M or L pin. This pin acts as a constant voltage source at approximately 2V dc and sinks current from the external circuit, which R_{LS} supplies. The internal current-sense thresholds for undervoltage and overvoltage detection are roughly 50 µA

DIs Inside

▶What are your design problems and solutions? Publish them here and receive $150! Send your Design Ideas to edndesignideas@reedbusiness.com.

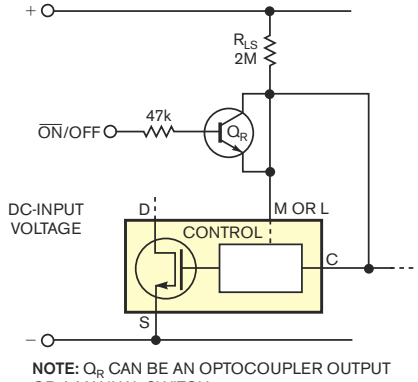

Figure 1 Adding transistor Q_R to the L pin of a Topswitch switching power controller enables an on/off-control feature.

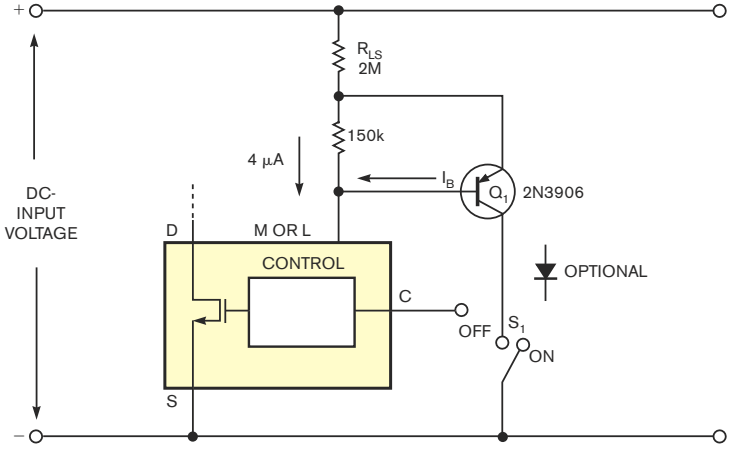

Figure 2 Instead of using just the external transistor to switch a Topswitch on or off, a simple on/off switch provides manual on/off control.

with 30 µA of hysteresis for undervoltage and 225 µA for overvoltage. That is, when the current into the M or L pin is less than 20 µA, or 50−30 µA, the regulator output switches off because of undervoltage. When the current into the M or L pin exceeds 225 µA, the regulator output switches off because of overvoltage. When the current into the M or L pin is 50 to 225 µA, the output is enabled.

The circuit of **Figure 2** works as follows: When the switch in the collector lead of Q_1 is open, Q_1 functions as a simple diode with a 0.6V drop from emitter to base. All the current that R_{LS} supplies flows into the M or L pin through the base-emitter junction of Q_1 and the 150-kΩ resistor. In this mode, the Topswitch IC senses the undervoltage and overvoltage thresholds. However, when the switch to ground closes, Q_1 functions as a nonsaturated transistor with high gain. The circuit siphons off most of the current through R_{LS} to ground as the collector current of Q_1. Only a small base current from Q_1 plus 4 µA through the 150-kΩ resistor flows into the M or L pin. For the values in **Figure 2**, this base current is less than 3.8 µA, even when Q_1 has minimum gain and input voltage is at a maximum of 450V dc. Therefore, 3.8+4 µA, or 7.8 µA, flows into the M or L pin. This low current flowing into the pin "fools" the regulator into "thinking" that the input voltage is undervoltage, and the regulator output switches off.

If another voltage or current source is present, you could replace S_1 with an open-collector switch that sinks current only. If the remote on/off driver can source and sink current, as the output of a logic gate can, then you should insert a diode in the collector lead of Q_1, and the driver must drive the cathode of that diode above 2V dc to turn off the regulator (optional in **Figure 2**). The M pin also allows current-limit-threshold adjustment. EDN

RC lowpass filter expands microcomputer's output port

Rex Niven, Forty Trout Electronics, Eltham, Victoria, Australia

It's almost a corollary to Moore's Law: Next year, microcomputers will have more features, and the software team will have bigger ideas. Unfortunately, though, the number of output pins will stay the same. Finding even one spare output for diagnostics, test, or even standard I/O can be a tussle. The single-pin "bus" in **Figure 1** can provide an unlimited number of parallel outputs with simple additional hardware. A microcomputer output with an RC lowpass filter controls serial-to-parallel converter HC164. To enter data into the serial-to-parallel converter, each bit consists of a one-to-zero-to-one transition, which alters the length of the low state. If the low state is longer than the lowpass filter's time constant, a zero shifts into the register. If the low state is short, then a one shifts into the register. The clock and data signals thus combine into one signal. A lowpass filter separates the clock and data signals (**Figure 2**).

Listing 1, a simple "Whip" routine,

LISTING 1 WHIP-ROUTINE OUTPUT FUNCTION

```
Whip                                        ;  the data to transmit is in W
        MOVWF       My_Data
        MOVLW       8
        MOVWF       Bit_Counter             ; set up for eight bits
        BSF         My_Port,  My_Bit        ; ensure output is initially high
B1:
        RLF         My_Data,  F             ; data to send is in CC
        BTFSS       STATUS, CC
        BCF         My_Port,  My_Bit        ; zero, so falling edge is early
        CALL        Delay_10us              ; if a one, pin stays high for 10us
        BCF         My_Port,  My_Bit        ; if a one, edge falls here
                    NOP                     ; ensures output pin is low for 0.2us min
        BSF         My_Port,  My_Bit        ; rising edge here clocks data into HC164
        DECFSZ      Bit_Counter
        GOTO        B1

        RETURN
```

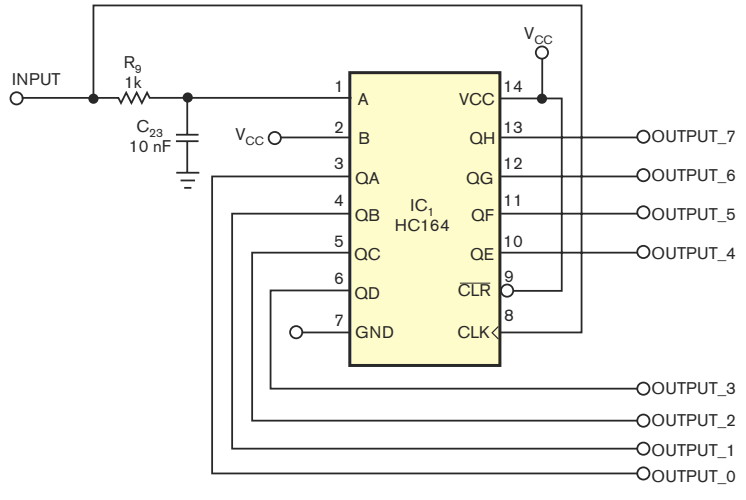

Figure 1 This single-pin "bus" can provide an unlimited number of parallel outputs with simple additional hardware.

50mA AUTOMOTIVE LDO REGULATOR IS IDEAL FOR ALWAYS-ON APPLICATIONS

IDEAL FOR
- TIRE-PRESSURE MONITORING
- REMOTE KEYLESS ENTRY
- CAR ALARMS
- OTHER ALWAYS-ON APPLICATIONS

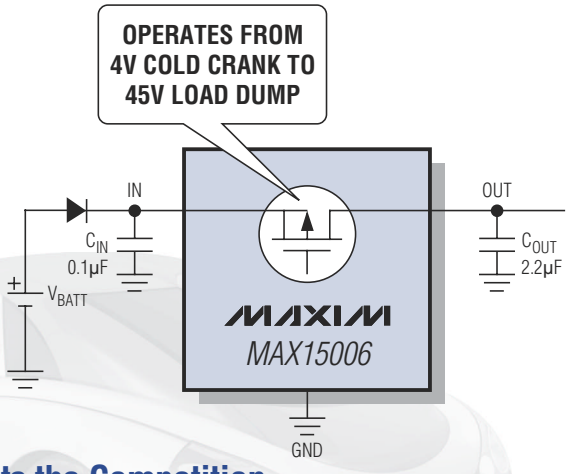

Maxim Beats the Competition

- ♦ 33% Lower Quiescent Current
  - ♦ 10µA at No Load
- ♦ 46% Lower Junction-to-Case Thermal Resistance (θ_{JC} = 8.5°C/W)
- ♦ 40% Smaller Package
  - ♦ 3mm x 3mm, 6-Pin TDFN (Also Available in 8-Pin SO with Exposed Pad)

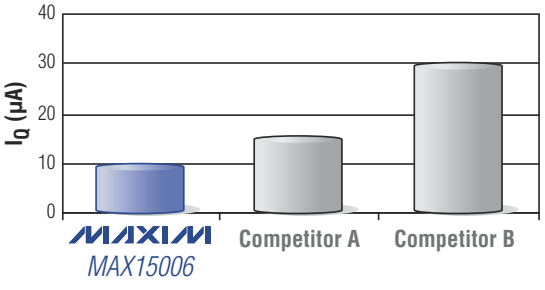

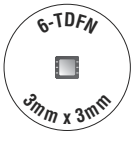

Additional Features

- ♦ Wide 4V to 40V Input-Voltage Range
- ♦ Guaranteed 50mA Output Current
- ♦ Thermal and Short-Circuit Protection
- ♦ 300mV Low Dropout Voltage at 50mA
- ♦ Stable Operation with Tiny, 2.2µF Output Capacitor
- ♦ Prices Start at $0.59†

†1000-up recommended resale. Prices provided are for design guidance and are FOB USA. International prices will differ due to local duties, taxes, and exchange rates. Not all packages are offered in 1k increments, and some may require minimum order quantities.

MAXIM

www.maxim-ic.com/MAX15006-info

FREE* Automotive Design Guide—*Sent Within 24 Hours!

CALL TOLL FREE 1-800-998-8800 (7:00 a.m.–5:00 p.m. PT) for a Design Guide or Free Sample

MAXIM/DALLAS DIRECT! DISTRIBUTION 1-888-MAXIM-IC

ARROW ELECTRONICS, INC. 1-800-777-2776

AVNET electronics marketing 1-800-332-8638

Distributed by Maxim/Dallas Direct!, Arrow, Avnet Electronics Marketing, Digi-Key, and Newark.

performs the output function for eight bits. Assume that the RC time constant is 3 µsec, and the instruction time should be 1 µsec or less at a crystal frequency of 4 MHz or greater. The routine uses bitwise manipulation of output My_Bit of port My_Port.

Although the circuit in **Figure 1** can control slow-reacting devices, such as relays or LCDs, using it with LEDs can give an annoying flicker when the HC164 is writing. To address that problem, the circuit in **Figure 3** uses another serial-in/parallel-out register, the 4094, which has a strobe input to allow simultaneous updates of all outputs without temporary levels. A twin monostable circuit supplies the data and strobe signals. This circuit should be able to control parallel devices, such as display modules based on HD44780 devices. EDN

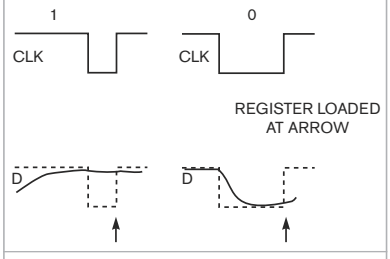

Figure 2 The clock and data signals combine into one signal.

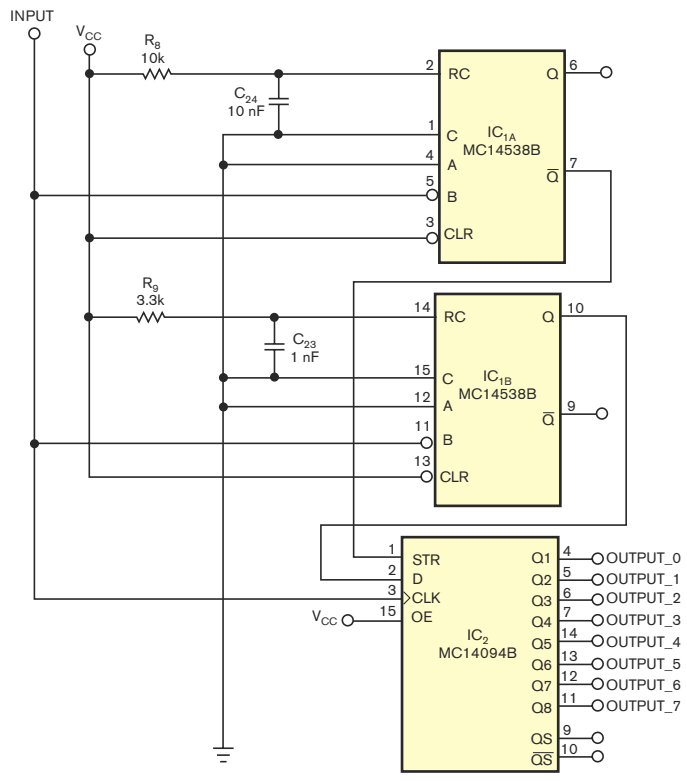

Figure 3 This circuit uses another serial-in/parallel-out register, the 4094, which has a strobe input to allow simultaneous updates of all outputs without temporary levels.

Simple dual constant-current load tests low-current power supplies

John S Lo Giudice, STMicroelectronics, Schaumburg, IL

Today's small electronic appliances, such as washers, dryers, and stoves, use switched-mode power supplies to replace bulky, heavy, linear-power supplies. The engineer testing these power supplies, which range in current from 50 mA to 1A, typically uses resistors or standard off-the-shelf electronic loads. An engineer would employ a variety of high-wattage resistors to verify multiple loading conditions to satisfy a proper design. Most off-the-shelf electronic loads target an average of 300W. When measuring 50 to 300 mA, a display is inaccurate; most of them display 0.1A, but accuracy is questionable at that low range. You can alternatively use the simple dual constant-current-load design in **Figure 1**, which you can build with inexpensive, common parts.

The load current passes through a MOSFET and a 1%, 1Ω sense resistor, R_6. Pin 2 of IC_{1A} compares the voltage drop in the resistor to a reference voltage. IC_1, an LM358 op amp, compares the two inputs and adjusts its output accordingly. The reference voltage at Pin 3 of IC_{1A} comes from a voltage-divider potentiometer, R_2 or R_3, which derives from a TS431 1.25V 1% reference. Because the maximum voltage can be 1.25V and the sense resistor's value is 1Ω, the maximum current per channel can reach 1.25A.

R_2 and R_3 are 15-turn, 1-kΩ potentiometers, which you can finely adjust to the desired load. One can set a minimum current, and the other can set a maximum current. Switch S_1 can then switch between minimum load, no load in the middle position, and maximum load. Furthermore, by attaching a standard DMM (digital multimeter) across R_6, you can directly read the current and adjust it to the proper level.

Input-voltage change does not affect the DMM's reading because it monitors the constant current through sense resistor R_6. The second channel is a duplicate of the first. Each chan-

SIMPLIFY DISTRIBUTED TEMPERATURE SENSING

Low-Voltage Digital Thermometer Allows Up to 27 Bus Addresses on the Same I²C Bus

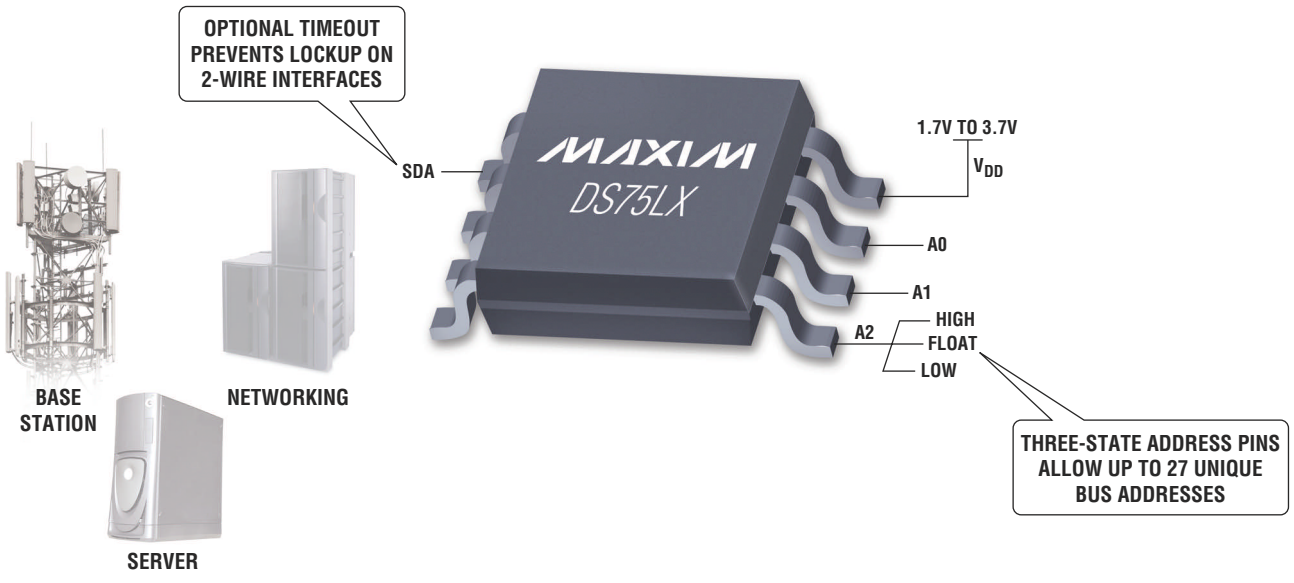

- ♦ Pin/Software Compatible with the LM75
- ♦ I²C/SMBus™-Compatible Serial Interface
- ♦ User-Selectable 9- to 12-Bit Resolution
- ♦ ±2.0°C Accuracy
- ♦ 1.7V to 3.7V Supply Range
- ♦ Frees Space for Other I²C Components
- ♦ Thermostat Output with User-Defined Trip Points
- ♦ -55°C to +125°C Operating Range
- ♦ Prices start at $0.75†

TINY SOT23 AVAILABLE FOR SPACE-CONSTRAINED DESIGNS

Part	Supply Voltage (V)	No. of Addresses	Optional Timeout	Package
DS75LX	1.7 to 3.7	27	✔	8-µSOP, SO
DS75LV	1.7 to 3.7	8	✔	8-µSOP, SO
DS1775	2.7 to 5.5	1		5-SOT23
DS75	2.7 to 5.5	8		8-µSOP, SO

SMBus is a trademark of Intel Corporation.
†1000-up recommended resale. Prices provided are for design guidance and are FOB USA. International prices

MAXIM

www.maxim-ic.com/DS75LX-info

FREE* Temp Sensors Design Guide—*Sent Within 24 Hours!

CALL TOLL FREE 1-800-998-8800 (7:00 a.m.–5:00 p.m. PT) for a Design Guide or Free Sample

MAXIM/DALLAS DIRECT! DISTRIBUTION 1-888-MAXIM-IC

ARROW ELECTRONICS, INC. 1-800-777-2776

AVNET electronics marketing 1-800-332-8638

Distributed by Maxim/Dallas Direct, Arrow, Avnet Electronics Marketing, Digi-Key, and Newark.

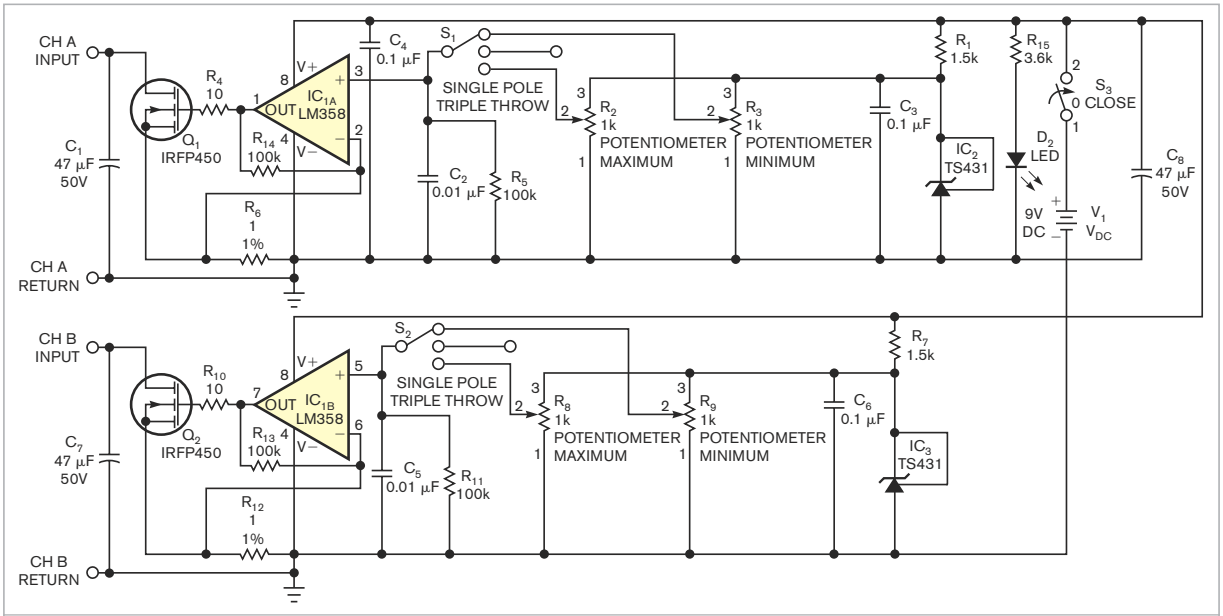

Figure 1 This dual constant-current load can measure the performance of dual power supplies supplying 0 to 1.25A per channel.

nel can control 0 to 1.25A and can handle a voltage of 3 to 50V. The capacitor input and the MOSFET set the upper limit. The two inputs can be in parallel to a load of 2.5A. For a two-output power supply, you can set the minimum and maximum current by precisely reading the level on a multimeter and then quickly testing a matrix of no load, minimum load, and maximum load. A 9V battery powers the unit. EDN

Stepper-motor motion controller and driver fit into a CPLD/FPGA

Stephan Roche, Santa Rosa, CA

This Design Idea further develops a previous one integrating a stepper-motor driver in a CPLD (**Reference 1**). However, this idea integrates not only the driver, but also a simple one-axis stepper-motor motion controller. Depending on the size of the target CPLD, you can implement multiple motion controllers into a single device. For example, a single-axis motion controller fits into a Xilinx (www.xilinx.com) XC95108 using 68 of, or 63% of, the available macrocells. The motion controller rotates the stepper motor clockwise or counterclockwise a given number of steps with a given speed profile versus time. When a motion begins, the controller accelerates until it reaches the cruise speed and then decelerates before stopping (**Figure 1**).

The controller can adjust the motor speed to 16 values, $V=V_{MAX}\times speed/16$, where speed is an integer with a value of one to 16. During the acceleration phase, the speed ramps up by increasing from one to 16; during the cruise phase, speed stays at 16; finally, during the deceleration phase, speed ramps down to one before stopping. If there are insufficient steps for the controller to reach the cruise phase, the controller goes directly from the acceleration phase to the deceleration phase. You can adjust the acceleration/deceleration rate in the program, which you can find at www.edn.com/070621di1 by the constant "accel," which can be one to 255. A high value of accel results in a slow acceleration/deceleration, and a low value results in a fast acceleration/deceleration. The inputs of the CPLD stepper-motor controller are clock, direction, full/half-step, reset, Nstep, start, and stop.

The clock input is active on the positive edge of the clock pulse. The maximum motor speed is one step every 16 clocks. The direction input determines the motor's rotational direction. The motor runs clockwise or counter-

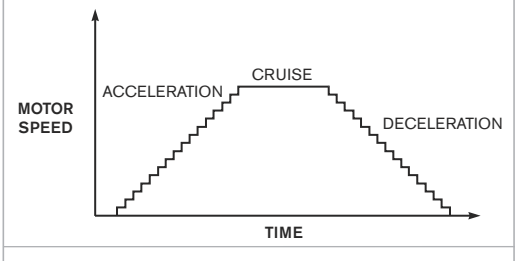

Figure 1 This design can control the motor with 16 speeds. The maximum speed in the cruise phase is such that the motor makes a step or half-step every 16 clock cycles.

DUAL CARD INTERFACE SIMPLIFIES SMART-CARD APPLICATIONS

Provides ±6kV ESD-Protected Connections to Two 5V, 3V, or 1.8V IC Cards with Few External Components

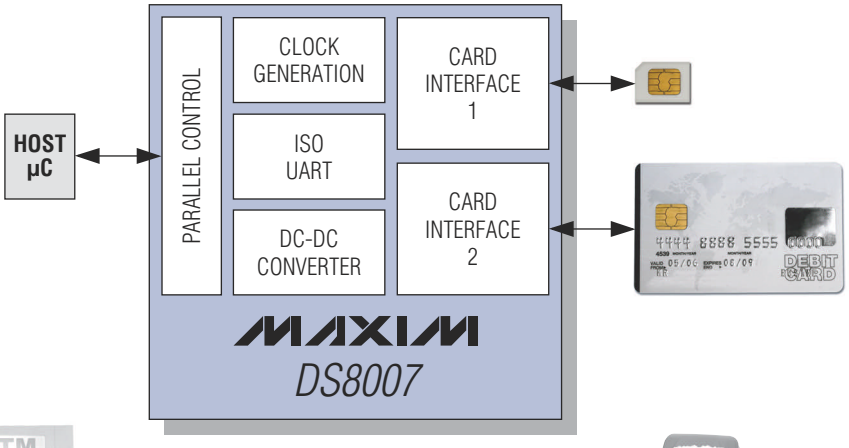

ATMs

SET-TOP BOXES

POINT-OF-SALE TERMINALS

- ±6kV ESD Protection on Card-Interface Pins
- Integrated ISO-7816 UART Provides Communication and Timing
- Short-Circuit and High-Temperature Protection
- Supplies Power to 5V, 3V, and 1.8V Cards
- Parallel Interface to Host µC
- Designed for EMV and GSM 11.11 Applications

Part	Temp Range (°C)	Card Voltages Supported (V)	Supply Voltage (V)	Package	Price† ($)
DS8007+ENG	-40 to +85	1.8, 3, 5	2.7 to 6.0	48-LQFP	2.65

†1000-up recommended resale. Prices provided are for design guidance and are FOB USA. International prices

MAXIM

www.maxim-ic.com/DS8007-info

FREE* A/D Converters Design Guide—*Sent Within 24 Hours!

CALL TOLL FREE 1-800-998-8800 (7:00 a.m.–5:00 p.m. PT) for a Design Guide or Free Sample

MAXIM/DALLAS DIRECT DISTRIBUTION 1-888-MAXIM-IC

ARROW ELECTRONICS, INC. 1-800-777-2776

AVNET electronics marketing 1-800-332-8638

Distributed by Maxim/Dallas Direct, Arrow, Avnet Electronics Marketing, Digi-Key, and Newark.

clockwise, depending on the level of this input and the motor connections. That value is latched at the first rising clock edge after start goes high. The full/half-step input determines the angular rotation of the motor for each clock pulse. In the low state, the motor makes a full step for each applied clock pulse, and, in the high state, the motor makes a half-step. A high level on the reset input sets the motor in a defined state. The motor ignores any clock pulse when reset is high. The 16-bit Nstep value defines the number of steps the next motion will perform. That value is latched at the first rising clock edge after start goes high. A high level on the start input starts the motion, and a high level on the stop input stops the motion, aborting the current motion.

The outputs of the CPLD stepper-motor driver are A, A_N, B, and B_N (**Figure 2**). The A and A_N outputs control one of the motor's coils through power drivers, and the B and B_N outputs control the motor's second coil through power drivers.

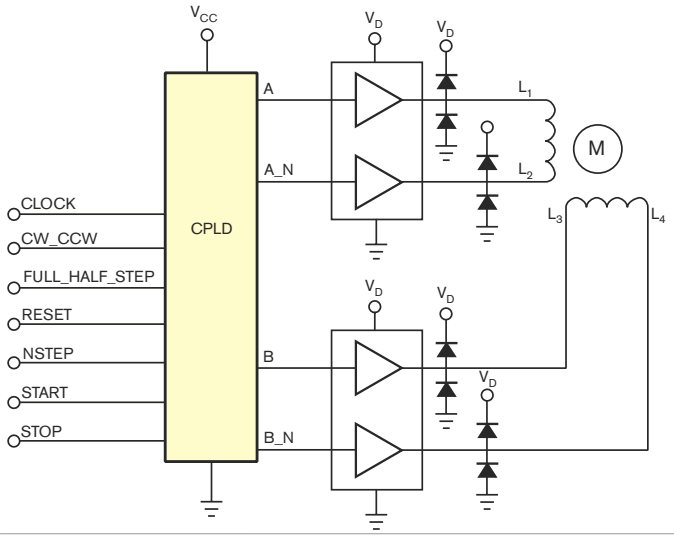

Figure 2 The FPGA/CPLD requires external drivers.

The CPLD/FPGA cannot directly drive the motor, so it requires external drivers. The driver must arrive at the motor's nominal voltage. The Schottky diodes at the output of each driver allow current freewheeling in the motor coils. If you use MOSFET drivers, external Schottky diodes should be unnecessary because MOSFETs have built-in diodes; the Microchip (www.microchip.com) TC4424A dual driver can drive motor coils to 18V and 3A. EDN

REFERENCE

1 Roche, Stephan, "Implement a stepper-motor driver in a CPLD," *EDN*, Feb 15, 2007, pg 90, www.edn.com/article/CA6413791.

World's First Single-Chip Clock Conditioner with Jitter Performance as Low as 200 fs

National Integrates the PLL, VCO, and Distribution Circuitry to Deliver the Industry's Smallest Solution

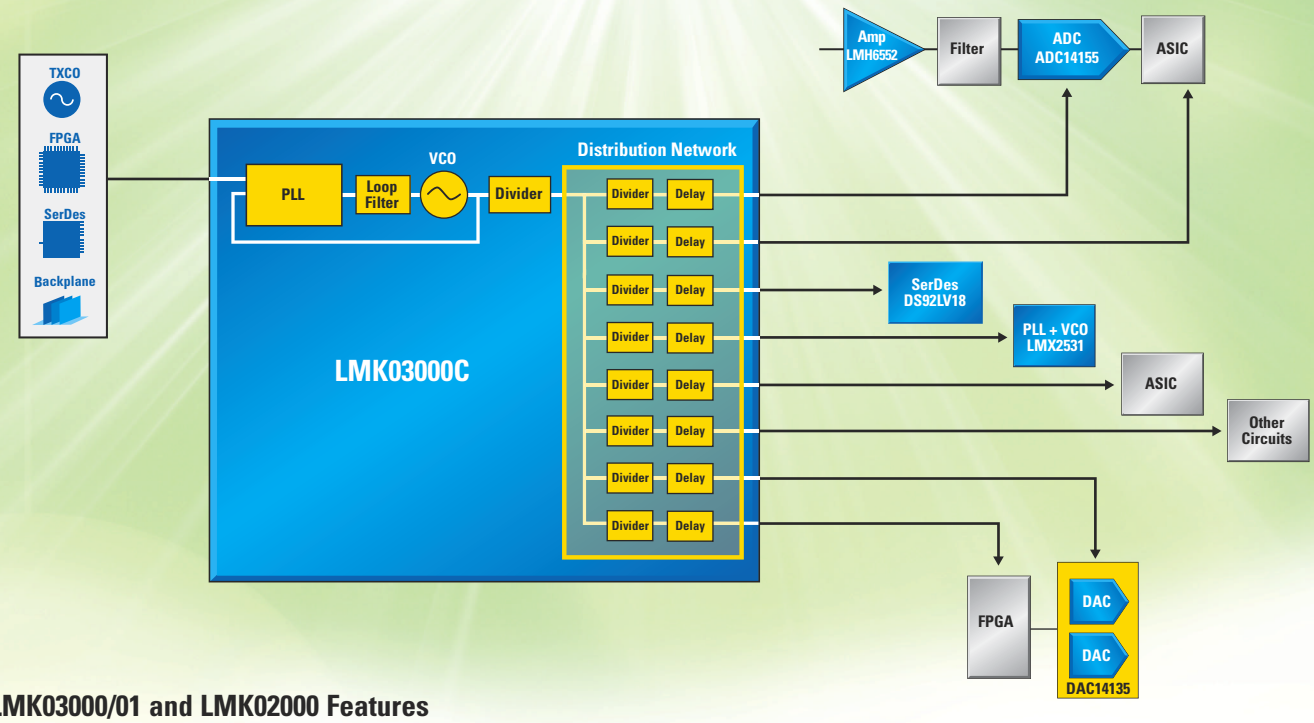

LMK03000/01 and LMK02000 Features

- Fully integrated VCO option delivers unprecedented jitter performance, reducing board space and risk
- Can be configured as jitter cleaner or clock generator
- Available in three performance grades for clocking various high performance applications with diverse jitter requirements
- Footprint compatibility between performance grades
- Three LVDS and five LVPECL clock outputs with dedicated divider and delay blocks simplifies distribution architecture
- Wide clock output frequency range of 1 to 785 MHz
- Small form factor minimizes PCB space by 70%

Performance Grade Table

Performance	Jitter (RMS Typ)
LMK02000	200 fs
LMK03000C/LMK03001C	400 fs
LMK03000/LMK03001	800 fs

Ideal for use in 2G/3G basestations, data converter clocking, networking, medical equipment, instrumentation, military, and aerospace applications

For FREE samples, datasheets, Signal Path Designer, and more information on the LMK03000/01 and LMK02000, contact us today at www.national.com/see/timing

Or call 1-800-272-9959

National Semiconductor
The Sight & Sound of Information

productroundup

OPTOELECTRONICS/DISPLAYS

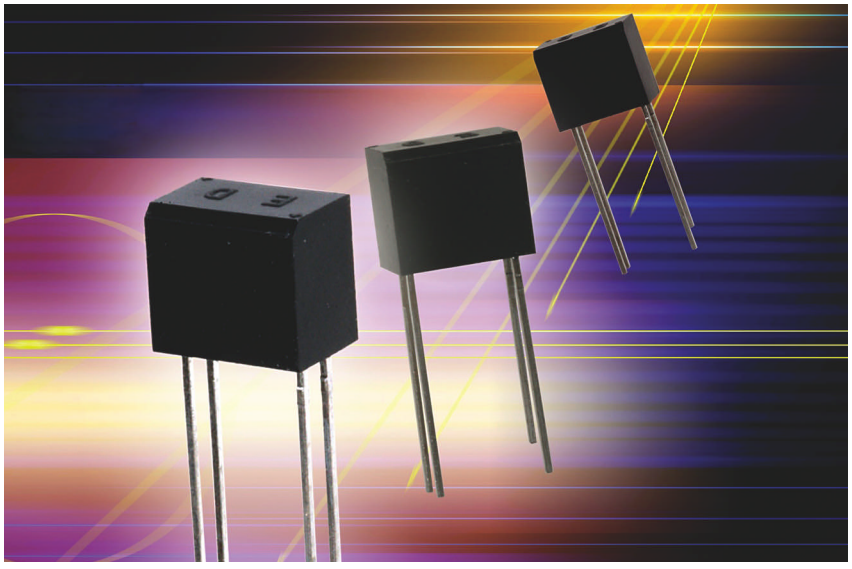

Optically coupled isolators suit 6-kV isolation

The optically coupled OPI7000 isolator family includes an infrared-emitting diode coupled to an NPN silicon phototransistor. A reversed emitter-and-collector pinout provides direct pin-for-pin replacement of lower-isolation-voltage, four-pin, dual inline optocouplers. Features include a ±6-kV input-to-output isolation voltage and a 100-mW power-dissipation rating. The OPI7002 and OPI7010 have an 890-nm LED peak wavelength, and the OPI7320 and OPI7340 isolators have an 890- or 935-nm LED peak wavelength. The OPI7000 series costs $3.25 (1000).

Optek Technology, www.optekinc.com

High-voltage LED driver withstands harsh operating environments

Targeting automotive and general-lighting applications, the MAX16831 high-voltage, high-power, constant-current LED driver features analog and PWM (pulse-width-modulation) dimming. The device integrates a floating LED current-sensing amplifier and a dimming MOSFET driver, reducing the component count and providing high reliability in applications using high-brightness LEDs. Suiting harsh operating environments, the device has a 5.4 to 76V operating input-voltage range, allowing cold start and an 80V load-dump compliance. Aiming at front-light automotive designs, the device suits high- and low-beam assemblies, adaptive front-light systems, DRLs (daytime running lights), and fog lamps. Available in a TQFN-32 package, the MAX16831 costs $2.33 (1000).

Maxim Integrated Products, www.maxim-ic.com

Color-sensing chip comes with an embedded-system interface

Aiming at automatic LCD-backlight-control devices, such as CCFLs (cold-cathode fluorescent lamps) or white LEDs, the VM6101 color-sensing chip provides color- and ambient-light sensing with an embedded-system interface requiring no external components. Using a light-to-frequency-pixel architecture, the device provides a 30 mil-

lion- to 30,000-lux dynamic range. Additional features include 1-mA power consumption in active mode and 1-µA power consumption in standby mode. Available in a 3×3×0.7-mm MLPD-8 package, the VM6101 costs 77 cents (10,000).

STMicroelectronics, www.st.com

Photomicrosensors accurately detect motion and position

Available in a fixed U-shaped housing and carrying both the emitter and the detector, the nonamplified EE-SX1 slotted phototransistor series provides accurate alignment and aiming. Able to detect the presence of an object without direct contact, these slotted, microminiature, transmissive photomicrosensors accurately detect motion and position. Measuring 3×3×3.4 mm with slots measuring 1 to 5 mm, the device provides high resolution through a 0.15- to 0.5-mm aperture width. The EE-SX1 is available in SMD, through-hole-PCB, and panel-mount-with-connector mounting styles, and prices range from 50 cents to $1.10 (10,000).

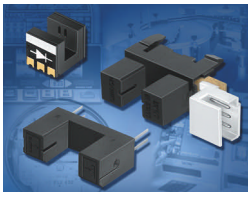

Omron Electronic Components, www.components.omron.com

Ultrabright white SMD LEDs use InGaN technology

The ultrabright white VLMW1100 SMD LEDs have a 0603 footprint with InGaN (indium-gallium-nitride) technology. Features include an 80° angle of half-intensity, a 140- to 280-mcd luminous-intensity range, and a 5500K typical color temperature. Measuring 1.6×0.8×0.6 mm, the VLMW1100 SMD LEDs cost 25 cents.

Vishay Intertechnology, www.vishay.com

Vacuum-fluorescent display features multicolored-LED backlights

Using multicolored LEDs functioning as backlights allows this vacuum-fluorescent display to provide a bright-colored or multicolored background. The display features a 5V power supply and size compatibility with the U-version LCD module. Price for the device is $36 to $80, depending on quantity.

Noritake, www.noritake.com

INTEGRATED CIRCUITS

I²C real-time clocks integrate nonvolatile FRAM

Adding nonvolatile FRAM (ferroelectric-random-access memory) to a real-time clock, the 2-kbit×8-bit DS32B35 and 8-kbit×8-bit DS32C35 devices keep time within one minute per year over the 0 to 70°C temperature range. Suiting GPS (global-positioning systems), utility-power metering, and server applications, the devices use an I²C interface to access the real-time clock and the FRAM. The FRAM provides true nonvolatility without requiring an energy source for data retention without the write-cycle or block-erase limitations of EEPROM or flash memory. Designers can size a separate battery or capacitor for the backup requirements of the real-time clock, providing flexibility in layout and sourcing of components. The devices also feature an integrated accurate TCXO (temperature-compensated crystal oscillator) and a crystal. The devices come in a 0 to 70°C commercial-temperature range or a −40 to +85°C industrial-temperature range. Available in a 300-mil SO-20 package, the DS32B35 costs $4, and the DS32C35 costs $4.80 (1000).

Dallas Semiconductor, www.maxim-ic.com

Step-down dc/dc converter provides 100-, dual 200-, and 300-mA outputs

Using a constant-frequency current-mode architecture, the 2.25-MHz LTC3544 quad-synchronous buck regulator generates a 0.8V output voltage, targeting use in low-voltage DSPs and microcontrollers. This step-down dc/dc converter provides 100-, dual 200-, and 300-mA outputs and has a 2.25 to 5.5V input-voltage range, suiting single-cell lithium-ion/polymer or multicell alkaline/NiCD/NiMH (nickel-cadmium/nickel-metal-hydride) applications. Burst-mode operation provides a 70-µA quiescent current for four channels under no-load conditions; the LTC3544B achieves a lower output ripple using a pulse-skipping mode instead of burst mode. Additional features include 1-µA shutdown current; 95% efficiency; low-dropout, 100%-duty-cycle operation; short-circuit protection; and overtemperature protection. Each channel has independent enable pins and internal soft-start. Available in a 3×3-mm, lead QFN-16 package, the LTC3544 and the LTC3544B cost $2.95 each (1000).

Linear Technology, www.linear.com

EDN ADVERTISER INDEX

Company	Page
Advanced Interconnections	83
Agilent Technologies	2
Analog Devices Inc	27
	29, 43
austriamicrosystems AG	80
Digi-Key Corp	1
Fairchild Semiconductor	15
Freescale Semiconductor	25
International Rectifier Corp	33
Intersil	51
	53, 55
Ironwood Electronics	85
IXYS Corp	68
Keithley Instruments Inc	41
Kemet Electronics Corp	18-19
LeCroy Corp	52
	10
Linear Technology Corp	67
	69-70
	71, 72
Luminary Micro	39
Maxim Integrated Products	75
	77, 79
Micrel Semiconductor	37
Microchip Technology	59
	61
Morsun Guangzhou Science	82
National Instruments	16A-16B
	42
National Semiconductor	6
	47, 81
NEC Electronics	31
NEC Tokin Corp	38
Performance Motion Devices	17
Pico Electronics	40
	54
Power-One Inc	63
Radicom Research	85
Renesas Technology Corp	49
Rohde & Schwarz GmbH & Co KG	12-13
Samsung Electro-Mechanics	34
Samsung Semiconductor	C-3
Samtec USA	66
STMicroelectronics	C-4
Summit Microelectronics	20
Tern	85
Texas Instruments	C-2
	3, 4
	5, 8
Toshiba America	56
Vicor Corp	65
WinSystems	23
XP Power PLC	84

This index is provided as an additional service. The publisher assumes no liability for errors or omissions.

EDN productmart

This advertising is for new and current products.

EDITED BY RON WILSON

scope

CHART YOUR COURSE

LOOKING AHEAD

TO MEMCON SAN JOSE

A growing number of system architects say that memory is one of the hardest problems they face. They must decide which kind of memory structures to use, where to put them, how to organize and implement them, how to model the resulting system performance, and how to deal with the impact of memory choices on system reliability and other issues. Surprisingly, the only conference to deal explicitly with these semiconductor-memory issues is vendor-organized MemCon, which Denali Software runs. The 2007 edition of MemCon, which takes place July 17 through 19 at the Hyatt Regency in Santa Clara, CA, will emphasize memory in the consumer-electronics industry. Three days of technical and strategy tracks, including one on high-speed-interface design for memory systems, explore the design and business aspects of the memory world.

LOOKING BACK

AT GLOBAL COMMUNICATIONS, '50s-STYLE

A single-sideband transmitter designed jointly by the US Army Signal Engineering Labs and Continental Electronics can reach any point on earth even through severe interference, according to the designers. Using a beam antenna, the unit reaches an effective power of 24 MW. The device's filters remove the carrier and one sideband from the modulated signal, directing all power into the remaining sideband. The transmitter can operate across the short-wave band from 4 to 30 megacycles at any of 10 previously set frequencies—selected from a single switch—and carries up to 64 teletype or four voice signals simultaneously. The transmitter occupies only a 50-square-foot room. The final stage is a 150-lb RCA vacuum tube operating at 18 kV.—*Electrical Design News*, June 1957

LOOKING AROUND

FOR THE DISAPPEARING ASIC

Where have all the ASICs gone? Design starts are down, and vendors are quietly de-emphasizing or even dropping out of the business. Probably several factors are at work. One is the increasing capacity of chips. Designs that used to require a half-dozen 100,000-gate ASICs now fit easily into one chip. But the advance of process technology has had less obvious effects, too. That one chip today is likely to be an entire system, including specialized accelerators and multiple CPU or DSP cores. These SOCs (systems on chips) are often beyond the skills of the typical system OEM, and so design activity is shifting to a handful of SOC specialist houses and fabless-chip vendors, which may sell the result as an ASSP (application-specific standard product), not an ASIC. Instead of doing its own chip design, the OEM may customize or order an ASSP. In China, for financial reasons, manufacturers are spinning off ASIC design to small, fabless ASSP start-ups, further distorting the picture as China grows in design activity. The designs are still there, but they are increasingly difficult to count.

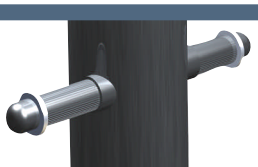

Samsung Memory and Windows Vista… performance worth celebrating

Windows Vista represents a new era in computing. To take full advantage of all the new operating system features, Samsung offers the key hardware components your PC will need. Our full range of DRAM, flash memory devices, graphics memory, and hybrid hard drives, makes Samsung the first stop for Vista-optimizing hardware. Sourcing your parts from Samsung means you're getting quality products from one of the industry's most trusted brands. Learn more about our Vista-enhancing components at:
www.samsung.com/semi/vista

SAMSUNG

Embedded Systems

Solutions

Capture the Power of ARM9 in Flash STR9 Microcontrollers.

Ethernet, USB, CAN, DSP Plus Embedded Control... No Problem

Innovative products for multi-segment application systems

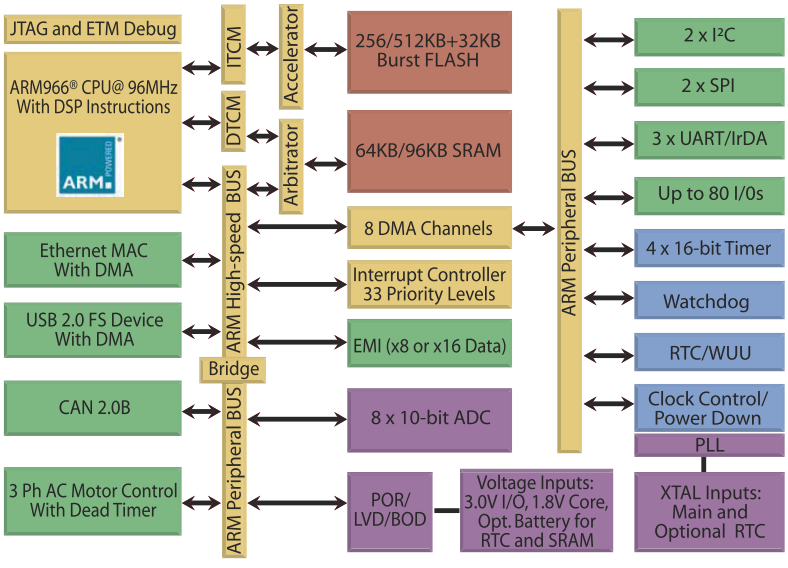

It's All Inside: *STR9 MCUs minimize the need for external components and add innovative features on the inside such as a 1µA real-time clock and a full-function system supervisor.* ***Rapid Transit:*** *96MHz operation, optimized internal data and instruction paths, and 9 DMA channels ensure data moves where you want it, when you want it.* ***No Crowded Spaces:*** *Up to 96KB of SRAM and 544KB of Flash memory accommodate complex applications, TCP/IP, real-time operating systems and multinational products.*

World's first ARM9-based general-purpose Flash microcontroller enables new generation of networked embedded applications

Part Number	Flash Memory (KB)	RAM (KB)	A/D Inputs	Timer Functions (IC/OC/ PWM)	Timer Functions Other	Serial Interface	GPIO (HI Current)	Package	Supply Voltage	Special Features
STR910FM32X	256 + 32	64	8x10-bit	7x16-bit (8,8,7)	RTC WDG	2xSPI 2xI²C 3xUART w/IrDA	40 (16)	LQFP80	Core: 1.8V I/O: 2.7 to 3.6V	CAN
STR910FW32X	256 + 32	64	8x10-bit				80 (16)	LQFP128		CAN, EMI
STR911FM42X	256 + 32	96	8x10-bit				40 (16)	LQFP80		USB, CAN
STR911FM44X	512 + 32	96	8x10-bit				40 (16)	LQFP80		USB, CAN
STR912FW42X	256 + 32	96	8x10-bit				80 (16)	LQFP128		Ethernet, USB, CAN, EMI
STR912FW44X	512 + 32	96	8x10-bit				80 (16)	LQFP128		Ethernet, USB, CAN, EMI

For further information, datasheets, and application notes, visit www.st.com/str9

Order free samples online at: www.st.com

Solutions @

DEVICE AVAILABLE FROM THESE DISTRIBUTORS: ARROW • AVNET • DIGI-KEY • FUTURE • MOUSER • Newark InOne • NU-HORIZONS