Reed Electronics Group

EDN

VOICE OF THE ENGINEER

JULY 19
Issue 15/2007
www.edn.com

Supply Chain: Linking design and resources Pg 91
Signal Integrity Pg 34
Analog Domain Pg 36
Tapeout Pg 38
Tales from the Cube Pg 40
Prying Eyes Pg 42
Design Ideas Pg 77

Is FPGA a simpler puzzle for ASIC designers?

Page 58

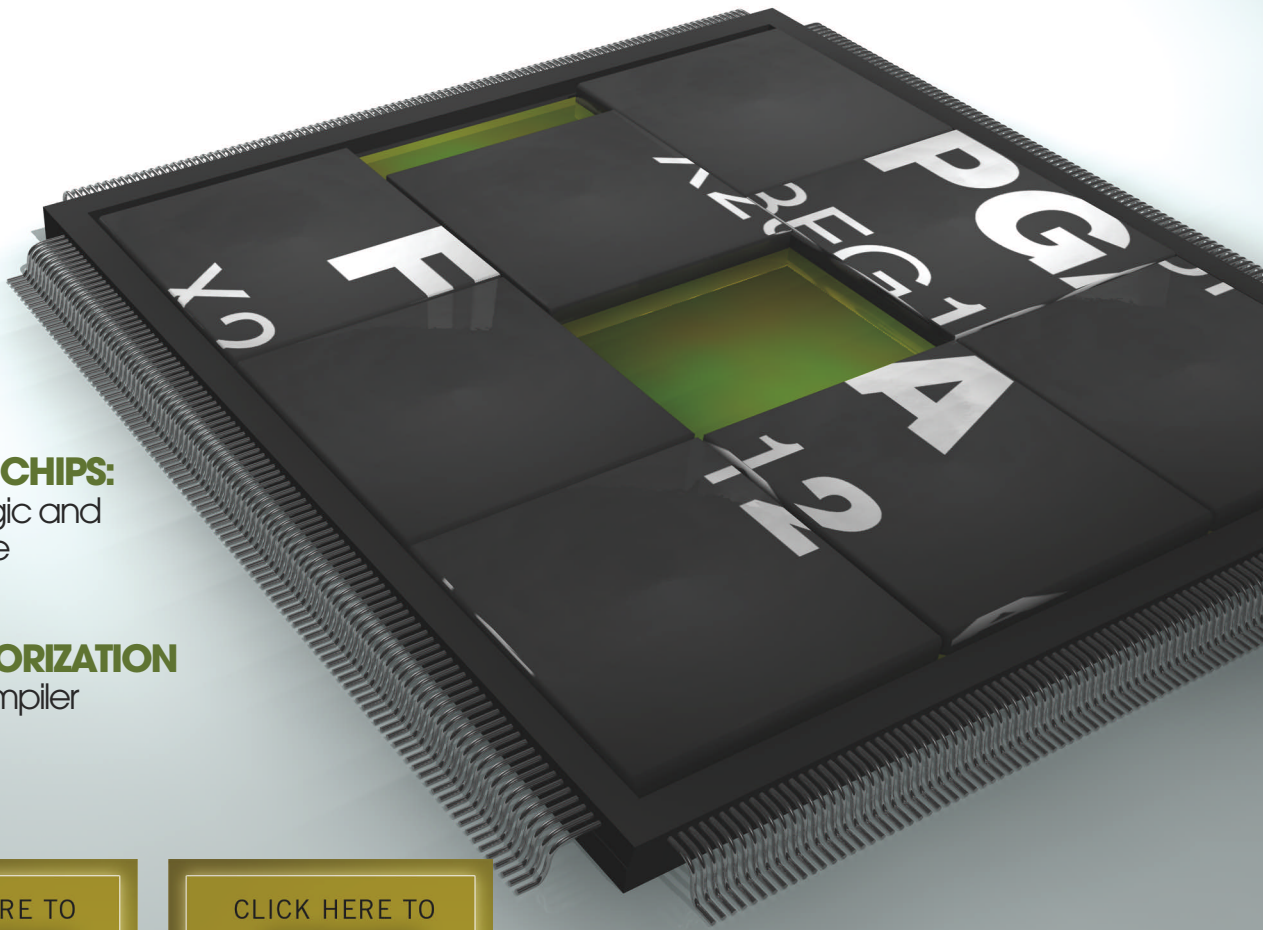

INTERFACE CHIPS: between logic and a hard place
Page 45

AUTOVECTORIZATION for GCC compiler
Page 69

Locate Your ZigBee Solution

First Location System-on-Chip for ZigBee

HIGH-PERFORMANCE ANALOG >> YOUR WAY™

PRODUCTS APPLICATIONS SUPPORT

The **CC2431** from Texas Instruments is the first SoC with a hardware location engine targeting low-power ZigBee® wireless sensor networking applications, such as asset tracking, patient monitoring, inventory control, security and commissioning networks. It features a powerful RSSI-based location engine, which reduces network traffic compared to centralized location systems, and is supported by TI's Z-Stack™ protocol stack.

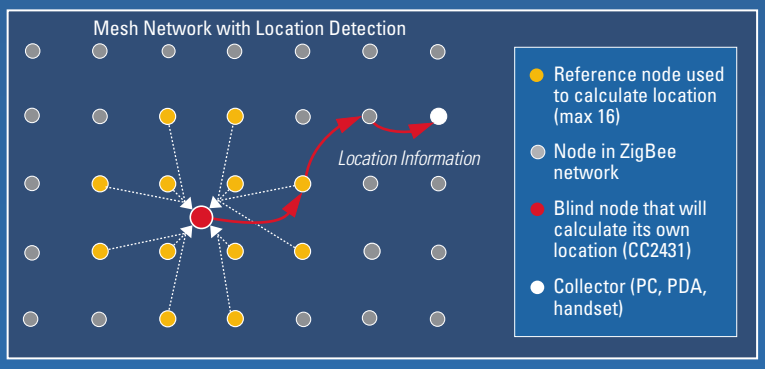

High-Performance Analog >>Your Way

For development kits, samples and datasheets, visit >>

www.ti.com/cc2431

1-800-477-8924 ext 2431

Technology for Innovators™

Texas Instruments

RAQ's

Rarely Asked Questions

Strange but true stories from the call logs of Analog Devices

Amplifier plumbing 101: Are leaky amps sinking your design?

Q. *My amplifier "leaks" a low level signal at the output with the power supplies off. What gives…is my circuit sunk?*

A. Ah yes, I've come across this "plumbing" problem once or twice before, it truly is one of those "rarely asked questions." Upon further discussion with the customer, I found out the amplifier is connected as a voltage follower, the power supplies are off and there is a 2-V p-p signal at the amplifier input… a key piece to the puzzle. The problem appears to be output related, but the source of the leak can be traced back to the input. More specifically, the input protection circuitry for the differential input stage.

Processes used to fabricate high-speed (>50 MHz) amplifiers utilize transistors with very high ft. While these devices operate at very high frequencies, their breakdown voltages can be quite low, hence the need for protection. Input protection can be as simple as a few series diodes connected in anti-parallel between the non-inverting and inverting input of an amplifier. The diodes limit the amount of voltage applied across the base-emitter junctions of the input differential pair and prevent reverse breakdown. The number of series diodes determines the differential input voltage rating. Voltage ratings can range anywhere from ±0.8 V to $\pm V_S$. The value is easily determined by checking the Max Ratings Table in the datasheet, under Differential Input Voltage.

This particular customer was using an amplifier with a ±0.8-V differential input rating. What happened was the protection diodes became forward biased when the 2-V p-p signal was applied to the input. The input signal then "leaked" through the protection diodes to the output via the feedback resistor. Fortunately, you don't need to call a plumber to fix this type of leak. Instead, reduce the input signal level to below ±0.8 V p-p, or simply use an amplifier with a higher max differential input voltage rating.

So, next time you come across a "leaky" amplifier, check your differential input voltage rating. And, remember, if you can't lower the input voltage to solve the problem, then, just like a plumber who needs to replace a washer to stop a leak you might have to replace the amplifier.

To learn more about high-speed amplifiers,
Go to: http://rbi.ims.ca/5397-101

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.

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

High Performance x 4

Quad, 125 MSPS, Low-Power, Small Form-Factor ADC

HIGH-PERFORMANCE ANALOG >> YOUR WAY™

PRODUCTS APPLICATIONS SUPPORT

The new **ADS6425** from Texas Instruments is the first in a family of four-channel, low-power, 12- and 14-bit pin-compatible ADCs. Consuming just 410 mW/channel, the 125 MSPS ADS6425 features serialized LVDS outputs that allow four highly programmable ADCs to fit in a small 9x9mm, 64-pin QFN package, making it ideal for a wide variety of high-density, low-power applications.

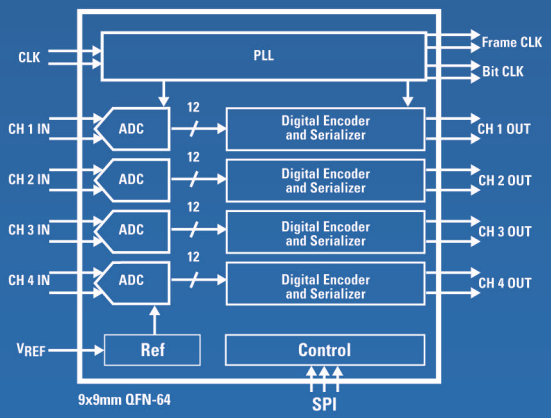

High-Performance Analog >>Your Way

For Data Sheets, Samples, Evaluation Modules and Deserializer Solution, Visit >>
www.ti.com/ads6425

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

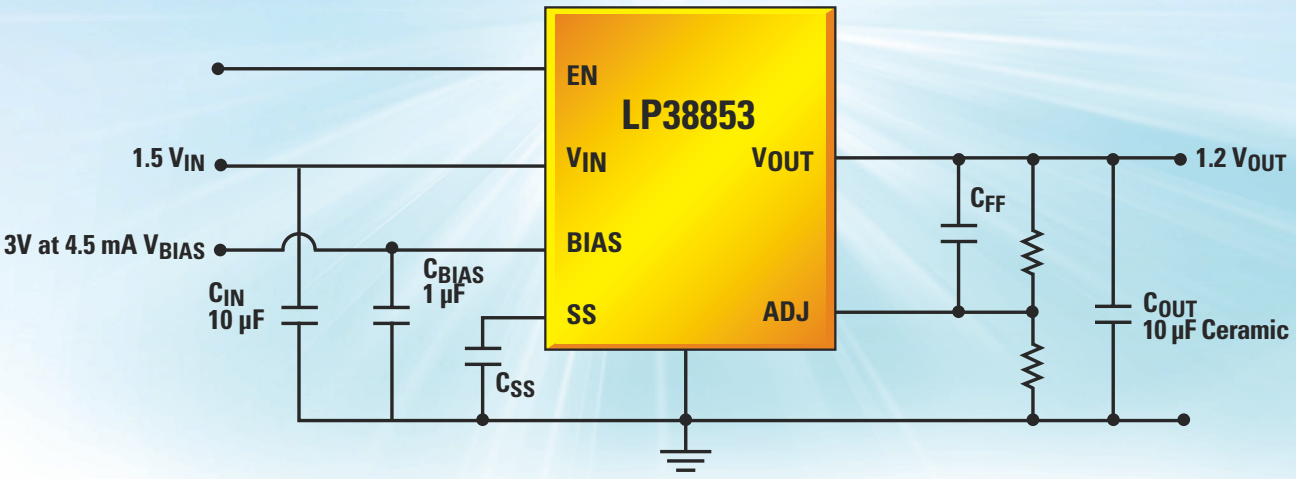

LDOs for Powering Digital ICs

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

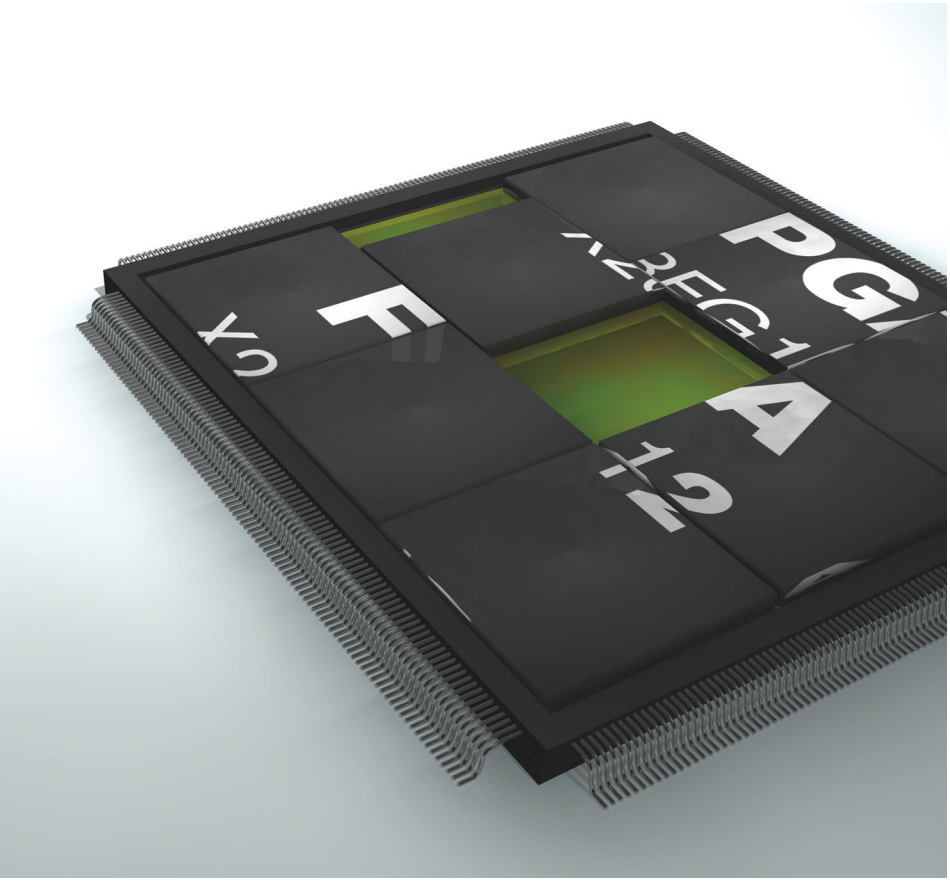

EDN contents

7.19.07

Autovectorization for GCC compiler

69 You can use industry-standard benchmarks to improve compiler performance.

by Markus Levy, EEMBC, and Ron Olson, IBM

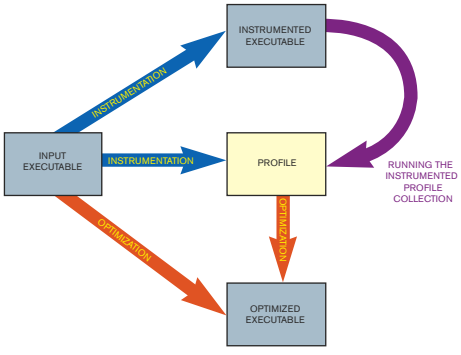

Is FPGA a simpler puzzle for ASIC designers?

58 With rising mask costs, complexity, and tool expenses to develop ASICs and SOCs, many design groups today are opting to implement their production designs in FPGAs. But, before designers make the leap, there are several factors—good and bad—they should consider.

by Michael Santarini, Senior Editor

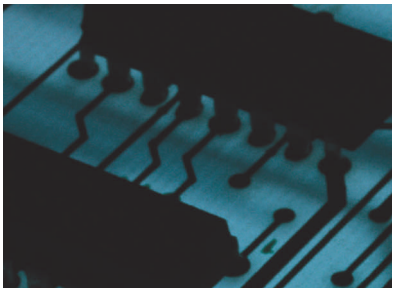

Interface chips: between logic and a hard place

45 Interface chips must mediate between ever-lower logic voltages and real-world loads. Here are the circuits and techniques you need to know.

by Paul Rako, Technical Editor

DESIGNIDEAS

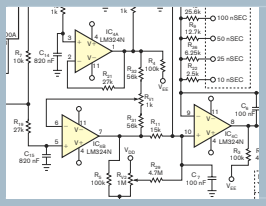

77 Analyzer tests reverse-recovery behavior of diodes

78 High-power LED drivers require no external switches

82 Perform PSRR testing with analyzers having no dc-bias ports

▶ Send your Design Ideas to edndesignideas@reedbusiness.com.

M-LVDS – The Bus Standard

Signal Integrity for Backplanes and Cables

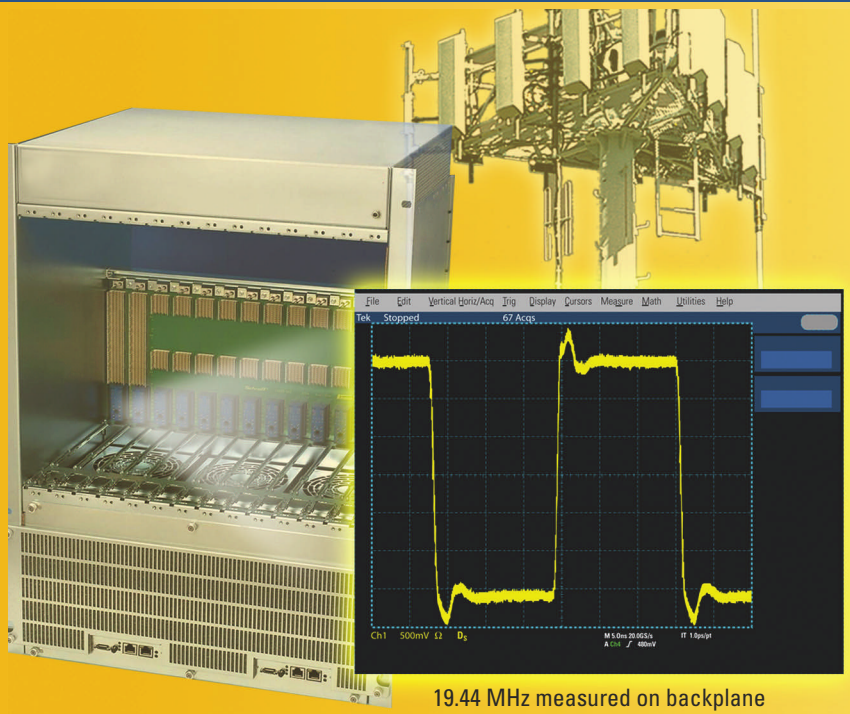

19.44 MHz measured on backplane

Applications

- AdvancedTCA (ACTA) and MicroTCA
- Clock, data, sync and control
- Cellular base stations
- Central-office switches
- Test and measurement

Features

- Meets M-LVDS (TIA/EIA-899)
- 11.3mA driver output current
- 50mV receiver thresholds
- Contention provisions:
  - Driver short-circuit limited to 43mA
  - Devices limit bus voltage from 0 to 2.4V
  - Drivers tested with 32 contending nodes
- Failsafe support

HIGH-PERFORMANCE ANALOG >> YOUR WAY™

PRODUCTS APPLICATIONS SUPPORT

Texas Instruments provides a broad range of innovative Multipoint LVDS (M-LVDS) devices that deliver the speed, power and robustness required for backplane and multipoint bus applications. With contention features for data and control signals, single- and eight-channel transceivers, a quad driver, single-channel receivers, and clock buffer devices for 1:8 and dual 1:4 fan-outs, TI's M-LVDS products are ideal for distributing clock and synchronization signals.

Device	Signaling Rate (Mbps)	No. of Tx / Rx	Input	Output	ESD HBM (kV)	Function
SN65MLVD047A	200	4/0	LVTTL	M-LVDS	9	Quad Driver
SN65MLVD080	250	8/8	LVTTL, M-LVDS	M-LVDS, LVTTL	8	Octal Half-Duplex Transceiver
SN65MLVD128	250	8/1	LVTTL	M-LVDS	9	Repeater
SN65MLVD201	200	1/1	LVTTL, M-LVDS	M-LVDS, LVTTL	8	Half-Duplex M-LVDS Transceiver
SN65MLVD207	200	1/1	LVTTL, M-LVDS	M-LVDS, LVTTL	8	Full-Duplex M-LVDS Transceiver
SN65MLVD2	250	0/1	M-LVDS	LVTTL	8	Single Channel M-LVDS Receiver
SN65MLVD3	250	0/1	M-LVDS	LVTTL	8	Single Channel M-LVDS Type-2 Receiver

New products listed in red

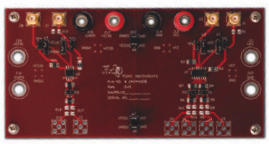

For evaluation modules, data sheets and samples visit >> **www.ti.com/mlvds**

800. 477. 8924 ext.13841

Technology for Innovators™

Texas Instruments

contents 7.19.07

pulse

Dilbert 24

23 Mezzanine card features dual cores

23 Oscillator offers one part/trillion drift/day

24 Sierra releases IC-place-and-route tools for 45-nm era

26 MIPS introduces 1-GHz 74K processor core

26 COM Express adopts multicore architecture

28 Small, 1-GHz-bandwidth, two-channel PCI digitizer offers low power consumption

28 Quad-core computer features five PCIe lanes

30 Cadence offers low-power-methodology kit

30 Video DSPs double performance, reduce system-BOM cost

32 **Research Update:** Dense nanotubes allow efficient interconnect; Circuit-design technique could make carbon nanotubes useful; Thermoacoustics speaks up at ASA convention

40

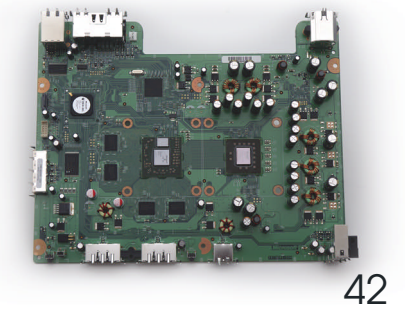

42

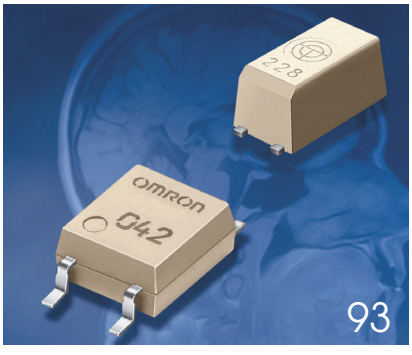

93

DEPARTMENTS & COLUMNS

10 **EDN.comment:** Qualcomm-IP quagmire launches supply-chain page

34 **Signal Integrity:** Uncertainty principle

36 **Analog Domain:** From thin air

38 **Tapeout:** Where is the BOM in IC design?

40 **Tales from the Cube:** Avalanche!

42 **Prying Eyes:** Eviscerating the Xbox 360 Elite

91 **Supply Chain:** Linking design and resources

98 **Scope:** Hot Chips, computerized weapons systems, and missing cell phones

PRODUCT ROUNDUP

93 **Discrete Semiconductors:** MOSFET relays, eye-safe diode lasers, DirectFET MOSFET chip sets, overvoltage-protection controllers, Schottky rectifiers and TVS devices, and more

94 **Embedded Systems:** Synchro/resolver-to-digital converters, SATA-to-CF adapters, digital-I/O PCI cards, USB digital-I/O devices, and more

95 **EDA Tools:** Multimode simulation platforms

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 15 (Printed in USA).

LeCroy invites you to unusual events.

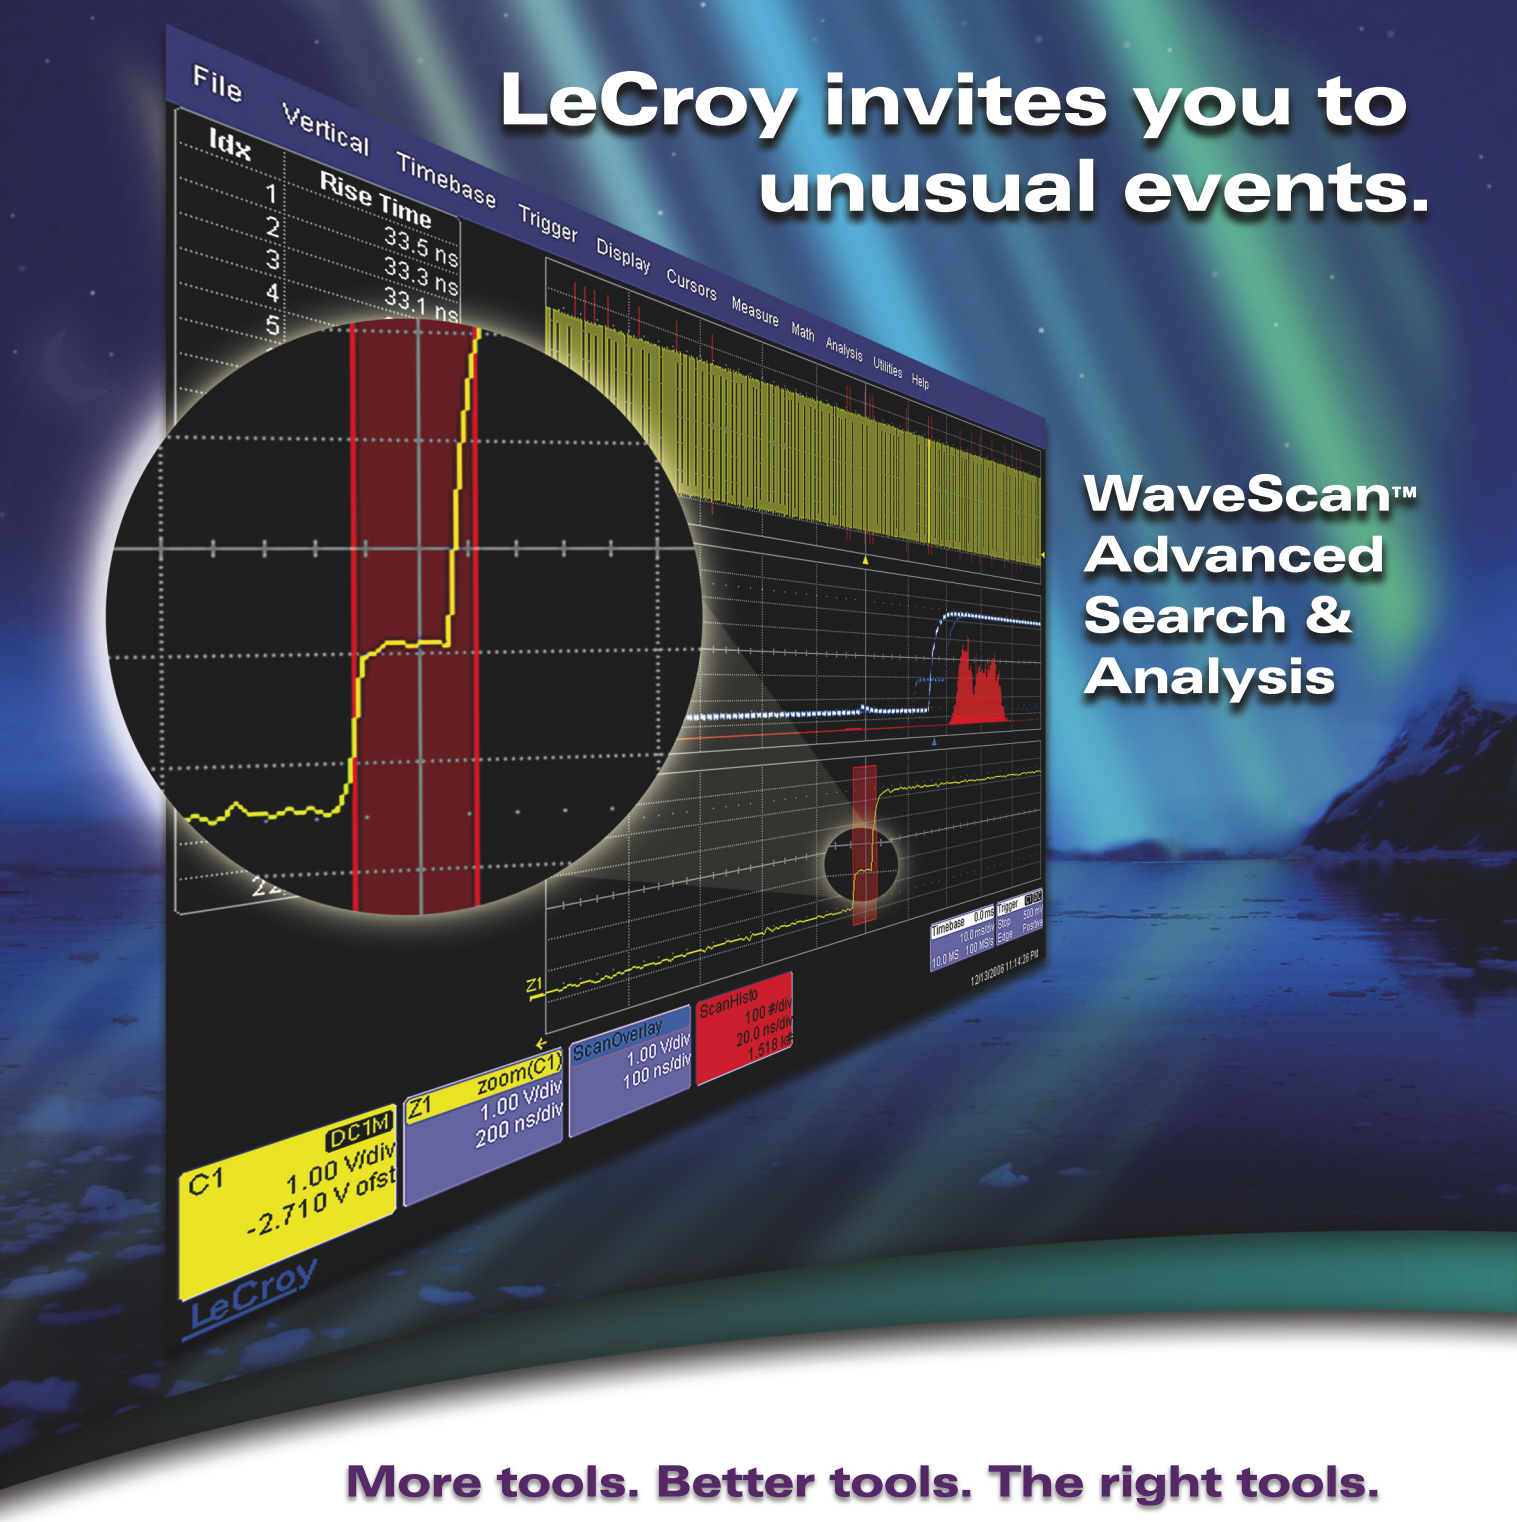

WaveScan™ Advanced Search & Analysis

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

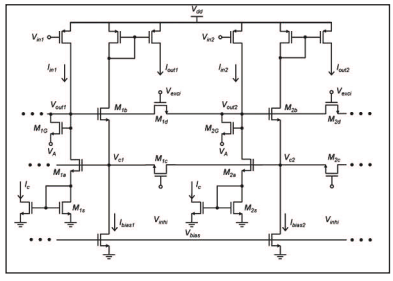

ONLINE ONLY

Check out these Web-exclusive articles:

CMOS winner-take-all circuits: A tutorial
The human-vision-processing system is built of numerous complex neural layers that communicate with one another by means of feedforward and feedback neural connections.
→www.edn.com/article/CA6454765

Audio, video chips enhance handheld devices' features, shrink power needs
→www.edn.com/article/CA6456165

Refining multicore concepts (parts 1 and 2)
→www.edn.com/article/CA6454564
→www.edn.com/article/CA6456728

Designing "with" test instead of "for"
→www.edn.com/article/CA6454314

EDA's ESL vendors: Unite and take over! ... or at least get organized
→www.edn.com/070719toc3

iPhone's A2DP Absence: What's the defense?
→www.edn.com/070719toc1

HOTweezer thermal wire strippers and Harleys
→www.edn.com/070719toc2

The real story of hard drives versus solid-state drives
→www.edn.com/070719toc4

Beyond vanilla CMOS: Are III/V compounds coming into the mainstream?
→www.edn.com/070719toc5

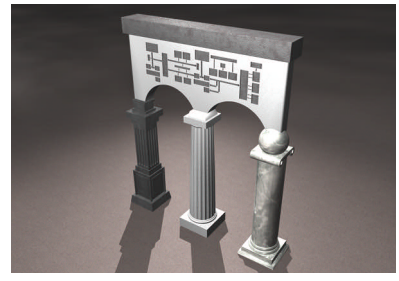

READERS' CHOICE

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

Exploring memory architectures: pillars of processing performance
→www.edn.com/article/CA6451245

Free PCB layout manual
→www.edn.com/070719toc7

Design for debugging: the unspoken imperative in chip design
→www.edn.com/article/CA6451246

Kodak Moment: A revolutionary color filter?
→www.edn.com/070719toc6

RC lowpass filter expands microcomputer's output port
→www.edn.com/article/CA6451248

Nice engineering demo of wireless power transmission
→www.edn.com/070719toc8

Harvesting energy from vibrations
→www.edn.com/article/CA6454543

Add a grounded-switch feature for Topswitch on/off control
→www.edn.com/article/CA6451247

CEO Roundtable: The brightest ideas in electronics
→www.edn.com/article/CA6445798

EZ-Color controller family simplifies HB-LED-lighting-system designs
→www.edn.com/article/CA6448812

IBM Inside: Interview with Bernard Meyerson
→www.edn.com/article/CA6452442

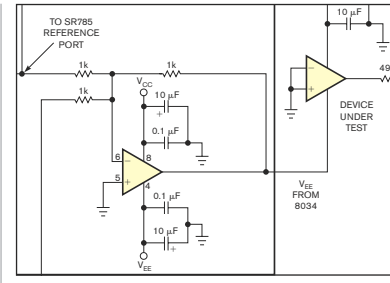

DESIGN IDEAS

Need more Design Ideas?
Visit the Design Ideas home page for all the latest engineer-submitted circuit designs, plus Design Ideas editor Charles H Small's commentaries, a Design Ideas topic index, and information on submitting your own clever concept.
→www.edn.com/designideas

MOVERS and SHAKERS
THE DEFINITIVE GUIDE TO LEADERSHIP IN GLOBAL ELECTRONICS 2007

Movers and Shakers
Check out our annual look at the people, companies, and technologies driving the tech industry.
→www.edn.com/moversandshakers

PART SEARCH

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

BY MAURY WRIGHT, EDITORIAL DIRECTOR

Qualcomm-IP quagmire launches supply-chain section

This issue of *EDN* marks the debut of a supply-chain section (pg 91) that will initially appear in every other issue and, we hope, eventually in every issue. In the new section, we will cover breaking news in the pricing, inventory, and product-availability areas. We'll include regular updates on green and environmental issues that continue to change the way engineers and vendors design and manufacture products worldwide.

Finally, we'll cover the distribution front, supply-chain-oriented events, and pertinent market statistics. The first instance of this section leads with a story on the recent legal IP (intellectual-property) battles between Qualcomm and Broadcom and the impending North American banishment of mobile handsets that use certain Qualcomm ICs.

The supply chain is nothing new to *EDN*. We've sporadically covered such stories in our Pulse section and in feature articles. Beginning with the relaunch of EDN.com last April, we introduced an online channel dedicated to supply-chain issues in the business center on the Web site. Suzanne Deffree is our resident supply-chain expert as well as the managing editor of our news center. She has been driving content on our online supply-chain channel, and she will edit the print section, as well.

Please take a look at the new section and let me know what you think. I believe the information is critical to design engineers. The ability to meet market windows means that designers must use efficient and quality design techniques and make supplier and component choices that will not impact the move to volume production.

A lack of cooperation will slow the promulgation of every compelling new technology.

I find the Qualcomm story especially interesting because I live in the San Diego area, where Qualcomm is based and where we get daily news of the legal skirmish. Qualcomm has long been among the most hawkish of companies when it comes to profiting from IP, and now the predator has become the prey. I think that Qualcomm will soon realize that it had better pay up. Others, such as Nokia, are ready to pile on with their own claims that Qualcomm has illegally used their IP. The threat of escalating royalties that Qualcomm might owe, combined with the lack of revenue from North American phone shipments, should nudge the company toward compromise. The ramifications go beyond Qualcomm. The handsets in question also carry other vendors' components. Moreover, these skirmishes could significantly affect consumers.

I recently bought a nifty Samsung Sync handset for the AT&T/Cingular network. The GSM (global-system-for-mobile-communications) network hasn't in the past been Qualcomm's market. Lately, however, the company has begun to compete vigorously with market leader Texas Instruments and other vendors in the GSM market. My Sync carries a small sticker on the back that notes that the phone uses Qualcomm HSDPA (high-speed-downlink-packet-access) technology—the 3G technology favored by the GSM universe. My understanding is that the recent legal ruling would ban such phones.

I think that the Qualcomm-versus-Broadcom battle indicates a growing problem that might quash innovation. Companies in the semiconductor industry have simply forgotten how to work together. Once-prevalent cross-licensing is now rare. Without more cooperation and cross-licensing, standards battles will continue to escalate, just as in the IEEE 802.11n wireless-LAN effort. In effect, a lack of cooperation will slow the promulgation of every compelling new technology.

Companies—especially public companies—need to make money and deliver value to shareholders. Somehow, they did just that during the era of broad cross-licensing and cooperation that started around 1960 and continued for most of the rest of the last century. Perhaps we need the industry to step back before moving forward. Consumers can't drive the industry by buying compelling new products when the courts are holding up introduction of those products. EDN

Contact me at mgwright@edn.com.

MORE AT EDN.COM

Go to www.edn.com/070719ed and click on Feedback Loop to post a comment on this column.

FAIRCHILD SEMICONDUCTOR®

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

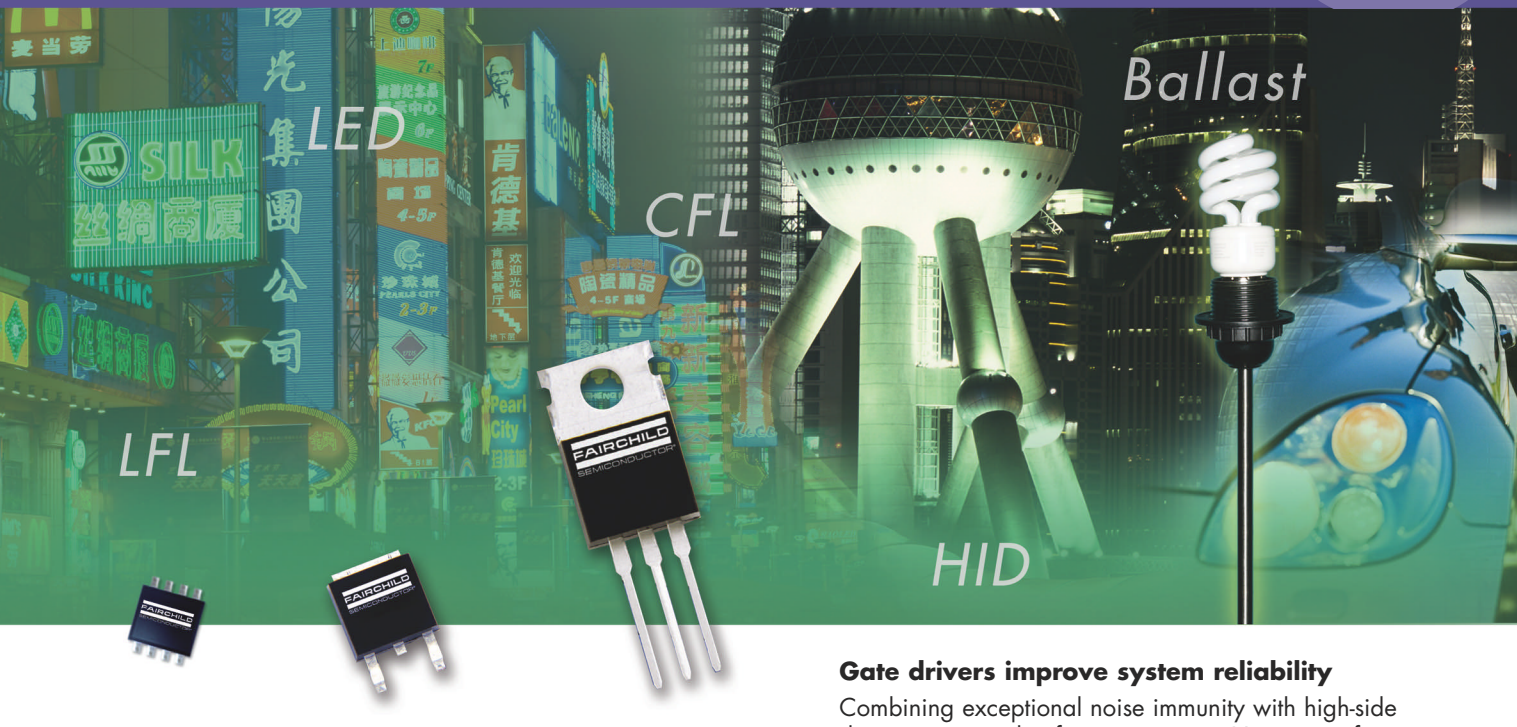

Integrated ballast control ICs simplify designs

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.

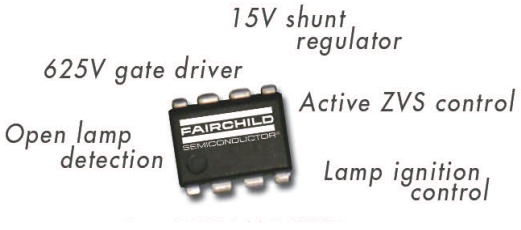

FAN7711 Electronic Ballast Control IC

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.

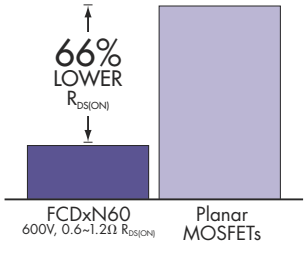

To learn more about our complete lighting solutions—including PFC ICs, IGBTs and transistors—please visit www.fairchildsemi.com/lighting.

FAIRCHILD SEMICONDUCTOR CELEBRATING 50 10 YEARS

www.fairchildsemi.com

the power franchise™

SuperH Flash Microcontroller Reaches Speeds up to 160MHz

Superscalar performance,
high-speed on-chip flash memory access, and much more

Renesas Technology

No.1* supplier of microcontrollers in the world

proudly introduces the SuperH family of superscalar flash devices. With 160MHz superscalar performance, high-speed access to on-chip flash memory and up to 200MHz CPU operation from SH-2A devices, the SuperH family allows you to incorporate all of your brilliant ideas into super high performance applications. The SuperH RISC engine and the SH-2A core are establishing new performance standards in the industry. The SH-2A core is geared towards systems that demand real-time, high-precision control and require a combination of high performance CPU and high-speed flash.

SuperH MCU Roadmap

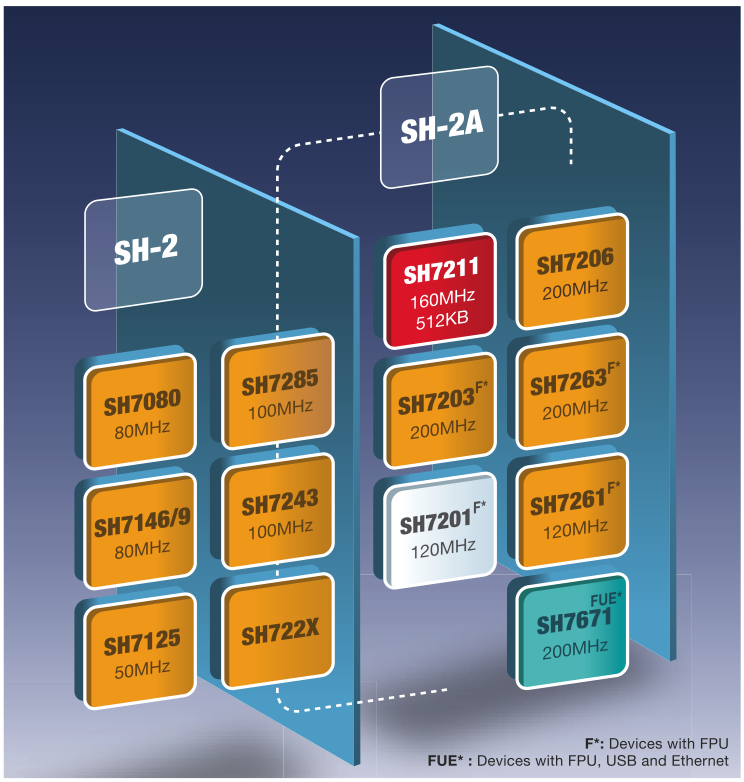

HOT Products **SH7211F**

(R5F72115D160FPV)

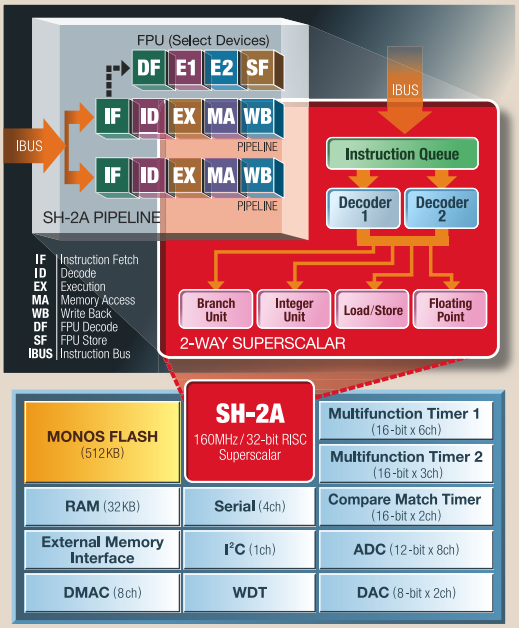

Top Reasons to Select SuperH

480 Dhrystone MIPS at **200** MHz

- **Superscalar Performance**
  - Two instructions are executed per cycle at 160MHz
  - World's fastest embedded FLASH with 12.5ns read access time.
- **Fast Real-Time Control**
  - Register-Banks enable 37.5nsec (6 cycles) interrupt latency time.
- **High Integration**
  - Large on-chip flash and on-chip RAM
  - Advanced 16-bit PWM timers to drive two motors simultaneously
  - 1.25μsec 12-bit A/D conversion with 3 sample & hold circuits

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

#1 MCU REACH FURTHER

Get Started Today -

Register to qualify for a SuperH MCU Kit

www.america.renesas.com/ReachSH/a

Starter Kit for SH7211F:
R0K572115S000BE
USB Emulator:
HS0005KCU11H

Renesas Technology Corp.

Everywhere you imagine. RENESAS

SIGNAL PATH designer®

Tips, tricks, and techniques from the analog signal-path experts

No. 110

Feature Article....1-7

LVDS Buffers 2

High-Speed Differential Amps and ADCs 4

Video Cable Equalizer 6

Design Tools 8

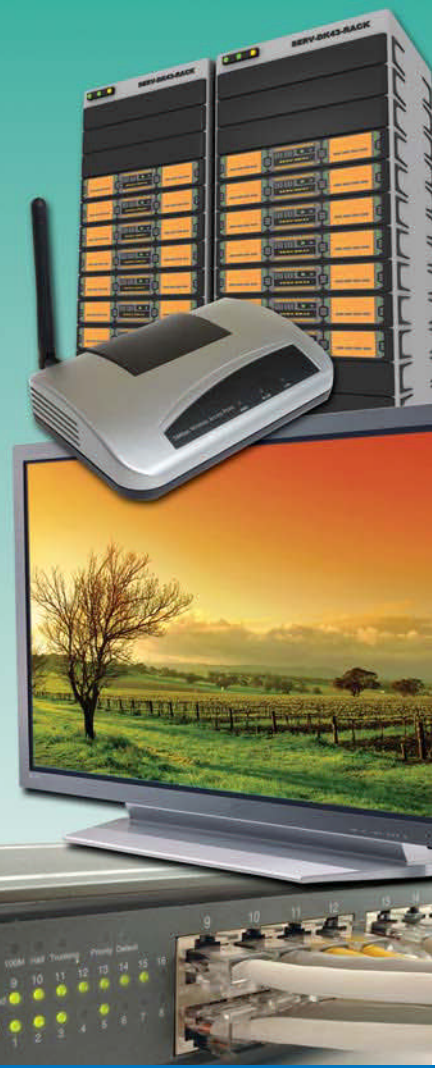

Extending the Signal Path Over Data Transmission Lines

— By Lee Sledjeski, Applications Engineer

After using an ADC to convert an RF signal to the digital domain, this signal often needs to be sent across a backplane or cable to an FPGA for additional processing. As sampling speeds continue to increase, the challenge becomes driving this signal more than a few dozen centimeters without bit errors. This article examines how to solve the challenges of extending high-speed signal paths over FR-4 traces and copper cables.

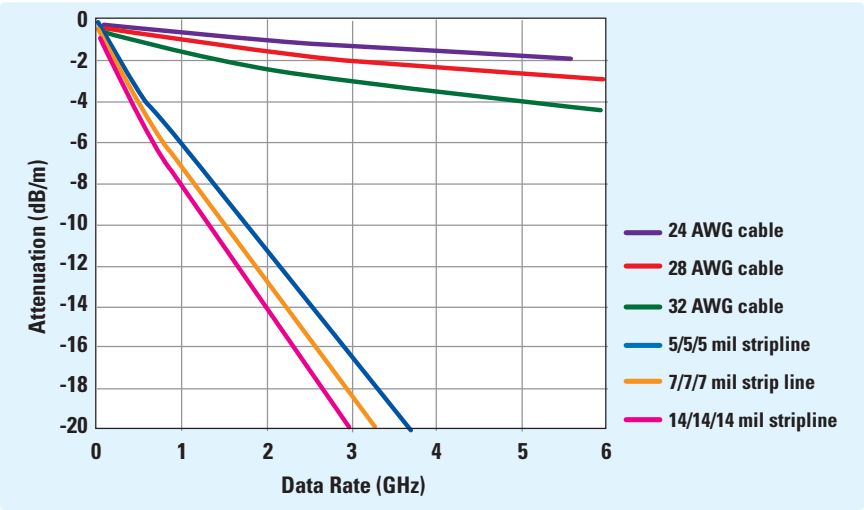

Figure 1. Transmission Media Loss (Differential) Shows the Low Pass Filter Effect

As an example of these high-speed signal challenges, consider the graph in ***Figure 1***. When a 2.5 Gbps NRZ signal transitions every bit cycle it generates the equivalent of a 1.25 GHz clock signal. This component of the 2.5 Gbps signal sees up to 10 dB loss after just 1 meter of 100Ω differential FR-4 stripline. It takes as much as 10 meters of twinaxial cable to generate a similar media loss. This transmission loss creates jitter that eventually closes the signal "eye" or sampling window as seen in ***Figure 2a***.

For high-speed signals, a low pass filter (***Figure 1***) represents the transmission media. The highest frequency that passes through the media with less than 3 dB of attenuation defines the filter bandwidth. Signal components within the

National Semiconductor
The Sight & Sound of Information

LVDS Plus Signal Conditioning Solves Your Signal Integrity Problems

3.125 Gbps LVDS Buffers with Pre-Emphasis and Equalization

DS25BR100/110/120 Features

- 17 ps typ jitter (2.5 Gbps, 2^{23}-1 pattern)
- Programmable pre-emphasis and equalization
- Wide range input accepts LVDS, CML, and LVPECL
- Integrated 100Ω input and output termination
- 7K ESD protection on LVDS pins

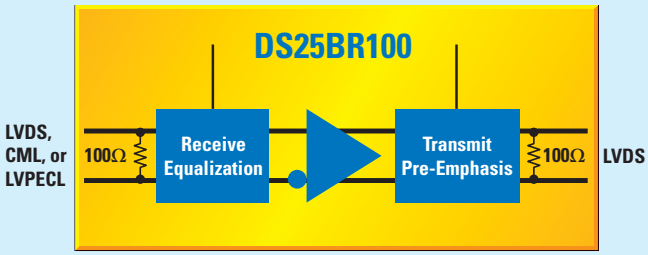

Signal Conditioning

Product ID	LVDS Buffer Function	Inputs	Outputs	Input Compatibility	Output	Pre-emphasis (dB)	Receive Equalization (dB)	Max Speed/Ch (Mbps)	Packaging	Comments
DS25BR100	Single	1	1	LVDS/LVPECL/CML	LVDS	0/6	3/6	3125	LLP-8	Int termination, 8 kV ESD
DS25BR110	Single	1	1	LVDS/LVPECL/CML	LVDS	—	0/3/6/9	3125	LLP-8	Int termination, 8 kV ESD
DS25BR120	Single	1	1	LVDS/LVPECL/CML	LVDS	0/3/6/9	—	3125	LLP-8	Int termination, 8 kV ESD
DS90LV804	Quad	4	4	LVDS/LVPECL/CML	LVDS	—	—	800	LLP-32	Int termination, 15 kV ESD
DS90LV004	Quad	4	4	LVDS/LVPECL/CML	LVDS	0/2/4/6	—	1500	TQFP-48	Int termination, 15 kV ESD
DS15BR400	Quad	4	4	LVDS/LVPECL/CML	LVDS	0/6	—	2000	LLP-32, TQFP-48	Int termination, 15 kV ESD
DS15BR401	Quad	4	4	LVDS/LVPECL/CML	LVDS	0/6	—	2000	LLP-32, TQFP-48	15 kV ESD

Ideal for use in routers, switches, storage, medical imaging, video security, and image capture and processing

Repeater

Equalizer

Pre-emphasis Buffer

3 x 3 x 0.8 mm LLP-8
(Shown actual size)

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

lvds.national.com

Extending the Signal Path Over Data Transmission Lines

media bandwidth pass through with minimal attenuation. Furthermore, the media attenuates the amplitude of signal components and harmonics beyond the -3 dB point or transmission line bandwidth. Data-dependent jitter due to Inter-Symbol Interference (ISI) results from this non-linear, frequency-dependent loss. Therefore, signal conditioning refers to the compensation techniques used to mitigate the effects of high-frequency transmission losses. Conditioning input and output signals enhance the performance and extend the signal-path distance.

Input and Output Signal Conditioning

Measured at the Equalizer Input

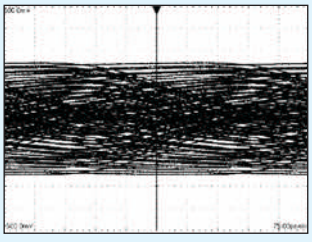

2.5 Gbps NRZ PRBS-7 after 70" Differential FR-4
Stripline V:100 mV/DIV, H:75 ps/DIV

Figure 2a. 2.5 Gbps NRZ Signal After 70 Inches of FR-4 Without Equalization

Signal integrity designers use equalizers to condition input signals. The equalizer circuit goal is to reverse the signal losses incurred during transmission by "flattening" the frequency response of the system and thereby reducing the distortion or "smearing" that the signal encounters. The equalizer flattens the transmission frequency response by acting as a high-pass filter that approximately complements the low-pass effect of the transmission medium within the signal's frequency band. As a general rule of thumb, the equalization should flatten the frequency response up to the frequency produced by the highest transition density data pattern possible. For example, a 2.5 Gbps data rate signal in a 1-0-1-0 NRZ pattern would require a flat frequency response out to 2.5 GHz (***Figure 1***). When properly tuned, equalization can significantly reduce the ISI effects from the transmission media as shown in ***Figure 2b***. Even for 29.5" stripline, attenuation

Measured at the Equalizer Output

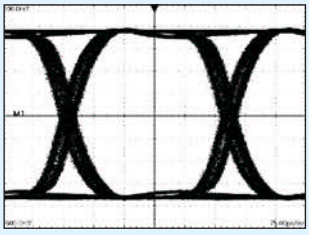

Equalized 2.5 Gbps NRZ PRBS-7 after 70" Differential FR-4
Stripline V:100 mV/DIV, H:75 ps/DIV

Figure 2b. 2.5 Gbps NRZ Signal After 70 Inches of FR-4 With Equalization (b)

approaches -6 dB at 2.5 Gbps, and it is safe to say that reliable communication across the transmission link requires some equalization.

High-speed devices such as the DS25BR110 feature a receiver input equalization circuit to reduce the effects of frequency dependent losses caused by the transmission medium (***Figure 2***). Four levels of EQ control ranging from 0 to 16 dB allow for easy optimization of signal quality across a broad range of typical transmission media lengths.

Transmitter output signal conditioning is used to produce a similar overall net effect. While input equalization acts to cancel frequency-dependent losses, output Pre-Emphasis (PE) alters the frequency content of a clean, unattenuated signal with the expectation that the transmission media attached to the driver will attenuate the signal. This results in a clean signal at the receiver positioned on the far end of the transmission line. By applying PE to the output waveform, the highest frequency components of the signal are emphasized at the driving device. Because signal conditioning requirements increase with higher data rates, multiple levels from 0 to 9 dB can be selected to optimize the signal integrity at the receiver for most common PCB and copper cable transmission distances (***Figure 3***).

Setting Pre-Emphasis and Equalization

Signal conditioning is designed to compensate for the low-pass filter effect of the transmission medium. As a starting point, adjust the level to match the loss of

High-Speed Comparator and High-Bandwidth A/D Converters with LVDS Outputs

Low-Power Comparator and ADCs for High-IF Sampling Systems

LMH7220 Features

- 2.9 ns propagation delay (100 mV overdrive)
- 0.6 ns rise/fall times
- Supply voltage: 2.7V to 12V
- Supply current: 6.8 mA
- Temperature range: -40°C to 125°C
- LVDS output
- Available in TSOT-6 packaging

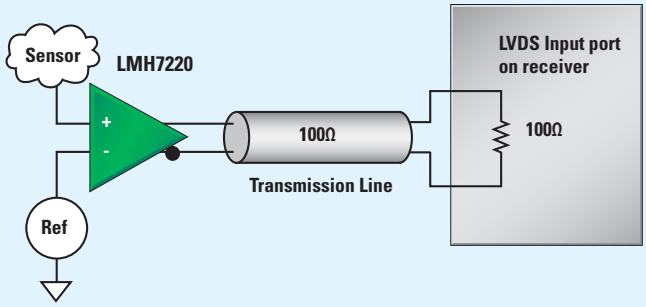

ADC12V170 and ADC14V155 Features

- 1.1 GHz full power bandwidth
- Dual 3.3V, 1.8V supply operation
- Dual data rate parallel LVDS outputs
- Pinout equivalent to ADC12C170 and ADC14155 with CMOS outputs
- Clock duty cycle stabilizer
- Internal sample-and-hold circuit
- Internal precision 1.0V reference
- Available in LLP-48 packaging (7 x 7 x 0.8mm)

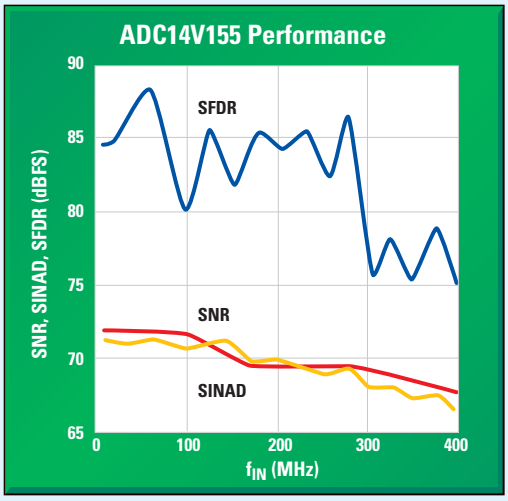

Product ID	Resolution (Bits)	Speed (MSPS)	Output Format	Power (mW)	SNR (dB)	SFDR (dB)	Packaging
ADC14V155	14	155	DDR Parallel LVDS	951	71.7	86.9	LLP-48
ADC12V170	12	170	DDR Parallel LVDS	781	67.2	85.8	LLP-48
ADC14155	14	155	CMOS	967	71.3	87.0	LLP-48
ADC12C170	12	170	CMOS	715	67.2	85.4	LLP-48

For FREE samples, datasheets, and more information, visit:
signalpath.national.com

Extending the Signal Path Over Data Transmission Lines

a.

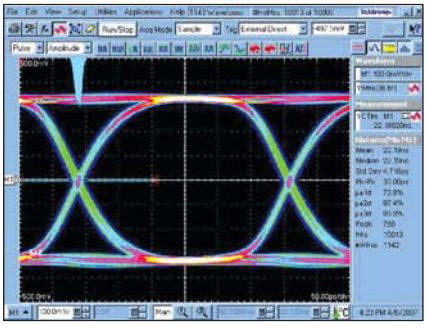

DS25BR120 Pre-Emphasis: OFF
3.125 Gbps with a PRBS7 pattern
Output load: 4" FR-4 100Ω differential trace

b.

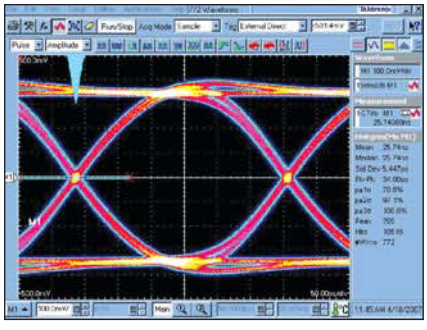

DS25BR120 Pre-Emphasis: ON (6 dB setting)
3.125 Gbps with a PRBS7 pattern
Output load: 20" FR-4 100Ω differential trace

Figure 3. 3.125 Gbps LVDS Output Without Pre-Emphasis After 4" FR-4 (a), and With PE After 20" FR-4 (b)

your transmission line at the primary frequency of interest, tweaking one level up and down to ensure optimal quality of the received waveform eye pattern. When considering signal conditioning, the primary frequency of interest is normally the data rate divided by 2 (e.g., for a 2.5 Gbps signal use 2.5÷2=1.25 GHz, for a 1 Gbps signal use 500 MHz). This is the frequency of the data pattern with the highest level of attenuation. A network analyzer is the easiest way to generate a loss graph like that in ***Figure 1***. If a network analyzer is unavailable, it is also possible to send a sine wave at the frequency of interest across your transmission line and use the attenuation value as your loss value.

Input EQ, output signal conditioning, or a combination of both can be used to compensate for lossy transmission lines. Each method has advantages and disadvantages and the choice of which one to use depends on the application and personal preference. Output PE has the advantage of being easily measurable. Its effect is visible at the receiver and can be easily monitored and adjusted. Receive EQ is often used in crosstalk-sensitive and low-power applications since it does not add extra energy to the transmission line. In very lossy applications where EQ or PE alone is not sufficient, both input and output signal conditioning may need to be used in combination, e.g. 9 dB of pre-emphasis with 9 dB of receive equalization.

Should I use LVDS or CML?

Discrete LVDS implementations can be used effectively to data rates in excess of 3 Gbps. Current Mode Logic (CML) is the I/O of choice at speeds of 3.5 Gbps (***Figure 4***) and beyond. LVDS generally has lower power and less EMI and is well-defined as an interface standard. CML is capable of higher speeds and typically has higher drive strength than LVDS. Either signaling technology can often be used at rates between 1 and 3.125 Gbps so a good translation strategy is important for optimal signal integrity.

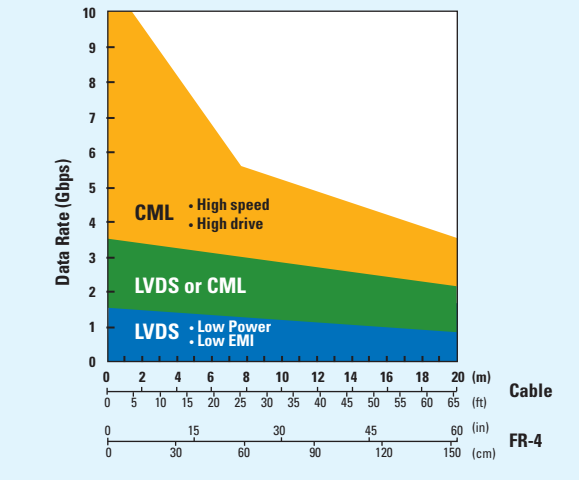

Figure 4: Typical LVDS and CML Applications

For successful level translation, the driver's differential output voltage (V_{OD}) and common mode voltage (V_{CM}) must fall within the receiver's input range. For LVDS the V_{OD} (as defined in the EIA/TIA-644A standard) is the voltage difference across the driver outputs with a 100Ω resistive load.

Industry's Lowest Jitter Video Cable Equalizer

DS16EV5110 Enables Superior Picture Quality for 1080p Deep-Color Video Applications and Extends HDMI and DVI Cable Reach Beyond 40 Meters

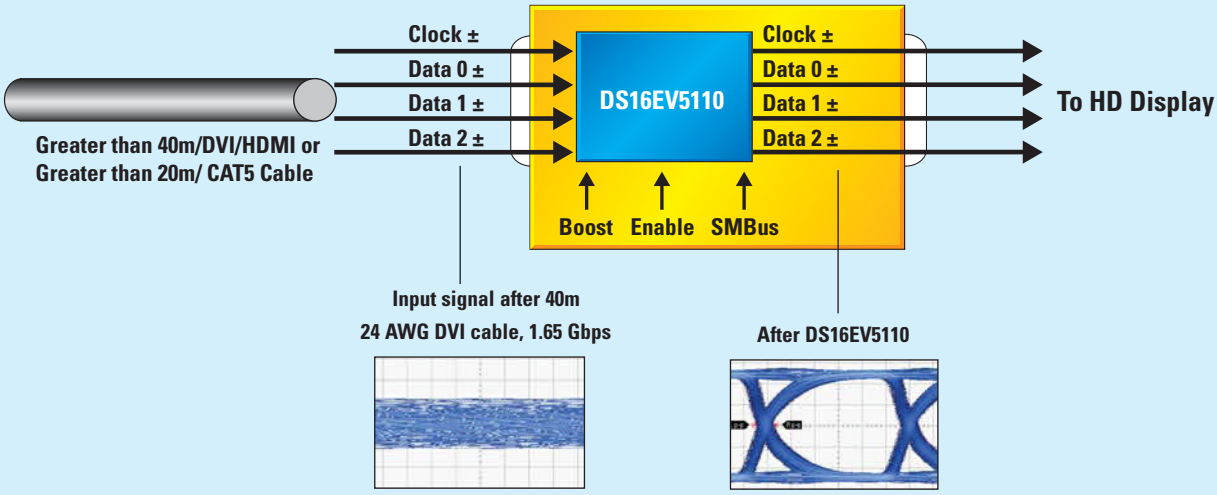

DS16EV5110 Features

- Significantly extends reach of HDMI, DVI, and CAT5 cables
- Settable boost for optimal equalization
- Low total jitter: 0.13 UI on 30m 24 AWG DVI cable
- Low power: 475 mW
- Clock channel LOS and output enable for power savings
- Serial Management Bus (SMBus) provides extra design flexibility
- Available in tiny 7 mm x 7 mm LLP-48 packaging

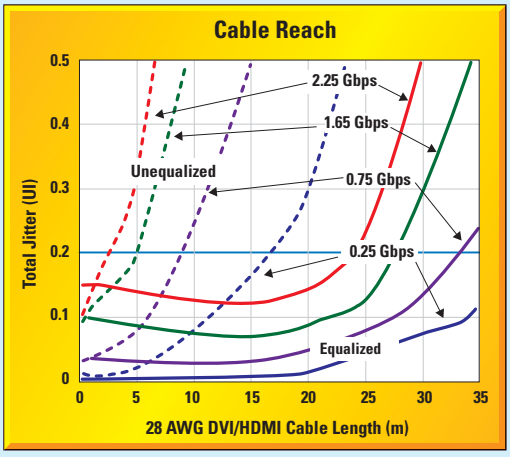

Ideal for use in HDMI, DVI, and CAT 5 cable extenders, switchers, digital routers and switches, projectors, and high-definition displays

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

lvds.national.com

Analog by Design Show

"Designing Long-Reach Applications with DVI HDMI and PCI Express Cable Standards"
www.national.com/nationaltv

Extending the Signal Path Over Data Transmission Lines

Other differential standards specify the output voltage amplitude as a peak to peak number. As an example, a 400 mV LVDS V_{OD} is exactly equal in amplitude to an 800 mV peak-to-peak CML output, but located at a different offset voltage with respect to GND.

LVDS receiver devices offer the greatest flexibility of any differential technology because their wide input common mode range easily accepts the signal swings of 1.2V/1.5V/2.5V CML and LVPECL differential I/O. This allows for a direct DC connection between most differential outputs and LVDS inputs, minimizing the added PCB real estate and cost of multiple AC-coupling capacitors.

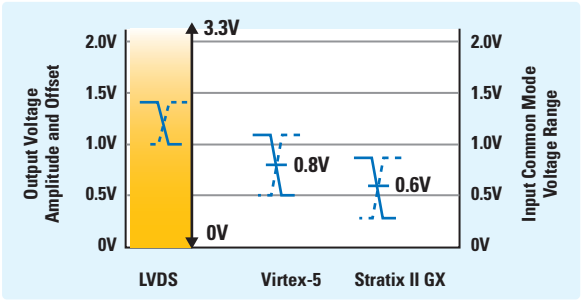

Figure 5. Standard and FPGA Differential I/O Signaling Levels

The LVDS transmitter output V_{OD} is specified to be 250 mV minimum (500 mV peak-to-peak) for high performance with low EMI and low-power consumption. Adding output signal conditioning enables LVDS to drive extended lengths of cable or large backplanes. The LVDS output rides on a 1.2V common mode voltage developed from an internal bandgap reference and can be DC coupled to many LVPECL inputs. CML inputs with a limited common mode range, however, require an AC-coupled interface because the LVDS output voltage swing does not meet the minimum common mode requirements for most CML inputs.

Many LVDS and CML buffer chips offer pre-emphasis and/or receive equalization to boost signals coming from ADCs, DACs, FPGAs, and DSPs, and some have multiplexing functions for switching/redundancy applications. A low-power solution for redundancy, multiplexing, and signal distribution can be achieved with a combination of high-performance LVDS crosspoint switches and high-speed FPGA CML I/O. Programming CML outputs for 600 mV to 800 mV will reduce transmit power expended by the FPGA integrated SerDes, lower the overall EMI signature of the interface, and provide the optimum signal for LVDS inputs with EQ.

The latest generation FPGAs with CML I/O have common-mode output voltages lower than the nominal 1.2V stated in the LVDS standard (***Figure 5***). The extended input common mode range allows the high-speed FPGA I/O to interoperate with LVDS devices.

Driving Over 20 Meters

LVDS and CML signal conditioners are capable of extending signal paths up to a maximum of 20m at multi-Gbps speeds. For longer distances, Ethernet can be used but requires additional protocol and timing overhead to fit the signal path raw data into Ethernet packets. Alternatively, a cable driver and adaptive equalizer can be used to extend serialized data up to 100s of meters. The DS15BA101/EA101 is a 1.5 Gbps cable extension chipset for 100Ω twisted pair and 50Ω coaxial cables up to hundreds of meters. Unlike many other signal conditioners with equalization, this chipset automatically compensates for various cable lengths and types. Adaptive equalization is critical to the success of long distance cable driving. The adaptive nature of the equalizer ensures minimal added signal noise and jitter caused during the amplification of high-frequency energy. An evaluation reference design, number DriveCable02EVK, is available for quick evaluation and design implementation.

Conclusion

Today many signal paths contain separate signal acquisition and processing modules. As sampling speeds increase, it becomes harder to transfer the signal-path data between modules due to the lossy effects of the transmission medium. In fact, even if the transmission distance is constant, losses and therefore jitter increases as signaling speed increases. Luckily, it is easy to estimate and overcome these loss effects through the use of signal conditioners with pre-emphasis and equalization. ■

Design Tools

WEBENCH® Signal Path Designer® Tools

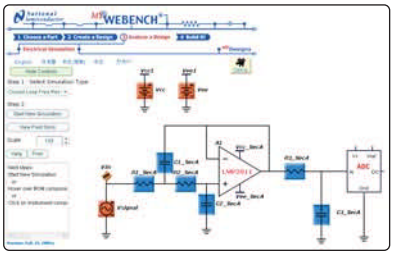

Design, simulate, and optimize amplifier circuits in this FREE online design and prototyping environment allowing you to:

- Synthesize an anti-alias filter
- Select the best amplifier/ADC combo for your system specs
- Make trade-offs based on SNR, SFDR, supply voltage
- Simulate real-world operating conditions using SPICE
- Receive samples in 24 hours

webench.national.com

WaveVision 4.1 Evaluation Board

Test and evaluate A/D converters with National's easy-to use WaveVision 4.1 evaluation board. Each evaluation board comes complete with USB cable and support software.

Features and benefits:

- Plug-n-play ADC evaluation board
- USB 2.0 interface to PC
- PC-based data capture
- Easy data capture and evaluation
- Highlighted harmonic and SFDR frequencies
- Easy waveform examination
- Produces and displays FFT plots
- Dynamic performance parameter readout with FFT
- Produces and displays histograms

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

Mailing address:
PO Box 58090
Santa Clara, CA 95052

Visit our website at:
signalpath.national.com

For more information, send email to:
new.feedback@nsc.com

Don't miss a single issue!

Subscribe now to receive email alerts when new issues of Signal Path Designer® are available:

signalpath.national.com/designer

Also, be sure to check out our Power Designer! View online today at:

power.national.com/designer

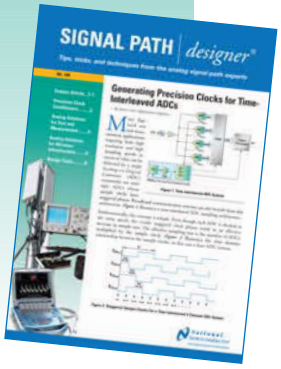

National Semiconductor
The Sight & Sound of Information

EDN

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

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

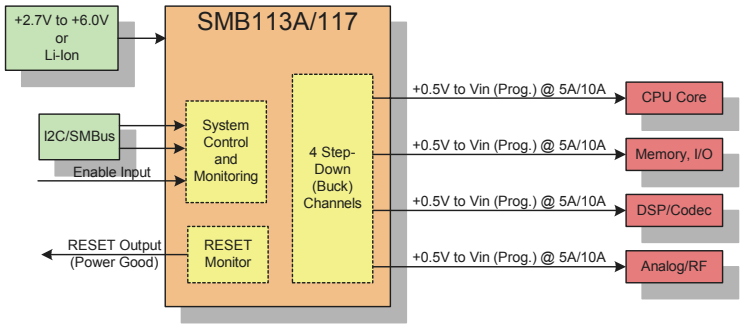

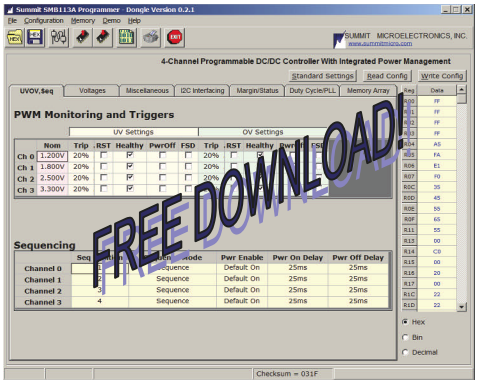

Develop application with convenient programming tool and upload HEX file
Summit ships pre-programmed "custom" units in high volume

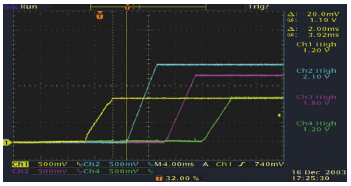

Sequencing

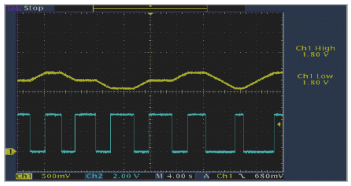

Dynamic Output Voltage

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

Mezzanine card features dual cores

With low-power-embedded-telecom, network-centric-warfare, medical-imaging, surveillance, and industrial-automation applications in mind, Motorola recently announced the PrAMC-7210, an AMC (advanced mezzanine card) that the company based on the Intel (www.intel.com) Core 2 Duo processor core and Intel 3100 server-class chip set. With as much as 4 Gbytes of DDR SDRAM running at 400 MHz and an onboard flash disk, the PrAMC-7210 targets high-performance operation and reliable-storage applications. One eight-lane PCI Express link supports autonegotiation to lower lane widths and reduce port bifurcation. The module also includes two GbE (gigabit-Ethernet) links for offboard communications and two SATA interfaces for local hard-disk storage.

The PrAMC-7210 is available with either the MontaVista (www.mvista.com) Carrier Grade Linux or the Wind River (www.windriver.com) Platform for Network Equipment, Linux Edition operating software. The first PrAMC-7210 modules will debut in the third quarter of 2007 with prices of less than $2000 for volume orders.

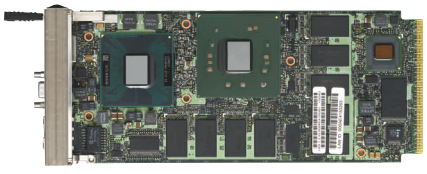

Motorola's latest AMC offers dual-core performance for AdvancedTCA (Telecommunications Computing Architecture)- or MicroTCA-communications applications.

—by Warren Webb

▷**Motorola Inc**, www.motorola.com.

FEEDBACK LOOP

"This is the most brilliant idea I have seen recently. It is getting rare to see ideas that use fewer components and a lot of intellectual ingenuity."

—V Ramprakash, in *EDN*'s Feedback Loop, at www.edn.com/article/CA6451248. Read that Design Idea and add your comments.

Oscillator offers one part/trillion drift/day

Jackson Labs has announced the availability of Fury, a 10-MHz GPSDO (global-positioning-system-disciplined oscillator) that requires no calibration, delivers greater stability than typical rubidium-clock references, and meets the stability requirements of a Stratum 1 frequency standard.

The unit offers less than 10^{-12}, or one part per trillion, frequency drift per day; 4.5W power consumption; a phase-noise floor of −155 dBc/Hz, and extremely low spurs and jitter. The device generates a one-pulse-per-second output, phase-synchronized to UTC (Universal Time Coordinated) and adjustable to UTC in 1-nsec steps. You can control the unit through an RS-232 port using industry-standard SCPI commands (standard commands for programmable instruments). Optional GPSCon software provides SNTP (Simple Network Time Protocol) timing.

The manufacturer's co-founder and president, Said Jackson, says that Fury has generated excitement not just because of its high performance, but also because of its competitive price: "The devices' exceptional performance, low cost, and calibration-free operation have enabled our customers to use the PCB [printed-circuit-board] version as a frequency reference within OEM products and the fully enclosed modular version to synchronize groups of laboratory instruments." The company is currently shipping the units. The PCB version sells for $960 (one) and $750 (20 to 99), with lower prices for larger quantities.**—by Dan Strassberg**

▷**Jackson Labs**, www.jackson-labs.com.

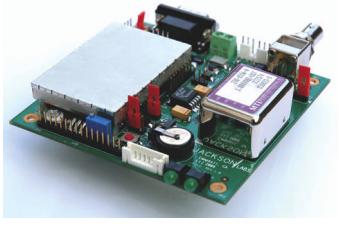

The Fury GPS-disciplined, 10-MHz oscillator provides one part-per-trillion-per-day drift. This PCB version targets OEM applications, but you can also obtain the unit fully enclosed for laboratory use.

Sierra releases IC-place-and-route tools for 45-nm era

Sierra Design Automation has released a new version of its Olympus physical-design suite for 45-nm-IC design with three new tools for layering and buffer swapping, multicorner-CTS (clock-tree synthesis), and shape-based DRC (design-rules-checking)-accurate routing.

At fine process geometries, manufacturing closure is the problem, and addressing it requires a DFM (design-for-manufacturing)-driven place-and-route system, according to Sudhakar Jilla, director of marketing at Sierra. The company built the Olympus suite for the new challenges of the 65- and 45-nm-design era.

At 2003's DAC (Design Automation Conference), the company introduced a much speedier router than competing tools. At 2004's DAC, the company introduced multimode design and analysis for power savings; 2005's DAC brought Sierra's lithography-driven routing engine for 65-nm design and close coupling with Mentor's (www.mentor.com) Calibre physical-verification lineup. At last month's DAC in San Diego, Sierra introduced the Olympus SOC (system on chip) for the challenges of 45-nm design.

Jilla says that the problems now emerging for 45-nm design are resistance spikes and variability. "The resistance-per-unit length has doubled from 65 to 45 nm, while capacitance has stayed fairly stable," he says. "As a result, the RC-unit length has significantly jumped from 90 to 65 to 45 nm. On top of that, there is now a large variation in resistance across a die ... which makes the clock tree and the performance fluctuate. Thus, for 45 nm, you need technologies to compensate for these problems." At the 45-nm node, resistance in vias is two to three times greater than resistance in vias at the 65-nm node. This increase means that designers must try to limit the number of vias they use.

> Sierra has added correct-by-construction features to the router.

Traditionally, layout groups have dealt with increased resistance and capacitance by adding buffering—breaking down long wires into a series of smaller wires. But that technique doesn't work well at smaller process nodes. "Now that you have resistance becoming such a big factor in determining a design's performance, it is no longer sufficient to simply break up a wire into little pieces," says Jilla. "You now have to think about what layer that wire is going onto." Higher layers in the layer stack tend to have lower resistance than middle layers, but you can't put an entire design into just the upper layers. Sierra has also introduced the FalconGR engine, which dynamically makes trade-offs between layer assignments and buffering.

To deal with resistance variability, Sierra has developed a multicorner-CTS tool that works with other tools in the flow. "Resistance fluctuates so widely across a die that you really can't pin down your performance unless you measure all the corner cases, [accounting for] resistance fluctuation," says Jilla.

Sierra's multicorner-CTS tool takes into account intercorner and intracorner skew. A sample design with a conventional one-corner CTS shows a 145-psec skew, but, with the multicorner-CTS engine, a design with nine corners has only 87-psec skew—a 40% decrease. The tool reduces the hold-buffer area by 54% and total area by 18%. And, if you build a clock tree with more corners, each new corner you add further decreases the amount of fluctuation. If you add enough corners, the skew becomes minimal to nonexistent.

Sierra has also tuned up its routing engine to better handle the large rule decks for the 45-nm process and has added support for model-based DRC. With the advent of finer process geometries, the number of DRC rules increases. In addition, today's routers now must account for extensive lists of DFM rules. Jilla notes that routing is traditionally a four-step process: global routing, assigning tracks, detailed routing, and postprocessing. Postprocessing has become the longest step in the 65- and 45-nm-design era and often requires users to run several iterations between DRC and the physical-implementation flow to ensure that their designs comply. To minimize iteration and shorten the postprocessing step, Sierra has added correct-by-construction features to the global-routing and assigning-track steps of the router. The company has also added a shape-based DRC engine to all its router engines to ensure that the tool creates DRC-compliant layouts.

The multithreaded shape-based router has an average run speed of 260,000 polygons/sec and routes 5 million nets—123.5 million polygons—in a 90-nm, metal-layer process with a DRC check that takes only 25 minutes. It routes 20,000 nets, or 5.59 million polygons, in a 65-nm, metal-layer process with a DRC check that takes 2 minutes. The company provides no runtime numbers for the 45-nm process. The base configuration of Olympus SOC sells for $1.6 million per seat.

—by Michael Santarini

▷**Sierra Design Automation**, www.sierra-da.com.

DILBERT By Scott Adams

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

Applications such as robotics, transportation, pipelines, MIL/COTS, medical, security, machine control, and industrial automation that must work in harsh, demanding environments need WinSystems' EBC-855.

This x86, PC-compatible SBC supports Linux and Windows® XP embedded and other popular RTOS along with popular video and wired and wireless network standards.

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

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

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

Ask about our 30-day product evaluation

WinSystems®

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

MIPS introduces 1-GHz 74K processor core

At May's Microprocessor Forum in San Jose, CA, MIPS extended its family of 32-bit cores with the 74K. The company's highest-performance, single-threaded, fully synthesizable core, the 74K achieves an operating frequency of greater than 1 GHz in a 65-nm general-purpose process. The 74K's updated microarchitecture is compatible with the software and system interfaces of the 24K, 24KE, and 34K processors. The 74K's 17-stage, asymmetric, limited, dual-issue (ALU and address-generation) pipeline supports out-of-order instruction dispatch and completion (eight-instruction-wide window per pipeline) that can improve the performance and efficiency of legacy binary code without recompiling. The branch-prediction logic includes three 256-entry branch-history tables and an eight-entry return-prediction stack.

The 74K core implements the DSP ASE (application-specific extension) Revision 2, which is a superset of the DSP ASE Revision 1 that the 24KE and 34K cores implement. The DSP ASE Revision 2 includes 27 new instructions that enable automatic compiler vectorization and that optimize video and image processing, VOIP (voice-over-Internet Protocol), and Viterbi algorithms. The 74K supports L2 cache with the MIPS SOC-it L2-cache controller.

The 74K core family comprises the 74Kc integer core and the 74Kf integer core with high-performance, floating-point support that fully complies with the IEEE 754 specification. Both 74K cores support CorExtend, which allows designers to add their own proprietary instructions and tightly coupled hardware. The 74K core family is available now for general licensing; the core deliverables include RTL, application notes, simulation testbenches, and EDA-design flows for Magma (www.magma-da.com), Cadence (www.cadence.com), and Synopsys (www.synopsys.com). Software-development and debugging tools from MIPS include the Eclipse-based IDE (integrated development environment) for the MIPS SDE (software-development environment), a full GNU tool chain, and a free version of SDE-Lite. Simulation support includes the full processor simulator MIPSsim and a bus-functional model. Hardware-development tools include the Malta FPGA-based hardware-development platform.**—by Robert Cravotta**

> Software-development and debugging tools include the Eclipse-based IDE for the MIPS SDE, a full GNU tool chain, and a free version of SDE-Lite.

▷**MIPS**, www.mips.com.

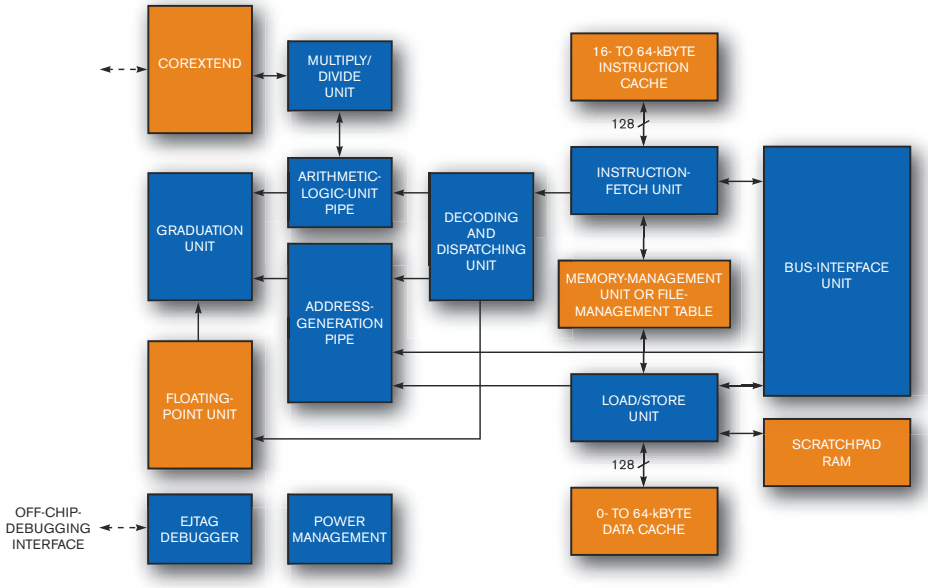

The 74K core family has a 17-stage, asymmetric, dual-issue pipeline.

COM EXPRESS ADOPTS MULTICORE ARCHITECTURE

07.19.07

Targeting embedded-system applications ranging from point-of-sale kiosks and mobile data-acquisition systems to medical 3-D-image scanners, automation systems, and other multidisplay applications, Kontron recently announced the ETXexpress-MC COM (Computer-On-Module) Express board computer.

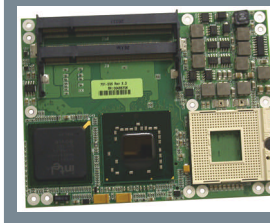

The ETXexpress-MC is the first embedded module to integrate the Intel Core 2 Duo processor and the GM965 chip set into the COM Express form factor.

The module features the Intel (www.intel.com) Core 2 Duo processor, the X3100 graphics-media accelerator, and support for as much as 4 Gbytes of dual-channel memory through two DDR2 SO-DIMM sockets on the top of the module. Standard features include five PCI Express lanes, three serial-ATA ports, and eight USB 2.0 ports, along with GbE (gigabit Ethernet). Sample units of the Kontron ETXexpress will be available in the third quarter, with prices starting at $830 (small quantities).–by Warren Webb

▶**Kontron America, www.kontron.com.**

Design with confidence.

Introducing the new Stratix III family.

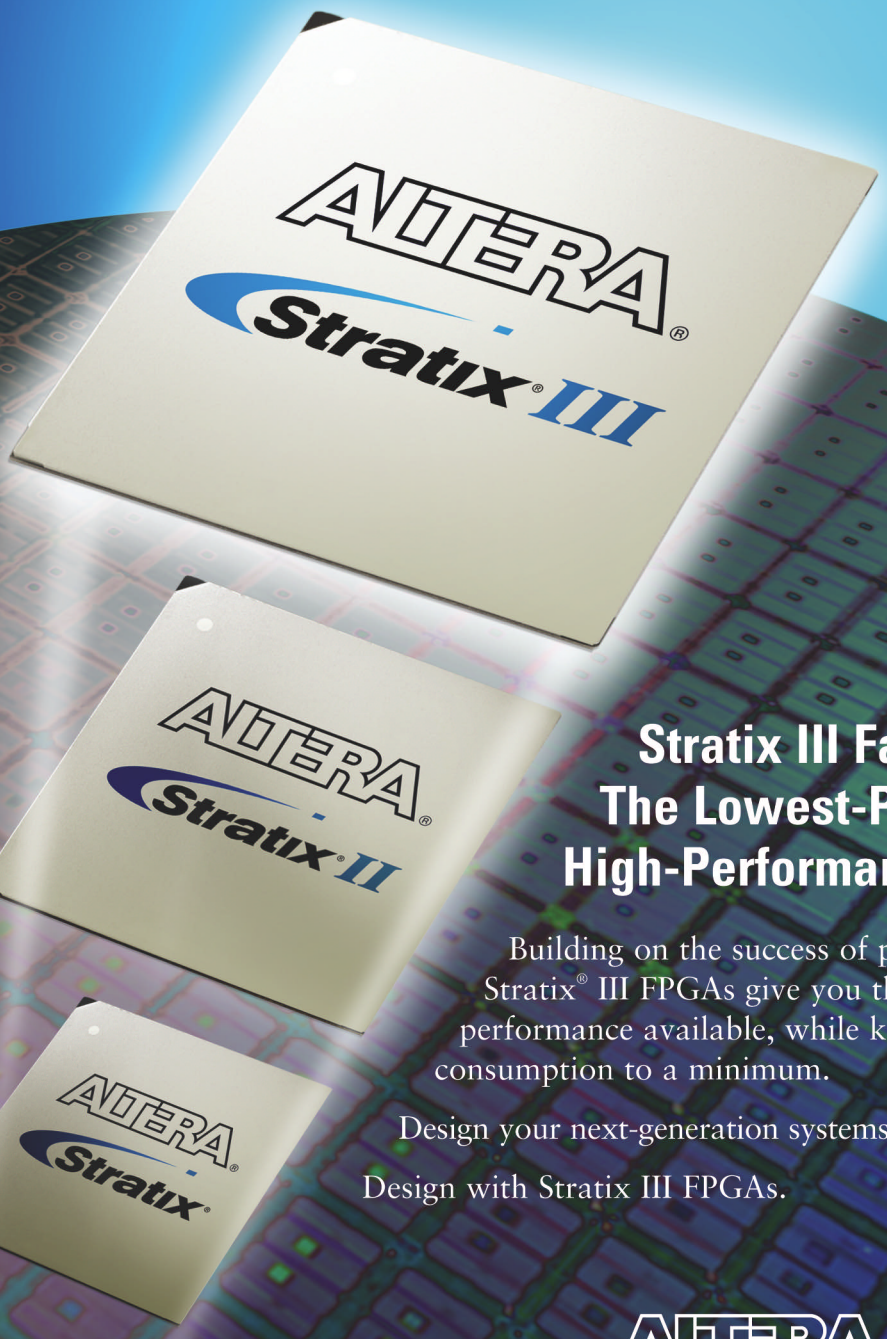

Stratix III Family: The Lowest-Power High-Performance FPGAs

Building on the success of prior generations, Stratix® III FPGAs give you the highest performance available, while keeping power consumption to a minimum.

Design your next-generation systems with confidence.

Design with Stratix III FPGAs.

Supported by

QUARTUS II

ALTERA®

The Programmable Solutions Company®

www.altera.com/confidence

Small, 1-GHz-bandwidth, two-channel PCI digitizer offers low power consumption

The low-power, two-channel Acqiris DP1400 digitizer provides guaranteed 1-GHz-minimum −3-dB bandwidth on each channel and can capture 2G 8-bit samples/sec in the single-channel mode. You may notice that the panel bears both the Acqiris and Agilent names because Acqiris has been part of Agilent since 2006. Then, note that the DP1400 is not a CompactPCI board; the short format can fool you. The board plugs into the standard 66-MHz, 32-bit PCI bus, so it works in standard desktop PCs. Most other boards in this class are full-length, even though dense PC packaging is now making full-length slots increasingly rare. Those large boards also tend to be power-hungry; some use nearly 60W, and several draw 30W. The DP1400 uses less than 15W.

The DP1400 offers an unusual optional multibuffer-simultaneous-acquisition-and-readout mode. This option lets you divide the board's internal 256-kbyte (128-kbyte/channel) memory into as few as three or as many as 960 segments, make conversions with a maximum dead time of 350 nsec, and rapidly send the data to the host PC while acquisition continues. This mode does not prevent the board from acquiring data faster than the PCI bus can remove it from the buffers. However, should all of the buffers momentarily become full, onboard logic temporarily suspends conversions whose results would otherwise overwrite previously acquired data.

For those interested in implementing random repetitive sampling, a mode normally associated with digital oscilloscopes but not with data-acquisition boards, the DP1400 contains a TDC (time-to-digital converter), part of the hardware you need to sample repetitive waveforms at an equivalent rate of 40G samples/sec.

The unit offers seven full-scale input ranges from 50 mV to 5V in a 1-2-5 sequence with offsets adjustable to ±2V on ranges to 500 mV and ±5V on less sensitive ranges. Each ADC samples at rates from 100 to 1G samples/sec in a 1-2-5 sequence, and you can interleave the two converters to double these rates. Besides the full-bandwidth mode, the board provides internal 20-, 200-, and 700-MHz lowpass filters for its analog inputs. Timebase error is less than ±2 ppm, and trigger jitter is less than 2 psec rms on a 10-µsec record. The unit offers a variety of trigger modes and can pretrigger to as much as the full record duration. The board also incorporates a second-generation autosynchronous-bus system that can distribute trigger and clock signals to multiple modules. Prices start at $9490; the simultaneous-acquisition-and-readout option adds $1000.**—by Dan Strassberg**

▷**Agilent Technologies**, www.agilent.com.

Despite its short PCI form factor and less-than-15W power consumption, the DP1400 digitizes two channels and offers 1-GHz guaranteed −3-dB bandwidth. In the single-channel mode, the maximum conversion rate is 2G samples/sec, which is adequate to prevent aliasing of 1-GHz signals.

QUAD-CORE COMPUTER FEATURES FIVE PCIe LANES

Although many embedded-system designers have for years relied on the standard desktop computer and its PCI bus as project starting points, rising performance expectations are changing the landscape. Parallel-bus-bandwidth limitations and high-performance-processor technologies have forced system designers to adopt the latest serial-board-to-board-communications strategies. The PICMG (PCI Industrial Computer Manufacturers Group) 1.3 specification solves the bandwidth problem by replacing the backplane-parallel-bus interfaces with high-speed PCIe (PCI Express) serial links.

To meet these new bandwidth requirements, Adlink Technology recently announced the PICMG 1.3-compliant NuPro-965 system-host board, which supports the Intel (www.intel.com) Core Quad and Core 2 Duo family of processors. Four onboard DIMM slots support as much as 8 Gbytes of dual-channel DDR2-800 RAM with a peak transfer rate of 12.8 Gbytes/sec. With a compatible backplane, the board supports 10-Gbps data transfers to and from external I/O cards and memory. The NuPro-965 integrates a 3-D-graphics engine and a 16-lane PCIe interface for additional high-end-graphics options. The device also has four additional one-lane PCIe interfaces for high-bandwidth-I/O applications, two Ethernet connectors for redundancy, support for four 3-Gbps SATA storage devices, four USB 2.0 ports, and an UltraATA IDE interface. Prices for the NuPro-965 start at $460.

—by Warren Webb

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

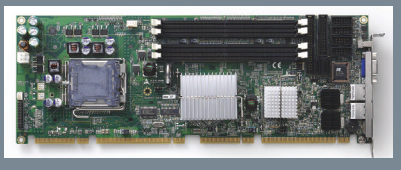

The PICMG 1.3-compliant, full-size, quad-core NuPro-965 single-board computer features five PCI Express lanes and 10-Gbps data transfers.

07.19.07

Reset your thinking about clocks.
In precision timing, analog is everywhere.

AD9516 14-Channel Clock Generator with Integrated VCO

- Two reference inputs
- Integer-N PLL frequency synthesizer
- Selectable VCO frequencies up to 2.95 GHz
- 14 outputs with total jitter < 0.5 ps rms
- Programmable dividers, up to 1024
- Four adjustable delay lines
- Selectable LVPECL, LVDS/CMOS logic
- 64-lead LFCSP
- 5k price: $11.47

AD9549 Network Clock Generator/Synchronizer

- Two flexible reference inputs
- Automatic or manual holdover/switchover modes
- Digital PLL core
- Programmable loop bandwidth < 1 Hz
- Programmable output frequencies from 8 kHz to 750 MHz
- Differential HSTL clock output
- 64-lead LFCSP
- 5k price: $11.39

Increase performance and reliability while cutting your system clock costs by 50%

With Analog Devices' newest timing innovations, you can develop efficient and effective clocking designs.

The AD9516 combines low phase noise clock generation with 14-channel clock distribution at jitter levels below 500 fs. Integration of all critical timing functions into a single IC not only reduces costs and board space, but also diminishes reliability risks associated with large discrete oscillators. By implementing the AD9516, designers save time and reduce problems related to common debug issues when dealing with multiple components. Five versions with on-chip VCOs tuned from 1.45 GHz to 2.95 GHz make them ideal for communications and instrumentation designs.

The AD9549 is a two input clock generator with extended holdover functionality. It has a programmable loop filter to clean jitter from input clocks and features flexible output frequency options. In addition, automatic redundant reference clock switchover with user-selectable rate of phase adjustment helps ensure system stability and maximizes network uptime.

Get more performance and greater value for your next clocking design. For more information on all our timing ICs, please call 1-800-AnalogD or visit *www.analog.com/clock-timing-ad*.

analog is everywhere.™

www.analog.com/clock-timing-ad

ANALOG DEVICES

Cadence offers low-power-methodology kit

Building on its thrust into the low-power-design area and leveraging its CPF (Common Power Format), Cadence Design Systems is now announcing a low-power-methodology kit. Until recently, low-power design has been the responsibility mainly of specialists and low-power-design gurus. However, starting at the 130-nm-process node, the transistors that fabs employed to increase the performance of their silicon processes came with a nasty side effect: static leakage. That is, high-performance transistors leak power even when they are not in use, which wastes power in wired-system applications and more quickly drains battery life in wireless-system applications. Therefore, most designers working with new processes must employ low-power-design techniques. Cadence has for the last couple of years been working on low-power approaches, says Neil Hand, director of vertical-solutions marketing at the company.

Two years ago, the company started developing CPF with partner companies and ensuring that most of its digital-tool

> The company is trying to increase adoption of its low-power format and help mainstream designers deal with low power.

flow could use CPF to help digital-IC-design customers better create low-power designs. This year, Cadence introduced some of the fruits of that labor when it officially released its Low Power Platform, in which its Encounter RTL-to-GDSII (register-transfer-level-to-Graphic Design System II) flow and its Incisive verification tools all now support low-power design through the CPF. Now, the company is trying to increase adoption of its low-power format and help mainstream designers deal with low power by offering the Low Power Methodology kit—essentially, a bundling of tools, IP (intellectual property), reference designs, and training. Cadence also offers similar kits for AMS (analog-mixed-signal), verification, and RF-SIP (system-in-package) design.

Cadence has developed a segment-representative reference design to illustrate the blocks and segments in a design that require low-power optimization. The company has built a modularized design-flow structure that illustrates the tools required during each phase of the design, including logic synthesis, functional verification, design for test and ATPG (automatic-test-pattern generation), physical design, formal-implementation verification, and power-grid sign-off. The Low Power Methodology kit also includes a list of IP and all the tools that Cadence offers for a low-power flow.

The tool flow in the kit is modular, so customers can incorporate non-Cadence tools if they wish. Using a pure-Cadence flow yields greater benefits in interoperability, however. The starting price for the kit is $30,000.**—by Michael Santarini**

▷**Cadence Design Systems Inc**, www.cadence.com.

VIDEO DSPs DOUBLE PERFORMANCE, REDUCE SYSTEM-BOM COST

Texas Instruments' new TMS320DM647 and 648 DaVinci technology-based DSPs deliver twice the processing performance of the previous-generation DSP and reduce system-BOM (bill-of-materials) cost for multichannel-video-security and infrastructure applications. The increase in processing performance is the result of using the C64X+ core, which improves cycle performance over the C64x core by 20%; an integrated VICP (video-and-imaging coprocessor); EDMA (enhanced-direct-memory-access) 3.0 support that provides a 400% increase in bandwidth; and the ability to perform twice as many 16-bit MMAC operations per cycle as predecessor devices, such as the DM642. These devices are 100% code-compatible with the C64x core.

To reduce system-BOM costs, the DM648 includes two extra video ports to remove the need for a video multiplexer and an integrated SGMII (serial-gigabit-media-independent-interface) port with a switch that supports as many as four devices without the need for a PCI bridge. In addition, the VICP eliminates the need for an FPGA to perform video preprocessing, such as color-space conversion, resizing, and deinterlacing. The instruction-based VICP for front-end preprocessing works in parallel with the main DSP core. The VICP functional library will be available with the release of the beta DM648 DVDP (digital-video-development platform) in the third quarter of this year. The DVDP also includes TI's TVP5154 video decoder.

The new processors support dual streaming that allows both a monitoring and a recording stream for every video channel. The DM648, with 512 kbytes of L2 cache and two switched Ethernet media-access controllers, targets security-digital-video-recording and network-video-server applications. The DM647, with 256 kbytes of L2 cache and a single Ethernet media-access-controller port, targets machine-vision and high-performance-imaging applications. The devices include five 16-bit, dual-channel video ports; a McASP (multichannel-audio-serial port); an I^2C; an SPI; and two 64-bit timers. The 720-MHz versions of the $39.95 (10,000) DM647 and the $49.95 (10,000) DM648 are available for sampling now, and production quantities will become available in the fourth quarter of this year.

TI and members of its DSP Third Party Network offer support for video analytics, such as object recognition and tracking. They also provide support for customizing codecs to implement security-specific modifications, such as adjusting bit rate and quality or providing codec information for security identification. Ittiam Systems (www.ittiam.com) offers codecs and a video-security-media system with multichannel, high-resolution, and low-latency features.

—by Robert Cravotta

▶**Texas Instruments, www.ti.com.**

07.19.07

HIGH PERFORMANCE, CHANNEL COUNT, SPEED, AND ACCURACY

fast ■ flexible ■ multi-function measurement solutions

INTEGRATED MULTIMETER/DATA ACQUISITION SYSTEMS

Do you need:

High precision measurements with high throughput?
Multiple channels with signal conditioning and low cost per channel?

We have exactly what you need.

Request a FREE *Simplified Test Toolkit* CD at www.keithley.com/at/305.

KEITHLEY

A GREATER MEASURE OF CONFIDENCE

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

RESEARCH UPDATE

BY RON WILSON

Dense nanotubes allow efficient interconnect

Because of the physics of electron motion in their structures, carbon nanotubes look like the perfect conductors for replacing copper-damascene interconnections on ICs. The tubes bring with them several minor problems, however. One is that, after you grow the tubes, they may not turn out to be conductors. Another issue is that the only reliable way to grow the tubes is in a forest: You must anchor them at the bottom and vertically grow them in a soup of available carbon atoms. That limitation is not helpful for IC interconnect, in which tiny bundles of tubes would have to follow routing channels from one contact to another. Yet another issue is that a single tube is of little use; you need a bundle of them to carry useful current.

At least the last problem is yielding to solutions—pun intended. Researcher James Jiam-Qiang Lu, associate professor of physics and electrical engineering at Rensselaer Polytechnic Institute (Troy, NY), led a small team investigating the bundling of the nanotubes. He presented the work at the International Interconnect Technology Conference, which took place last month in Burlingame, CA. The tubes normally grow in loose bunches, rather like hay. Lu found that, when his team immersed the bunch of nanotubes in isopropyl alcohol and then let it dry, capillary action pulled the tubes together into a much tighter bundle, typically 25 times denser than the original bunch. This density is still insufficient to compete with high-quality metallic copper for current density, but it is getting closer. And the problems of fabricating interconnect lines from copper are growing with each new process node.

▷**Rensselaer Polytechnic Institute**, www.rensselaer.edu.

▷**International Interconnect Technology Conference**, www.his.com/~iitc.

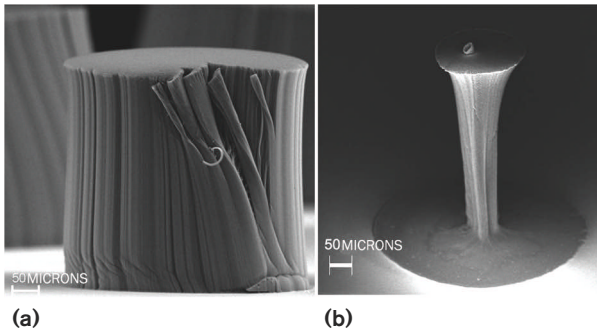

(a) (b)

Researchers at Rensselaer developed a process that makes carbon-nanotube bundles denser, thus increasing their conductivity. The researchers pulled a bunch of tubes (a) into a tight bundle 25 times denser than the original bunch (b).

CIRCUIT-DESIGN TECHNIQUE COULD MAKE CARBON NANOTUBES USEFUL

To be of any use to designers, carbon nanotubes must align with circuit patterns, and an elegant approach to that necessity may now have emerged. Subhasish Mitra, a Stanford University (Palo Alto, CA) assistant professor of electrical engineering and computer science, presented a paper at the Design Automation Conference in San Diego last month that suggests a solution to the problem. Mitra and his team devised a way to form functioning logic circuits from a random tangle of nanotubes on the surface of a die—not by sorting and reconnecting them but by simply trimming away undesirable ones. Mitra's work assumes that the tubes form as semiconductors with transistor functions, not as conductors. If you have a random tangle of transistors, you can make a circuit of it by saving the transistors that connect some points and eliminating the transistors that go elsewhere.

Toward this end, Mitra's team superimposed a grid over the layout of standard logic circuits. Then, for each type of circuit, team members identified the cells through which semiconducting nanotubes should pass if they are to be useful and the cells through which they should not pass. Using simulations, Mitra's team showed that etching away tubes in the undesirable cells could produce working logic circuits, even though the researchers did not know the exact arrangement of nanotubes in the grid. They produced a program that can automatically generate the pattern of good and bad cell locations in the grid. Problems still exist, of course. There's no way of knowing whether a given nanotube will be a transistor, a resistor, or a "pretty-good wire." And there's no guarantee that a sufficient number of semiconducting nanotubes will connect the areas that must connect in a circuit.

▷**Stanford University**, www.stanford.edu.

▷**Design Automation Conference**, www.dac.com/44th.

THERMOACOUSTICS SPEAKS UP AT ASA CONVENTION

Thermoacoustic devices convert heat directly into intense sound waves with relatively high efficiency–reportedly 40% of Carnot-cycle efficiency. (A Carnot cycle employs the principles of the Carnot Theorem, which states that a cycle continuously operating between a low temperature and a high temperature can be no more efficient than a reversible cycle operating between the same temperatures.) Physicists can use the thermoacoustic effect simply to cool a system by converting its wasted heat to sound, which they can then convert to electrical energy by focusing the sound on a piezoelectric device.

A number of papers at last month's Acoustical Society of America Conference in Salt Lake City suggested progress in all of these areas. Papers described the construction and optimization of the thermal-sound generators and the energy-capture devices that went with them.

▶Acoustical Society of America, http://asa.aip.org/saltlakecity/Fridayam.pdf under session 5aPA.

07.19.07

From day one, the new Tyco Electronics brings many advantages to our customers, investors, employees and communities. As a $12.8 billion company, with a 60-year head start, we've been at the forefront of the explosive growth in electronics worldwide. With nearly 100,000 employees in more than 45 countries, and a peerless reputation in engineering, our approximately 500,000 products include electronic components, network solutions and wireless systems. Tyco Electronics is committed to the technologies that bring advantages to everyone's life every day.

tycoelectronics.com

Tyco Electronics

Our commitment. Your advantage.

BY HOWARD JOHNSON, PhD

Uncertainty principle

The uncertainty principle that German physicist Werner Heisenberg made famous in 1927 asserts the improbability of knowing with complete precision the position and momentum of any subatomic particle. According to Heisenberg, you may in some cases know the position quite well, or, in other cases, know the momentum—that is, the velocity—quite well, but there are limits to the accuracy of simultaneously knowing these two independent properties. Heisenberg's principle slammed shut the door on older, particle-oriented, deterministic interpretations of physics, making way for a new generation of quantum-mechanical understanding.

In the language of quantum-mechanical-wave functions, author Enrico Persico puts the principle like this: "The smaller the extent of a wave packet in space, the wider must be the distribution of propagation vectors of the wave trains composing it" (**Reference 1**). In that form, you may recognize the uncertainty principle as a relation between the time domain (physical length) and the frequency domain (spectral width) of a wave packet.

The proof of the uncertainty principle rests on properties of the Fourier transform that apply just as well to electrical engineers as to physicists. In the high-speed-digital world, the uncertainty principle translates as, "The shorter the duration of an event in time, the wider must be the spread of frequencies associated with it."

The principle stipulates an inescapable relation between time and frequency, but the large number of ways to define the terms "signal duration" and "signal bandwidth" makes that relation somewhat fuzzy.

Regarding measures of rising-edge duration, you might define the rise time of a digital signal at its 20 to 80% points, whereas your buddy uses the 10 to 90% points. Other persons might draw a line tangent to the rising edge at some point (the midpoint?) and interpret the rise time from that measurement. A mathematically inclined individual might differentiate the signal (turning each step edge into a brief pulse) and then compute the rms duration of each pulse amplitude. A physicist would square her signal, forming a power waveform, and then use the rms duration of the power waveform.

Regarding bandwidth, you can choose between the −3-dB and the −6-dB point, or perhaps you can choose the "equivalent noise-power bandwidth," which is the bandwidth required to pass a specified amount of white-noise power. A physicist computes the Fourier transform of her signal, squares it (making a power spectrum), and then evaluates the rms width of the squared power spectrum.

Each of these definitions—and this is not a complete list by any means—was created to solve a specific problem. For example, the combination of the rms duration of a power waveform and the rms width of the associated squared-power spectrum forms the basis of Heisenberg's principle.

Each measure is slightly different, yet they all exhibit a common behavior: Regardless of the model, shrinking the signal duration in time always increases its bandwidth.

If I could, I would give you an easy means of translating from one measurement standard to another. Unfortunately, a precise translation exists only if you precisely specify the signal shape. For gaussian shapes, the translations between forms of measurements all just involve simple scaling constants. (For example, the product of 10 to 90% rise time and 3-dB bandwidth is 0.338 Hz.) That simplicity explains, in part, why mathematicians love gaussian shapes; the math works out so neatly. For other signal shapes, the translations differ, sometimes markedly (**Reference 2**).

Small details of measurement technique can influence the result. If your data sheet calls out a 20 to 80% rise/fall time, does that mean 20 to 80% of full-scale or just 20 to 80% of the effective signal swing under loaded conditions? What are the loaded conditions? With what bandwidth probe was the signal measured?

With attention to details of measurement technique, I hope you can improve the consistency and reduce the uncertainty in your design process. EDN

REFERENCES

1 Persico, Enrico, *Fundamentals of Quantum Mechanics*, Prentice-Hall, 1950, pg 118.
2 Johnson, Howard, *High-Speed Digital Design: A Handbook of Black Magic*, Prentice-Hall, 1993, pg 399.

Howard Johnson, PhD, of Signal Consulting, conducts workshops for digital engineers at Oxford University and other sites worldwide. Contact him at www.sigcon.com or howie03@sigcon.com.

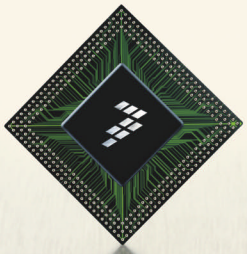

It helped Datang Mobile create a TD-SCDMA network that supports even the world's most populated areas.

Datang Mobile Base Station

Multimedia experiences are on the move. Working with Freescale, Datang Mobile created the world's first tower-to-handset solution for TD-SCDMA—making 3G networks an exciting reality for China. Optimized for high-bandwidth and compute-intensive applications, Freescale quad-core DSPs are designed to deliver enhanced performance while saving power and reducing system cost—an optimal solution for 3G wireless base stations. By working together, Freescale and Datang Mobile make even more of the world a more connected place. freescale.com/smarter

freescale™
semiconductor

BY JOSHUA ISRAELSOHN, CONTRIBUTING TECHNICAL EDITOR

From thin air

After my last column, I expected to be done with the topic of professional conferences for a time (**Reference 1**). Since that column, however, the Darnell Group presented the first annual NanoPower Forum with an outstanding and unusual program that I recommend to your attention (http://nanopower.darnell.com).

The NanoPower Forum pulled together a range of topics that, at first glance, do not appear to relate readily to one another. However, the topics jigsaw together to form a tightly coupled study of certain low-power applications. In all examples, the design goal was to implement a remote function with no signal or power leads connecting the function's remote location and the rest of the system.

Key among the presentation topics was energy harvesting by various means, including piezoelectric devices, thermopiles, and magnetoelectric devices. In many presentations, the energy these devices made available powered sensors or autonomous communications nodes, often both acting in concert. The applications for such devices cover a noteworthy range, including human-body implants, industrial-equipment monitoring, and sensor networks for commercial aircraft.

With such a variety of technologies, applications, and environments, one challenge the segment must face is the lack of clearly defined descriptors and specifications. For example, some presenters used the phrase "energy harvesting" to mean recapturing waste energy from the application under discussion. They used the phrase "energy scavenging" to mean capturing energy available from sources unrelated to the subject application. Other presenters and some attendees, however, seem to use the terms interchangeably or, at least, with less rigorous distinction. This difference is more than nit-picking: If a function is harvesting energy from the application in which it operates, then the designer can characterize the available energy with greater certainty than if the function scavenges energy from unrelated sources.

One challenge the segment must face is the lack of clearly defined descriptors and specifications.

For example, Boeing reported using a thermopile to power remote stress and corrosion sensors. The thermopile operates between the fuselage and the warm side of a nearby insulation blanket. During flight, the fuselage temperature is roughly −30°C, while the plane maintains its interior at or near 20°C for human comfort. The thermopile can power the sensors and their communication links from the 50°C in-flight temperature differential, and the application can tolerate the sensors being nonoperational when the plane is on the ground. By contrast, other presentations reported on scavenging energy from ambient-RF fields—an energy source that, as described, is uncorrelated to the subject application.

Beyond energy harvesting and scavenging, designers of such wireless, remote functions must contend with a variety of other distinctions. For example, does the remote function need to operate continuously or only intermittently? If intermittently, what is the duty cycle? What are the peak and average energy needs? Is the function critical or peripheral to the application? If critical, does it require multiple energy sources to ensure sufficient power under all operating conditions?

Is the energy source narrowband, broadband, or quasistatic? For example, vibration energy from the cases of rotating machinery is often a narrowband source. Ambient-RF energy is a broadband phenomenon. Thermal sources tend to be quasistatic. In each case, the bandwidth of the harvesting or scavenging mechanism must reasonably match that of its energy source. EDN

REFERENCE

1 Israelsohn, Joshua, "Confer for continuing education," *EDN*, May 24, 2007, pg 32, www.edn.com/article/CA6442440.

Contributing Technical Editor Joshua Israelsohn writes about analog and power applications and technologies. Contact him at jisraelsohn@ieee.org.

MORE AT EDN.COM

Go to www.edn.com/070719ji and click on Feedback Loop to post a comment on this column.

Amplifiers | Power Management | Processors | DSP | RF | MEMS | Converters

DIGITAL-TO-ANALOG CONVERTER ICs

Volume 7, Issue 7

YOUR SEMICONDUCTOR SOLUTIONS RESOURCE

FEATURED ARTICLES

Low Voltage, Buffered, Quad DAC Offers 1 LSB INL at 16-Bits 1

Adjusting DAC Voltage Span 2

Dual, 1 GSPS, Interpolating DACs Deliver Speed, Performance, and Consume Low Power 2

Super-Nyquist Operation Delivering High RF Output 3

Design Tips and Tricks: Measuring Ground Impedance 4

Low Noise DAC for Current Source Control 5

128-Position, I²C-Compatible Digital Potentiometer for Electronic Adjustment Functionality 5

Clock IC with Internal VCO Enables Low Jitter 6

Industry's Smallest, Most Flexible Multichannel DACs 7

New Low Cost Alternatives to Standard 12-Bit Bipolar Devices 7

Lens Driver for VCM Autofocus Camera Modules 8

Industry's First Low Voltage, Buffered, Quad DAC Offers 1 LSB INL at 16 Bits for Combined Performance and Flexibility

As PLC and distributed control system designers struggle to accommodate the multiple drive requirements (different voltage and current levels) of their customers, they simultaneously require increased performance and minimized board area in their systems, and design compromises are often necessary.

Solution

ADI's AD506x family, the latest innovation resulting from the continued development of the industry-leading *nano*DAC™ digital-to-analog converter portfolio, allows designers to avoid design compromises by combining both flexibility and performance in a space-saving package.

The AD5064 is a quad, 16-bit buffered digital-to-analog converter with 16-bit accuracy over a 0 V to 5 V output span—delivering a total unadjusted error of less than ±2 mV in 14-lead TSSOP. It is also available in 16-lead TSSOP with individual reference pins allowing each DAC the flexibility to be driven with a different reference voltage, if necessary. In addition, the devices are extremely flexible and provide many convenient control features such as pin-selectable reset/power-on to mid- or zero-scale, hardware LDAC and CLR pins, and a programmable CLR function to allow the user to select clear options. Other features include LDAC override to allow only selected channels to be synchronously updated, an SDO pin to allow daisy-chaining of multiple devices, and a per-channel power-down feature.

The pin-compatible 12-bit quad DAC AD5024, and 14-bit quad DAC AD5044, ensure easy upgrade or downgrade of a design. The family also includes dual 12-bit, 14-bit, and 16-bit devices, and all are available in 14-lead and 16-lead TSSOP packaging, and are controlled using an SPI interface. For more information, visit *www.analog.com/nanoDAC*.

- Pin-compatibility:
  - AD5024: 12-bit
  - AD5044: 14-bit
  - AD5064: 16-bit, 1 LSB, quad DAC
- Pin-selectable functions provide very flexible control

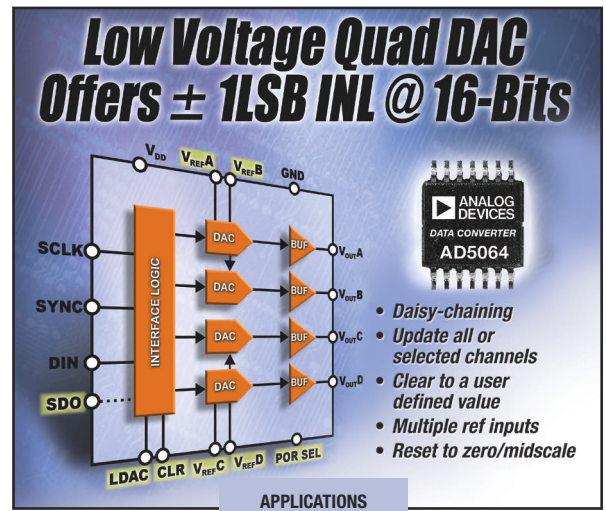

Knowledge Base Resource

Explore ADI's extensive list of free online DAC design tutorials. Visit *www.analog.com/DACtutorials*.

APPLICATIONS

- Process control
- PLC and DCS systems
- Gain and offset calibration

AD5064	$11.50
AD5044	$8.99
AD5024	$5.10

AVNET electronics marketing

All prices in this bulletin are in USD in quantities greater than 1000 (unless otherwise noted), recommended lowest grade resale, FOB U.S.A.

analog is everywhere™

www.analog.com/V7DAC

ANALOG DEVICES

How to Adjust a DAC Voltage Span Using Internal Calibration Registers

The AD5360 is a highly integrated 16-channel, ±10 V, serial input, voltage output 16-bit DAC, in an 8 mm × 8 mm, 56-lead LFCSP package. It offers a nominal output span of 4 × V_{REF}. If, for example, a design requires an output range of –8 V to +8 V, this implies a nonindustry-standard 4 V reference, which will not take into account the zero-scale and full-scale errors of the DAC and may affect the output range.

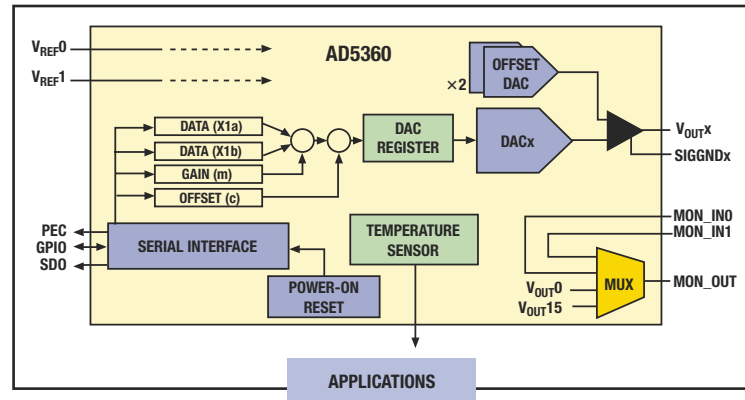

APPLICATIONS

- Industrial control systems
- Instrumentation
- Level setting at ATE

AD5360 $29.95

To overcome this challenge, the solution is to select a reference that will give a higher span than desired and use the internal gain (m) and offset (c) registers to independently trim each channel output to the desired range.

- Use a 4.096 V reference to give an approximate output span of –8.192 V to +8.192 V (including zero-scale and full-scale errors).
- Measure zero-scale and full-scale voltages, e.g., –8.193 V and 8.195 V, respectively.
- Calculate the LSB size, i.e., 8.195 V – (–8.193 V)/65536 = 250 μV.
- To move the zero-scale voltage from –8.193 V to –8 V requires a step of 0.193 V/250 μV = 772 LSBs.
- The full-scale voltage is now 8.195 V + 0.193 V = 8.388 V. To reduce this to +8 V requires a step of 0.388 V/250 μV = 1552 LSBs.
- You can now program these values to the offset (c) and gain (m) registers for the DAC channel.
- The default value for the offset register is 32768. Reprogram this with 32768 + 772 = 33540 to give a zero-scale voltage of –8 V.
- The default value for the gain register is 65535. Reprogram this with 65535 – 1552 = 63983 to give a zero-scale voltage of +8 V.

Dual, 1 GSPS, Interpolating DACs Deliver Unparalleled Speed and Performance with Low Power Consumption

Many modern wireless base stations have adopted a complex I/Q approach for the transmitter chain in either a zero IF (intermediate frequency) or digital IF configuration. While there are multiple reasons for this dual DAC approach, some of the major ones are centered on the increased bandwidth and the image rejection enabled by this architecture.

To achieve optimal performance from a high speed DAC, setup and hold times must be adhered to on the data input. At input data rates in excess of 100 MSPS, determining the timing budget and completing system verification over temperature is not a trivial task. The AD9779A TxDAC® transmit DAC solves this challenge by adding an automatic timing mode. If the user uses the DAC data_clock_out (DCO) in combination with the autotiming mode, the part will select the optimum setup and hold times for the data input and eliminate any effect due to drift over temperature. This feature also improves the interfacing capability to ASICs, or devices, that do not have output timing flexibility. This procedure is automated, thereby decreasing manufacturing time and improving the overall timing margin.

Additionally, the AD9779A DAC contains innovation within the PLL-based clock multiplier circuit. Its dual 16-bit, 1 GSPS parent DAC device, the AD9779, requires the user to manually determine its PLL lock band settings on a part-by-part basis. The AD9779A also features an autosearch function that, when enabled, eases the requirement by determining the best PLL settings for a given frequency and thereby eliminating the trial and error associated with the manual selection.

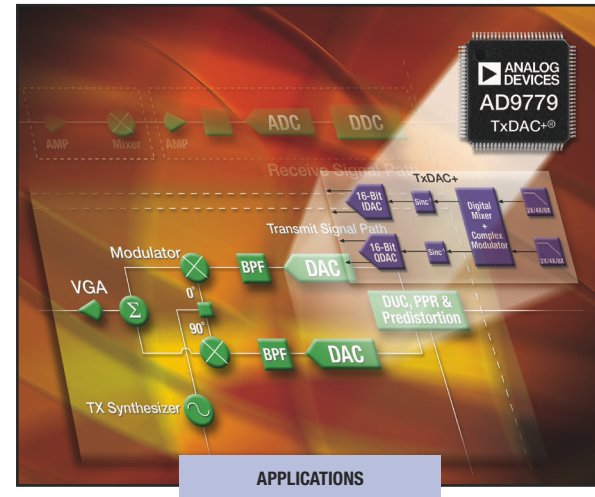

APPLICATIONS

- Wireless infrastructure
- Direct conversion transmit architectures
- Wideband communications

AD9779 $28.95
AD9779A $28.95

Super-Nyquist Operation of the AD9957 Quadrature Digital Upconverter Yields a High RF Output Signal

One often frustrating challenge a designer may face involves frequency limitations of Nyquist devices. Because these devices use a sampling clock to generate output signals, the outputs are typically confined to a frequency range between dc and half the sampling clock frequency. While this frequency limitation is often interpreted as a shortcoming, a Nyquist device actually produces useful signals (spectral images) that extend well beyond the frequency of the sampling clock—thus enabling very high frequency applications.

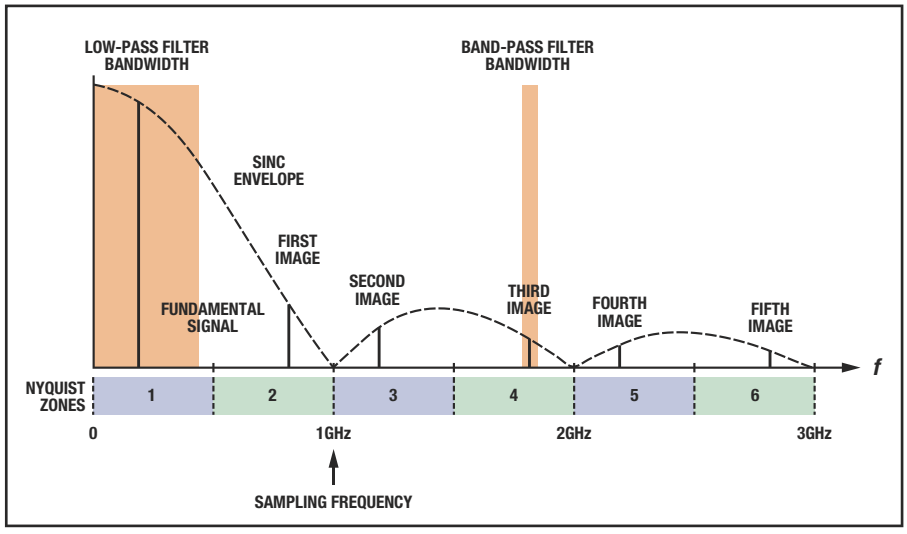

AD9957 $21.00

The AD9957 is a DDS-based quadrature digital upconverter (QDUC) with an integrated 14-bit DAC. It is classified as a Nyquist device, so its 1 GHz sampling clock can generate a fundamental output signal over a frequency range of dc to 500 MHz. Outside of this normal output frequency range, the output signal contains useable images of the fundamental signal out to several GHz. This process is shown graphically in the diagram at left.

Normally, the images would be suppressed via a low-pass reconstruction filter with a bandwidth of approximately 400 MHz. However, for super-Nyquist operation, a band-pass filter is used to select the desired image and reject the unwanted images along with the fundamental signal. A representation of a typical low-pass and band-pass filter bandwidth is shown in the above diagram for clarity.

In the same diagram, also notice that the spectral signals exhibit a frequency dependent attenuation, which is mostly attributed to the well-known sinc response (shown in the diagram with linear magnitude scaling rather than logarithmic, or decibel, scaling). The sinc response is a consequence of the DAC's digital-to-analog conversion process. In addition to attenuation distortion, images in the even Nyquist zones are spectrally inverted relative to the fundamental signal, which is of little consequence for a sinusoidal output signal. However, applications using a modulated carrier are affected by spectral inversion, that is, the equivalent of inverting the Q component of an I/Q symbol constellation. Another point of consideration for applications that use modulation is the effect of the sinc envelope on the transmitted signal. The frequency dependent attenuation can have a negative impact on modulation schemes that rely on a flat response over the transmission bandwidth.

Although the images do suffer from attenuation due to the sinc response, use in high RF applications is not precluded, especially in applications that can tolerate the reduced SNR and SFDR associated with the use of an image. The key to optimizing super-Nyquist operation is to choose a sampling frequency, fundamental frequency, and Nyquist zone that place the image relatively close to the peak of the sinc response curve while keeping other spurious artifacts (and images) in the stop band of the band-pass filter. A proper frequency plan can make very high RF applications possible in spite of the limitations often attributed to Nyquist devices. For more information, visit *www.analog.com/DDS*.

The frequency synthesizer is a critical element in many systems ranging from communications LO generation to chirped radar systems. Explore the design issues, considerations, and trade-offs of using PLL and DDS frequency synthesizer solutions in our on-demand technical seminar: "Options and Solutions in Frequency Synthesis" available at *www.analog.com/onlineseminars*.

Design Tips and Tricks: Measuring Ground Impedance

The AD5933 is a high precision impedance converter system that combines an on-board frequency generator with a 12-bit, 1 MSPS, analog-to-digital converter. In applications where the unknown impedance/sensor being measured has one connection to ground, with no signal return path, the AD5933 can be used to analyze the impedance or sensor.

The AD5933 output excitation voltage is developed across the sense resistor, which is in series with the grounded load. The resulting current through the grounded load is monitored by the AD8220 high precision in-amp, which measures the difference in voltage across R_{SENSE}. The AD8220 sine wave output signal is symmetric about a set bias value of $V_{DD}/2$ and connected back into the AD5933 for digital signal processing.

The AD8220 output voltage is then connected to a series resistor that is coupled with an equal value feedback resistor (RFB). This solution results in an inverted unity gain voltage being generated at the output of the receive side current-to-voltage (I-to-V) amplifier within the AD5933.

- Impedance range: 100 Ω to 10 MΩ
- Phase measurement capability
- Programmable frequency sweep up to a maximum of 100 kHz

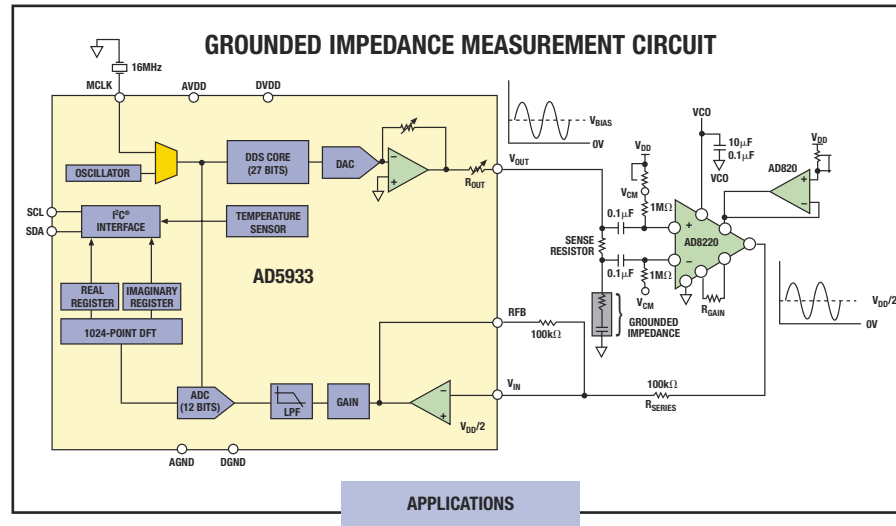

APPLICATIONS

- Impedance spectroscopy
- Complex impedance measurement
- Material property analysis

AD5933	$5.85
AD8220	$2.29

Part Number	Impedance Range	System Accuracy	Power Supply (V)	Temperature Range (°C)	Sampling Speed	On-Chip Oscillator	Temperature Sensor	Price ($U.S.)
AD5933	100 Ω to 10 MΩ	0.5% typical	3, 5	–40 to +125	1 MSPS	Yes	Yes	6.65
AD5934	100 Ω to 10 MΩ	0.5% typical	3, 5	–40 to +125	250 kSPS	No	No	4.35

For more information, visit *www.analog.com/AN-847.*

Online Seminar Series

"Leveraging Advanced Converter Architectures for Impedance and Capacitance Sensors" at *www.analog.com/onlineseminars.*

Low Noise DAC Is Ideal for Programmable Current Source Control

The AD5441 is a high accuracy, 12-bit, single-channel, multiplying DAC specified for programmable current source applications. The input offset voltage of the op amp is multiplied by the variable noise gain of the DAC. A change in noise gain between adjacent codes can produce a step change in the output voltage due to the amplifier's input voltage, which directly leads to degraded linearity performance.

The AD5441's external amplifier allows users to tailor circuit design to suit a variety of application needs—a perfect choice for the construction of an arbitrary signal in applications that require low noise spectral density and greater standards of performance across a wide range of applications. It offers much improved performance over the previous generation DAC8043A product, as well as the associated benefits such as smaller package and lower cost. The AD5441 is available in an 8-lead MSOP and LFCSP. The AD8628 and AD8510 are recommended, complementary low noise amplifiers.

AD5441 IN A VERSATILE V-I CONVERSION CIRCUIT USING AN IMPROVED HOWLAND CURRENT PUMP

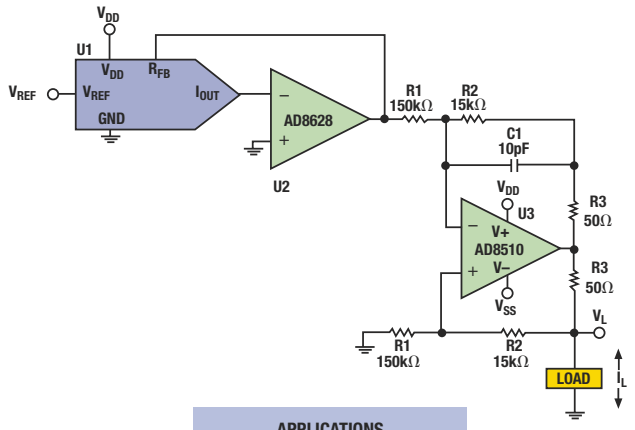

APPLICATIONS

- PLC applications in industrial control
- Programmable amplifiers and attenuators
- Digitally controlled calibration and filters
- Motion control systems

AD5441 $2.47

128-Position, I²C-Compatible Digital Potentiometer Performs the Same Electronic Adjustment as a Mechanical Potentiometer, in an Ultracompact Package

As requirements to reduce board space and size continue to grow, particularly in portable applications, customers demand smaller and smaller packages. However, the control signal chain often requires many different voltage levels to be managed simultaneously. The challenge inherent in these applications is to have multiple compact devices that can be addressed from a single communication bus.

Solution

The AD5247 is a single-channel, 128-tap digital potentiometer with an I²C interface in an ultracompact SC70 package. It has a low tempco of 45 ppm/°C and is available in 5 kΩ, 10 kΩ, 50 kΩ, and 100 kΩ options. Most significantly, the device is available with multiple unique addresses allowing customers to implement various devices placed around the board from a single bus.

- I²C-compatible interface with multiple address options
- Low temperature coefficient: 45 ppm/°C
- Single supply: 2.7 V to 5.5 V
- Low power: I_{DD} = 3 μA typical
- Wide operating temperature range: –40°C to +125°C
- Package: 6-lead, 2 mm × 2.1 mm SC70

MULTIPLE AD5247 DEVICES IN A SINGLE BUS

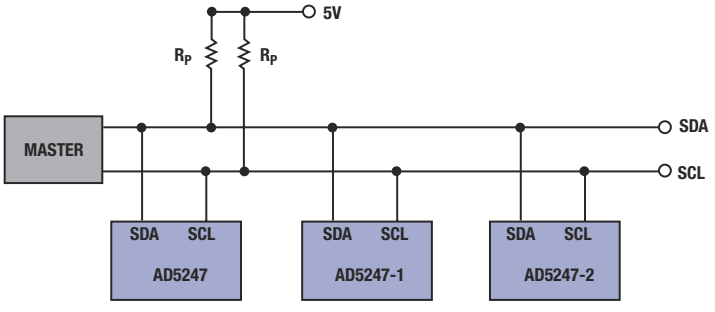

APPLICATIONS

- Mechanical potentiometer replacement in new designs
- Transducer adjustment of pressure, temperature, position, chemical, and optical sensors
- RF amplifier biasing

AD5247 $0.52

For more information, visit *www.analog.com/digitalpotentiometers.*

PLL-Based Clock with Internal VCO Supports Low Jitter Clock Generation Applications

The AD9510 is a low jitter clock distribution IC with on-chip PLL that can be used with an external VCO to create and distribute a low jitter clock signal for clocking DACs and ADCs. The AD9516 is a similar IC that offers the advantage over the AD9510 of not requiring an external VCO. Both devices require an external R-C loop filter to complete the phase-locked loop.

The performance of a clock generator depends on several factors that include, most importantly, the phase noise of the VCO, as well as the phase noise of the reference oscillator, and the loop bandwidth of the PLL. Some optimizations can be made by selecting an optimal loop bandwidth for the combination of the VCO and reference to minimize the overall contribution of the phase noise to the resulting time jitter.

The reference oscillator for both the AD9510 and the AD9516 will be external—making the two parts essentially equivalent in terms of choice of reference source. However, the AD9510 requires the selection and use of an external VCO; there are a number of VCOs available for a given frequency requirement, but some offer better phase noise characteristics than others. A clock generator design using the AD9510 allows for the selection of the highest performance VCO obtainable.

The plot in Figure 1 is an ADI*sim*PLL simulation that shows the AD9510 generating a 491.52 MHz clock using a VCO (Sirenza VCO190-1455T) at 1474.56 MHz, with a reference of 15.36 MHz, and a loop bandwidth of 100 kHz. The equivalent time jitter of this clock generator calculated by the integration of the phase noise from 12 kHz to 20 MHz is 338 fs rms.

The AD9516 features an on-chip VCO with very good phase noise performance, but not as good as the best standalone VCOs available for the same frequency ranges. However, a very high performance clock generator can be achieved using the AD9516.

Figure 2 shows the absolute phase noise of an AD9516-4 generating a 491.52 MHz clock signal at its LVPECL output. The VCO frequency is 1474.56 MHz and the loop bandwidth is 100 kHz. The reference input is an Agilent SMA100A signal generator at 15.36 MHz.

The equivalent time jitter of this clock generator calculated by the integration of the phase noise from 12 kHz to 20 MHz is 275 fs rms.

As the chart below illustrates, the AD9516 family of clock ICs with internal VCOs is capable of generating a very high quality low jitter clock signal, with performance that mitigates the expense, board space, and reliability headaches of implementing an external VCO.

For more information, visit *www.analog.com/clocks*.

FIGURE 1: AD9510 PHASE NOISE PERFORMANCE 491.52 MHz OUTPUT CLOCK

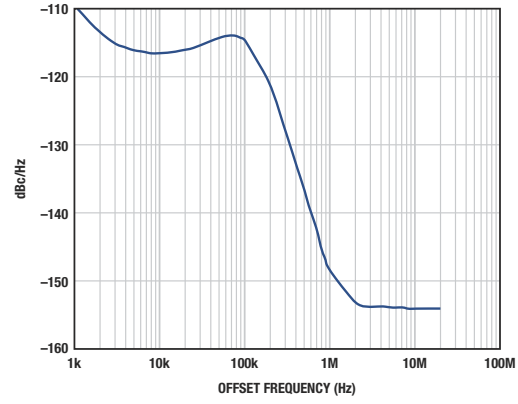

AD9510*	$10.99
AD9516*	$11.47

*5k unit pricing

FIGURE 2: AD9516 PHASE NOISE PERFORMANCE 491.52 MHz OUTPUT CLOCK

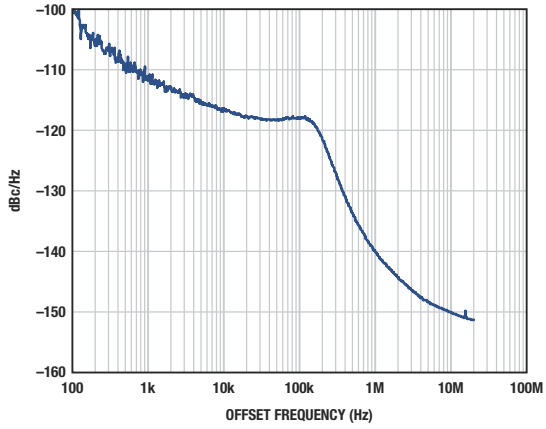

Online Seminar Series

At higher frequencies and bit resolutions, the quality of the sampling clock becomes an increasingly critical factor in the ultimate level of dynamic performance achieved by a given data converter application.

Considerations include: the relationship between jitter and phase noise, methods for measuring subpicosecond jitter, and analyzing clock jitter impact on converter performance.

Learn from the experts by viewing our on-demand, two-part technical seminar series: “Using Low Jitter Clocks to Enhance Data Converter Performance” and “Performance Clocks: Demystifying Jitter.”

Available at *www.analog.com/onlineseminars*.

Industry's Smallest, Most Flexible, Multichannel DACs Ease Board Design and Layout Challenges

Designers need flexibility when specifying converters for high voltage applications, such as factory process control, instrumentation, and dc setpoint-control applications. These requirements include choice of resolution, number of channels, voltage range, power supply, and more. In addition, board design and layout challenges present integration and package size needs.

Solution

The AD575x, AD573x, and AD572x DAC families provide a wide range of 12-bit to 16-bit, dual and quad channel options operating over a wide voltage range on either single or dual power supplies. The 16-bit AD5754R features 8 LSB INL, 0.1% total unadjusted error (TUE) and an on-board 5 ppm/°C reference. In addition, the highly configurable DACs feature software-programmable unipolar and bipolar voltage output ranges that can be independently set per channel providing designers the option to eliminate a number of discrete components, such as precision resistors, switches, and external amplifiers. Available in 24-lead TSSOP, the AD5754R packs uncompromised performance and features in a small package size. Combined with the elimination of discrete components, the benefits translate into board space saving of as much as 70%. In addition, the 16-bit performance, integrated features, and excellent TUE reduce cost and calibration time. For more information, visit *www.analog.com/industrial*.

APPLICATIONS

- Industrial automation
- Closed-loop process control
- Programmable logic controllers

AD5754R $10.05

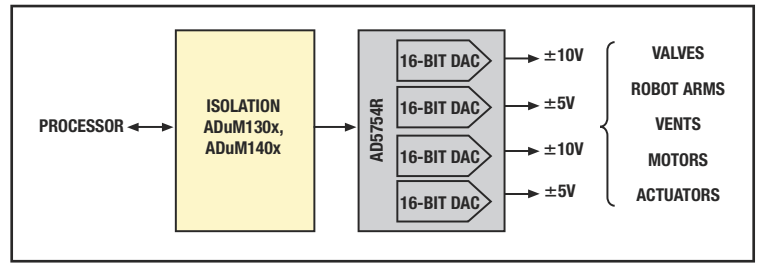

- AD572x: 12-bit dual and quad devices
- AD573x: 14-bit dual and quad devices
- AD575x: 16-bit dual and quad devices

Online Seminar Series

"Data Conversion Applications and Solutions for Precision Process Control" at *www.analog.com/onlineseminars*.

Introducing Two New, Low Cost Alternatives to Industry-Standard 12-Bit Bipolar Devices

Using technological advancements enabled by ADI's *i*CMOS™ industrial process technology, the AD5726 and AD5725 offer low cost, pin-for-pin and functionally equivalent alternatives to the industry-standard DAC84xx, DAC76xx, and DAC77xx series. These parallel (AD5725) and serial (AD5726) input quad 12-bit bipolar voltage out DACs operate from a wide variety of supply (+5 V to ±15 V) and reference voltages (+2.5 V to ±10 V). In addition, they also feature 12-bit resolution, 0.5 LSB INL (max), and are guaranteed monotonic over the industrial temperature range.

The devices integrate features essential to reducing design time and overall system cost, including a low headroom/wide swing amplifier, power-on reset, output control (asynchronous clear to zero-scale, midscale), and I/O lines. The AD5725 and AD5726 are ideal for closed-loop servo control, process control, and dc setpoint control applications. These parts have JEDEC-compatible 3 V and 5 V logic that allows easy interface to mixed-voltage systems, while eliminating the need for additional level shifting circuitry. Both generics are now in production. For more information, visit *www.analog.com/dac*.

Part Number	Description	Interface	Alternative/Upgrade	Price ($U.S.)
AD5726	12-bit quad bipolar	Serial	DAC8420,* DAC7714,* DAC7715, DAC7614,* DAC7615	6.95
AD5725, AD5725-1	12-bit quad bipolar	Parallel	DAC8412, DAC8413, DAC7724, DAC7725, DAC7624, DAC7625	6.95

*Denotes pin-for-pin replacement

For samples and data sheets, visit www.analog.com/V7DAC

ANALOG DEVICES

Lens Driver Optimized for VCM Autofocus Camera Modules

New, low voltage digital cores of image sensors and signal processors require a small, low voltage I²C-compatible interface. The AD5821 is a single 10-bit DAC with 120 mA output sink current capability. It is designed to meet the requirements of autofocus, image stabilization, and optical zoom applications in camera phones, digital still cameras, and camcorders. The IC operates from a single 2.7 V to 5.5 V supply, and is controlled via a 2-wire (I²C compatible) serial interface that operates at clock rates up to 400 kHz, and has 1.8 V-compatible logic thresholds compatible with new digital cores. It is available in a 1.5 mm × 1.7 mm wafer level CSP (WLCSP) package.

- 1.8 V, I²C interface
- All dc operation = no noise
- Package: 1.5 mm × 1.7 mm WLCSP

BLOCK DIAGRAM SHOWING CONNECTION TO VOICE COIL

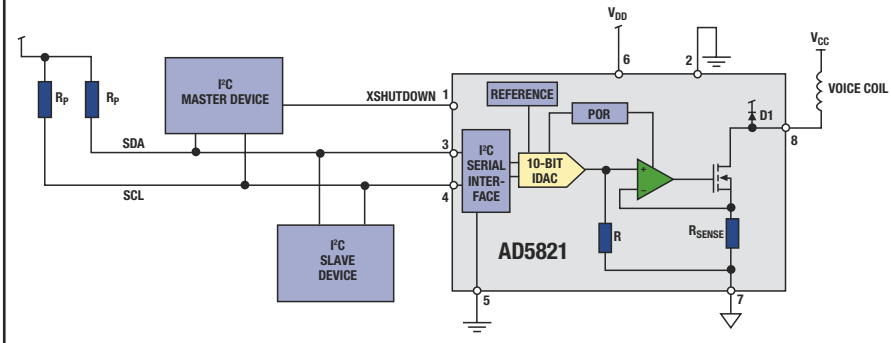

APPLICATIONS

- AF using VCM in camera modules
- Camera phones
- Web and PC cameras

AD5821

Contact ADI for pricing.

Part Number	Interface Voltage (V)	Size (mm)	Operating Voltage (V)	Resolution (Bits)	Output Current (mA)	Interface	POR	PD
AD5398	2.7 to 5.5	1.5 × 1.7	2.7 to 5.5	10	120	I²C	Yes	Yes
AD5821	1.8	1.5 × 1.7	2.7 to 5.5	10	120	I²C	Yes	Yes

Analog Devices, Inc. Worldwide Headquarters
Analog Devices, Inc.
One Technology Way
P.O. Box 9106
Norwood, MA 02062-9106
U.S.A.
Tel: 781.329.4700
(800.262.5643,
U.S.A. only)
Fax: 781.461.3113

Analog Devices, Inc. Europe Headquarters
Wilhelm-Wagenfeld-Str. 6
80807 Munich
Germany
Tel: 49.89.76903.0
Fax: 49.89.76903.157

Analog Devices, Inc. Japan Headquarters
Analog Devices, KK
New Pier Takeshiba
South Tower Building
1-16-1 Kaigan, Minato-ku,
Tokyo, 105-6891
Japan
Tel: 813.5402.8200
Fax: 813.5402.1064

Analog Devices, Inc. Southeast Asia Headquarters
Analog Devices
22/F One Corporate Avenue
222 Hu Bin Road
Shanghai, 200021
China
Tel: 86.21.5150.3000
Fax: 86.21.5150.3222

Printed in the U.S.A. SB06864-90-6/07

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

BY PALLAB CHATTERJEE, CONTRIBUTING TECHNICAL EDITOR

Where is the BOM in IC design?

As ICs have progressed to becoming full systems, such as SOCs (systems on chips), they have been mirroring historical development and design issues that occur in PCB (printed-circuit-board) designs of complex systems. These issues include interconnect and power-supply noise, signal coupling between layers, inductance of the interconnect, limitations in size and reliability of interlayer vias, and stacking issues on multiple layers of interconnect.

EDA tools have faced similar issues, with netlists, physical design, software/hardware-cosimulation requirements, mixed-mode (analog/digital)-simulation requirements, I/O simulation and loading, and autoplacement and routing tools. The industry has optimized the magnitude of devices and the algorithms for each market, but, on a high level, they are similar in function.

For the most part, even the design methodologies have a great deal of overlap in circuit design, design verification, physical design, physical-design verification, and application-interface verification and testing. Traditionally, there has been one major discrepancy in the flows that the industry should correct. BOM (bill-of-materials) costs, along with approved vendor and provider lists, drive PCB-design tools and flows; IC-design flows have no such automated drivers.

Current SOC designs comprise mostly premade parts or component IP (intellectual property). This assembly technique is similar to PCB design for full-system assembly. Rather than a real part, a "phantom" part, such as a symbol or an outline of the part that indicates its size, pinout, and functional model, manages the IP in physical design in both areas. Some of the custom-IC parts go down to the individual-device level. For example, rather than have formal outlines, automated routines, such as pcells, create passive components and diodes at physical-design runtime. The IC-design world, however, has no mechanism for visual inspection or validation of the BOM to ensure the replacement of phantoms with the correct real part.

The PCB community tends to lack the level of trust in tools and third-party operators.

According to a number of high-volume board-design and -manufacturing/assembly companies, the incidence of incorrect vendor parts or missing parts on the board is less than 0.01%. Engineers find most of these problems by optical inspection before the expense of electrical testing and debugging. In stark comparison, more than 20 major IC houses and a large number of major start-ups report the problem of incorrect component placement or missing devices in a physical design more than once per division, with occurrence rates of greater than 15%. On the IC side, because of the large number of devices, optical checking is impossible; instead, engineers use automated LVS (layout-versus-schematic) tools. The LVS environment has two major drawbacks. First, the tools cannot check nondevice objects that form critical parts of the design. Second, LVS uses black-box methodologies that verify and sign off a design based on an implementation of a phantom that the vendor provides after sign-off.

For some reason, officials at the new generation of semiconductor companies think it is acceptable to design a project, which may cost $10 million to $100 million and use $1 million to $50 million in tools; work on it for years; and then allow some blind, automated tool the companies don't control to finish the design. The PCB community tends to lack the level of trust in tools and third-party operators.

On a positive note, most of the major IC companies have implemented mandatory *written* sign-off procedures to validate correct linking to approved IP blocks and to make sure that generated devices instead of big shiny spots of bare silicon actually appear on the wafers. These procedures, with rudimentary automation, have resulted in an average schedule reduction of 22%. Most smaller companies have not yet implemented such procedures. One can only estimate the schedule reduction and yield improvement that would be possible if automation methods for the PCB community were available to the IC world. Unless customers require tool and IP providers to support this feature, we remain in a semiconductor chain that maximizes profit for the foundries, test vendors, and tool providers, rather than for the design and end-customer-support side. EDN

Contact me at pallabc@siliconmap.net.

Intersil Medical Products

High Performance Analog

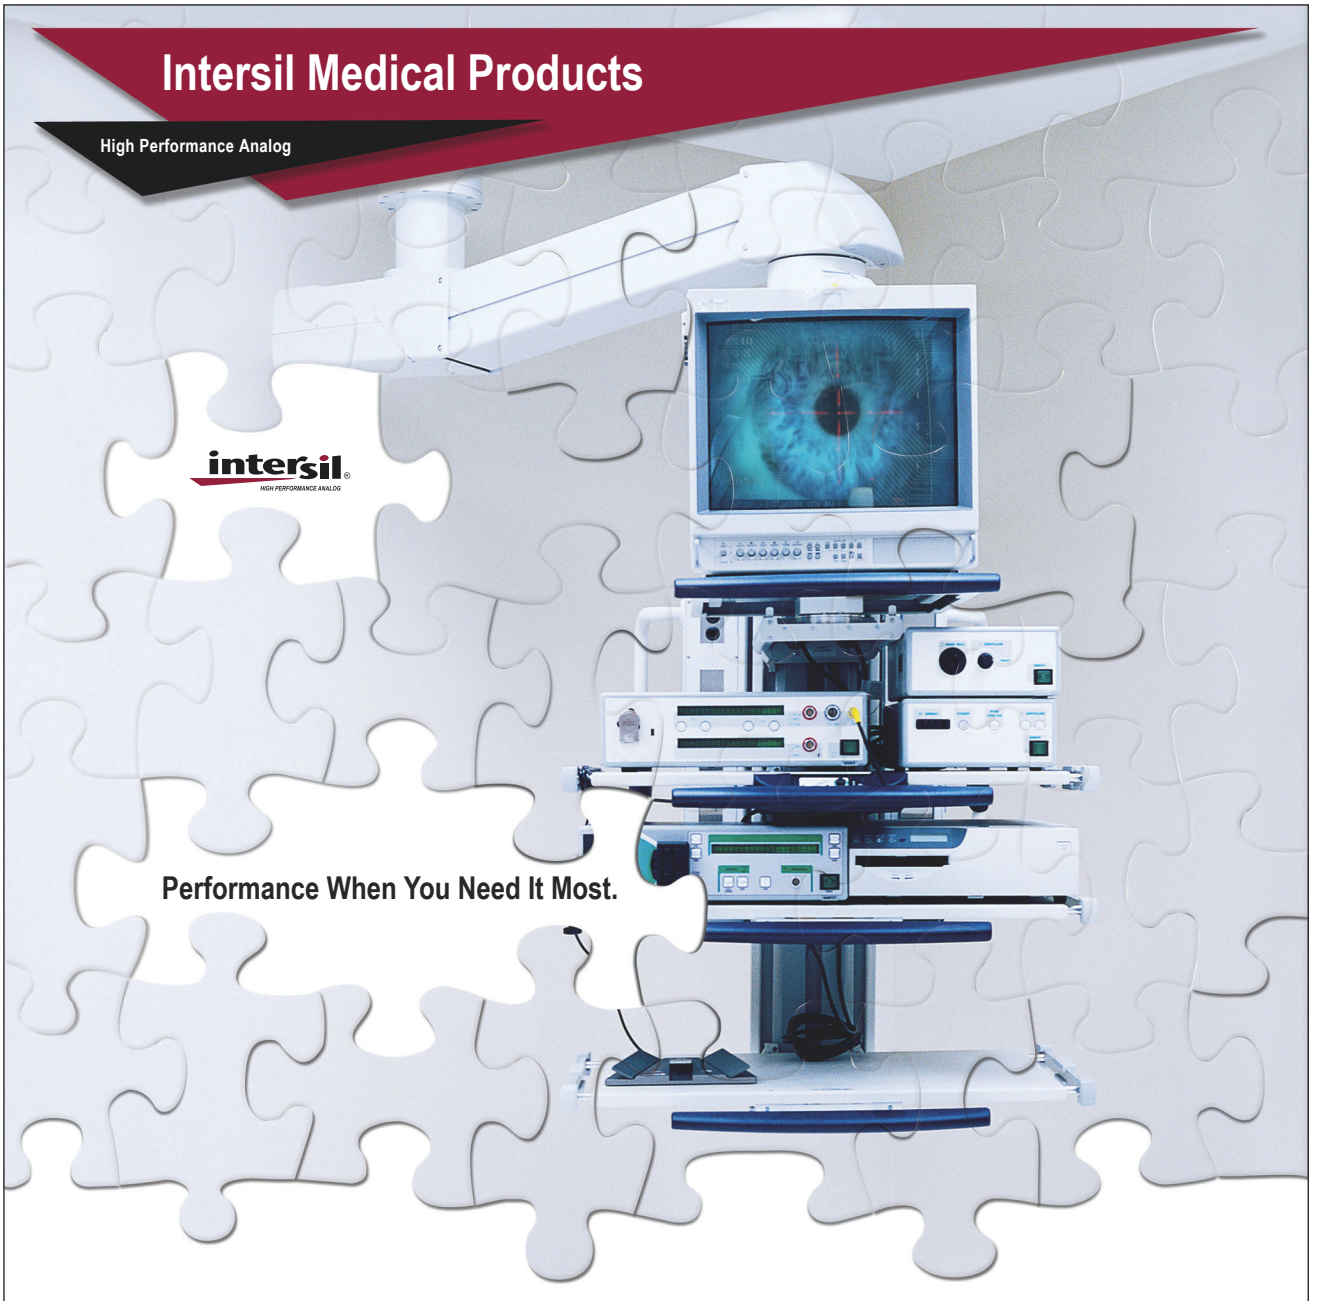

Performance When You Need It Most.

Get the performance you need with Intersil's high performance industrial and medical analog ICs.

Analog Mixed Signal:

- DCPs
- Data Converters
- Interface
- Instrumentation Amps
- Op-Amps
- Real-Time Clocks
- Switches/MUXes
- Voltage References

Power Management:

- Backlight Drivers
- Battery Authentication
- Battery Chargers
- Fuel Gauges
- Integrated FET Regulators
- LCD Display Power
- LDOs
- Memory Power Management
- Overvoltage and Overcurrent Protection
- PWM Controllers
- Voltage Monitors

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

Intersil – An industry leader in Switching Regulators and Amplifiers.

intersil®
HIGH PERFORMANCE ANALOG

Avalanche!

Customers often returned process controllers to us for repair, because the field operators could not adjust them correctly. For example, consider a gas pipeline that is to operate at 3000 psi. When measured by a pressure transducer calibrated for 0 to 5000 psi (with an output of 1 to 5V dc), 3000 psi produces an output of 3.4V, which is the process variable. The operator will have adjusted the setpoint to 3.4V also for operation at 3000 psi. Now the process variable is equal to the setpoint, and the control loop should "sit still"—that is, produce no output change.

Often, the field operator could adjust the controllers for constant output when the process variable and setpoint were both equal to 1V but not when they were both equal to 5V. In this situation, the CMRR (common-mode rejection ratio) of the input differential amplifier was unsatisfactory.

Consider a typical process-control loop, which contains a 45V power supply, a wire-line pair (or equivalent) from the control room to the measurement site, a current transducer, and a 250Ω "range resistor," all in series. The current transducer, which measures the process variable, produces a current of 4 to 20 mA in the wire-line pair, representing 0 to 100% of the value of the process variable. This current in the 250Ω-range resistor produces an input for the controller of 1 to 5V.

But sometimes a 1 to 5V input signal is already available and the range resistor is unnecessary. Then the installer does not plug it in at the controller. However, if the controller is to work with a 4- to 20-mA signal, and the installer forgets the range resistor, the input to the controller is about 45V. Not good!

The process-controller circuits included resistors between the input connector and the inputs of the first operational amplifier (configured as a differential amplifier) that were greater than 100 kΩ, so an input voltage that exceeded the amplifier's maximum common-mode rating would not cause damage, would it? Could a current great enough to change the characteristics of the op amp come from a 45V source through a 100-kΩ resistor?

I asked an engineer leading a Motorola seminar whether applying 45V to a 100-kΩ resistor connected to an op-amp input, operating with a ±15V power supply, could harm the amp.

"Yes," he said. "The op-amp input devices will *avalanche*. The amp may not fail, but the operating characteristics will change."

Avalanche? What did *that* mean?

Semiconductor junctions have a "reverse breakdown voltage" at which a reverse-biased junction begins to conduct. Zener diodes behave in this way, which makes them useful as voltage regulators. But this situation is unhealthy for sensitive op-amp input transistors. With a circuit like the one described above, this phenomenon is not usually fatal, but it will change the behavior of the op amp. *It will spoil the balance of the input differential pair.*

So, we had to keep the input voltage from exceeding the supply voltages (±15V). We placed diodes at the controller inputs (cathodes to the 15V supply; anodes to the connector-input pins) that would conduct if the controller-input voltage exceeded the positive supply voltage. To protect against input voltages more negative than −15V (in case of a wiring error), we also placed diodes between the controller inputs and the −15V supply (anodes to the −15V supply; cathodes to the connector-input pins).

These diodes had to be good; leakage had to be low at the highest operating temperature. At a 0% reading (4 mA, 1V), 40 μA represents an error of 1%.

After we added the diodes, customers no longer returned the modified controllers. EDN

Walter Lindenbach started and operated Calgary Controls Ltd from 1970 to 1990, at which point he retired.

DANIEL VASCONCELLOS

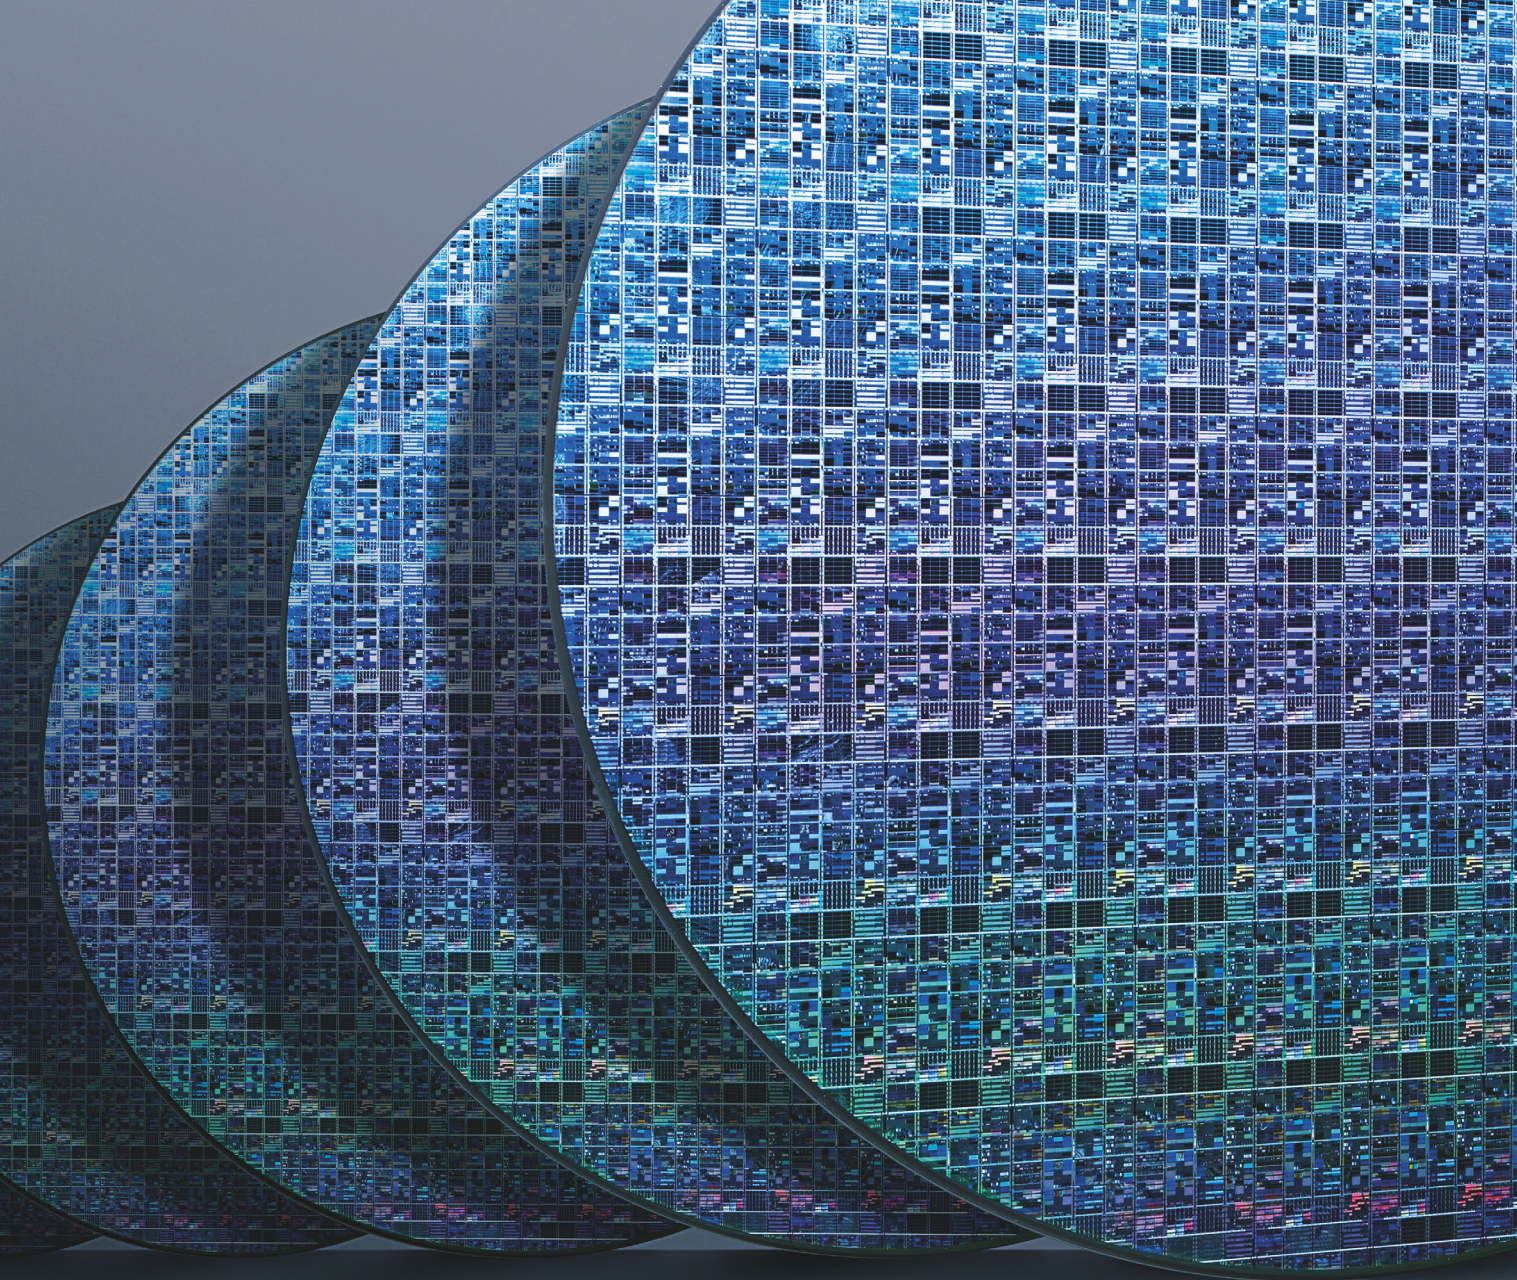

A CODEC CHIP THAT ACHIEVED 50% COMPRESSION IN DEVELOPMENT TIME.

THAT'S MODEL-BASED DESIGN.

To be first to market with a next-generation mixed-signal chip, the IC design team at Realtek used system models to continuously verify their design. The result: 50% less time to completion and a 50% market share. To learn more and to request the IBS study on Simulink ROI in Electronic Design, visit mathworks.com/mbd

The MathWorks

Accelerating the pace of engineering and science

MATLAB® & SIMULINK®

PRY FURTHER AT EDN.COM

Visit www.edn.com/070719pry for more in-depth coverage of the Xbox 360, including internal photographs of standard and Elite systems.

Eviscerating the Xbox 360 Elite

In late April, Microsoft refreshed its nearly 1½-year-old Xbox 360 product line with the high-end Elite variant, touting an upgraded-capacity 120-Gbyte HDD and a Version 1.2 HDMI port. When you crack open the Elite's sleek black case, what—if any—alterations to the initial console design do you discover beyond the HDMI augmentation?

The Xbox 360's hardware scaler represents the single biggest evolution in the system-design transition from the Core and Premium Xbox 360 models to the Elite version. This scaler is also a notable differentiation from the hardware-scaler-deficient Sony PlayStation 3. Microsoft uses the scaler to increase and decrease the resolution of, as well as to interlace and deinterlace, game and video content to match the connected display's desired attributes. The Core and Premium systems' Ana scaler IC embeds the necessary DACs for various analog-video connections; the Elite's Hana scaler (H=HDMI?) presumably also integrates the HDMI transmitter and therefore has access to both the digital-video and -audio data coming from other system ICs. Missing from the Elite motherboard is a Cypress clock-generator IC that was just to the left of Ana; a migration from TQFP to BGA packaging for the scaler IC also marks the Ana-to-Hana transition.

Four of the Xbox 360 ICs carry Microsoft labels. In contrast to the first-generation Xbox, for which Microsoft depended on the supply-versus-demand whims of its IC partners, the company this time acquired IP from other companies and directly manages chip manufacturing. The system's southbridge IC contains circuits from SiS; this chip's functions include SATA controllers for the system HDD and DVD drive, USB2 transceivers, audio DACs, and the Ethernet MAC.

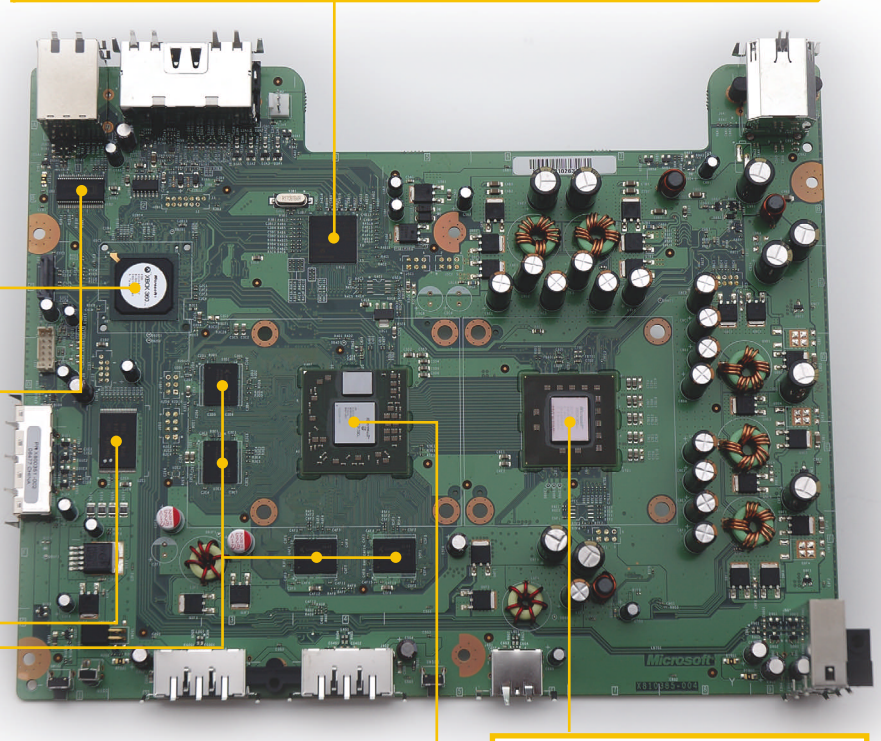

The Ethernet PHY comes from ICS; Broadcom is an alternative source. A transceiver and antenna residing on a separate daughterboard (not shown) handle the 2.4-GHz wireless communication between the console and its game controllers.

The bulk of system memory is 512 Mbytes of GDDR3 SDRAM. Four of the 512-Mbit ICs, which are multisourced from Qimonda and Samsung, reside on the PCB topside; the other four are mirror-image-mounted on the motherboard backside. Because the Xbox 360 comes in a Core variant with no HDD, a Hynix 128-Mbit NAND-flash memory serves as an alternative storage location for operating-system patches and other code and data updates. Some versions of the motherboard (not shown) also include a 2-kbit Atmel EEPROM directly below the CPU.

Under the heat sink, the GPU is a multidie module comprising the graphics processor and a 10-Mbyte frame-buffer-memory IC. Analysis by Semiconductor Insights suggests that neither the CPU nor the GPU has yet migrated from a 90- to a 65-nm lithography; single-IC GPU integration is one possible outcome of the 65-nm conversion slated for later this year.

The Xbox 360 has no discrete boot memory IC; Microsoft learned a lesson from the hacker community on the first-generation Xbox, whose exposed traces between the CPU and flash memory were its Achilles' heel. This time, Microsoft embedded the system firmware within the CPU. The boot code is in-system-upgradable.

COURTESY LLAMMA.COM

MILL-MAX®

Maximum Interconnect Solutions

Mill-Max Spring-loaded Connectors

MINIMIZE Noise.

When the power is on, you don't want shock or vibration to create spurious signals. Mill-Max Spring-loaded Connectors provide a reliable electrical connection in the most rigorous environments.

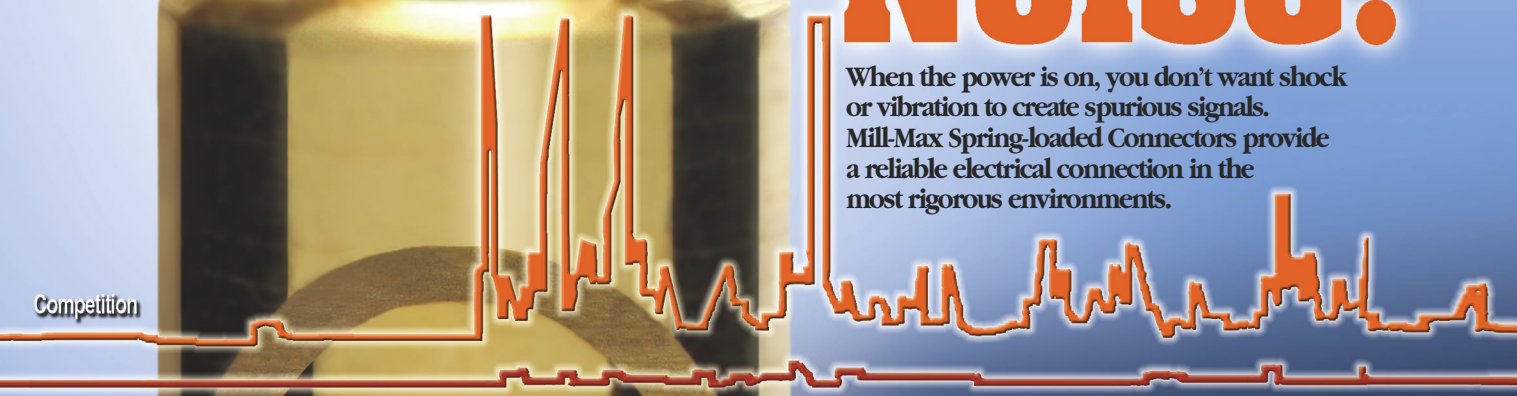

Maximum Continuity: Precision-machined gold-plated components and a low-resistance spring maintain a consistent electrical path.

Maximum Stability: Tested to a minimum of 50G shock and 10G vibration with no spikes >1µs and >1.15V with 0.5A applied.

Maximum Endurance: 1,000,000 cycles and still electrically silent.

Maximum Range: Surface mount and thru-hole configurations with various profiles and multiple stroke lengths.

Mill-Max Spring-loaded Connectors are typically used as the battery charging contacts in portable instruments, or as a rugged interconnection between circuit boards.

For information and our Free Design Guide,
go to **www.mill-max.com/respond**
Response Code: **EDN560**
Phone: 516-922-6000

AC TO IC

Any Voltage • Any Current
Any Power Architecture

RoHS Compliant

ZM7300 Digital Power Managers

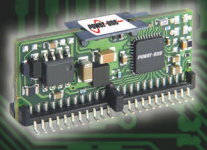

Z-Series Digital POLs

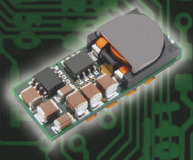

Y-Series Analog POLs

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

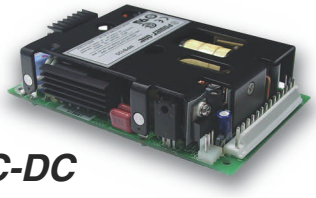

AC-DC Embedded

DIN Rail

Cassette

CompactPCI

power-one®
Changing the Shape of Power

www.power-one.com

INTERFACE CHIPS MUST MEDIATE BETWEEN EVER-LOWER LOGIC VOLTAGES AND REAL-WORLD LOADS. HERE ARE THE CIRCUITS AND TECHNIQUES YOU NEED TO KNOW.

INTERFACE CHIPS: between logic and a hard place

BY PAUL RAKO • TECHNICAL EDITOR

Logic signals have fallen from the 15V of the old CMOS 4000 series to 5V TTL (transistor-to-transistor logic) to modern CMOS levels of 3.3, 2.7, and 1.8V. Advanced processes use digital logic that operates at 1V or lower, although they offer higher I/O voltages. These ever-dropping logic voltages still must drive solenoids, twisted-pair wiring, and discrete-semiconductor-power stages, so they require external interface circuits. You need to understand these circuits, which run the gamut from FET drivers to open-collector solenoid drivers to LVDS (low-voltage-differential-signaling) and isolated drivers.

Interfacing between circuits is difficult because it connects signals of different voltage levels or across different impedances. Another hardship is the brutality of the real world. Anyone who has tried to pass CE (Conformité Européenne) immunity tests can confirm that any device or signal that connects to the real world must face voltage surges and spikes to survive. In addition to these difficulties, many interface ICs are operating at higher frequencies. The high speeds of these new interface ICs add a whole new set of challenges.

Let's begin our journey through the perilous waters of interface with one of the most basic devices, the bipolar transistor. Because silicon transistors turn on with a base voltage of 0.6V, you can, in principle, drive a transistor from almost any logic family. One approach is to put a series resistor in the base connection (**Figure 1**). You can drive the transistor directly from logic, but it is poor design practice to rely on the internal resistances in the logic IC to limit the current into the transistor. Because the base of a bipolar transistor stays at 0.6V when you turn it on, the logic IC experiences a short circuit on its output.

So, even something as prosaic as hanging a transistor onto the output of a gate requires some engineering. Look up the transistor's beta over the temperatures at which you expect the circuit to operate. Use that figure to determine the base current you need to turn on the transistor at the highest expected load. Remember that the load current

may also be a function of temperature. Then, look at the data-sheet charts of the logic part you are using. Once again, you may have to factor in temperature to see the actual drive current available. The drive current depends on the voltage drop. The transistor base is at 0.6V, whereas a CMOS-logic IC tries to swing the output to the power-supply rail, and a bipolar-logic family can get to within a diode drop, or 0.6V, of the power rail. Knowing the voltage drop and the current that the transistor needs, you can work out a nominal value for the resistance. It is usually advisable to halve the value of the resistor—from 150 to 75Ω, for example—to ensure that the transistor will fully turn on. Then, you must look at the power that the logic IC loses and make sure that it does not exceed any power-dissipation specs. Using all eight outputs of an octal buffer can cause the part to burn up, especially if the part comes in a TSSOP or another tiny package that cannot dissipate much heat. Also, the logic part has its own temperature derating of power dissipation, so you may have to factor that figure in to ensure that you are not overstressing the part. Remember: If you use a base resistor with too low a value, or if you entirely omit the resistor, then you will drive the transistor into heavy saturation, and it will take longer to turn off. This situation occurs because all the excess hole-electron pairs you have injected into the base-emitter junction must recombine before the collector-emitter current can go to zero.

All of these calculations apply to the use of a bipolar transistor to buffer a logic output. Using a MOSFET creates another set of criteria. FETs require no current to turn on; a voltage on the gate causes the source and drain to become low-impedance. The benefit of FETs is that they can conduct current with lower losses, and some parts, such as those from Supertex, can handle very high voltages. Older FETs needed 10V to turn on. Process engineers used ion implantation to lower the gate-to-source threshold voltage. Be aware, however, that developers designed these logic-level FETs for 5V logic. You must examine the data sheet to ensure that a part will turn on with 3.3V or lower logic. In addition to the drive-voltage problem, the digital-logic part that drives the FET gate sees that gate as a capacitor. The bigger the FET, the larger the capacitance the digital part sees. Because both CMOS and bipolar logic can drive only so much current, the effect of the FET-gate capacitance slows the FET's turn-on. Again, you need to review the worst-case specs in the logic's data sheet, based on operating the logic at the lowest expected power-supply voltage to ensure that the FET will turn on and off fast enough to meet your design intent.

Concerns such as these have given rise to integrated ICs, such as the Microchip TC4468 and the Texas Instruments SN75372 dual-MOSFET driver, which interface to MOSFETs. You can also use amplifiers to drive large MOSFETs, but be careful. Driving large FETs from amplifiers may cause the output to oscillate due to the capacitive load that the FET presents to the amplifier.

On top of all these concerns, you must also consider Miller capacitance: a small stray capacitance between the gate and the drain in a FET or between the base and the collector in a bipolar transistor. Raising the gate causes the drain's voltage to drop, and this drop "fights" the signal that is trying to raise the voltage on the gate (**Reference 1**). Thus, you must make sure that the transistor switches as desired under real-world operating conditions. If you need to control more than one device, you can use an integrated transistor array rather than a discrete power transistor. Allegro and other companies offer arrays of transistors on single chips. Typing "transistor array" into Digi-Key's search box returns a list of 16 companies.

It's remarkable that you need to put all this work, study, and calculation into the simple act of buffering a logic output with a transistor. But this work illustrates the difference between a hobbyist and an engineer: An engineer uses specifications and calculations to ensure that a circuit will work as he intends—not only in the lab, but also through an entire production run. For example, engineer John Massa, currently owner of a service bureau, has worked at Houston Instruments, Teledyne, and Quantum Disk Drives. Back then, he would counsel his fellow engineers that driving an LED with a 74LS244 octal buffer would not work. Many of these engineers were incredulous, saying that they had successfully done so many times. Massa would only smile and advise them to read the data sheets. The data sheet for the Texas Instruments SN74LS244 and a typical 20-mA LED in use at the time shows why Massa would take this position (**references 2** and **3**). The Texas Instruments part shows a guaranteed output voltage of 2.4V at only 3-mA output current. The Chicago Miniature Lamp red-LED 5306H1 has a maximum forward voltage of 3V at 10 mA. If you

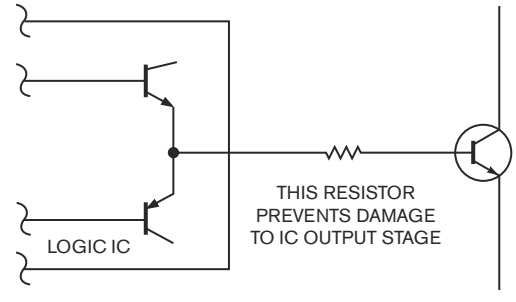

Figure 1 Be sure to use a series resistor when connecting logic outputs to a transistor. Large FETs with substantial input capacitance may also need a series resistor to protect the logic-output stage. The resistor slows FET switching, but the resistive damping stabilizes the circuit.

AT A GLANCE

- Interfaces can be as simple as a transistor or as complex as a FireWire chip.
- Using interface ICs alleviates many engineering problems.
- RS-232, RS-422, and RS-485 are older legacy interfaces.
- USB (Universal Serial Bus), FireWire, SATA (Serial Advanced Technology Attachment), and PCIe (Peripheral Component Interconnect Express) are modern interfaces.
- New ICs ease the design of isolated interfaces.

Micrel's Family of Sequencing LDOs Gets You Back on Track

MIC68xxx Solves your Tracking, Sequencing and Ramp Control™ Requirements

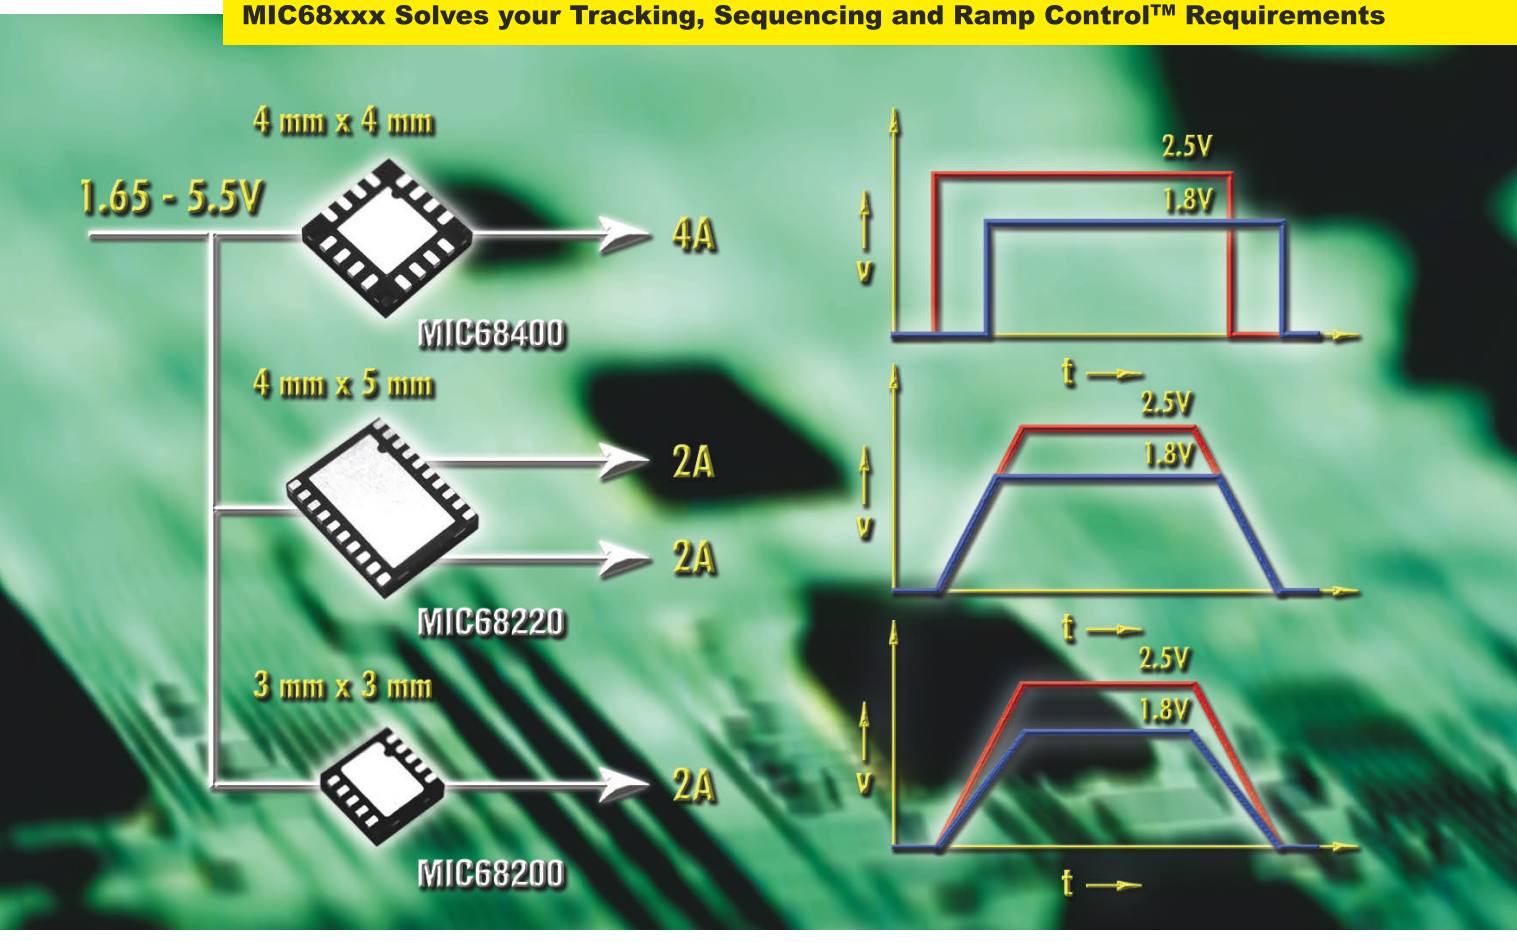

ASIC, FPGA and CPU manufacturers require complex and unique start-up protocols for loads on the system board. Micrel's MIC68xxx family of regulators are designed specifically to address those protocols and make the task of the system designer simpler.

All members of MIC68xxx family operate from a wide input range of 1.65V to 5.5V, which includes all of the main supply voltages commonly available today. They are designed to drive digital circuits requiring low voltage at high currents. The chips incorporate a delay pin (Delay) for control of power on reset output (POR) at turn-on and power-down delay at turn-off.

Additionally, the MIC68xxx family offers a ramp control pin (RC) for tracking applications or output voltage slew rate adjustment at turn-on. This is an important feature in applications where the load is highly capacitive and in-rush currents can cause supply voltages to fail and microprocessors or other complex logic chips to hang up. All power-up protocols, sequencing, window sequencing, tracking and ratiometric tracking can be implemented with minimal components, thereby eliminating the need for complex power management ICs.

For more information, contact your local Micrel sales representative or visit Micrel at: www.micrel.com/ad/mic68xxx.
Literature 1 (408) 435-2452 Information 1 (408) 944-0800

The Good Stuff:

- Input voltage range: 1.65V to 5.5V
- Output voltage as low as 0.5V
- Stable with small ceramic capacitor
- 1.0% initial output tolerance
- Output currents from 2A to 4A
- Power-on Reset (POR) supervisor with programmable delay time
- Programmable Ramp Control for in-rush current limiting and slew rate control of the output voltage
- Tracking on turn-on and turn-off with pin strapping
- Timing Controlled Sequencing On/Off
- Single Master can control multiple Slave regulators
- Daisy chain any combination and quantity of MIC68xxx regulators for synchronized, well controlled voltage supply
- Compatible with Z-One® Digital Power Architecture

MICREL®
Innovation Through Technology™

www.micrel.com

expect the LED to shine at full intensity based on a 20-mA current, this circuit is not guaranteed to work, just as Massa warned. Now, interface is an analog phenomenon, and, yes, you can usually drive an LED from an LS244 octal-bus driver. This approach may be acceptable for a consumer product, but not for a life-critical medical device or a military application; the part's specified limits do not guarantee that the LED will shine at required brightness—or any brightness at all.

Moving up the evolutionary chain from the simple interface transistor, you come to ICs that provide interface capability much like the aforementioned MOSFET-driver ICs. The 4000 series CMOS-logic hex-inverting CD4009 or noninverting CD4010 logic-buffer IC is one of the oldest of these interface ICs. Texas Instruments still makes these venerable parts. The 4000 series has the benefit of working at 15V, providing for remarkable noise immunity and assisting in directly driving interface signals. One benefit of CMOS logic, whether the 15V 4000 series or any low-voltage CMOS, is that you can tie multiple outputs together to increase the current drive of the circuit (**Figure 2**). The FET transistors are resistive when you turn them on, so they share current, whereas bipolar transistors in this configuration may hog current—that is, one bipolar output transistor becomes hot, causing it to conduct more current, in turn becoming even hotter, and finally becoming damaged. To prevent this situation, you can insert 5 or 10Ω resistors in series with each of the outputs, but this step is more costly and complex than using a CMOS chip.

The need for driving solenoids or other difficult loads has spawned entire series of interface ICs. Some of the most ubiquitous of these are the SN and DS series from Texas Instruments and other vendors. These parts range in variety from the SN55462 dual-high-voltage peripheral drivers to the DS3680 quad-telephone-relay driver. Also notable for this class of ICs is Freescale, a spin-off of Motorola, whose name implies the automotive radios that gave the company its start. As you would expect, Freescale makes interface ICs that can take the higher temperatures and voltage surges of the automotive environment. Freescale offers low-side switches; high-side switches; and H-bridge, predriver, power-train, and squib ICs. (Military designers will recognize squibs as the explosive bolts that blow to allow rocket stages to separate. In this context, a "squib" is the explosive canister that deploys automotive air bags.)

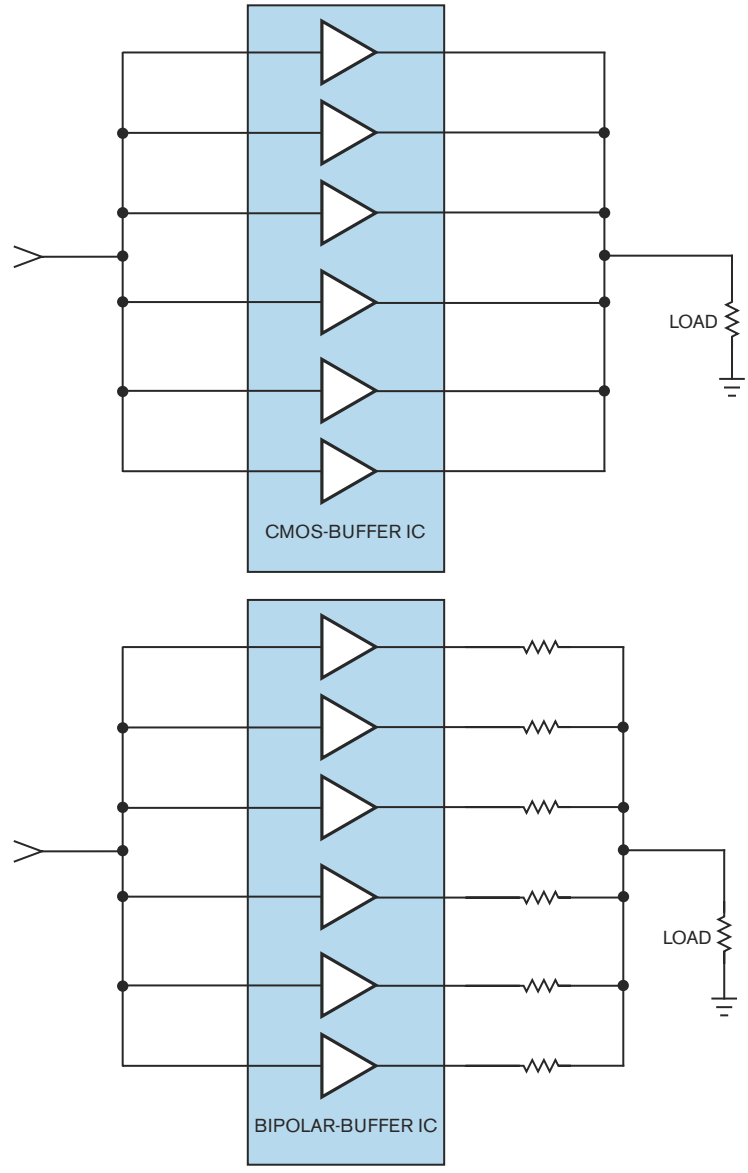

Figure 2 CMOS-family logic, such as CD4000 and 74HCxxx, can have connected outputs because MOSFETS share current. Bipolar-logic families would need low-value resistors in each output leg to prevent current from hogging the bipolar-output stages.

Speaking of high-side and low-side switches, it might seem logical to put a semiconductor switch on the low side of the load. The problem with low-side switching, though, is that when you turn off the switch, you leave the load floating at the positive-power-supply voltage. In a world that has been using simple mechanical switches for more than a century, you would expect to be able to simply switch off the power, rather than interrupt the ground side of the load. It may also be beneficial to keep the load grounded for noise reasons or to monitor currents, voltages, or temperatures. This problem has given rise to the use of high-side ICs.

One way to design a high-side switch

Create any kind of waveform

With the industry's best-selling function generators.

Get a FREE $444 hardware upgrade.

Agilent 33220A Function Generator

	33220A *20 MHz Sine & Square*	***33250A*** *80 MHz Sine & Square*
Basic Waveforms	*11 Built-in Low Distortion*	
Pulse	*Variable Edge 5 MHz*	*Variable Edge 50 MHz*
Arbitrary Waveforms	*14-bits, 64K points*	*12-bits, 64K points*
Modulation *(Internal / External)*	*AM, FM, PM, FSK & PWM*	*AM, FM & FSK*
Open Connectivity	*USB, GPIB & LAN*	*GPIB & RS232*
Price *(USD, Subject to change)*	*$1,853.00*	*$4,553.00*

From now through September 30, 2007:

Purchase an Agilent 33220A from an authorized distributor and get a free hardware upgrade worth $444. For a complete list of authorized distributors, visit www.agilent.com/find/functiongens

Authorized distributors

METRICTEST™
6,000 instruments. One source.
866-436-0887

Electro Rent Corporation
800-360-5688

Techni-Tool
800-832-4866

www.agilent.com/find/functiongens

Agilent Technologies

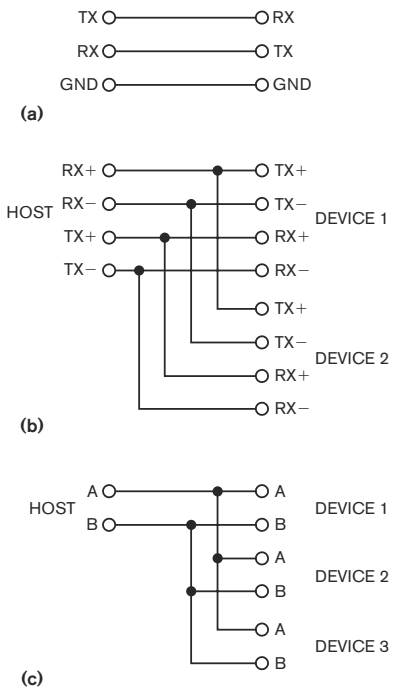

Figure 3 RS-232 started as a telecom standard to connect modems to computers (a). It used all 25 pins of a DB-25 connector. It has since evolved to a simple three-wire connection. The multidrop RS-422 standard can have 10 receivers (b), and RS-485 is a bidirectional bus with as many as 32 receivers and transmitters (c) (courtesy Omega Engineering).

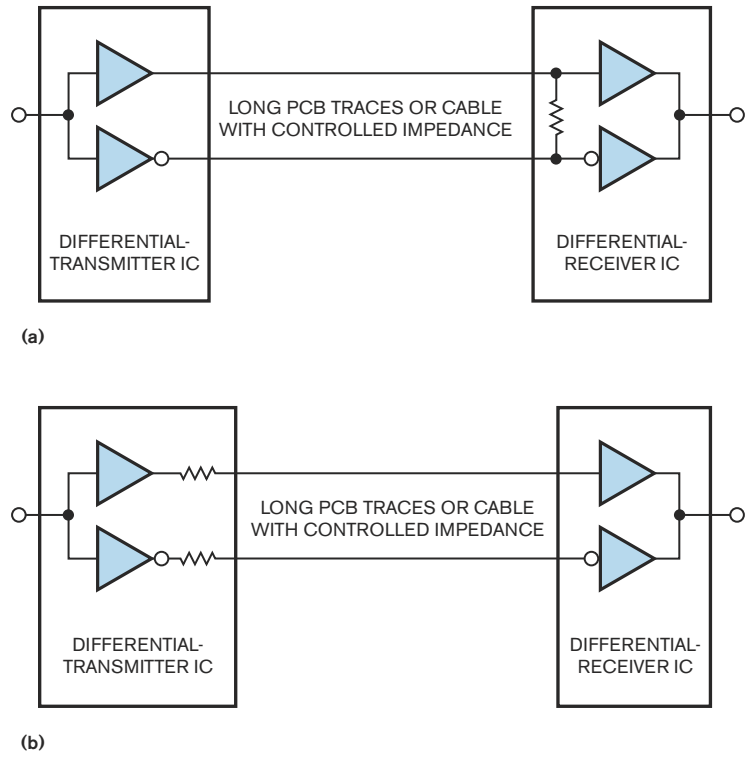

Figure 4 Differential signaling uses a pair of wires operating on opposite polarities. You can use a 100Ω shunt termination at the receiver or two 50Ω terminations in series to eliminate reflections (a). The series termination does not dissipate dc power (b). Differential signaling is the basis for LVDS, USB, SCSI, SATA, PCIe, CameraLink, and many other interfaces. The transmitter does not start and stop its operating current; it just switches it between the two wires, making decoupling the power pins less critical.

is to use PNP or P-channel FETs and live with the greater losses. Because electron mobilities are higher in N-type material, however, NPN-bipolar and N-channel MOSFETs can conduct current with lower losses. A more sophisticated technique uses a charge pump to boost voltage, which then turns on an N-channel MOSFET through a level-shifting circuit. One caution here involves gate capacitance. Data sheets may show a typical turn-on time for a charge-pump high-side driver, but manufacturers base that time on a specific load capacitance. If you try to control a large MOSFET, it takes a long time to charge the huge gate capacitance. As a result, the part may linger too long in the linear mode, causing excessive power dissipation. Just as seriously, the part may not quickly enough connect a power source, causing the system that the high-side switch is feeding to become unstable or to reboot.

One of the best known brands of high-side FETs is NXP's TopFET. This part integrates a power FET along with protection features for overcurrent, overtemperature, and overvoltage conditions. In addition, the part offers reverse-polarity protection, making it ideal for replacing both switches and fuses. NXP has minimized charge-pump noise to make the part suitable for a wide range of applications. Infineon makes the ProFET smart high-side switch with similarly robust features. This unit complements Infineon's line of low-side switches and motor controls. STMicroelectronics is also involved in difficult interface applications in the automotive environment. It has pioneered work involving minimizing electromigration and improving the thermal fatigue of high-side switches. This work brings designers ever closer to the ideal of a device with the robustness of silver contacts and the intelligence, fault protection, and monitoring of semiconductor devices.

SIGNAL INTERFACES

The term "interface" has grown to mean more than driving solenoids and other loads from logic voltages. It also refers to the several standards to connect digital systems. RS-232 was one of the earliest of these standards. Other early serial-interface ICs supported RS-422 and RS-485 (**Figure 3** and **references 4** and **5**). Whereas RS-232 has one wire for transmitting and one for receiving, RS-422 uses differential pairs for communications and needs two pairs

Why settle for capacitors when you can get a capacitance company?

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

21V, 8A, 600KHz
Synchronous Buck Regulator

42mΩ Internal Power MOSFET Switch

Only 8 External Components!

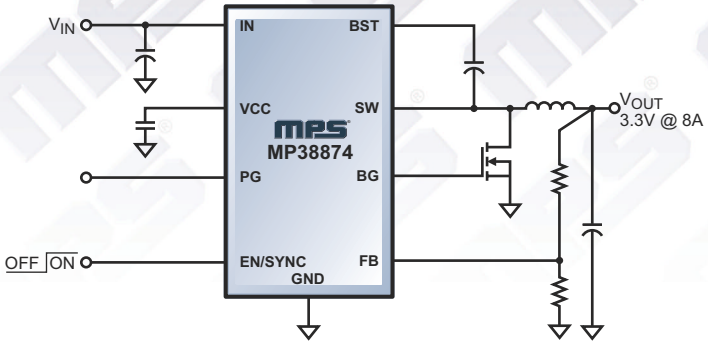

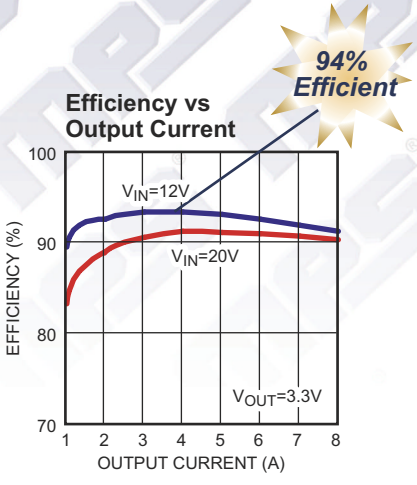

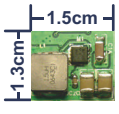

Actual Double Sided Footprint Solution Size

- *Up to 8A for Today's High Performance Digital Core*
- *Stable with Low ESR Output Ceramic Capacitors*
- *Internally Compensated*
- *Integrated Soft Start*

Featured (High-Current) Synchronous Bucks

Part	Frequency	V_{IN} (V)	V_{OUT} (V)	I_{OUT} (A)	Package
MP38872	600KHz	4.5 - 21	0.8 - 15	6	QFN 14 (3mm x 4mm)
MP38891	420KHz	4.5 - 30	0.8 - 15	6	
MP38874	600KHz	4.5 - 21	0.8 - 15	8	

DC to DC Converters CCFL / LED Drivers Class D Audio Amplifiers Linear ICs

ISO9001:2001 Certified

LEAD FREE RoHS COMPLIANT

Monolithic Power Systems®
The Future of Analog IC Technology®
www.monolithicpower.com

AVNET Memec

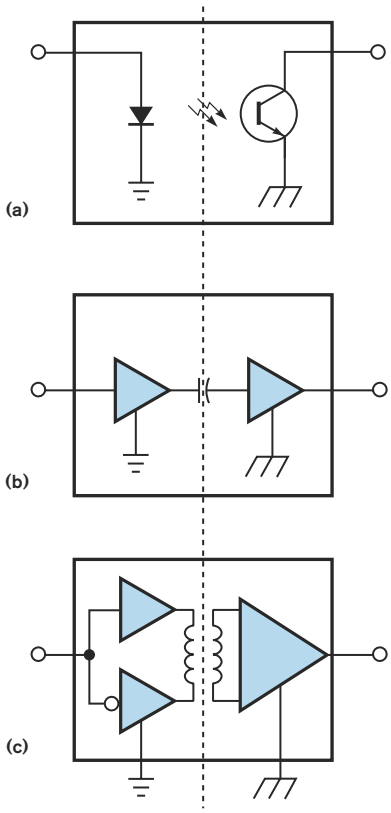

Figure 5 You can achieve digital isolation using optical means with chips such as these from Avago (a), capacitive means as in these chips from Texas Instruments (b), or inductive means as in these chips from Analog Devices (c).

to send and receive. The advantage is that it can work over longer distances—even kilometers—and can have as many as 10 receivers listening to one transmitter. The differential RS-485 interface combines transmission and reception on the same pair of wires. It can support 32 transmitters and receivers on the same bus.

One problem with older standards, such as RS-232, is that designers do not rigorously follow them. This lack of compliance applies to both the pins in the connectors and the voltage levels in the transmission. RS-232 started as a standard for connecting modems to computers. For this reason, it does not need such signals as RD (ring detect) to connect two computers or a computer to a peripheral drive. All the potential pin, cable, and voltage configurations have often caused designers to call the RS-232 standard the nonstandard standard.

In this regard, engineer Massa points out, many people design RS-232 interfaces for 5V operation rather than for the voltage the standard dictates: ±12V. The input-structure circuit diagram of a common RS-232 receiver, he notes, takes only one diode drop—0.6V—to turn on the input, so designers believe that 5V would do this job as well as 12V would. When the signal line goes to 0V instead of −12V, they reason, the input would turn off. Massa disagrees, however. "Remember: You have violated the standard by running the interface at 5V," he says. "You will have far worse noise immunity, and the length of cable runs you can achieve will be far less. Nevertheless, if all you are doing is trying to get a connection to a PC over two feet of wire, a 5V signal can function on an RS-232 line." Maxim and other companies provide a benefit by selling RS-232-interface chips that operate at 5V, use internal charge pumps to generate ±10V or higher, and provide 15-kV ESD (electrostatic-discharge) protection on the line. Having a serial interface that conforms to the standard and having fault protection to boot can provide considerable peace of mind to a design engineer.

You can gauge the tremendous popularity and proliferation of these early interface standards by examining Texas Instruments' offerings in this area. The company offers 63 RS-422 parts ranging from dual-line drivers selling for approximately 25 cents to parts with three differential-transceiver channels that cost several dollars. TI's RS-232 line comprises 95 parts, starting with drivers costing less than a quarter, to complete interface ICs, including charge pumps and ESD protection, selling for more than $4. For the RS-485 standard, TI offers 101 parts, with a broad range of capabilities, all with commensurate prices.

Faster standards, such as USB (Universal Serial Bus) and Apple-developed FireWire, now augment these older RS standards. These bus standards operate at higher speeds and with more sophistication than the RS-485-bus standard. USB and FireWire tout not only standardized electrical interfaces, but also standardized higher level protocols. The old serial-communication port still ex-

SIMULATION SOFTWARE
FOR HIGH-PERFORMANCE ELECTRONIC DESIGN

ANSOFT

ists as a protocol in USB, even though the electrical spec differs. Despite the modernity of these standards, they face the same environmental challenges as the older RS standards face. The fact that a long cable is extending from your digital system creates the same potential for havoc. The cable can receive EMI (electromagnetic interference) and RFI (radio-frequency interference) and may cause your system to malfunction, or the cable can radiate digital noise from your system. ESD from people or other systems also can damage your system unless the interface IC protects against this phenomenon. As you would expect with such popular standards, a surfeit of vendors makes USB- and FireWire-interface chips. USB also delivers power, and vendors offer interface circuits to address the need to manage and protect the power lines in the USB port. For example, Raychem and Fairchild offer fuses and switches, respectively, to handle this problem.

The need to transmit faster data rates between chips and subsystems has caused designers to consider differential signaling (**Figure 4**). By using differential pairs, designers can maintain impedance control and balanced currents over large distances. The SCSI (Small Computer System Interface) disk-drive standard, which Shugart Associates developed in 1979, exemplifies this technique. SCSI became a standard in 1986, and Apple Computer used it in the Apple II (**Reference 6**). Texas Instruments lists SCSI chips under the "interface" section of its Web site. The utility and advantages of differential signaling have caused developers to adopt it in SATA (Serial Advanced Technology Attachment) disk-drive interfaces and the PCIe (Peripheral Component Interconnect Express)-bus interface. Many of these differential interfaces use a variation of LVDS. The industry first applied this term to the signals in the ribbon cable that connects laptop-LCD panels to the video chip. The video interface had to reside on a ribbon cable because the laptop screen had to tilt and close. A further benefit of differential signaling is that it radiates less electrical noise because the signals are differential and closely coupled. This feature allowed the fast transmission of digital video over a ribbon cable in a laptop product that had to pass strict Federal Communications Commission standards for radiated noise in consumer electronics. National Semiconductor was an early champion of LVDS and also used it in the CameraLink standard, which connects industrial cameras to video-capture hardware (**Reference 7**). SCSI, SATA, PCIe, and LVDS do not connect logic to a high-voltage peripheral, instead connecting two digital systems at high speeds over long distances. CameraLink can work over distances of meters, and LVDS cables can run for hundreds of meters if you properly shield and design them. National Semiconductor has an impressive portfolio of LVDS chips, and many other vendors, such as TI and STMicroelectronics, also make these products.

BY USING DIFFERENTIAL PAIRS, DESIGNERS CAN MAINTAIN IMPEDANCE CONTROL AND BALANCED CURRENTS OVER LARGE DISTANCES.

A special type of LVDS interface, SERDES (serializer/deserializer) serializes parallel data, such as a computer bus or several bytes of video data, into a high-speed LVDS pair. These chips often operate at speeds in excess of 1 Gbps. Once the SERDES serializes the parallel data and sends it over the pair of differential wires, the deserializer extracts a clock from the serial-bit stream and converts the data back into a parallel bus. National Semiconductor, TI, Intersil, Fairchild, and a host of other companies make these types of chips. SERDES chips are finding applications in cell phones in which screen and camera data must traverse the hinge of the phone. This task is easier with one pair of wires as opposed to a parallel bus.

Another class of interface chip supports isolated communications using optical, capacitive, or magnetic isolation to allow digital signals to communicate to systems without conducting ground loops or hazardous voltages (**Figure 5**). The

Intersil Real-Time Clocks

High Performance Analog

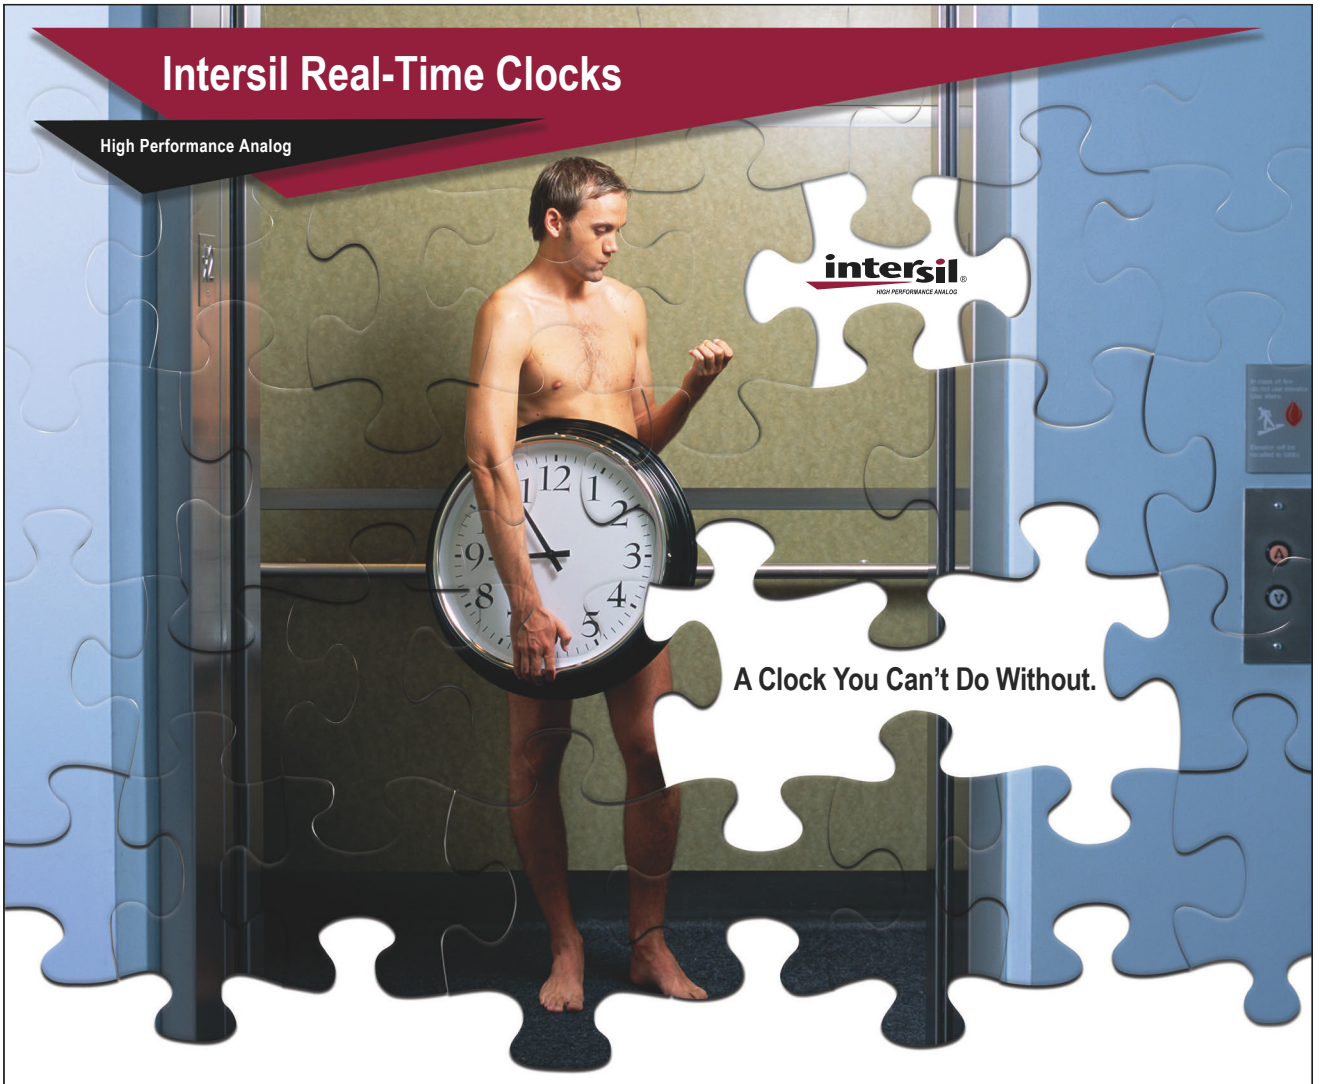

Real-Time Clock with unique ID saves cost and board space.

Intersil's ISL12025 is the only RTC with 4k-bits of EEPROM and embedded 64-bit unique ID that combines on-chip oscillator compensation and CPU supervisory functions.

Unique ID

- Factory-programmed 64-bit unique ID that can be read via I²C. Tamper-proof with no write access.
- Ideal for security-related and warranty-tracking applications.

Real-Time Clock/Calendar

- Tracks time in hours, minutes, and seconds
- Day of the week, day, month, and year

Automatic backup to battery or SuperCap

On-Chip Oscillator Compensation

- -64ppm to +110ppm frequency trimming range
- Internal feedback resistor and compensation capacitors
- 64-position digitally controlled trim capacitor
- Six digital frequency adjustment settings to ±30ppm

CPU Supervisory Function (ISL12025)

- Five selectable V_{TRIP} levels from 2.63V to 4.64V

Integrated 512 x 8-bit

- 16-byte page write modes (32 total pages)
- Eight modes of BlockLock™ protection
- >2,000,000 write cycle and 50 years endurance reliability

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

Intersil – Amplify your performance with advanced signal processing.

intersil®
HIGH PERFORMANCE ANALOG

MORE AT EDN.COM

Go to www.edn.com/070719df and click on Feedback Loop to post a comment on this article.

developers of the MIDI (Musical Instrument Digital Interface) standard for musical instruments based the technology on optical isolation. Avago Technologies offers optical-isolation chips that withstand thousands of volts and can transmit megabits per second of data. Both Texas Instruments and Analog Devices have chips that can provide thousands of volts of isolation and transmit signals in the hundreds of megahertz. The availability of these remarkable chips is changing the nature of some analog design. Previously, it was more common to achieve isolation in the analog section, such as with the Linear Technology LTC1531 isolated comparator. With the availability of these digital-isolator chips, it is now common practice to use an isolated regulator to power the analog section and employ a data converter on the isolated side. You can isolate the digital signals from this converter with a digital-isolation chip.

Interface issues are important; pay attention to them. A plethora of approaches is available from a large number of suppliers. If you consider how complex it is just to hook a transistor to a logic output, you can appreciate the need and desirability of using a specialized interface chip to do the job. If you are sending gigabit signals over the length of a large PCB (printed-circuit board) or onto a cable, you must use some of these sophisticated interface chips. The combination of capabilities and protection features is essential for most applications. It is a cruel world out there, and these interface chips can make your design far more resistant to many indignities. EDN

REFERENCES

1. Israelsohn, Joshua, "Miller on edge: The role of Miller capacitance in nonlinear circuits," *EDN*, March 29, 2007, pg 38, www.edn.com/article/CA6426883.
2. "SN54LS240, SN54LS241, SN54LS244, SN54S240, SN54S241, SN54S244, SN74LS240, SN74LS241, SN74LS244, SN74S240, SN74S241, SN74S244 Octal buffers and line drivers with 3-state outputs," Texas Instruments 74LS244 data sheet, February 2002, http://focus.ti.com/lit/ds/symlink/sn74ls244.pdf.
3. "5306 Series T-1 3/4 Right Angle LED Assembly," Chicago Miniature Lamp Inc, http://rocky.digikey.com/WebLib/Chicago%20Miniature/Web%20Data/5306%20Series.pdf.
4. "RS-232, RS-422, RS-485—Differences and Typical Applications," Omega Engineering, www.omega.com/techref/das/rs-232-422-485.html.
5. "Quick reference for RS485, RS422, RS232, and RS423," www.rs485.com/rs485spec.html.
6. SCSI, http://en.wikipedia.org/wiki/SCSI.
7. "Camera link," www.siliconimaging.com/ARTICLES/CameraLink.htm.

FOR MORE INFORMATION

Allegro MicroSystems
www.allegromicro.com

Analog Devices
www.analog.com

Apple
www.apple.com

Avago Technologies
www.avagotech.com

Chicago Miniature Lamp Inc
www.cml-it.com

Digi-Key
www.digikey.com

Fairchild Semiconductor
www.fairchildsemi.com

Freescale Semiconductor
www.freescale.com

Houston Instruments
http://en.wikipedia.org/wiki/Houston_Instruments

Infineon Technologies
www.infineon.com

Intersil
www.intersil.com

Linear Technology
www.linear.com

Massa Enterprises
www.massaenterprises.com/

Maxim Integrated Products
www.maxim-ic.com

Microchip Technology
www.microchip.com

Motorola
www.motorola.com

National Semiconductor
www.national.com

NXP Semiconductors
www.nxp.com

Omega Engineering
www.omega.com

Quantum
www.quantum.com

Raychem
www.raychem.com

Shugart Associates
http://en.wikipedia.org/wiki/Shugart_Associates

STMicroelectronics
www.st.com

Supertex
www.supertex.com

Teledyne Technologies
www.teledyne.com

Texas Instruments
www.ti.com

You can reach Technical Editor Paul Rako at 1-408-745-1994 and paul.rako@reedbusiness.com.

BY MICHAEL SANTARINI • SENIOR EDITOR

Is FPGA a simpler puzzle for ASIC designers?

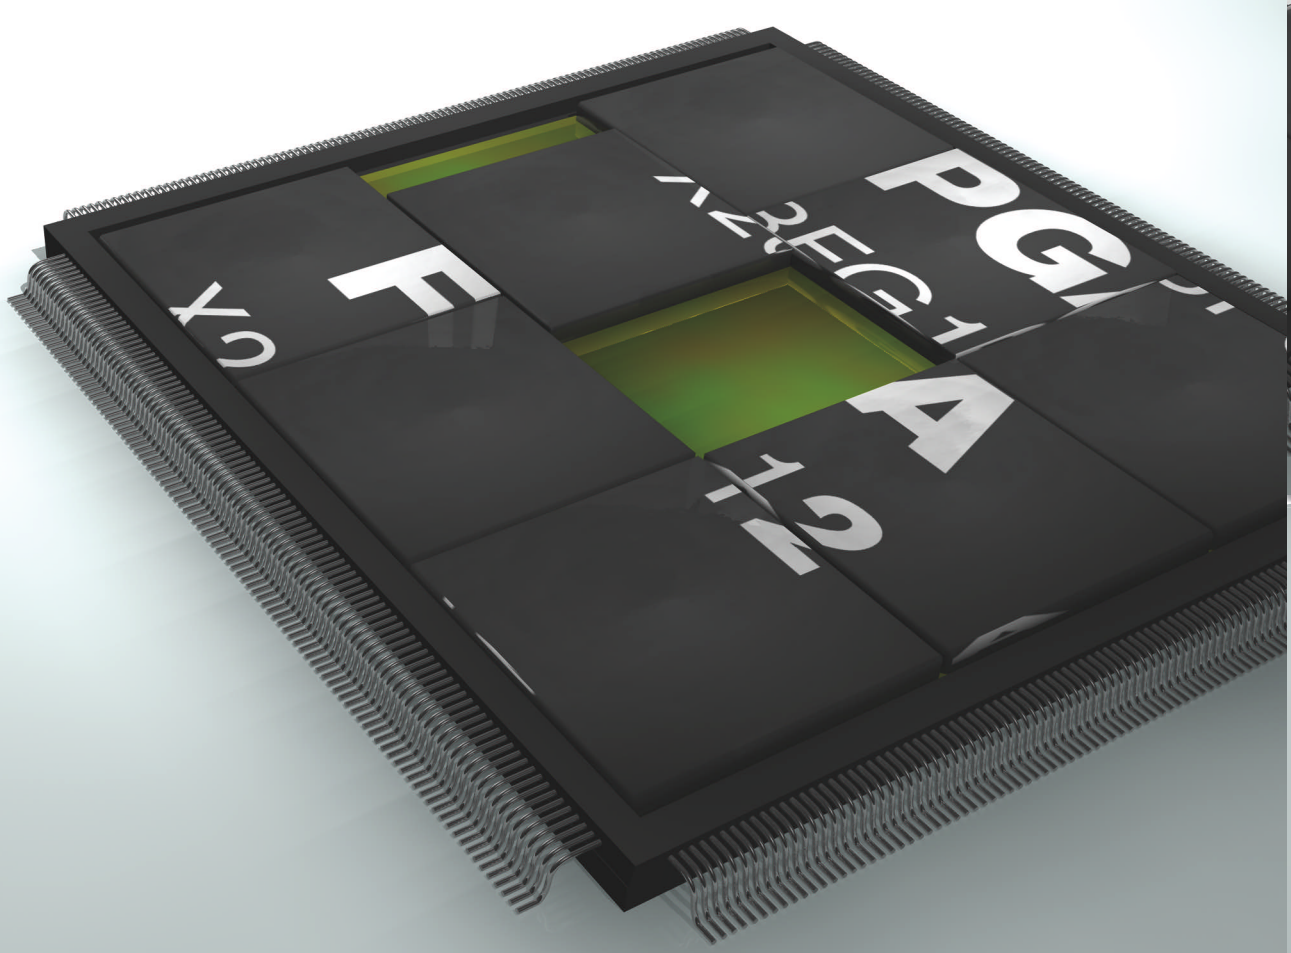

WITH RISING MASK COSTS, COMPLEXITY, AND TOOL EXPENSES TO DEVELOP ASICs AND SOCs, MANY DESIGN GROUPS TODAY ARE OPTING TO IMPLEMENT THEIR PRODUCTION DESIGNS IN FPGAs. BUT, BEFORE DESIGNERS MAKE THE LEAP, THERE ARE SEVERAL FACTORS—GOOD AND BAD—THEY SHOULD CONSIDER.

Over the last 10 years, FPGA vendors have made great strides in overcoming the shortcomings of FPGAs and taking share from the ASIC market. In the late 1990s, FPGA vendors increased the capacity of their devices to rival midsized ASICs. Then, circa 2001, FPGA vendors improved the performance of their devices to compete with midsized ASICs. Though FPGAs still consume much more power than ASICs of comparable densities and performance, last year, FPGA vendors made great strides to stabilize the amount of power that FPGAs consume (**Reference 1**).

On top of these device-attribute gains, FPGAs have also come down in price. Vendors such as Actel, Altera, Lattice, Quicklogic, and Xilinx offer a wide swath of devices ranging from pennies-per-part CPLDs (complex programmable-logic devices); to secure, nonvolatile FPGAs; to high-performance, high-LUT (look-up-table)-count, SRAM-based juggernaut FPGAs that cost thousands of dollars each.

In the early days of the FPGA industry, designers would use the most expensive and highest grade FPGAs mainly to functionally prototype code that they planned to implement in an ASIC or to do proofs of concept for system designs. They would create the logic for their ASICs, run verification, partition their ASIC designs, and then program those partitions onto several FPGAs residing on a prototyping board (**Reference 2**). Designers today still use this method, but, because FPGAs have improved on all fronts, many designers are using even the highest grade and most expensive FPGAs for production parts.

It's fairly easy to get a marketing executive at an FPGA vendor to tout the glory of FPGAs and how they are taking over ASIC real estate—even for production use. Designers are also stepping forward to say that FPGAs are truly viable production vehicles and that designers should no longer dismiss them simply as ASIC-prototyping tools.

Sanjay Singh, the technical lead in ASIC/FPGA design at Hewlett-Packard's nonstop-computer division, has over the course of his career designed 10 FPGAs and 25 to 30 ASICs. He started designing ASICs with Tandem Computer, which Compaq acquired in 1996; HP in turn in 2002 acquired Compaq. "I started at 0.5 microns doing ASICs for Toshiba nonstop computers," says Singh. Currently, his group is designing ASICs at 110- and 90-nm nodes, and, when the situation calls for it, his group uses the highest grade SRAM-based FPGAs—Altera Stratix- and Xilinx Virtex-class devices—for server applications.

"Our systems are based on Intel Itanium-server chips, and the FPGAs we design have to talk to memories, I/Os, and processors," says Singh. "Our value added is in hardware, so we have to do communications, porting functions, data-integrity functions, and duplicated functions. The FPGAs we design typically go on a

processor board, and the end systems will run upward of a million dollars."

On the other end of the design spectrum, Ranjit Rozario is a senior design engineer at communications start-up Sonoa Systems. Rozario is one of a few hardware designers in a 100-person company that comprises mostly software engineers. A long-time ASIC designer, Rozario recently took his first foray into FPGA design, ultimately choosing a Virtex-5 LX 220.

Both Singh and Rozario say they are more frequently going with FPGAs for various reasons but note there are advantages and disadvantages of designing with FPGAs. So, when making the move, ASIC designers must take into account several factors, including design size, performance and power budgets, PCB (printed-circuit-board) requirements, design and verification requirements, and FPGA-tool limitations. SRAM FPGAs also introduce new challenges, such as soft errors, which are more commonplace in SRAM fabric than in standard-cell material.

WHY MOVE TO FPGA?

There are several reasons that designers choose FPGAs over ASICs: FPGAs are reprogrammable and field-upgradable, and the design cycles are shorter than ASIC-design cycles (**Figure 1**); FPGAs have better prices for high-cost, low-volume applications; and they are relatively stable, so you can avoid respins, mask costs, and buying DFM (design-for-manufacturing) tools.

AT A GLANCE

- Designers should heed the rule of 80%: Buy an FPGA with 20% more LUTs (look-up tables) if you expect to hit high-performance targets.
- The biggest FPGAs can run at 550 MHz.
- The biggest FPGAs have 330,000 logic cells, or roughly 12 million equivalent-ASIC gates.
- FPGA-vendor synthesis tools are less efficient than commercial EDA FPGA-synthesis tools.
- When selecting an FPGA, observe the layout and hard-wired macro configuration; it can present layout challenges that affect timing.

But Singh says that his group has two main reasons to use FPGAs. The first is that they allow his group to take a lot of functions off the PCB and integrate them onto an FPGA to speed performance and save PCB space. The second and most convincing reason to use an FPGA, according to Singh, is simply that ASICs' unit volumes sometimes don't justify the mask, design, and tool costs, and, most of all, the risk (**Figure 2**). "FPGAs have evolved, and, in many cases, they can meet your performance and density requirements," says Singh. "If you are designing a midrange, midclass ASIC, you should ask yourself, 'Do I want to spend $2 million to $3 million to do a 90- or 65-nm ASIC, or do I want to use a 90- or 65-nm FPGA technology with a comparable benefit?'"

Singh notes that his group tends to favor FPGAs over ASICs for designs that will have a new architecture and do not use a lot of blocks from previous designs. Because FPGAs are reprogrammable, the group can try new architectures and simply reprogram the FPGA when changes are necessary.

Similar to Singh's, Rozario's group also uses an FPGA rather than an ASIC, primarily for cost reasons. "When you're a start-up and money's tight, one of the first things to do is to look at an FPGA, because the development costs are a lot less and there aren't mask costs," says Rozario.

Rozario says that his company wanted to integrate and speed up some of the software functions in its product by implementing many of those functions in hardware on a single FPGA for its next-generation offering. "When we started, we really didn't know what functions we wanted to move to a chip and what functions needed to speed up," he says. He likes FPGAs because his group could add and subtract functions on the FPGA that they needed to speed up later in the process. "One of the great things about an FPGA is [that] you have the option of changes down the line."

Although the choice to go to an FPGA was fairly easy, however, Rozario and Singh say that figuring out how to design with an FPGA requires a bit of learning. Both engineers say that, after you decide to use an FPGA, the next big step is selecting the right one.

DETERMINE FPGA NEEDS

Designers need to look at what FPGA families are available from the various vendors and find the right mix of performance, power, and density. But buyer beware: One of the first things designers need to know when moving from an ASIC to an FPGA is that, to target a performance grade, you should buy a device with 20% more density than you need.

In most of the 10 FPGAs Singh has worked on so far, he says, he opted for

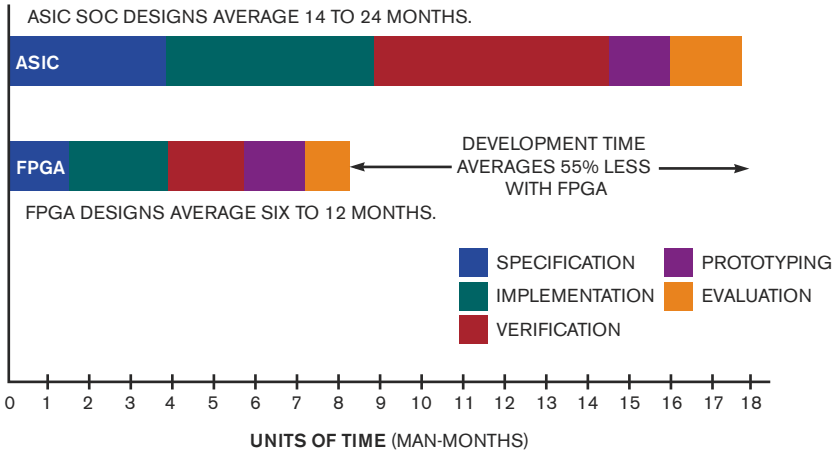

SOURCE: *EXPERIMENTATION MATTERS: UNLOCKING THE POTENTIAL OF NEW TECHNOLOGIES FOR INNOVATION,* STEFAN H THOMKE, HARVARD BUSINESS SCHOOL PRESS, 2003. BASED ON A HARVARD STUDY OF 400 DESIGNS.

Figure 1 One of the overlooked facets of FPGAs is their relatively fast development.

NI LabVIEW.

Limited Only by Your Imagination.

Build and program robots with LEGO® MINDSTORMS® NXT using software powered by NI LabVIEW

Communicate via multiple protocols including Bluetooth

Graphically program concurrent, real-time applications

Develop your human machine interface (HMI) display

Independently control multiple servo motors

Target 32-bit microprocessors and FPGAs

Real-Time and Embedded | Signal Processing | High-Performance Test | Industrial Control

PRODUCT PLATFORM

LabVIEW Real-Time Module

LabVIEW FPGA Module

LabVIEW Embedded Development Module

NI CompactRIO Embedded Hardware Platform

When the LEGO Group needed parallel programming and motor control tools intuitive enough for children, it selected graphical software powered by NI LabVIEW. With LabVIEW graphical system design, domain experts can quickly develop complex, embedded real-time systems with FPGAs, DSPs, and microprocessors.

>> Expand your imagination with technical resources at **ni.com/imagine** **866 337 5041**

NATIONAL INSTRUMENTS

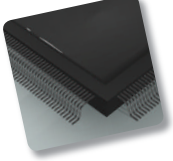

FPGAs because they fit his company's volume and performance needs. "In all those cases, we had to push the designs to the edge with respect to frequency, I/O time, and usage," he explains. A rule of thumb his group uses is that, if you are using 60 to 75% of the design's overall real estate, small changes from one compilation to another will give you a good chance of meeting your performance requirements and doing well in the lab. However, he notes, if you cross the 85%-usage mark, you may be unable to obtain the performance you need.

Singh says that his group has designed FPGAs with as much as 95% usage and reached performance goals, but doing so took a lot of work. "You have to be very familiar with the FPGA and how it works," says Singh. "In the ASIC world, you can write TCL (Tool Command Language) scripts and query the database to get the stuff you want, but, in the case of FPGAs, the tools are not as mature. They are typically GUI [graphical-user-interface]-based."

Rozario says that, when he first shopped around for FPGAs, the fact that vendors had improved the devices to meet most performance and density grades impressed him. However, FPGAs still don't reach the same top speeds or density grades as ASICs. The top-of-the-line SRAM-based FPGAs, when you scale them back in functions and optimize them to the fullest, top out at 550 MHz; ASICs can double that performance at top speeds. The biggest commercially available 65-nm FPGA, the Xilinx Virtex-5, has 330,000 logic cells, or roughly 12 million equivalent-ASIC gates, according to Xilinx.

Rozario was initially worried about FPGA-performance limitations. But the speed and density of the FPGAs surprised him. "If you are used to designing an ASIC, your expectations for performance will definitely have to come down in an FPGA," he cautions. Rozario notes that, in the first design, his group used 80% of the Virtex-5 LX 220 real estate and hit the design project's performance goals.

FPGA vendors today typically offer application-specific variants of their devices. A given FPGA-product family typically features a traditional sea-of-gates FPGA and variants that target specific markets. Some devices target network applications and contain hard-wired SERDES (serializer/deserializer) cores; other devices target communications applications and contain hard-wired DSP blocks. All contain fairly large amounts of memory. Xilinx, for example, offers the Virtex-5 LX for high-performance logic, the Virtex-5 LXT for high-performance logic with serial connectivity, the Virtex-5 SXT for high-performance DSP with serial connectivity, and the Virtex-5 FXT for embedded processing with serial connectivity.

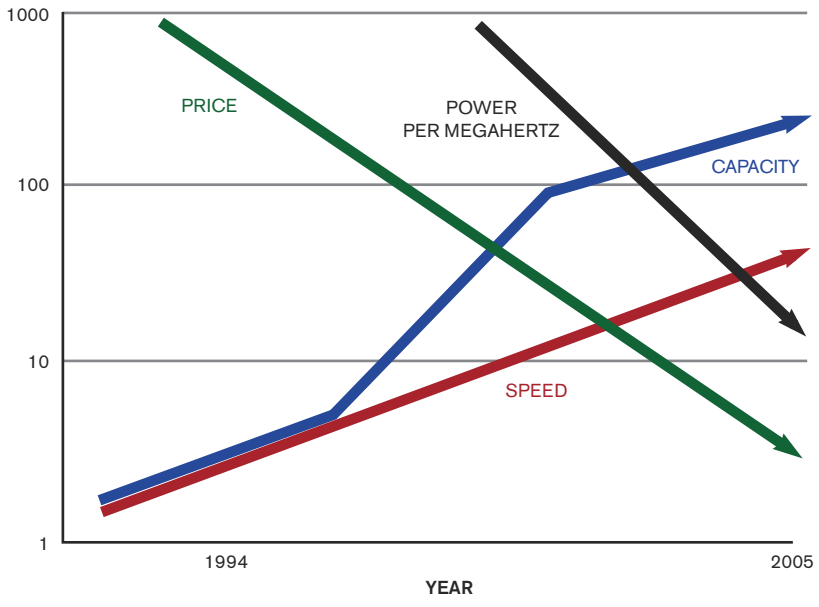

SOURCE: XILINX HISTORICAL DATA.

Figure 2 FPGAs have improved on all fronts: capacity, power, performance, and cost.

Singh and Rozario say that choosing the right device for your application is imperative, because FPGAs with unneeded hard-wired cores can consume real estate and may cause layout roadblocks later in the design cycle that could hinder you from achieving your performance goals.

Neither designer is targeting an application in which power is a primary concern, but Singh says that design groups designing for low power have to take the power-consumption issues of FPGAs into consideration, even though FPGA vendors have been making strides to control overall power consumption and leakage in the 90- and 65-nm nodes.

The only step Singh's group has taken to combat the power problem, he says, is shutting down hard-wired, 5-Gbps SERDES with the clock-gating technique. "The new FPGAs deal with low power fairly well, and there are a lot of techniques you can employ to lower power, but we haven't had to use them, yet," he says. Of course, that situation could change as server applications begin to rely on power savings as a selling point.

ASICs of the same performance and node generally consume much less power than FPGAs, though FPGA vendors are making progress on that front. Xilinx and Altera claim to have stabilized leakage power so that their 65-nm devices leak no more power than their 90-nm devices do.

FITTING SYSTEM NEEDS

In addition to ensuring that FPGAs meet performance, density, and power goals, designers also must consider the FPGA's impact on the IC package and the PCB. FPGA chips typically consume more area on a PCB, and their dense I/O often requires designers to add more layers to a PCB to handle the routing of this I/O. They also need more diligent

Looking for the world's largest inventory of Microchip products?

How much are you paying to program your PIC® microcontroller?

microchipDIRECT can also program code in your purchased Microchip devices. For details, visit our web site.

Sample Pricing

Example Device	Package Type	Quantity Programmed	microchipDIRECT Programming Charge
PIC10F202-I/OT	6-Lead SOT-23	5001+	$0.02 each
PIC16F630-E/SL	14-Lead SOIC	5001+	$0.08 each
PIC16F627A-I/SS	20-Lead SSOP	5001+	$0.10 each
PIC16F57-I/SP	28-Lead PDIP	5001+	$0.12 each
PIC18F4520-I/ML	44-Lead QFN	5001+	$0.20 each
PIC18F8722-I/PT	80-Lead TQFP	5001+	$0.23 each

microchip DIRECT
www.microchipdirect.com

Your most comprehensive resource for Microchip products.

MICROCHIP
www.microchip.com

Microcontrollers • Digital Signal Controllers • Analog • Serial EEPROMs

package and PCB-signal-integrity analysis and enough room to accommodate the extra power-supply circuitry to properly power the FPGA. All of these requirements can add costs to the design cycle and ultimately the end product.

Rozario notes that, for this design, the group had no other devices running at 1V, so, to accommodate the FPGA, it had to put an extra power brick on the PCB to power the device. This step was no problem, he notes, because the PCB has little more than the FPGA on it.

Beyond picking a part that meets performance, power, and density goals for your system, you should also look at what tools FPGA vendors and independent EDA vendors offer (**Figure 3**). Years ago, Altera encountered a significant setback when its Quartus development suite had usability issues. The company has since fixed those problems, but users must heed tool availability and quality.

TOOLS AND DESIGNS DIFFER

Both Singh and Rozario say that, although FPGA-vendor tools are fairly good for pushbutton use and typically free if you buy enough silicon, they are less sophisticated than ASIC tools. That is, you can't manipulate them as much as ASIC tools to perform custom tasks. In particular, Singh says, FPGA vendors offer decent compiler or synthesis technology, but the FPGA-vendor synthesis tools do not implement designs as well as or as efficiently as commercial FPGA-synthesis tools, such as those from Synplicity, Mentor Graphics, and Magma Design Automation. Some FPGA vendors act as OEMs for commercial EDA tools and offer them to customers in a "lite" configuration, targeting only that vendor's FPGAs, at little or no cost. In most cases, those tools are better than the FPGA vendor's tools but lack the features of the commercial, full-priced version of the EDA vendors' tools. Also, in the FPGA world, for the most part, users have no choice but to use their

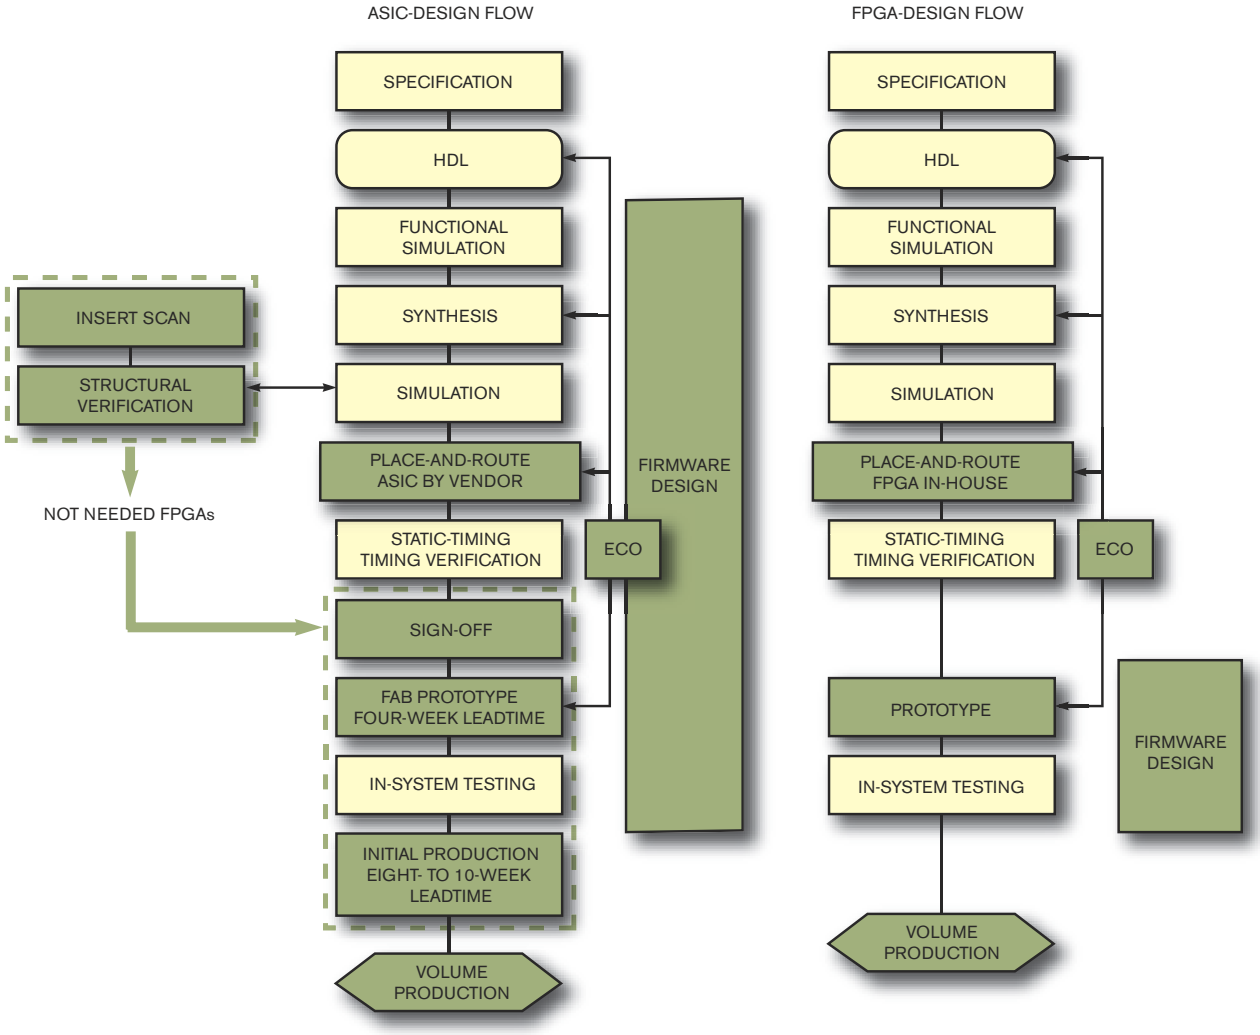

Figure 3 FPGA-development flows are far less complex and thus less expensive than ASIC flows.

Control Your Remote with Low-Cost 16-bit MCUs

Whether your design is a home theater remote or any other embedded application, the PIC24F delivers the peripherals, performance, development tools and software you need.

With 16 MIPS performance and an extensive peripheral set, Microchip's PIC24F microcontrollers are highly cost-effective solutions.

GET STARTED IN 3 EASY STEPS!

1. **Learn all about our 16-bit products in only 20 minutes!** Our *FREE* 16-bit web seminars highlight our 16-bit architecture and comprehensive support suite that includes low-cost development tools, a variety of software libraries (many free), application notes, reference designs and more!
2. **Take advantage of *FREE* resources for download!** Our MPLAB® IDE Integrated Development Environment and a full-featured trial of the MPLAB C30 C Compiler from our web site gets you up and running fast!
3. **Get hands-on technical training for *as low as* $49!** Attend a 16-bit Embedded Control Solutions class at a Microchip Regional Training Center to take your designs to the next level!

Purchase and program your 16-bit PIC24F devices and related development tools at...

microchip DIRECT
www.microchipdirect.com

Device	Pins	Flash (KB)	PIC24F Family Features
PIC24FJ64GA006	64	64	8 KB RAM Parallel Master Port 5 - 16-bit Timers 5 - Output Compare/PWM 5 - Input Captures Real-Tme Clock Calendar 2 - UART with IrDA® and LIN Protocols 2 - SPI, 2 - I²C™ 10-bit ADC, 16 Channels 2 Analog Comparators
PIC24FJ64GA008	80	64	
PIC24FJ64GA010	100	64	
PIC24FJ96GA006	64	96	
PIC24FJ96GA008	80	96	
PIC24FJ96GA010	100	96	
PIC24FJ128GA006	64	128	
PIC24FJ128GA008	80	128	
PIC24FJ128GA010	100	128	

Visit our web site for more information about additional 16-bit devices with higher performance and added features like DSP!

MICROCHIP

www.microchip.com/PIC24

Microcontrollers • Digital Signal Controllers • Analog • Serial EEPROMs

FPGA vendor's internally developed physical-design tools. FPGA vendors develop their own physical-design tools to help users take advantage of their FPGA architectures. But, Singh and Rozario say, like ASIC synthesis, the layout tools are rarely as sophisticated as equivalent-ASIC tools.

According to Singh and Rozario, layout of an FPGA can be tricky because many FPGAs have fixed macros—SERDES, RAM, PLLs (phase-locked loops), and DSP cores—and some have fixed microprocessor blocks.

Altera, for example, offers the Stratix and Stratix GX lines. Singh says that the two FPGAs are similar, but the GX includes hard-wired SERDES blocks. Once you locate the RAM, PLLs, and dedicated I/O blocks, he explains, you need to design from the bottom up: Obtain the pinout, package, and macros on the die, and then do your planning and implementation. Challenges with I/O technologies, SSTL (stub-series-terminated logic) for DDR, and HSTL (high-speed-transceiver logic) for PCI have taught Singh the importance of due diligence, knowing what each device offers and what you need to do to design with it.

Singh also notes that, although ASICs have multiple clock resources, FPGAs are more limited and typically have global clocking. "If you have an ASIC design and you are thinking of moving to FPGA but it contains multiple clock domains—particularly, large clock domains—then you have to work with the FPGA vendor [to determine] whether the device can handle your design," says Singh. Besides global clocking, he says, FPGAs have local clocks. But, Singh warns, those clocks are limited to certain quadrants, so be careful about your logic and its clocking. Singh's designs employ multiple clock domains, typically 15 to 20, which require a great deal of work to implement in FPGAs. "With ASICs, you can tune your I/Os," he explains. "FPGA I/O is complex, and tuning it can tax your overhead because it's difficult to tune slew rates, drive strength, and impedance." According to Singh, signal integrity and clocking can both be big issues if you fail to properly tune I/O from the beginning of the design.

Rozario's group experienced similar issues with layout. Because his design

"WE DO FUNCTIONAL SIMULATION, GATE-LEVEL SIMULATION, AND DYNAMIC SIMULATION, TAKING THE GATES ALONG WITH PARASITICS BACK INTO FUNCTIONAL SIMULATION."

includes relatively few clock domains, however, the clock structures were simpler and sufficient for his group's design. "FPGAs come with built-in clocks throughout the fabric, so clock balancing was a lot simpler than with an ASIC," Rozario says. "You don't have to worry about the H-tree signal integrity because the FPGA architecture has taken care of all the clocking, and you can have a large number of clocks per region."

Rozario warns that a poorly planned layout can result in timing problems. "Some of the larger FPGAs have functions that you don't necessarily need, and you have to work around them. The device we are using has hard-wired DSP functions. We aren't using a DSP in this design, but we may use it in the next one." Once his group established the optimum layout, the Xilinx ICE (in-circuit-emulator) tools did a "fairly good job" of adapting the design to the architecture. Initially, his group would floorplan every FPGA submodule and bring it up to the top level and then repeat that method for the entire chip. But ICE-application engineers would caution that the group's method was against their recommendation and urge the group to use automated tools. Rozario was skeptical but found that the results for both methods were similar.

Although FPGA tools are less customizable than ASIC-design tools, Singh encourages ASIC designers moving to FPGA design to take advantage of the free FPGA-vendor tools, IP (intellectual-property) portfolios, and customer support. He cites the vendors' good program management, architects, and design specialists that understand FPGAs, as well as field support. He finds the tools valuable for groundwork, such as floorplanning and signal integrity.

VERIFICATION ISSUES

When FPGAs first hit the market, one of their big selling points was that designers could program them and then directly test them on a board running in a mockup of the system, skipping simulation-based logic verification. Many FPGA veterans still use this approach. But Singh and Rozario say that, even midrange FPGAs are now generally too large and complex to simply program, run on a board, and debug during system test and bring-up.

Many veteran FPGA designers have told Rozario that they design the chip, skip logic verification, and test the chip in the lab. Rozario doesn't support that method, because, although the outsides of his chips have standard bus interfaces, the insides are complex, and debugging without visibility inside the chip is difficult. "It's a lot easier having waveforms, so we stuck with the ASIC-verification approach," he says. "We'll still debug in the lab, but we try to get most of the bugs out before we get to the lab."

Similarly, Singh says, because his group uses only the largest FPGAs, it employs the same verification methodology for FPGAs as it does for ASICs and runs each design through simulation with a testbench, debugging, and timing analysis.

"We do functional simulation, gate-

MORE AT EDN.COM

- For more on EDA topics, check out the following articles:
  - "Design verification and debug ging FPGA implementations," www.edn.com/article/CA333612.
  - "Low-cost kits: the new FPGA-designer trend," www.edn.com/article/CA6391429.
  - "Creating quasistatic, parameterized FPGA designs," www.edn.com/article/CA313053.
- Go to www.edn.com/070719cs and click on Feedback Loop to post a comment on this article.

level simulation, and dynamic simulation, taking the gates along with parasitics back into functional simulation. Then, we do a very thorough static-timing analysis over multiple corners and modes," says Singh. "Once we do a good job there, we then mount it on a board and power it up." He concedes that the first time you use this method, you have to go through a few iterations, but then it becomes a standardized process and easier on subsequent projects. The method typically brings Singh good results. FPGA vendors provide great simulation models for their hard and soft blocks, SERDES, PLLs, and other macros that help his group's verification.

SOFT ERRORS

In addition to subtle differences between ASIC and FPGA design, SRAM-based FPGAs also present users with a new challenge: soft errors. Soft errors occur when a random atmospheric neutron collides with an IC, causing a bit error or, in some cases, a false signal (**Reference 3**). Standard-cell devices are not typically susceptible to soft errors, but SRAM structures, logic, and other types of memory are. And companies such as Xilinx, Altera, and Lattice base their highest performing FPGAs on SRAM.

Singh says that his group has started to more frequently encounter soft errors. Thus, he warns, when picking a device, you need to ask for a reliability study. And, even if the vendor claims that the FPGA is reliable, he says, you need to add some amount of ECC (error-correction coding) to your design. An SRAM-based device needs an ECC configuration to make sure there are no false errors—flipped bits that a neutron strike introduces.

If you have control circuitry in your design that checks the integrity of your configuration RAM, and, every once in a while, the configuration of your RAM changes, you've most likely encountered a soft error. Soft errors are uncommon in ASIC designs, but they can affect all the SRAM-based segments of an FPGA, including SRAM-based logic structures and on-chip SRAM blocks. "With transistors getting smaller and [device] voltage dropping, you can't help it; you're always prone to getting hit by a neutron," says Singh.

Fortunately, vendors such as Actel and Lattice have answers to the soft-error problem as they provide flash-based, nonvolatile FPGAs. These devices are neither as fast nor as dense as SRAM-based devices, but they are soft-error-resistant. The nonvolatile devices are also gaining popularity because they offer better security for companies selling into geographic markets in which piracy controls are still questionable.

In the face of rising ASIC-mask, -development, and -tool costs, FPGA vendors are presenting designers with a viable option for quickly getting designs out the door. But, before designers take the leap, they need to research FPGAs and weigh each vendor's architecture, design tools, and technical support. Designers also must understand the limitations of FPGA architectures and augment design techniques accordingly. EDN

FOR MORE INFORMATION

Actel
www.actel.com

Altera
www.altera.com

Hewlett-Packard
www.hp.com

Intel
www.intel.com

Lattice Semiconductor
www.latticesemi.com

Magma Design Automation
www.magma-da.com

Mentor Graphics
www.mentor.com

Quicklogic
www.quicklogic.com

Sonoa Systems
www.sonoasystems.com

Synplicity
www.synplicity.com

Toshiba
www.toshiba.com

Xilinx
www.xilinx.com

REFERENCES

1. Santarini, Michael, "FPGAs balance lower power, smaller nodes drip by drip," *EDN*, June 8, 2006, pg 58, www.edn.com/article/CA6339245.
2. Santarini, Michael, "ASIC prototyping: make versus buy," *EDN*, Nov 21, 2005, pg 30, www.edn.com/article/CA6283831.
3. Santarini, Michael, "Cosmic radiation comes to ASIC and SOC design," *EDN*, May 12, 2005, pg 46, www.edn.com/article/CA529381.

ELECTRONIC COMPONENTS

system-level IP | SERDES | embedded memory | system in package | embedded microprocessor | mixed-signal IP

How do you balance the integration of design/IP/manufacturing in your custom SoC development?

Leading companies rely on Toshiba for end-to-end integration of custom SoC/ASIC design/IP/manufacturing to minimize design risk and speed time to production. By offering these critical capabilities under one roof, Toshiba is able to prove its IP in high-volume applications, then follow through in a seamless EDA environment to ensure your design is highly manufacturable. Finally, the Toshiba design-for-manufacturing methodology significantly reduces design defects. No other company can match Toshiba electronic components — nor the Toshiba commitment to your success. Together we can change the world.

Get a 2007 calendar poster FREE, plus qualified engineers can enter to win an iPod® nano at* ***www.SoCworld.toshiba.com****.*

TOSHIBA
Leading Innovation »»

 *Weekly drawings while supplies last.

BY MARKUS LEVY, EEMBC • AND RON OLSON, IBM CORP

Autovectorization for GCC compiler

YOU CAN USE INDUSTRY-STANDARD BENCHMARKS TO IMPROVE COMPILER PERFORMANCE.

You can't go out in public these days without seeing someone talking on the phone, listening to music, taking a digital photograph, or even watching a video. With all of these electronic devices available today, engineers are constantly seeking ways to make them smaller and more efficient. One way to accomplish this goal is to compress any data that must be stored or transmitted, but this task requires a powerful processor.

A common practice is to use a DSP to handle data compression and decompression and a general-purpose processor to handle everything else. The source code targeting the general-purpose processor is often written in C or C++, whereas the source code targeting the DSP is often written in assembler. But most OEMs would prefer not to have two or more processors in their products, so many contemporary, general-purpose processors include SIMD (single-instruction-multiple-data) capabilities. General-purpose processors with SIMD capabilities can offer improved performance when executing complex algorithms, such as those in portable-device applications, and can run general-purpose code, such as a Linux OS.

The basic premise behind a vector/SIMD unit is to allow the processor to simultaneously execute a mathematical operation on multiple pieces of data within special vector registers. The benefit is that a single instruction inside the CPU effects the parallel operation. To support more programmers writing in C, developers are improving the GCC (GNU Compiler Collection) compiler to more efficiently take advantage of general-purpose processors with these SIMD capabilities.

Executing industry-standard EEMBC (Embedded Microprocessor Benchmark Consortium, www.eembc.org) benchmarks on the IBM (www.ibm.com) PowerPC 970FX, compiled with the GCC compiler, illustrates the results of these compiler optimizations. Specifically, out-of-the-box EEMBC TeleBench scores for the 970FX PowerPC processor yielded more than 150% better results than those the consortium published more than a year ago using the same processor and platform and running at the same speed, but with a different compiler. You can attribute the improved scores to autovectorization and FDPR (feedback-directed program restructuring).

APPLYING AUTOVECTORIZATION TO GCC

Autovectorization improves the performance of programs compiled with GCC for processors that have vector/SIMD capabilities. You can download the GCC from the GNU Web site, http://gcc.gnu.org. For example, the IBM 970FX processor VMX (vector/SIMD-multimedia-extensions) unit provides 32 quad-word, 128-bit vector registers that can hold different-sized elements, such as signed and unsigned bytes, half-words, full words, and single-precision floating-point numbers. The VMX unit is a complex design, so using it involves penalties. For example, loading the vector registers causes overhead, and restrictions exist, mainly in the alignment of the data, to read into these registers. Until now, it has been difficult for compilers to automatically recognize when performance benefits you accrue with VMX outweigh the overhead.

GCC is a versatile and readily available compiler and finds wide use in server, desktop, and embedded-system settings. Typically, its optimization technology lags many commercially available compilers. For now, however, GCC is among the leaders in the compiler industry for autovectorization technology. IBM engineers contributed much of this new technology while working on the autovectorization-branch version of GCC, which is available at http://gcc.gnu.org/projects/tree-ssa/vectorization.html. These changes will be part of the official Version 4.3 release.

Autovectorization targets loops in the program that a vector/SIMD unit can parallelize. The first step is to analyze the loop to see whether it qualifies for parallelization. For analysis of the loop, autovectorization uses recently added loop-analysis features of GCC in the tree-SSA (static-single-assignment) facility, which is available at http://gcc.gnu.org/projects/

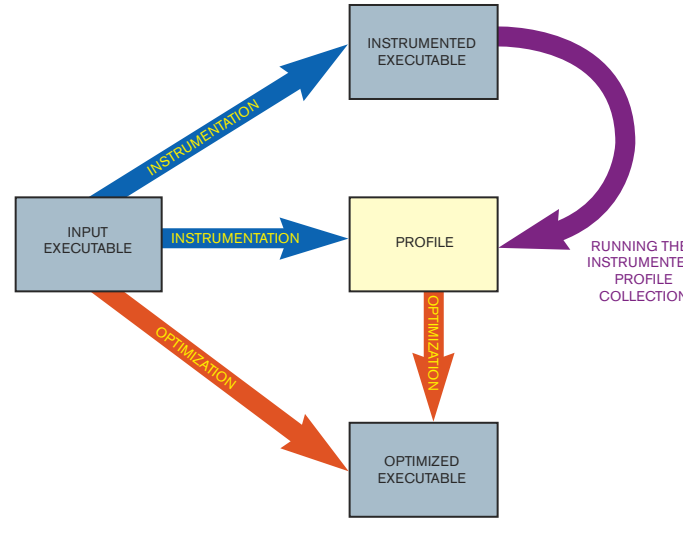

Figure 1 FDPR-Pro takes an executable program as input.

tree-ssa. If it qualifies, GCC will convert the body of the innermost loop to a number of parallel operations, reducing the number of iterations by that factor. For example, if GCC can modify the loop to do two simultaneous operations, that modification would halve the number of iterations. Similarly, if the GCC can modify the loop to do four simultaneous operations, that modification would reduce the loop by a factor of four.

Architectural differences pose the main difficulty with autovectorization. Currently available vector/SIMD units typically use special-purpose instructions and special vector registers; they have many features that a general-purpose compiler would not use. Recently, developers have significantly improved the autovectorizer of GCC. Some of these improvements include the ability to handle "strided," nonconsecutive accesses to memory; the addition of an idiom-recognition engine that includes a widening multiplication idiom; the ability to handle loops operating on multiple-sized data, including type conversion; the addition of enhanced "if" conversion and store sinking to replace conditional branches; and the addition of efficient loop versioning to resolve aliasing.

FEEDBACK-DIRECTED RESTRUCTURING

Another new technology to improve the performance on out-of-the-box code is FDPR-Pro. The FDPR-Pro postlink tool, available on the IBM alphaWorks Web site, www.alphaworks.ibm.com/tech/fdprpro, analyzes the behavior of the compiled program when it is running a typical workload and creates an executable program after modifying the code and data layout to increase performance. FDPR-Pro takes an executable program as input (**Figure 1**). This executable program works in the same way as any other program, except for the addition of one extra flag that tells the linker to retain relocation data; this extra flag results in a slightly larger file. The additional relocation information does not load to memory when the program is running.

FDPR-Pro first creates a new, instrumented, executable program by inserting additional code into the original program. This additional code is responsible for collecting profile information—typically, basic block- and control-flow edges' execution counts. Next, using a representative workload, it runs the instrumented program to create the profile data. Finally, it generates a new, optimized executable program, using the original program and the profile data.

FDPR-Pro can optimize any program or shared library by reordering the code to reduce cache and TLB (translation-look-aside-buffer) misses, reduce page faults and branch penalties, and improve branch prediction. It also removes unneeded NOP (no-operation) instructions, sets branch-prediction bits, and applies branch folding when beneficial. IBM engineers used FDPR-Pro to achieve an improvement of as much as 67% for single-kernel execution in the out-of-the-box EEMBC benchmark score for the 970FX.

TABLE 1 EEMBC OPTIMIZED-TELEMARK SCORES

Processor	Speed	Out-of-the-box score	Optimized-telemark score	Difference (%)
Freescale 7448	1.7 GHz	50.4	601.4	1093
Freescale 7447	1.3 GHz	37.7	507.8	1247
Freescale 7447	1.3 GHz	37.7	432.5	1047
Freescale 7447A	1.4 GHz	41.4	500.6	1109
Freescale 7455	1 GHz	28.3	121.6	330
IBM 970FX	2 GHz	56.1	1058.7	1787
LSI 402ZX	200 MHz	3.3	40.3	1121
TI TMS320C6203	300 MHz	6.8	44.6	556
TI TMS320C6203	300 MHz	6.8	68.5	907
TI TMS320C6413	500 MHz	13.5	263.3	1850
TI TMS320C6416	1 GHz	27.1	526.5	1843
TI TMS320C6416	1 GHz	27.1	873.1	3122
TI TMS320C6416	720 MHz	19.5	379.1	1844
TI TMS320C6416	720 MHz	19.5	628.6	3124
Average	NA	26.8	431.9	1512

TABLE 2 RESULTS OF RUNNING EACH KERNEL IN TESTBENCH SUITE

Benchmark	GCC autovectorization	GCC autovectorization FDPR-Pro	FDPR gain (%)	GCC	Autovectorization gain (%)
Autocorrelation/pulse	2,264,526	2,285,207	1	2,649,007	−15
Autocorrelation/sine	114,943	117,647	2	23,669	386
Autocorrelation/speech	107,527	110,867	3	24,691	335
Convolutional encoder/xk5r2dt	696,864	698,813	0	48,309	1343
Convolutional encoder/xk4r2dt	700,525	763,505	9	54,644	1182
Convolutional encoder/xk3r2dt	1,028,560	1,096,191	7	63,694	1515
Fixed-point bit allocation/typical	11,835	12,063	2	10,000	18
Fixed-point bit allocation/step	110,538	184,502	67	163,044	−32
Fixed-point bit allocation/pent	17,410	17,718	2	15,000	16
FFT/IFFT/pulse	63,966	65,076	2	60,976	5
FFT/IFFT/spn	63,966	65,076	2	60,976	5
FFT/IFFT/sine	63,966	65,076	2	60,976	5
Viterbi decoder/get	40,236	41,459	3	9740	313
Viterbi decoder/toggle	40,226	41,448	3	10,638	278
Viterbi decoder/ones	40,237	41,460	3	11,450	251
Viterbi decoder/zeros	40,170	41,277	3	14,286	181
Telemark	133.8	141.8	6	49.2	172

BENEFITS OF GCC

Implementing effective autovectorizing optimization for

Half the Size, Half the Loss
All the Power

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

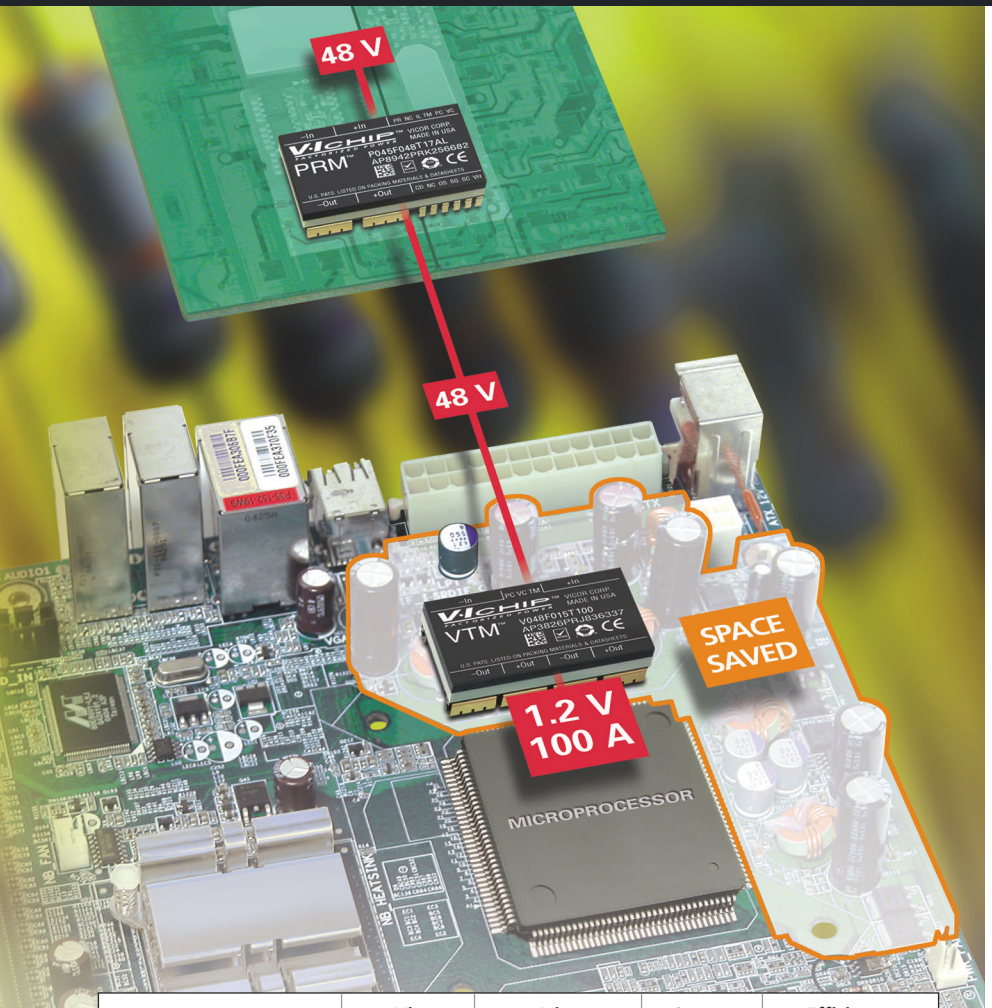

RoHS

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

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

a target processor is a difficult task. One of the many hurdles is the fact that test cases represent real programs. At this point, EEMBC telecommunication benchmarks come into play. IBM engineers running the EEMBC TeleBench kernels on the IBM PowerPC 970FX processor obtained data that they then used to improve the autovectorization capabilities of the GCC compiler. EEMBC's TeleBench suite comprises kernels that include autocorrelation, convolutional-encoder, bit-allocation, inverse-FFT (fast-Fourier-transform), FFT-benchmark, and Viterbi-decoder tests. These benchmarks represent tasks that can benefit from vector/SIMD execution. For example, the Viterbi encoder computes the most probable transmitted sequence of a convolutional coded sequence. The most computationally intensive part of Viterbi performs a maximization of a likelihood function through a sequence of add-compare-select operations.

EEMBC publishes two types of scores: out of the box and "full fury," or optimized. The organization obtains out-of-the-box scores by compiling unchanged source code, and it obtains full-fury scores by changing the source code to improve performance but still follow the EEMBC rules. In most cases, the changes engineers make to the code enable use of vector/SIMD instructions that compilers were unable to do automatically. The full-fury scores show an average improvement of more than 1500% over out-of-the-box scores for the same processor running at the same speed (**Table 1**).

You can make many optimizations by restructuring the code or rewriting it in assembler, but a compiler that recognizes when using vector/SIMD will be beneficial could automatically do a significant portion of these gains. To illustrate this point, compare the IBM 970FX optimized score of 1058.7 to the out-of-the-box score of 56.1 in June 2005 using Green Hills Software's (www.ghs.com) Multi compiler (**Table 2**). The percentage of improvement with hand-optimizing is on the same order of magnitude as many of the other processors' optimized improvements. The out-of-the-box score when using GCC with autovectorization is 141.8. This score is 153% better than the previous out-of-the-box score. Internal testing at IBM shows that the compiler from Green Hills Software is better than the GCC compiler without autovectorization. The tests also show that, when comparing results of autovectorization with results compiled using the previous version of GCC, Version 4.1.1, the gain is closer to 190%.

The first column of **Table 2** shows the EEMBC score in iterations per second of the test EEMBC compiled with autovectorization but before running FDPR-Pro. The next two columns show the results of running FDPR-Pro on the autovectorized executable and the resulting performance improvement expressed as a percentage. The next two columns compare the results of compiling the benchmarks with autovectorization and without autovectorization, ignoring gains from FDPR-Pro in both cases.

As you can see from the "autovectorization-gain" column, autovectorization provided more of a boost to some of

MORE AT EDN.COM

Go to www.edn.com/ms4249 and click on Feedback Loop to post a comment on this article.

Compact Power Supplies for FPGA-Based Systems

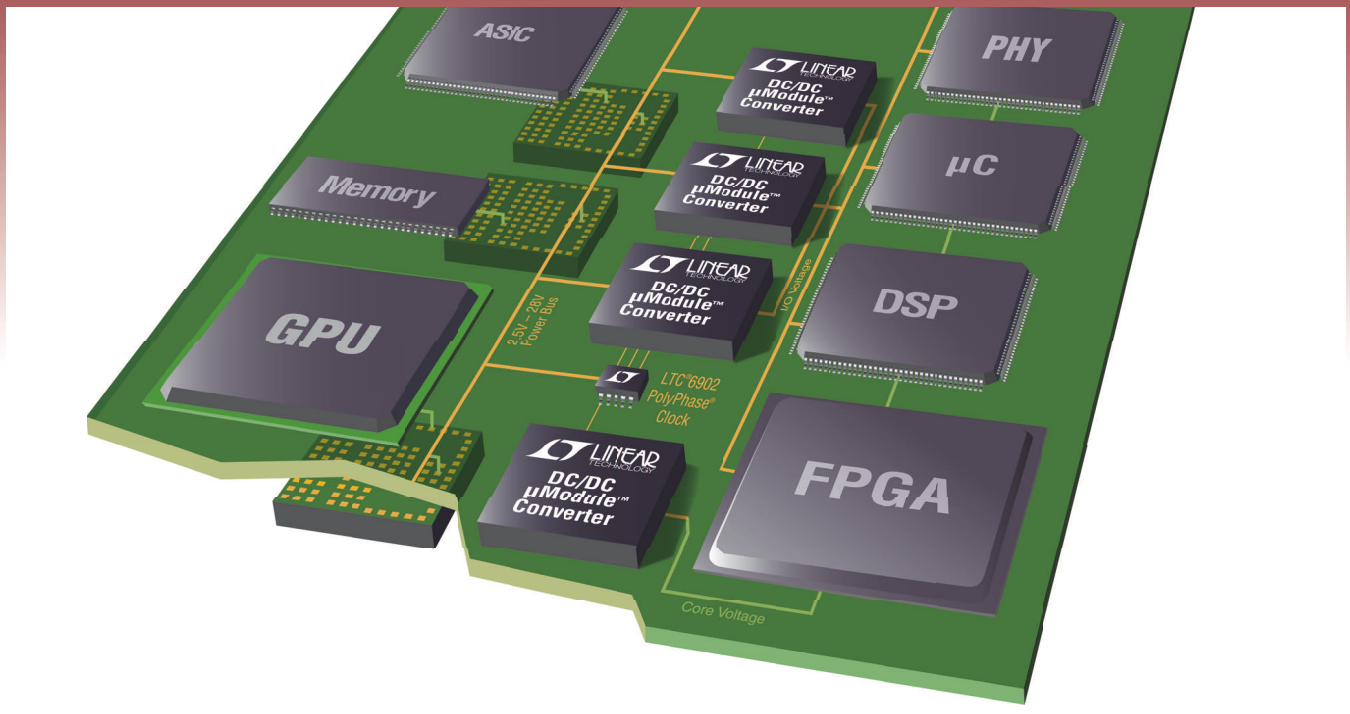

Tiny µModule™ Power Supplies Fit on Both Sides of PC Board

Our µModule DC/DC point-of-load power supply family is complete with built-in inductor, MOSFETs, bypass capacitors and compensation circuitry. At only 2.8mm height, these tiny, lightweight (1.7g) point-of-load regulators fit the tightest spaces on top and bottom of your board. Small size and impressive low thermal impedance allow high power conversion from a wide range of voltages. Our µModule DC/DC converters simplify the design of your FPGA-based system and are backed by rigorous testing and high product reliability.

µModule DC/DC Converters for Core, I/O, Clock & System Power

V_{IN}: 4.5V-28V; V_{OUT}: 0.6V-5V						LGA Package (15°C/W)	
Part No.	I_{OUT} (A)	Current Share	PLL	Track, Margin	Remote Sense	Height (mm)	Area (mm)
LTM®4602	6	Combine two for 12A to 24A or 4x LTM4601 for ≤48A				2.8	15x15
LTM4603	6		✔	✔	✔		
LTM4603-1	6		✔	✔			
LTM4600	10						
LTM4601	12		✔	✔	✔		
LTM4601-1	12		✔	✔			
V_{IN}: 2.25V-5.5V; V_{OUT}: 0.8V-3.3V							
LTM4604*	4	2x for 8A	✔			2.3	9x15

*Future Product

Info & Online Store

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

Application Notes
Available for Download

LINEAR TECHNOLOGY

the benchmarks, particularly Viterbi benchmarks. For other benchmarks, such as convolutional encoder, scalar replacement of aggregates, rather than autovectorization, improved scores. Vectorization can further boost the convolutional encoder, as the optimized-telemark version demonstrates, but it is more difficult to automate.

The benefit of improved compiler optimizations is not higher benchmark scores, but better performance in real products with less time and effort. Today, OEMs are demanding more performance from embedded processors, and processor vendors are responding by adding execution units that can do simultaneous operations in parallel, such as vector/SIMD engines. Adding capabilities to the CPU will do no good unless the software can take advantage of these features. Until now, the best choice for taking advantage of vector/SIMD functions has been to modify generic code to use these special features of the processor. This method has many disadvantages, including the fact that it binds you to a specific implementation. Now, with GCC 4.3 and beyond, you can benefit from vector/SIMD by simply recompiling your code. **EDN**

REFERENCES

1 "Auto-vectorization in GCC," http://gcc.gnu.org/projects/tree-ssa/vectorization.html.

2 "Feedback Directed Program Restructuring," www.haifa.il.ibm.com/projects/systems/cot/fdpr.

3 "PowerPC Microprocessor Family: Vector/SIMD Multimedia Extension Technology Programming Environments Manual," Aug 22, 2005, www-3.ibm.com/chips/techlib/techlib.nsf/techdocs/FBFA164F824370F987256D6A006F424D.

4 Nuzman, Dorit, Ira Rosen, and Ayal Zaks, "Auto-Vectorization of Interleaved Data for SIMD," PLDI 2006, June 12 to 14, 2004, Ottawa, ON, Canada, pg 132.

5 Nuzman, Dorit, and Ayal Zaks, "Autovectorization in GCC–two years later," GCC Developers' Summit 2006, June 28 to 30, 2006, Ottawa, ON, Canada.

6 Nuzman, Dorit, and Richard Henderson, "Multi-platform Auto-vectorization," Proceedings of the International Symposium on Code Generation and Optimization, March 26 to 29, 2006, New York.

AUTHORS' BIOGRAPHIES

Markus Levy is president of EEMBC, where he has worked for 10 years and where he is responsible for management, marketing, business development, and conference presentations. He has a bachelor's degree in electrical engineering from San Francisco State University, and his personal interests include exercising, mountain biking, skiing, and volunteering as a firefighter. You can reach him at markus@eembc.org.

Ron Olson is a software engineer at IBM, where he has worked for 23 years. He is responsible for PowerPC-software development, benchmarking, and enablement. He has a bachelor's degree in electrical engineering from DeVry University, and his personal interests include sailing, fishing, and swing dancing. You can reach him at ronolson@us.ibm.com.

16-Bit, 8-Channel ADC with Software-Selectable Input Span

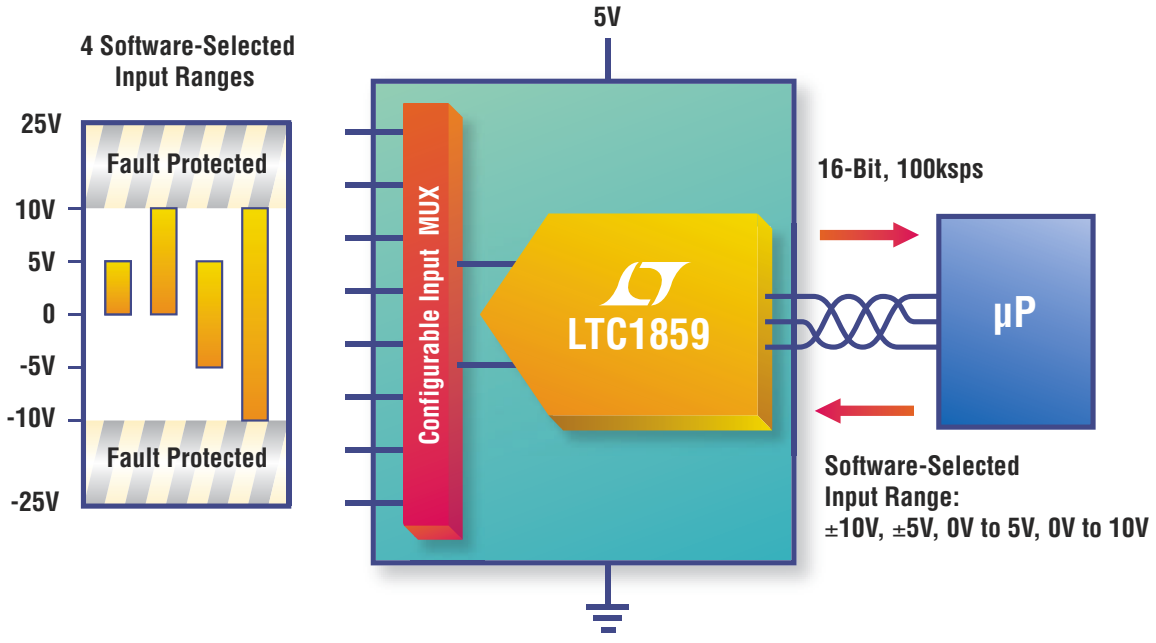

16-Bit, 100ksps and ±3LSB INL on 5V Supply

The extremely versatile input multiplexer on the LTC®1859 16-bit SoftSpan™ ADC is software-configurable for unipolar, bipolar, single-ended or differential inputs. The LTC1859 achieves ±3LSB INL, the 14-bit LTC1858 ±1.5LSB and the 12-bit LTC1857 ±1LSB. Each channel or differential pair can be independently programmed. Conversions on any channel are unaffected by fault conditions up to ±25V on any other channel.

Features

- 100ksps, 16-Bit Resolution
- Software-Selectable Input Ranges:
  - 0V to 5V
  - 0V to 10V
  - ±5V, ±10V
  - Single-Ended or Differential
- 8-Channel MUX with ±25V Protection
- Single 5V Supply, 40mW
- Serial Interface
- Compact 28-pin SSOP
- Pin-/Software-Compatible SoftSpan/ Fixed ±10V Family

Industrial ±10V ADCs

Resolution	SoftSpan	Fixed ±10V
16-bit	LTC1859	LTC1856
14-bit	LTC1858	LTC1855
12-bit	LTC1857	LTC1854

Info & Free Samples

www.linear.com/1859

1-800-4-LINEAR

Free Industrial Signal Chain Brochure

www.linear.com/indsolutions

LINEAR TECHNOLOGY

2GHz ADC Driver

Get All the Bits You Paid for

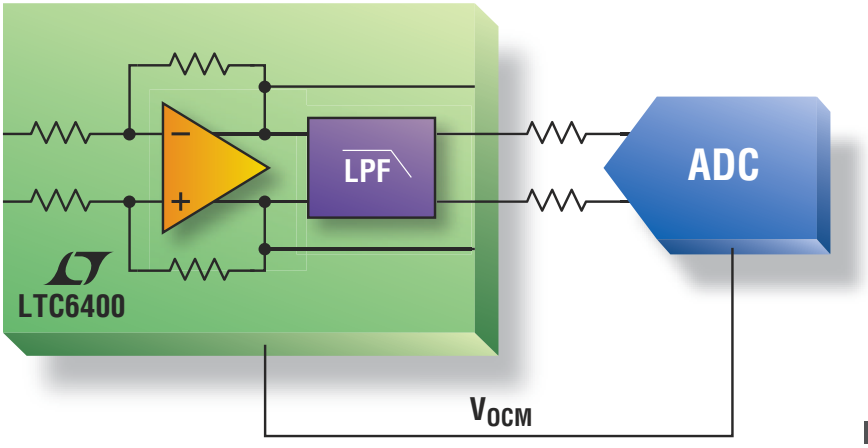

Actual Driver Size
3mm x 3mm QFN

-93dBc IMD3 at 140MHz

Introducing the LTC®6400, an ADC driver with performance to match today's most advanced high-speed ADCs. Featuring excellent distortion performance and 1.8 nV/rtHz input-referred noise for 300MHz inputs, the LTC6400 is optimized for driving 12-bit, 14-bit and 16-bit ADCs. And it does this all on a single 3V supply at 270mW, while providing a differential output swing of $4.4V_{P-P}$. The LTC6400 fixed gain architecture eliminates several external components, simplifies design and maximizes performance.

Features

- 1.8GHz -3dB Bandwidth
- Fixed Gain of 10V/V (20dB)
- -93dBc IMD_3 at 140MHz
- 6.2dB Noise Figure
- 2.85V to 3.5V Supply Voltage
- 90mA Supply Current (270mW)
- Other Gain Options and Lower Power LTC6401 Also Available

Third Order Intermodulation Distortion vs Frequency
V_{OUT} = $2V_{pp}$ Composite

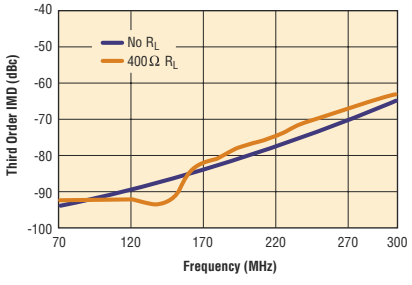

Info & Free Samples

www.linear.com/6400
1-800-4-LINEAR

Free High Speed ADC Brochure

www.linear.com/adcsolutions

LINEAR TECHNOLOGY

EDITED BY CHARLES H SMALL
AND FRAN GRANVILLE

designideas

READERS SOLVE DESIGN PROBLEMS

Analyzer tests reverse-recovery behavior of diodes

Louis Vlemincq, Belgacom, Evere, Belgium

Testing the reverse-recovery behavior of diodes normally requires complex testing gear. You must be able to establish the forward-conduction conditions, the blocking state, and the transition between the two. You also need a means of extracting the characteristics from the resulting waveform. In short, a specialist should handle this complex job; it is not something you routinely control in the field. This fact explains why engineers generally prefer to rely on published data.

Checking the reverse-recovery time yourself could be advantageous, however, if testing were simple and straightforward. Such a setup would enable you to compare devices from different manufacturers under identical conditions and test devices having no such specification, such as substrate diodes of driver ICs, zener diodes, and standard rectifiers. (Because of the number of combinations of the test parameters, a direct comparison of the data is rare-

DIs Inside

78 High-power LED drivers require no external switches

82 Perform PSRR testing with analyzers having no dc-bias ports

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

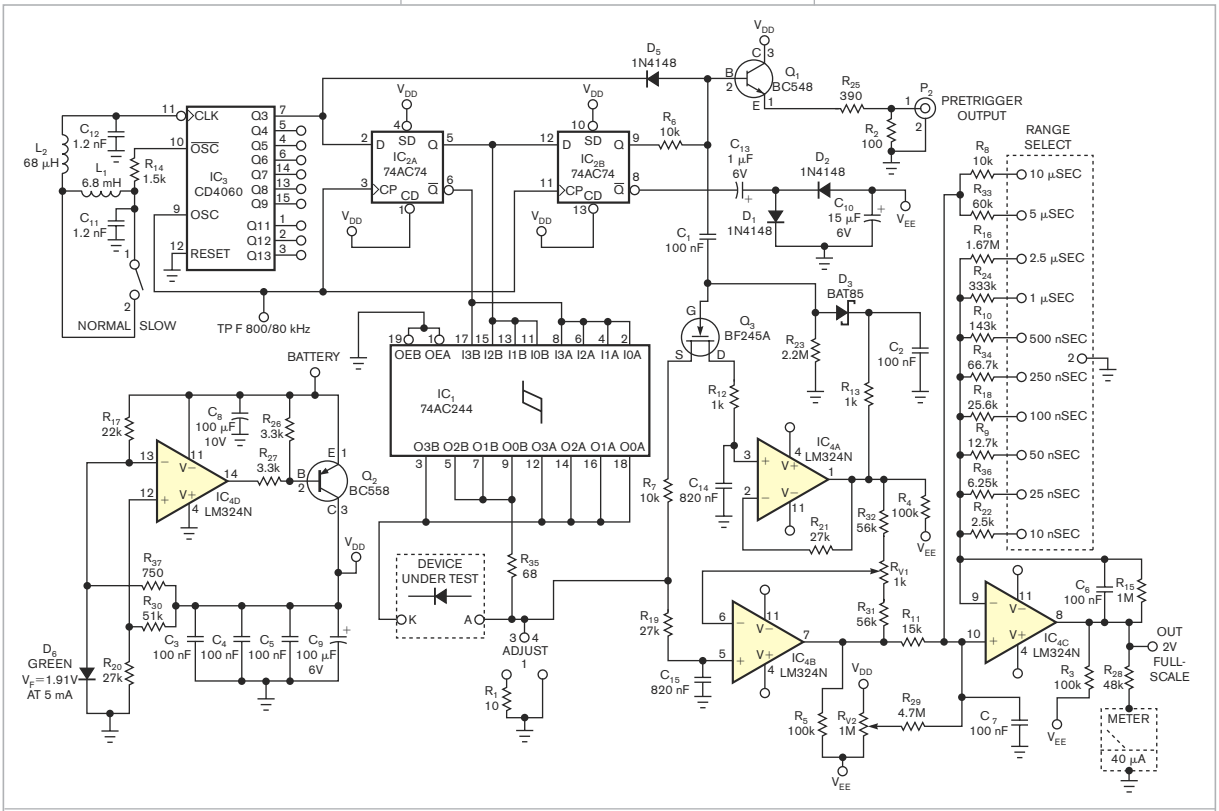

Figure 1 This diode-recovery test setup allows you to compare devices from different manufacturers under identical conditions.

ly possible.) Note that shorter reverse-recovery time is not necessarily better. Slow diodes can be useful, too. They can generate small dead times, improve the efficiency of converters, and provide other benefits (**Reference 1**).

This Design Idea presents a tester that, using only a handful of inexpensive, standard components, allows you to check reverse-recovery time. The test conditions are fixed for simplicity, to normalize the tests and to provide a common standard for comparison purposes. These conditions are compatible with 99% of the devices susceptible to test. The tester's forward current is just low enough to be safe with small switching diodes but high enough to overcome the capacitive effects in larger devices.

A diode-resistor AND gate lies at the heart of the circuit; the gate's diode is the DUT (device under test, **Figure 1**). IC_1 buffers flip-flop IC_{2A}, which derives the antiphase square waves that drive this gate. R_{35} sets the DUT's forward current to approximately 75 mA. With an ideal diode, the gate's output would always stay low, because one of the inputs is always low. But a real diode remains conductive after the transition, generating a positive pulse across R_{35}. Instead of using the brute-force approach of directly measuring this pulse width, the circuit uses a subtler scheme. The R_{19}/C_{15} network averages the pulse and amplifies and displays the resulting voltage. Because the measurement frequency is fixed at 50 kHz, a correct scaling factor is all that is necessary.

A real diode also has a forward voltage, which you would average with the result. Q_3 takes care of this problem by sampling this forward voltage through IC_{4A} and subtracting it from the output voltage through R_{32}. Varying the gain of amplifier IC_{4C} sets the various ranges. In this case, the ranges are in a 1, 2.5, 5 sequence, which suit the salvaged galvanometer this circuit uses as an indicating device. You could easily create other ranges by adapting the values of R_8 through R_{22}. The big advantage of this measuring method is that it handles only dc or low-frequency signals, requiring no fast comparators or samplers, yet it can resolve a few hundreds of picoseconds.

The built-in oscillator of IC_3 generates the clock. The clock frequency is 800 kHz and divides down to produce the 50-kHz reference at Q_3. An optional slow mode is available for those needing to test devices slower than 5 µsec. The insertion of coil L_1 decreases the clock frequency to 80 kHz and enables you to measure reverse-recovery time as fast as 50 µsec. IC_2 generates the test waveforms and shifts the 50-kHz signal at the clock rate. The leading and trailing states then exit through the D_5/R_6 AND gate to produce a sampling pulse that centers on the conduction period. Because the sampling occurs far from any transitions, it need not be particularly fast or accurate. C_1 transfers the sampling pulse and provides a convenient pretrigger signal, which Q_1 buffers. This option enables a comfortable observation of the waveform when you connect an oscilloscope to the anode of the DUT.

The unused output, Pin 8 of IC_{2B}, feeds a negative-voltage generator, serving as a bias source for the outputs of IC_4 to let them reach a true zero. The measurement circuits receive power from a 9V battery by a supply encompassing IC_{4D}. An LED serves as a reference to the 5.5V and provides some temperature compensation because the reverse-recovery time depends highly on ambient temperature.

You can make some adjustments to the circuit. For example, with no diode inserted, you can short the Adjust testpoint 1 to 4. In the 10- or 25-nsec range, R_{V2}, which is 0 nsec in the range, to get a midscale reading. Move the short to Adjust testpoint 3, R_1, and R_{V1}, thereby providing V_F cancellation, to read the same value. Repeat the procedure until the reading is independent of the position of the short. The adjustment interacts with the zero due to the offset of the amplifiers.

Now, you have eliminated the effect of V_F. You can adjust the 0 nsec by shorting Adjust testpoints 1 and 4 and adjusting R_{V2} to read zero on the 10-nsec range. This adjustment yields 0 nsec with a typical offset of 1 to 2 nsec in the positive direction. Residual skew in the timing and charge-injection effects cause this offset. Normally, this offset should not be a problem, because it is small, stable, and constant. If you need an absolute accuracy down to the picosecond, you have to test a known, ultrafast diode, such as an FD700 or a BAY82, and adjust the 0 nsec to read the actual value. If you lack access to such a diode, you can always arbitrarily shift the value by 1.5 nsec. This adjustment is normally sufficient to reach a ±500-psec accuracy. Schottky diodes are unsuitable. Despite their low recovery time, they generate a nonzero reading because of their relatively high capacitance and non-negligible leakage currents. Low-capacitance, mixer-type diodes are too fragile for this tester. EDN

REFERENCE

1 Vlemincq, Louis, "Slow diodes or handy timing devices?" *EDN*, Sept 16, 2004, pg 83, www.edn.com/article/CA450598.

High-power LED drivers require no external switches

Alfredo H Saab and Steve Logan, Maxim Integrated Products, Sunnyvale, CA

As the latest generation of new LEDs achieves higher levels of power and efficiency, use of these devices extends to new areas, such as flashlights and vehicular applications. High-power LEDs are finding use even in ambient lighting, long the sole province of incandescent bulbs and fluorescent tubes. A current source is the best way to power LEDs. Because most energy sources, including batteries, generators, and industrial mains, look more like voltage sources than current sources, LEDs require that you insert some

INDUSTRY'S MOST SECURE MICROCONTROLLER FOR FINANCIAL APPLICATIONS

Integrated Tamper-Reactive Sensors Provide Strongest Key Protection

Since 1993, Dallas Semiconductor (a wholly owned subsidiary of Maxim Integrated Products) has been the industry leader in secure microcontrollers that drive financial terminals. Our fourth-generation secure microcontroller, the DS5250 is ideal for embedded applications that require a high level of security.

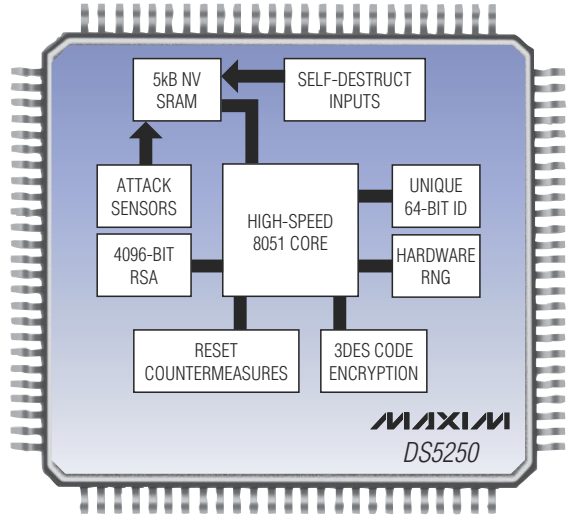

- ♦ Cryptographic Accelerators for RSA, DES and 3DES
- ♦ Voltage, Temperature, and Intrusion Sensors Detect Attack
- ♦ Hardware Random-Number Generator
- ♦ NV SRAM Wipes Instantly on Tamper
- ♦ Encrypted External Code and Data Memory

Part	Temperature (°C)	Package
DS5250F+125	0 to +75	100-MQFP
DS5250F+825	0 to +75	80-MQFP
DS5250F+1N5	-40 to +85	100-MQFP
DS5250F+8N5	-40 to +85	80-MQFP

MAXIM

www.maxim-ic.com

FREE* Authentication & Data Security 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 DIRECT™ 1-888-MAXIM-IC

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

AVNET electronics marketing 1-800-332-8638

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

electronic circuitry between them and the source of power. This circuitry can be as simple as a series resistor, but a better choice, considering energy efficiency and other factors, is a high-efficiency, voltage-fed current source. For LEDs with currents greater than 0.35A, an inductive switching regulator is usually the best choice.

This Design Idea presents a series of circuits based on single-power-IC switching regulators, with efficiency and miniaturization as the main objectives. The circuits' designers approach these objectives by minimizing the use of large components, such as external power transistors, switches, high-value capacitors, and current-sense resistors, and by maintaining regular operation by delivering constant, high-intensity light over as extended a range as possible.

The circuits in **figures 1** through **3** are suitable for applications in which the power source comprises three or four alkaline, NiMH (nickel-metal-hydride), or NiCd (nickel-cadmium) cells. Those in **figures 4** and **5** are for vehicular applications in which the nominal line voltage for the power-distribution system is 12, 24, or 42V. The circuits of **figures 4** and **5** are also useful in industrial systems that include a 24V distribution line for control and emergency subsystems and in telecom applications for which the system power is distributed as a −48V line.

The designers of these circuits based them on the same concept: a fully integrated, single-die-IC switching regulator and a micropower operational amplifier. The op amp drives the 1.25V feedback terminal on the IC. Although that node targets the topology of a standard voltage regulator, the op amp matches it to the much smaller current-sense voltage and the slightly different topology of a current regulator. None of the circuits requires the use of external power switches. The design eliminates the use of the large-valued filter capacitors you usually find in a switching regulator, because there is no need to smooth out high-frequency ripple in the LED current. Common to all circuits is the option of adding a dimming capability by introducing adjustable bias at an op-amp input through

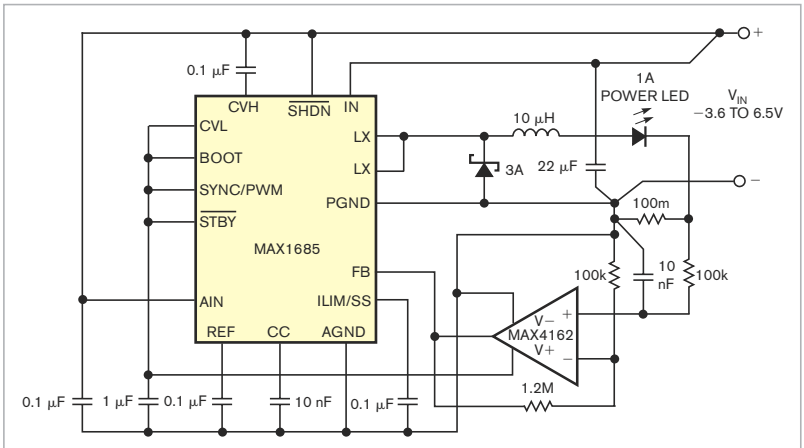

Figure 1 This miniature, 1A, high-power LED driver operates on 3.6 to 6.5V.

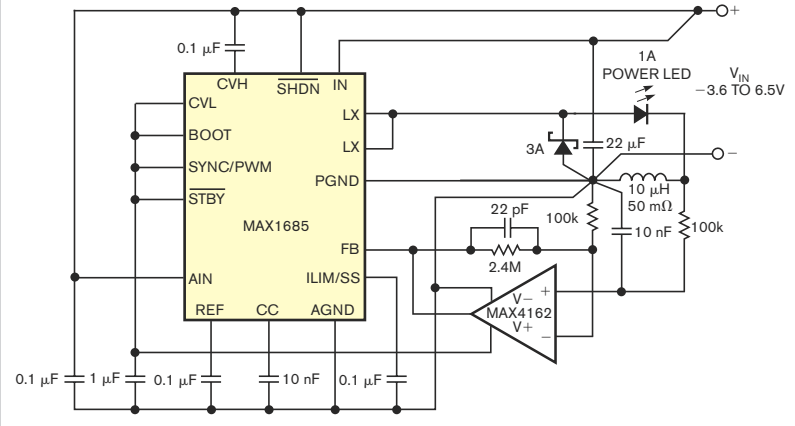

Figure 2 Similar to the circuit in Figure 1, this miniature, 1A, LED driver operates on 3.6 to 6.5V but requires no current-sense resistor.

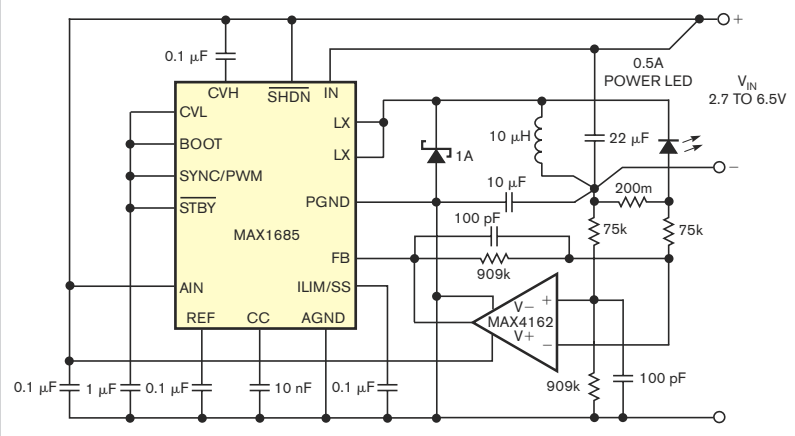

Figure 3 Another miniature, high-power LED driver delivers 0.5A and operates on 2.7 to 6.5V.

FIRST 3A/6A, STEP-DOWN REGULATORS WITH INTEGRATED SWITCHES AND VID INPUTS

Over 95% Efficiency, Up to 2.4MHz Switching Frequency, Low On-Resistance, and Sync Input—All in a Small 4mm x 4mm TQFN

The MAX8643/MAX8646 step-down regulators provide the highest integration with on-board low on-resistance switches and full overcurrent and thermal protection. Frequencies are programmable and synchronizable from 250kHz to 2.4MHz. A two-pin VID input provides one of nine output preset voltages that can be selected by the user. This eliminates the need for an external resistor-divider, thus providing higher overall accuracy. PWRGD and EN facilitate sequencing, where REFIN can be used for tracking and/or termination, such as in DDR applications. This set of features makes these parts easy to use and suitable for any power-supply application.

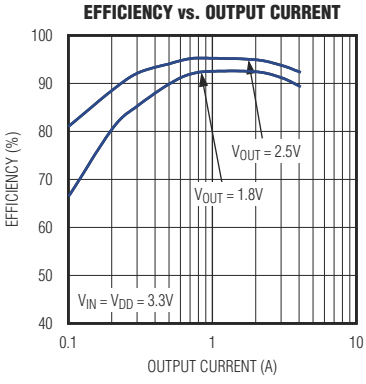

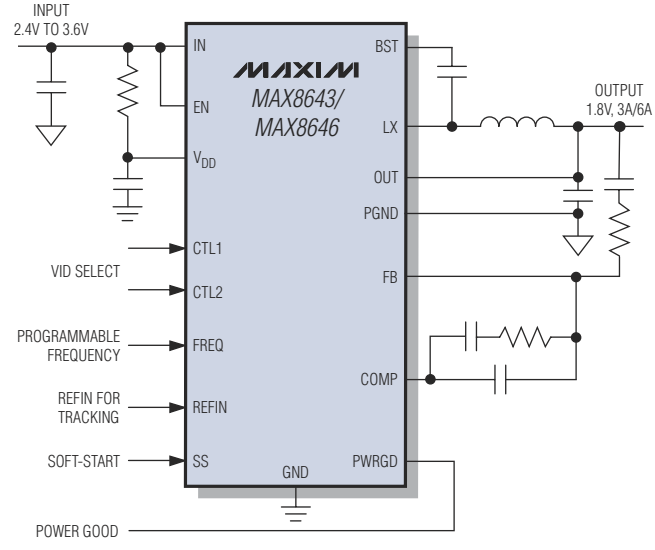

- VID-Set Output Voltages: 0.6V, 0.7V, 0.8V, 1.0V, 1.2V, 1.5V, 1.8V, 2.0V, and 2.5V
- Continuous 3A/6A Output (MAX8643/MAX8646)
- ±1% Output Accuracy Over Temperature
- Input Supply from 2.3V to 3.6V
- Soft-Start Reduces Inrush Supply Current
- 250kHz to 2.4MHz Adjustable Switching or Sync Input
- All-Ceramic Capacitor Design Reduces Space
- REFIN for Tracking and DDR Termination
- PWRGD and EN for Sequencing
- Available in 24-Pin TQFN Package with Exposed Pad
- Applications
  - ASIC/DSP/CPU Core Voltages
  - DDR Termination
  - Memory Power Supply
  - Telecom/Datacom Power Supply
  - Industrial Power Supply
  - Point-of-Load Power Supply

MAXIM

www.maxim-ic.com

FREE* Power Supplies 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 DIRECT 1-888-MAXIM-IC

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

AVNET electronics marketing 1-800-332-8638

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

a resistor and a potentiometer powered from the internal regulator—the VD or CVL terminal, depending on the IC.

A high-frequency switching regulator powers the basic regulator circuit for LEDs (**Figure 1**). It operates with input voltages of 3.6 to 6.5V, drives a single LED with currents as high as 1A, and uses a current-sense resistor to control the current-regulation loop. The circuit of **Figure 2** is similar, but, in place of a current-sense resistor, it employs the parasitic resistance of the inductor as a current-sensing element. Like the circuit in **Figure 1**, it operates with 3.6 to 6.5V inputs and drives one LED with currents as high as 1A.

For the single-LED circuit of **Figure 3**, the starting voltage of the MAX1685 defines the input range, which goes as low as 2.7V. Its maximum current capability is 0.5A versus 1A for the circuits in **figures 1** and **2**. The upper operating limit remains 6.5V. Once this circuit is operating, it maintains power to the LED even for input voltages as low as 1.7V. Applications for the circuits of **figures 1**, **2**, and **3** include headlights, flashlights, and any other portable lights powered by three or four alkaline primary cells, three or four NiMH/NiCd secondary cells, or a single lithium secondary cell.

The circuits of **figures 4** and **5** operate over 8 to 50V. Assuming a 12V system in which all the components are properly specified, these circuits can survive load dumps, thanks to the 76V absolute maximum rating for the IC's input-power terminal, V_{IN}. The maximum available current is 1A, and the circuits can drive as many as three LEDs in series, provided that you increase the lower limit of the operating range to 11.5V. These two circuits are similar, except for the use of the inductor resistance as a current sensor in **Figure 5**. The disadvantage of using the inductor resistance in this way is the resulting dependence of output current on temperature, due to the large temperature coefficient of copper resistivity. The inductor winding is made of copper, and its dc resistance has a first-order temperature coefficient of 3.9 parts/1000/°C. As a result, the regulated current decreases about 4% for each 10°C increase in operating temperature. EDN

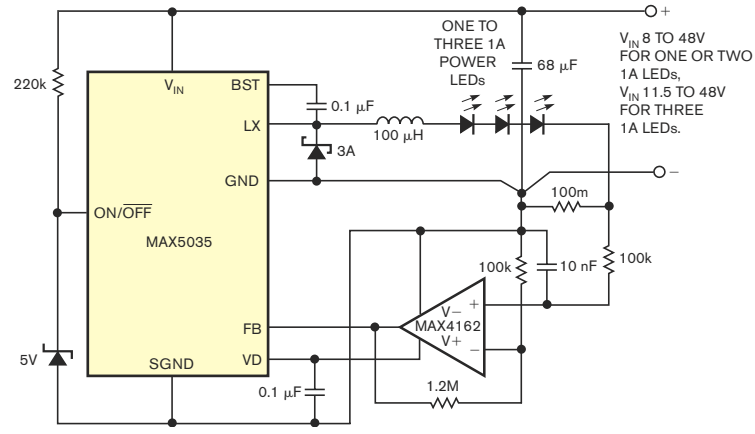

Figure 4 This miniature, 1A, LED-driver circuit operates on 8 to 50V and drives as many as three LEDs in series.

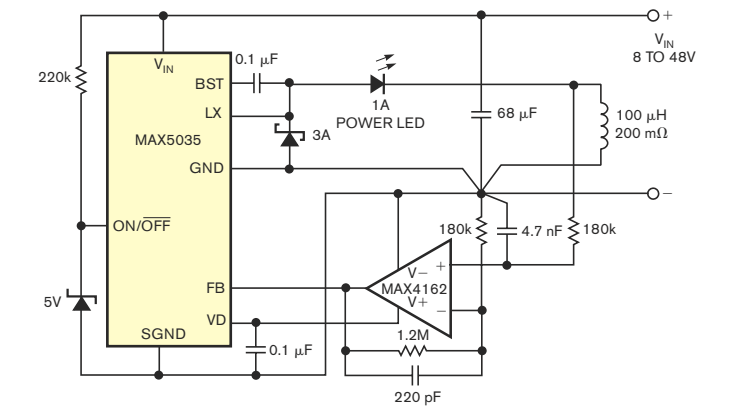

Figure 5 Otherwise similar to the circuit in Figure 4, this circuit requires no sense resistor.

Perform PSRR testing with analyzers having no dc-bias ports

David Karpaty, Analog Devices, Wilmington, MA

An amplifier's PSRR (power-supply-rejection ratio) is among the most commonly characterized parameters when analyzing the performance of an op amp. Examples of some noise sources on an amplifier's power-supply pins include parasitic supply-line traces, their interaction with currents that the amplifier draws, and the noise that switching circuits sharing the same supply create. Both sources produce voltage-amplitude variations reproduced as noise signals at the amplifier's input pins.

Characterizing PSRR over frequency

BROADEST LINE OF THERMAL-MANAGEMENT ICs FOR INDUSTRIAL PROCESS CONTROL

Over 100 Solutions for the Most Demanding Applications

Analog Temp Sensors

Over 30 Different Solutions

MAX6613—Perfect for Handheld Instruments

- 1.8V to 5V Supply Voltage Range
- ±1.3°C Accuracy
- Improved LM20

6-SC70
2.0mm x 2.1mm

Thermocouple to Digital

Many Thermocouple Solutions

MAX6675

- Reads Up to +1024°C
- 12-Bit, 0.25°C Resolution
- -20°C to +85°C Operation
- 3V to 5.5V Supply Voltage Range

8-SO
5mm x 6.2mm

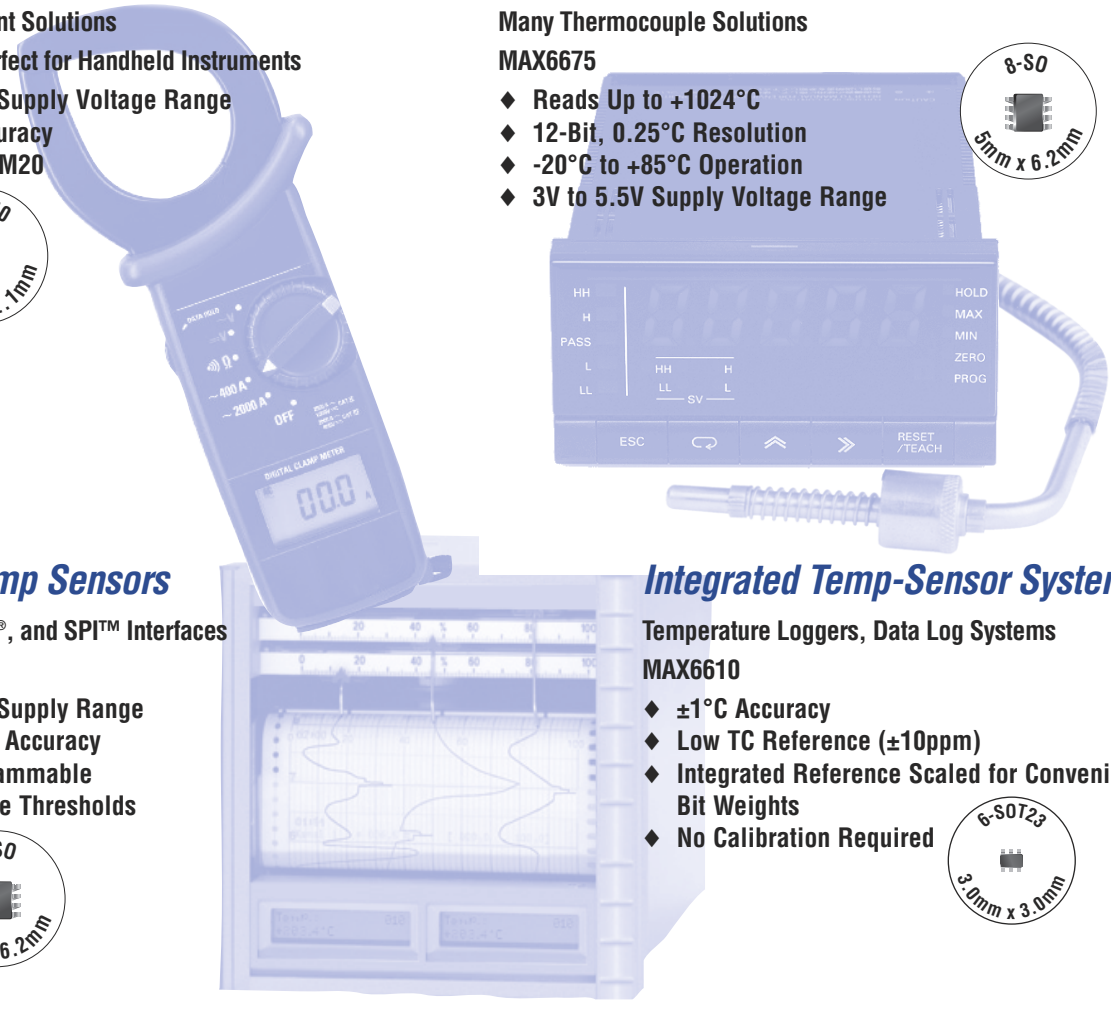

Digital Temp Sensors

2-Wire, 1-Wire®, and SPI™ Interfaces

MAX6634

- 3V to 5.5V Supply Range
- ±1°C (max) Accuracy
- User-Programmable Temperature Thresholds

8-SO
5mm x 6.2mm

Integrated Temp-Sensor Systems

Temperature Loggers, Data Log Systems

MAX6610

- ±1°C Accuracy
- Low TC Reference (±10ppm)
- Integrated Reference Scaled for Convenient Bit Weights
- No Calibration Required

6-SOT23
3.0mm x 3.0mm

SPI is a trademark of Motorola, Inc.
1-Wire is a registered trademark of Dallas Semiconductor Corp.

MAXIM

www.maxim-ic.com/industrial

FREE* Temp Sensor 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 DIRECT™ 1-888-MAXIM-IC

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

AVNET electronics marketing 1-800-332-8638

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

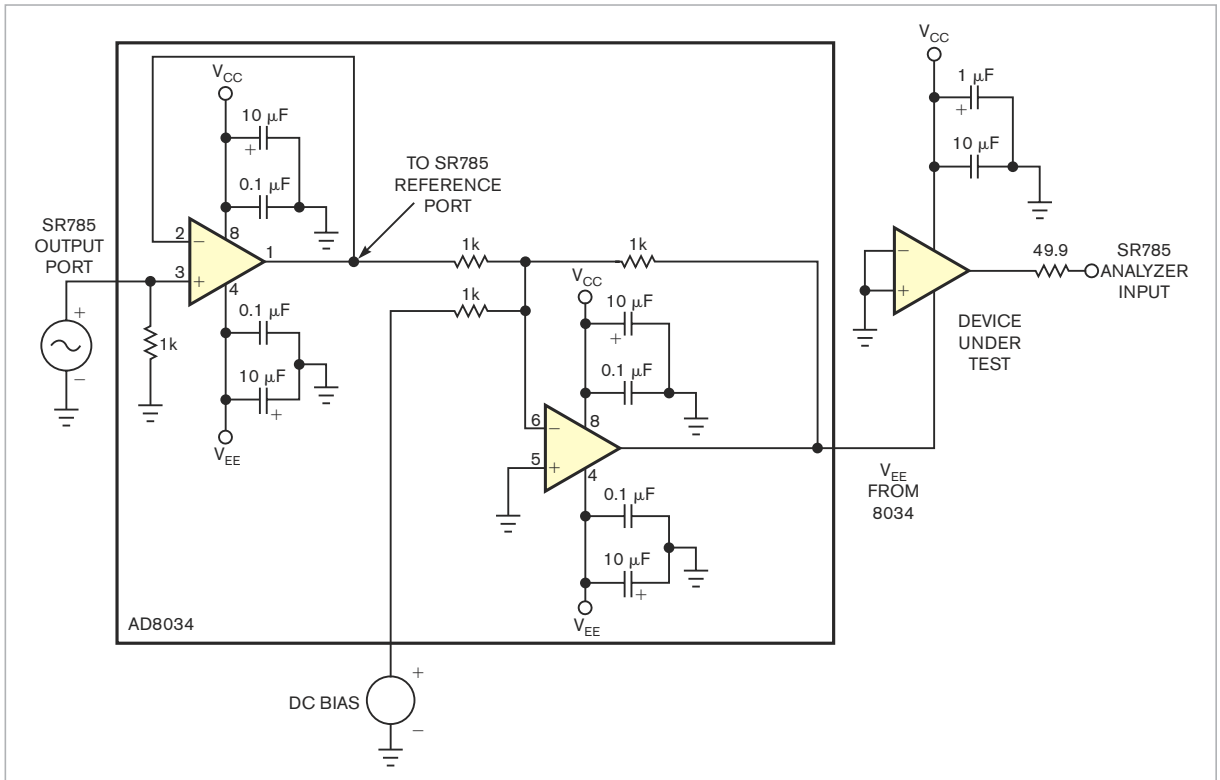

Figure 1 The AD8034 sums the analyzer's reference output and a dc bias, supplying the negative-supply voltage for the device under test to perform negative-PSRR analysis.

commonly involves the use of analyzers equipped with a dc-bias port, such as Agilent's (www.agilent.com) 8753. To measure negative PSRR, for example, the amplifier's $-V_S$ pin comes through Port 1, with the negative dc voltage through the bias port, of the 8753 with a superimposed sinusoid. To complete the measurement, you measure the amplifier's output on Port 2. Unfortunately, the 8753 doesn't measure frequencies below 30 kHz because of the limitations of the analyzer's internal bias, T. Additionally, most PSRR-versus-frequency plots begin at frequencies far below 30 kHz.

An alternative technique would involve the use of an analyzer that has no dc-bias port but that can characterize frequency response as low as 10 or even 1 Hz. One such analyzer is the Stanford Research Systems (www.thinksrs.com) SR785, which can make measurements better than −120 dB. One way of approaching this problem is to connect the output port of the SR785 to a buffer/inverting-summer circuit constructed with an Analog Devices (www.analog.com) AD8034.

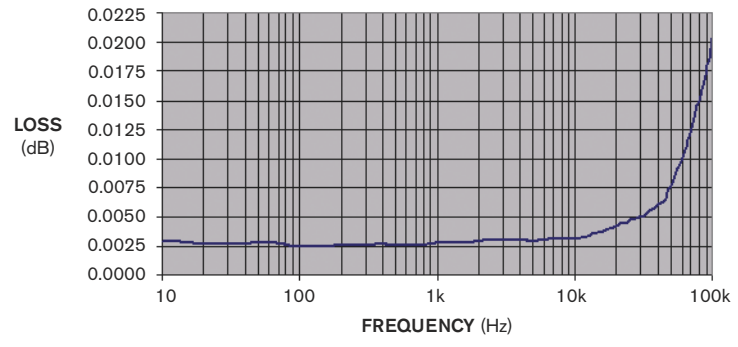

Figure 2 You can neglect any loss the AD8034 incurs. The response of the AD8034 buffer/inverting summer from 10 Hz to 10 kHz is approximately flat.

Figure 1 illustrates a negative-PSRR test-circuit configuration. Pin 3 connects to the SR785 source-output port. Pin 1, which is V_{OUT} of the buffer amplifier, connects to the reference port of the SR785. Here, the first amp isolates the output port of the SR785 from the dc bias and provides the sinusoidal output. The second amplifier within the AD8034 sums the dc bias and sinusoid, which it uses to feed the DUT's (device under test's) negative-supply pin. The **figure** omits all bypass capacitors at the DUT's negative-supply pin. A 1-kΩ resistor from Pin 3 to ground prevents the noninverting input from

LINEAR TECHNOLOGY

DESIGN NOTES

USB Power Solution Includes Switching Power Manager, Battery Charger, Three Synchronous Buck Regulators and LDO

Design Note 420

Brian Shaffer

Introduction

Linear Technology offers a variety of parts to simplify the task of extracting power from a battery or a USB cable. These devices seamlessly manage the power flow between an AC adapter, USB cable and Li-Ion battery, all while maintaining USB power specification compliance. As battery capacities rise, battery chargers must keep pace by steadily improving efficiency to minimize thermal concerns and charge times. A USB-based battery charger must squeeze as much power from the USB as possible, and do it efficiently to meet the stringent space and thermal constraints of today's power-intensive applications.

The LTC®3555 combines a USB switching power manager and battery charger with three synchronous buck regulators and an LDO to provide a complete power supply solution in one small (4mm × 5mm) package (Figure 1). The constant-current, constant-voltage Lithium-Ion/Polymer charger utilizes a Bat-Track™ feature to maximize the efficiency of the battery charger by generating an input voltage that automatically tracks the battery voltage (described below). An I²C serial interface affords the system designer complete control over the charger and the DC/DC bucks for ultimate adaptability to changing operating modes in a wide range of applications.

Switching PowerPath Controller Maximizes Available Power to the System Load

The LTC3555 improves over earlier generations of USB battery chargers with the addition of several new features. It uses a proprietary switching power manager to extract power from a current-limited USB port with the highest possible efficiency, while maintaining average input current compliance. It minimizes power lost in the linear charger with its Bat-Track feature.

First generation USB applications implemented a current-limited battery charger directly between the USB port and the battery, where the battery voltage powers the system. This is referred to as a battery-fed system. In a battery-fed system, the available system power is $I_{USB} \bullet V_{BAT}$ because V_{BAT} is the only voltage available to the system load. When the battery is low, nearly half of the available power is lost to heat within the linear battery charger element.

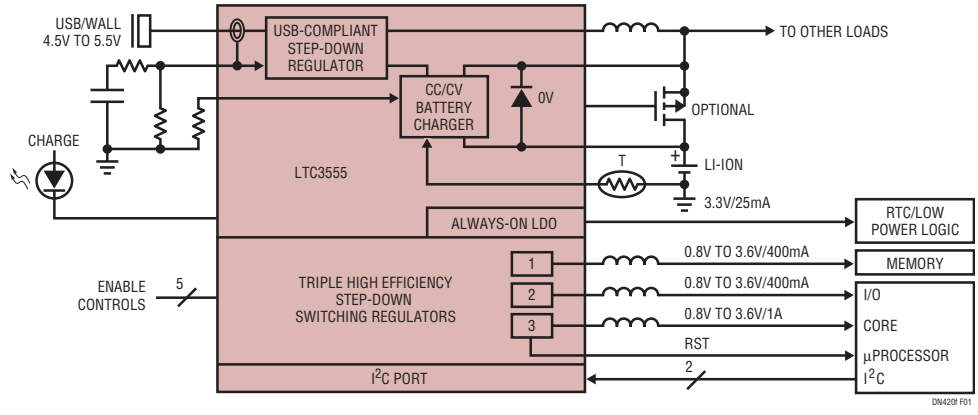

Figure 1. All-in-One USB Power Solution Includes Switching Power Manager, Battery Charger, Three Synchronous Buck Regulators and LDO

Second generation USB chargers developed an intermediate voltage between the USB port and the battery. This intermediate bus voltage topology is referred to as a PowerPath™ system. In PowerPath ICs, a current-limited switch is placed between the USB port and the intermediate voltage. The intermediate voltage, V_{OUT}, powers the linear battery charger and the system load. By using the intermediate bus voltage topology, the battery is decoupled from the system load and charging may be carried out opportunistically. PowerPath systems have the added benefit of being "instant-on" because the intermediate voltage is available for system loads as soon as power is applied to the circuit, independent of the state of the battery. In a PowerPath system, more of the 2.5W available from the USB port is made available to the system load as long as the input current limit has not been exceeded. PowerPath systems offer improvements over battery-fed systems, but significant power may still be lost in the linear battery charger element if the battery voltage is low.

The LTC3555 is the first IC in the third generation of USB PowerPath chargers. These PowerPath devices produce an intermediate bus voltage from a USB-compliant step-down regulator that is regulated to a fixed amount over the battery voltage (a Bat-Track feature). The regulated intermediate voltage is just high enough to allow proper charging through the linear charger. By tracking the battery voltage in this manner, power lost in the linear battery charger is minimized, efficiency increases and power available to the load is maximized.

Figure 2 provides an efficiency comparison and power savings between chargers with switching vs linear PowerPath systems. The amount of power saved while charging large batteries can make the difference between a device that runs in thermal limit and one that runs cool.

Complete Power Solution in a Single IC

The LTC3555 also contains three user configurable step-down DC/DC converters capable of delivering 0.4A, 0.4A and 1A. Regulator 1 has a fixed reference voltage of 0.8V while regulators 2 and 3 may have their reference voltage changed via the I²C interface between 0.8V and 0.425V. All of the converters operate at a switching frequency of 2.25MHz, allowing the use of small passive components while maintaining efficiencies up to 92% for output voltages greater than 1.8V (see Figure 3). All three regulators may be programmed to operate in pulse-skipping mode, Burst Mode® operation or LDO mode via the I²C port or through I/O pins. In Burst Mode operation the output ripple amplitude is slightly increased and the switching frequency varies with the load current to improve efficiency at light loads. If noise is a concern, all of the regulators may be set to operate in LDO mode or pulse-skipping mode. The device also provides an always-on 3.3V output capable of delivering 25mA for system needs such as a real-time clock or pushbutton monitor.

Conclusion

The LTC3555 is an advanced and complete power solution in a single chip. The third generation PowerPath management technology with its reduction in both heat generation and battery charge time is ideally suited for tomorrow's high density, feature rich battery-powered products. By integrating three I²C-controlled, highly efficient step-down DC/DC converters, the LTC3555 allows the system designer complete flexibility to adapt to changing demands and operating modes.

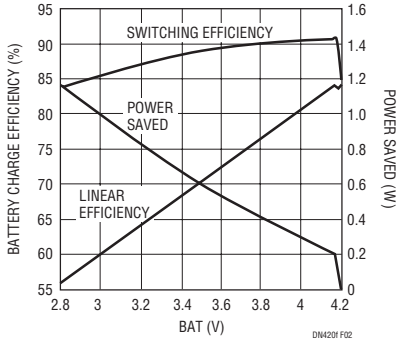

Figure 2. Switching PowerPath Battery Charger Efficiency and Power Savings Relative to a Linear Charger. (V_{BUS} = 5V, 5X MODE, R_{CLPROG} = 2.94k, R_{PROG} = 1k, I_{BAT} = 0.7A at V_{BAT} = 2.8V)

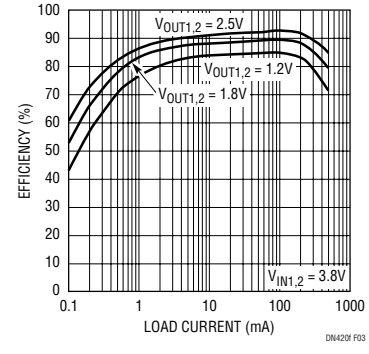

Figure 3. Efficiency of Switching Regulators 1 and 2 with Burst Mode Operation

Data Sheet Download

www.linear.com

For applications help,
call (978) 656-3755

floating. The positive terminal of the external dc-power supply feeds Pin 6 through a 1-kΩ resistor. Connecting the DUT's output to Channel 2A of the SR785 completes the test-circuit configuration.

Building the buffer/inverting summer with an AD8034 dual amplifier is a good choice because it has a supply range of 5 to 24V; a signal-frequency response well beyond 1 MHz; and a large capacitive-load-drive capability, allowing you to neglect the capacitance of test cables. Further, the AD8034 can deliver as much as 40 mA of load current.

To instill confidence that this buffer/inverting-summer configuration works, **Figure 2** proves that you can neglect any loss that the AD8034 incurs. The **figure** demonstrates that the response of the AD8034 buffer/inverting summer of 10 Hz to 10 kHz is approximately flat with a loss of only 0.0025 dB, and the loss from 10 to 100 kHz is approximately 0.024 dB. **Figure 3** shows negative-PSRR test results. The Hewlett-Packard (www.hp.com) HP8753 provides the PSRR-versus-frequency responses beyond 100 kHz. You can measure positive PSRR (**figures 4** and **5**) by connecting Pin 3 to the SR785's output port. Pin 1, V_{OUT} of the buffer amplifier, connects to the reference port of the SR785. Here, you use the first amp to isolate the output

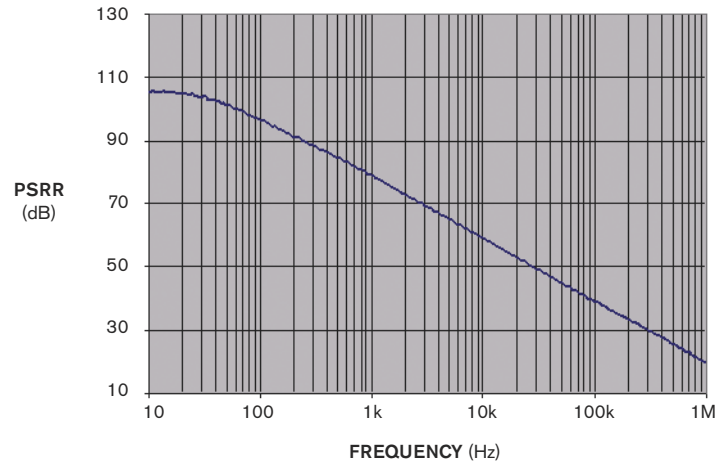

Figure 3 The negative-PSRR test results show responses beyond 100 kHz.

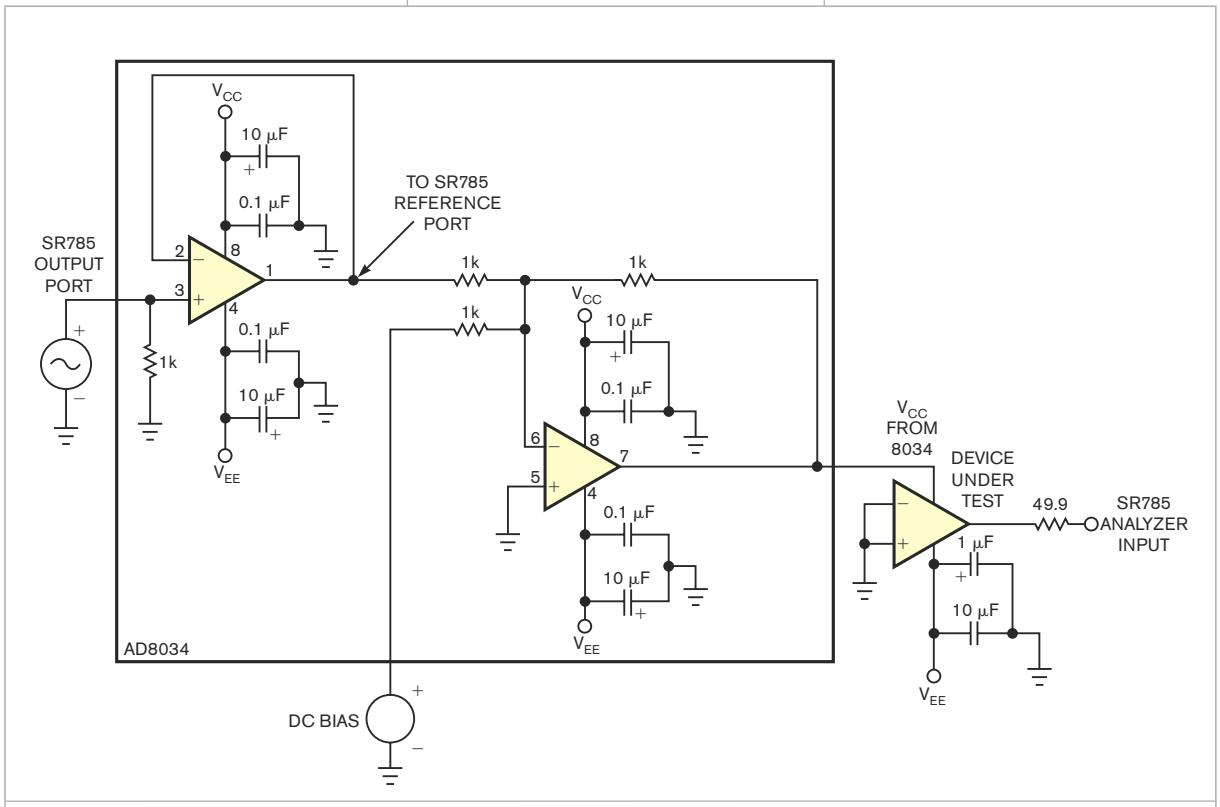

Figure 4 This test setup measures the positive PSRR. Note that bypass capacitors are absent from the device under test's positive-supply input.

port of the SR785 from the dc bias and provide the sinusoidal output. The second amplifier within the AD8034 sums the dc bias and sinusoid, which you use to feed the DUT's positive-supply pin. You must remove all bypass capacitors at the DUT's positive-supply pin. A 1-kΩ resistor from Pin 3 to ground prevents the noninverting input from floating. Feed the negative terminal of the external dc-power supply to Pin 6 through a 1-kΩ resistor. Connecting the DUT's output to Channel 2A of the SR785 completes the test-circuit configuration.

For the AD8034, assume that the DUT has a maximum supply voltage of ±15V, that you need to test negative PSRR, and that the DUT supplies ±10V. If you want to accommodate the maximum SR785 output of 5V peak, the first amplifier of the AD8034 needs enough head room to avoid clipping the ±5V signal from the output port of the SR785. In this case, a supply setting for the AD8034 of 6 and −16V is sufficient to prevent any problems.

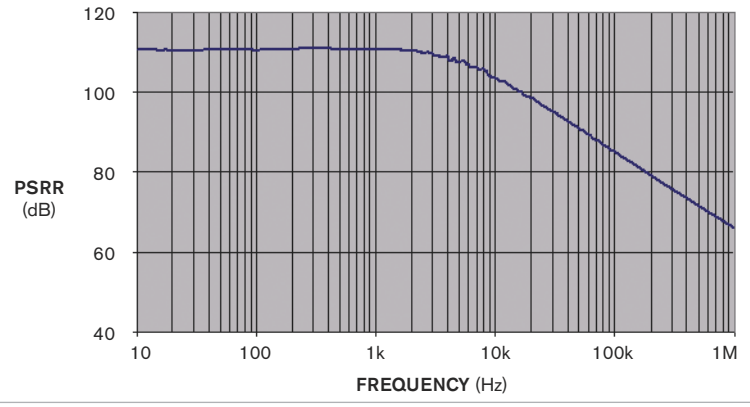

Figure 5 The positive-PSRR test results show responses beyond 100 dB.

This amount provides enough head room to accommodate the first amp of the AD8034, which handles a ±5V signal centered at ground. The −16V accommodates the dc bias of −10V and the ±5V signal centered at −10V at the output of the second amplifier of the AD8034. Positive PSRR is similar: Just set the AD8034 supplies to 16 and −6V for this example.

You might consider using separate power supplies for the DUT and the AD8034 to simplify matters. However, you can use the same dc-power supply for the DUT to provide the dc-bias voltage at Pin 6 of the AD8034 buffer/inverting summer. Choose the output voltage of the SR785 or whichever analyzer you use so that the DUT operates within its linear region of operation. You can apply this technique to other applications. EDN

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

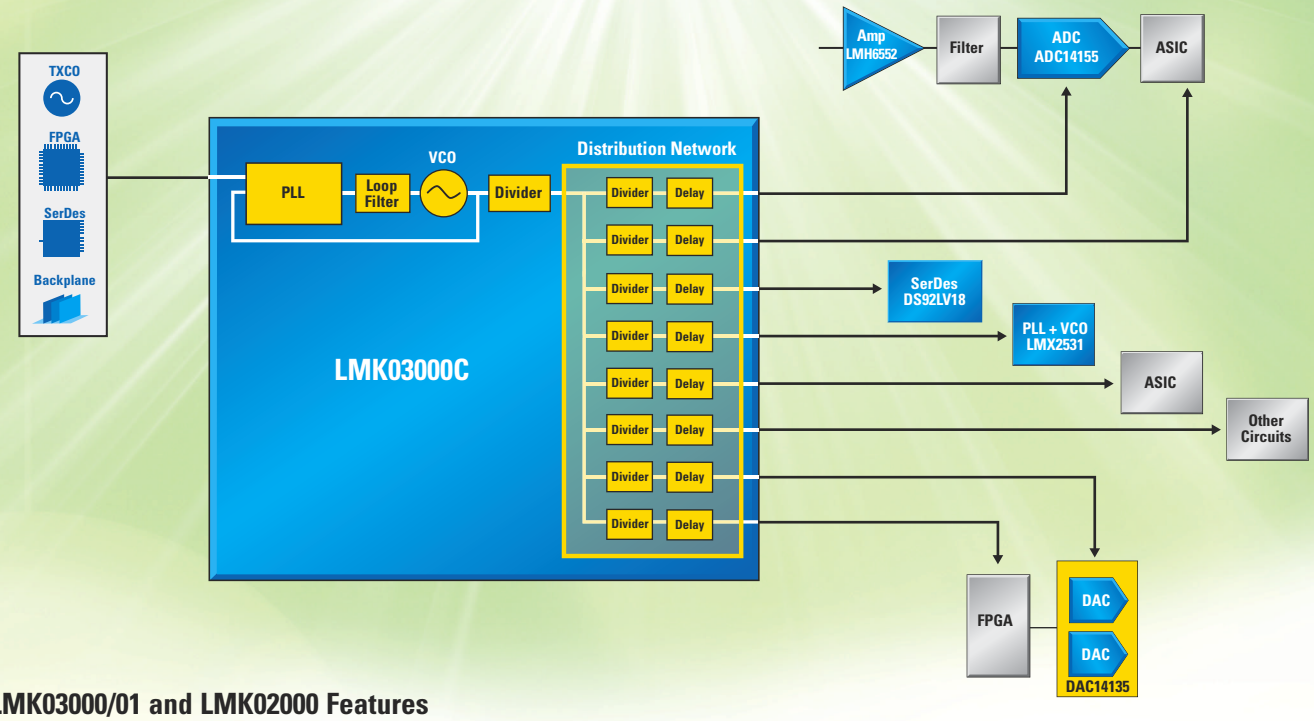

LMK03000/01 and LMK02000 Features

- 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

Accelerating the **Speed of Design.**™

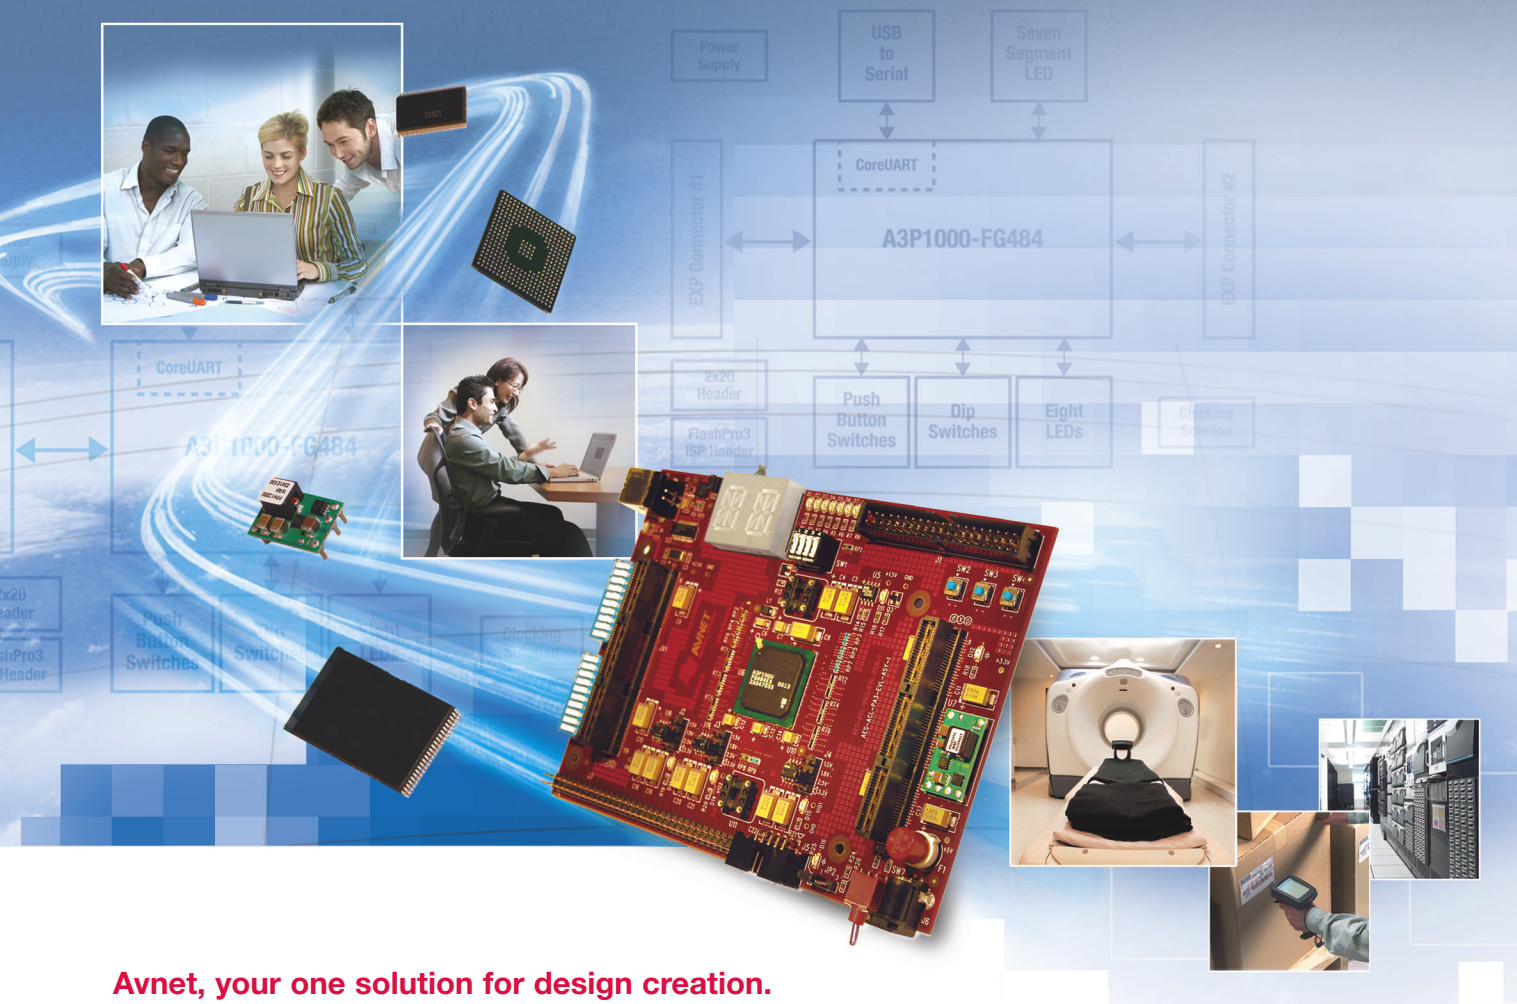

Avnet, your one solution for design creation.

Avnet Memec focuses on a select number of leading suppliers to provide original equipment manufacturers (OEMs) with the technical expertise you need in designing cutting edge products.

With factory trained and certified technical support, Avnet Memec specialists offer a level of technical depth that is unmatched in the industry. That translates into faster time to market for you, our valued customer. And it keeps getting better. **Avnet is ranked Best-In-Class* for seminars, technology, education and training, as well as design and engineering services. We deliver:**

> Proactive Engineering Support

> Factory Certified FAEs

Accelerate Time to Market.™
Visit the Avnet Design Resource Center™ at: www.em.avnet.com/drc

AVNET® Memec

1 800 408 8353
www.em.avnet.com

Avnet Green Initiative

NEDA Supplier Authorized Distributor

*As rated by Hearst Electronics Group, 2006 The Supplier Interface Study.

EDITED BY SUZANNE DEFFREE

supplychain

LINKING DESIGN AND RESOURCES

How much will the Qualcomm chip ban impact the global mobile-phone market?

The ITC's (International Trade Commission's) recent decision to ban US imports of some mobile phones that include certain Qualcomm (www.qualcomm.com) 3G chips will affect the industry beyond just the company itself, but to what extreme is questionable. The ITC's determination stems from a patent battle between Broadcom (www.broadcom.com) and Qualcomm; the commission has banned importation of infringing Qualcomm chips and future products, such as 3G cellular phones that use those chips. At press time, Qualcomm planned to appeal to President George W Bush to veto the ITC's decision.

The CTIA (Cellular Telecommunications Industry Association, www.ctia.org) backs Qualcomm's argument for a veto. "The ITC's importation ban, if implemented, will force the redesign of virtually all handsets that utilize the banned chips," Steve Largent, CTIA's chief executive officer, wrote in a letter to the president. "This redesign process requires a substantial amount of time (18 to 24 months) and many millions of dollars for each of the companies involved."

CTIA's chief executive officer, Steve Largent, is urging President Bush to veto the ITC's mobile-phone ban.

Meanwhile, researchers at iSuppli Corp estimate that the ban will affect 4.2 million shipments of EVDO (Evolution Data Optimized) and WCDMA mobile phones in 2007, representing only 4.4% of North American mobile-phone shipments and 3.2% of worldwide 3G mobile-phone shipments in the second half of 2007. According to analyst Tina Teng, the ban would affect only 11 mobile-phone models in 2007, or 0.9% of phone-model introductions for the year.

Analysts at iSuppli say that the ban will have the greatest impact on OEMs Samsung Electronics, LG Electronics, and Motorola. They do not expect the ban to reduce volume shipments of mobile phones overall this year.

GREEN UPDATE

"PC" MEANS "POWER CONSCIOUS" IN NEW ENERGY STAR REQUIREMENTS

Companies that use the EPA's (Environmental Protection Agency's) Energy Star logo on their products as marketing leverage will need to meet new, stricter design requirements. The first phase of Energy Star 4.0 goes into effect July 20 and includes new performance requirements to qualify for the rating on desktop, notebook, and tablet computers; workstations; integrated computers; desktop-derived servers; and game consoles.

Version 4.0 defines a set of testing criteria and power limits that could reduce the amount of energy that idle equipment consumes by an average of 45%, according to the EPA. To do so, Energy Star 4.0 requires, regardless of the system type, that the display enter a sleep state after 15 minutes of system inactivity; the platform enter a sleep state after 30 minutes of system inactivity; and the platform's GbE (Gigabit Ethernet) link switch to a lower rate mode when entering the sleep mode. In addition, systems for an enterprise market must include wake-on LAN during the sleep state. The specification includes certain additional power limits, which depend on the device type and power state.

The EPA has set a goal of 25% compliance for each of the platform categories. The agency plans a second phase of Version 4.0 in 2009.

OUTLOOK

ASPs DOWN, BUT DEMAND SOLID

As have many industry watchers, the SIA (Semiconductor Industry Association, www.sia-online.org) dramatically downgraded its 2007 global microchip-sales-growth forecast from its earlier 10% estimate to 1.8% in recent weeks. The SIA notes sharp declines in ASPs (average selling prices) for microchips in key segments, including microprocessors and DRAMs, as prime contributors to the slower growth.

The recently lowered expectations are for sales, and, when you take lower ASPs into consideration, the supply-chain picture is more optimistic than it might seem. The SIA reports that the end markets that drive sales of these products continue to be in line with previous forecasts, keeping unit demand for semiconductors strong.

SEMI (Semiconductor Equipment and Materials International, www.semi.org) concurs. "Unit growth for some key materials has been increasing since the early part of the first quarter. So, while revenue forecasts by some of the analysts have been downgraded, we're starting to see growth for units in materials, and that corresponds to growth in demand for semiconductors," said Jonathan Davis, a SEMI vice president, in the group's June Voices of the Industry audiocast. The new SIA forecast projects total sales of $252 billion in 2007.

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

productroundup

DISCRETE SEMICONDUCTORS

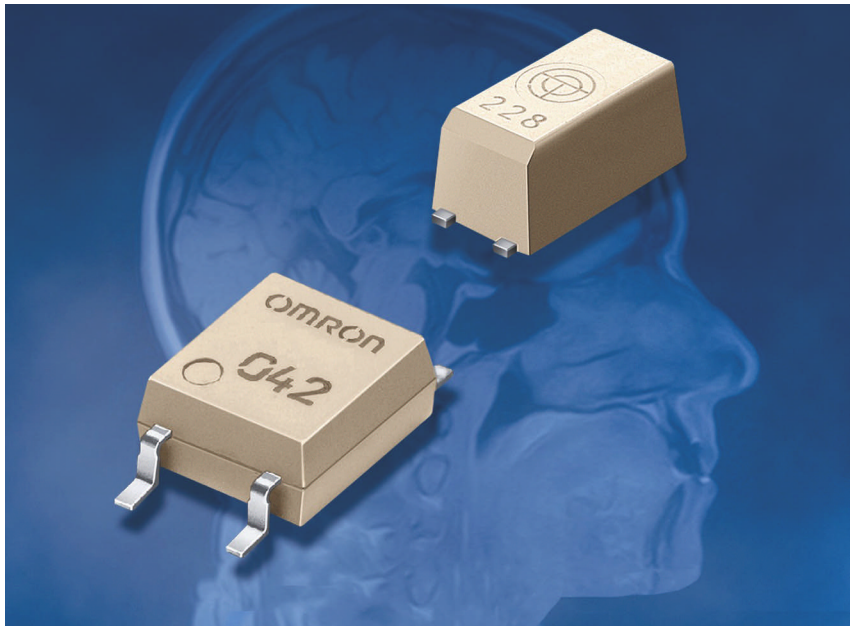

MOSFET relays claim a low range of capacitance, resistance values

Available in a 1-Form-A contact configuration, the G3VM-21LR and 41LR series of MOSFET relays come in SSOPs. The devices provide fast turn-on and -off times in the microsecond range and handle ac and dc loads. Suiting medical-imaging applications, the devices feature 1Ω on-resistance, 1-pF off-state capacitance, and 2.5-pFΩ range of capacitance and resistance values. The 21LR and 41LR series measure 2×3.8×1.8 mm and sell for $6.79 to $7.10 (10,000).

Omron Electronic Components LLC, www.components.omron.com

Eye-safe diode lasers come as single diodes and vertical stacks

The high-peak-power HPD Series 2400 eye-safe laser diodes feature 1550-nm pulsed lasers, targeting distance-measurement applications requiring eye-safe laser sources. Operating at 200-nm pulse widths and a 0.1% duty cycle, the lasers come in 150- to 380-micron source sizes. Available in single diodes and stacked arrays, the family comprises the HPD2410 at 5W and the HPD2424 at 12W diodes and the HPD2410-2S at 10W, the HPD2424-2S at 24W, and the HPD2424-4S at 48W vertical stacks. The four-chip vertical array provides 48W output power. Available in 5.6-mm, 9-mm, TO18 coaxial, and TO5-twin-lead packages, the HPD Series 2400 laser diodes cost $255.

Intense Ltd, www.intenseco.com

Power switch conserves energy using valley-switching technique

The Green FPS (Fairchild Power Switch) e-Series touts high energy efficiency and reliability and suits DVD players, set-top boxes, LCD monitors, and other 25W and lower power-supply designs. Using the vendor's proprietary "valley-switching" technique, the devices combine the functions of avalanche-rated SenseFETs, current-mode-PWM (pulse-width-modulator) ICs, and various protection features. The power switch has a 0.2W standby-power consumption at no-load conditions and less than 1W power consumption at a 0.5W load. Units in the Green FPS e-Series costs 84 cents each.

Fairchild Semiconductor, www.fairchildsemi.com

DirectFET MOSFET chip set has a low on-resistance

Suiting high-current dc/dc converters in notebooks, high-end desktops and servers, and advanced telecom and datacom systems, the DirectFET MOSFET chip set includes the IRF6712S control FET and IRF6716M synchronous FET. The IRF6716M provides a 1.2-mΩ typical on-resistance at a 10V gate-to-source voltage and 2-mΩ on-resistance at a 2V gate-to-source voltage. The IRF6712S achieves a 12-nC gate charge and a 4.4-nC gate-to-drain charge, as well as a 3.8-mΩ typical on-resistance at a 10V gate-to-source voltage and a 6.7-mΩ on-resistance at a 4.5V gate-to-source voltage. Available in an MX-pad outline with a 0.7-mm profile, the IRF6712S and IRF6716M cost 59 cents (10,000) and $1.27 (10,000), respectively.

International Rectifier, www.irf.com

Overvoltage-protection controllers require no integrated charge pump

Targeting cell phones, MP3 players, PDAs, and other devices using

productroundup

DISCRETE SEMICONDUCTORS

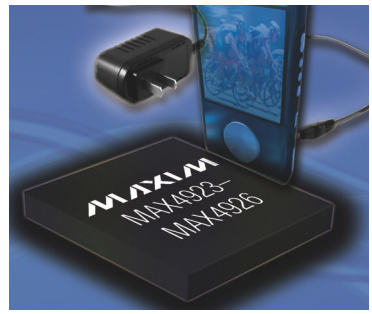

Schottky rectifiers and TVS devices come in micro-SMPs

Available in a micro-SMP (surface-mount package), these Schottky rectifiers and TVS (transient-voltage-suppressor) devices suit digital cameras, MP3 players, navigation systems, and cellular phones. The rectifiers provide a 30°C/W junction-to-lead thermal resistance, and the TVS devices provide a 20°C/W junction-to-lead thermal resistance. Features include a 25A forward-surge-current capability for an 8.3-msec pulse. The rectifiers operate at 1A at a 110°C lead temperature, and the two voltage suppressors provide a 100W peak pulse current at a 10/1000-sec waveform. Measuring 2.5×1.3×0.65 mm, the Schottky rectifiers and TVS devices cost 2.5 cents each (1000).

Vishay Intertechnology, www.vishay.com

battery charging from a wall adapter or a USB port, these overvoltage-protection controllers drive external PFETs, drawing a typical supply current of 13 µA. The controllers reduce excessive current drain by using an external PFET, eliminating the need for an integrated charge pump. Available in 1.5×1-mm micro-DFN-6 packages, the MAX4923, MAX4924, MAX4925, and MAX4926 cost 67 cents each (1000).

Maxim Integrated Products, www.maxim-ic.com

EMBEDDED SYSTEMS

Synchro/resolver-to-digital converter comes on a PC-104 card

The DSP-based, six-channel 73-SD4 synchro/resolver-to-digital converter on a PC-104 card features six independent, transformer-isolated, programmable synchro/resolver-tracking-converter-measurement channels. Each channel has 16-bit resolution, a ±1 arc-minute accuracy, a 150-rps tracking rate, an accurate digital-velocity output, A and B incremental encoder outputs, and a wraparound self-test. The channels are programmable for speed ratios of 1-to-1 and 255-to-1. Requiring a 5V-dc power supply, the unit operates over a 47-Hz to 10-kHz frequency range and has a 2 to 28V-rms autoranging-input range. Additional features include 16 programmable digital-I/O channels, a latch for simultaneously reading all channels, and an optional programmable excitation-reference supply. The 73SD4 costs $2500 (100).

North Atlantic Industries, www.naii.com

SATA-to-CF adapter suits through-panel mounting

Targeting use in embedded systems, the CFADPTCS-SATA SATA-to-CF (Serial Advanced Technology Attachment-to-CompactFlash) adapter connects a CF drive to any system using a SATA port. The device includes a CF-activity LED and an eject button for the CF connector. Requiring 5V power, the device suits through-panel mounting. Measuring 2.6×2.25 in., the CFADPTCS-SATA costs $29.

Mesa Electronics, www.mesanet.com

EMBEDDED SYSTEMS

XMC modules run at 1.5 and 3G samples/sec

The AD1500 and AD3000 high-speed analog-input XMC (express-mezzanine-card) modules feature the Xilinx Virtex-5 FPGA. The dual-channel AD1500 ADC runs at 1.5G samples/sec; the single-channel AD3000 ADC runs at 3G samples/sec. The modules' analog input uses an 8-bit National Semiconductor ADC083000 or an ADC081500 converter. The designs have an FPGA back end with a Virtex-5 SX95T or an LX110T FPGA option and have a direct connection to the analog input. The AD1500 and AD3000 cost $13,995 each.

VMetro, www.vmetro.com

Digital-I/O PCI card provides 128 opto-isolated channels

The high-density PCI-7442, -7443, and -7444 digital-I/O PCI cards provide fast card detection and troubleshooting in systems with multiple digital-I/O cards. Each device delivers 128 optoisolated channels. The PCI-7443 provides 128 identical nonpolarity channels of optoisolated digital inputs, and the PCI-7444 module has 128 channels of open-drain, power-MOSFET-drive digital outputs. At a 100% duty cycle, the PCI-7444 provides 250-mA sink current per channel with a 300-mA maximum load. The series supports Microsoft Windows and Linux operating systems and includes software drivers for Visual Basic, Visual C, BCB, Delphi, LabView, and Matlab. The PCI-7442, -7443, and -7444 cost $495 each.

Adlink Technology, www.adlinktech.com

Family of USB digital-I/O devices features locking USB connectors

The SeaDAC Lite family of USB digital-I/O devices provides a compact alternative for interfacing USB-equipped computers with real-world signals. The first two products of the family feature four optically isolated inputs and four Reed-relay outputs or Form C-relay outputs. The family includes SeaLatch locking USB connectors, preventing accidental cable disconnection. Operating at 0 to 70°C, the modules are also available in a −40 to 85°C temperature range. The SeaDAC costs $159.

Sealevel Systems, www.sealevel.com

Six-channel ADC integrates FPGA-based DSP

The Tarvos VXS combines a six-channel, 16-bit, 160M-sample/sec ADC with an FPGA-based DSP and a DAC-output channel in one slot. The ADCs link into a Xilinx FPGA with a 5-Gbyte double-data-rate-SDRAM architecture. The Tarvos VXS costs $19,900.

TEK Microsystems, www.tekmicro.com

EDA TOOLS

Multimode simulation platform upgrades all simulators for custom design

The Virtuoso MMsim (multimode-simulation) platform Version 6.2 provides custom-designed simulators with a common kernel and database of netlists and models, assisting designers in effectively simulating analog, RF, memory, or mixed-signal designs. The upgrades of the MMsim tools ex-

productroundup

EDA TOOLS

tend to the Virtuoso Spectre circuit simulator, UltraSim full-chip simulator, and AMS Designer. Improvements to the engine in the Spectre XL circuit simulator make the simulator three times faster than traditional SPICE simulators and provide 10-times-faster enhanced Monte Carlo analysis than the previous version. Integrating analog-, RF-, and high-speed-IC-simulation capabilities, Spectre XL features an enhanced frequency-domain, multirate, harmonic-balance engine; high dynamic range; and weakly nonlinear-RF circuits. Suiting highly nonlinear circuits, the simulator features a time-domain algorithm and a flow for analysis of analog noise and jitter in PLLs (phase-locked loops). The UltraSim XL full-chip simulator runs custom designs 10 times faster than the previous version. The EM and IR features now support electrical verification of memories and large analog- and mixed-signal devices. Improving the link between the AMS Designer and its custom simulator provides enhanced mixed-signal RF with integration to Spectre XL and improved performance with UltraSim XL for SOC (system-on-chip) verification. The Virtuoso MMsim Version 6.2 costs $24,000 for a one-year license.

Cadence Design Systems, www.cadence.com

EDN ADVERTISER INDEX

Company	Page
Agilent Technologies	49
Altera Corp	27
Analog Devices Inc	2, 29
	36A-36H
Ansoft Corp	54
austriamicrosystems AG	88
Avnet Electronics Marketing	90
Avtech Electrosystems Ltd	97
CML Microcircuits (UK) Ltd	96
Delta Electronics	97
Digi-Key Corp	1
Fairchild Semiconductor	11
Freescale Semiconductor	35
Intersil	37
	39, 56
Keithley Instruments Inc	31
Kemet Electronics Corp	51
LeCroy Corp	8
Linear Technology Corp	73
	75, 76
	85-86
Mathworks Inc	41
Maxim Integrated Products	79
	81, 83
Melexis Inc	74
Mentor Graphics	53
	55, 57
Micrel Semiconductor	47
Microchip Technology	63, 65
Mill Max Mfg Corp	43
Monolithic Power Systems	52
Morsun Guangzhou Science	96
National Instruments	21, 61
National Semiconductor	4
	13-20
	89
Performance Motion Devices	95
Pico Electronics	67, 94
Power-One Inc	44
Radicom Research	97
Renesas Technology Corp	12
Samtec USA	72
Speeding Edge	96
STMicroelectronics	C-4
Summit Microelectronics	22
Tern	97
Texas Instruments	C-2
	3, 6
Toshiba America	68
Tyco Electronics Corp	33
Vicor Corp	71
WinSystems	25
Xilinx Inc	C-3

EDN provides this index 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 HOT CHIPS

The IEEE Technical Committee on Microprocessors and Microcomputers again hosts the pre-eminent technical forum on high-performance processing ICs at Stanford University (Palo Alto, CA) Aug 19 to 21. Appropriately enough for the times, the conference will include three sessions on multicore architectures, one of which—a keynote—will examine the future of the x86 architecture. Recognizing the growing role of software in IC-architectural decisions, organizers have planned another session titled "Technology and Software Directions." The forum will devote one session to exploring the IBM Power6 architecture. Sunday tutorials will include "Approaches to System Design for the Working Engineer" and "Enterprise Power and Cooling: a Chip to Data Center Perspective." It's hard to find a more concentrated source of technical conversations on the chip architectures that are making the news.

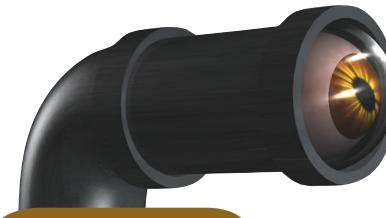

LOOKING BACK

AT EARLY COMPUTERIZED WEAPONS SYSTEMS

Performing 136,000 mathematical steps in under one second, the Multi-Weapon Automatic Target and Battery Evaluator assists antiaircraft-operations-center officers in organizing defense against approaching enemy aircraft. The MATABE contains more than 20 miles of wiring in seven cabinets. Its heart is a Burroughs Corp-developed real-time electronic digital computer, capable of 200,000 multidigit additions per second. That speed allows the MATABE to calculate a missile's time-to-burst point, the response time of the missile battery, the missile-intercept point, whether that point is within effective range, the relative threats of targets, and which target should be attacked first. The machine keeps a running account of the firing history on each target and records this information on punched paper tape.

—Electrical Design News, July 1957

LOOKING AROUND

FOR SOME MISSING CELL PHONES

There is an intimate connection between the health of the cell-phone-handset business and the glow in the eye of the electronics industry. Nothing can absorb that many chips—RF chips, SOCs (systems on chips), system controllers and power chips, memory chips, plus a lot more—without having influence. So, when the Semiconductor Industry Association starts talking about slower growth in handset sales, everyone takes notice. The trouble is that at least one indicator suggests that the growth isn't slower; it's gone altogether. One analyst recently stated that, for the first time, shipments of displays for handsets fell compared with the year-earlier period. Because virtually every handset has a display, and because this industry is not famous for letting inventories get out of balance, it sounds like handset shipments may be in for a decline. That news could be sobering in a lot of areas.

Low-Power Transceivers

Ultimate Connectivity . . .

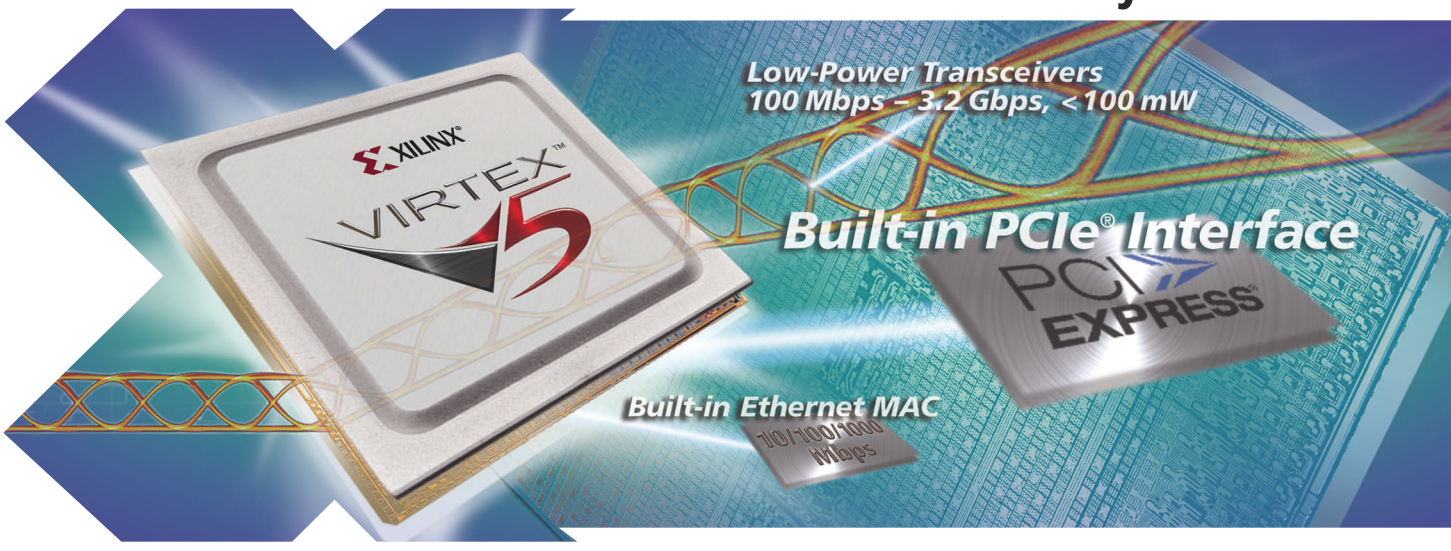

VIRTEX™ 5

Reduce serial I/O power, cost and complexity with the world's first 65nm FPGAs.

With a unique combination of up to 24 low-power transceivers, and built-in PCIe® and Ethernet MAC blocks, Virtex™-5 FPGAs get your system running fast. Whether you are an expert or just starting out, only Xilinx delivers everything you need to simplify high-speed serial design, including protocol packs and development kits.

Available Now!
Built-in x1-x8 PCIe Solution
$2,200 USD MSRP

Complete PCI Express® Development Kit

Lowest-power, most area-efficient serial I/O solution

RocketIO™ GTP transceivers deliver up to 3.2 Gbps connectivity at less than 100 mW to help you beat your power budget. The embedded PCI Express endpoint block saves 60% in power and 90% in area compared to other solutions. Virtex-5 chips and boards are on the PCI-SIG® Integrators List. Plus our UNH-tested Ethernet MAC blocks make connectivity easier than ever.

Visit our website to get the chips, boards, kits, and documentation you need to start designing today.

EDN innovation
HONORING EXCELLENCE IN ELECTRONICS
2006 WINNER

Delivering Benefits of 65nm FPGAs Since May 2006!

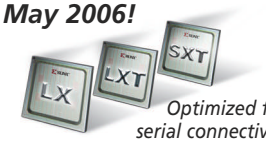

Optimized for logic, serial connectivity and DSP.

XILINX®

www.xilinx.com/virtex5

The Ultimate System Integration Platform

Embedded Systems

Solutions

Amazingly Small, Tremendously Useful ST7UltraLite 8-Pin Microcontroller.

The super small 8-bit MCU that offers more.

Innovative products for multi-segment application systems

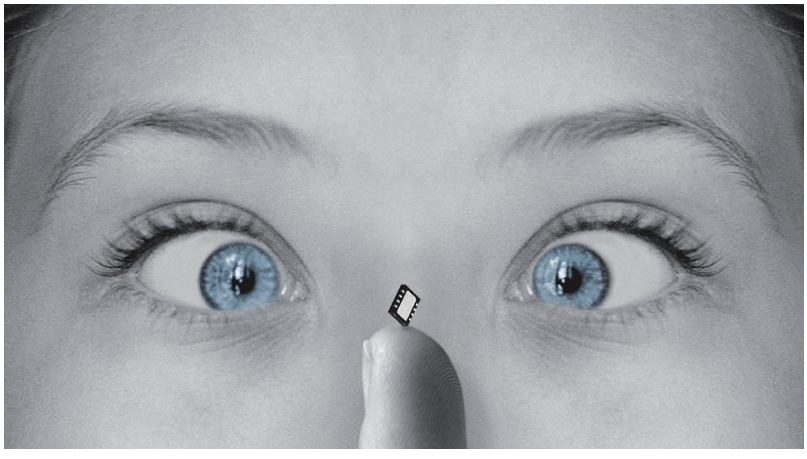

With a footprint of less than 16mm², the ST7UltraLite in an 8-bit DFN8 package includes 6 I/Os, a 10-bit A/D converter and a 12-bit auto-reload timer.

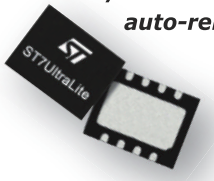

ST7UltraLite Primer complete low-cost mini application

ST's new series of 8-bit microcontrollers in tiny 8-pin packages, are intended to be used in simple, very cost-sensitive applications such as security, lighting, power control, small appliances, sensors and motor control. The new ST7FLITEUS series extends the embedded peripherals of the popular and easy-to-use ST7Lite Flash MCU family into very inexpensive, very small devices, which enable users to develop applications quickly and easily.

Key Features

- 1KB Flash or 2KB Flash with 128 bytes EEPROM
- Up to 8MHz Internal RC Oscillator, Low Power and Low Frequency Clock
- 8-bit Lite Timer, 12-bit Auto-reload Timer, 5 Input Channels, 10-bit ADC, 6 I/Os
- SO8, DIP8 and DFN8 Packages
- -40°C/125°C Temperature Range Qualified

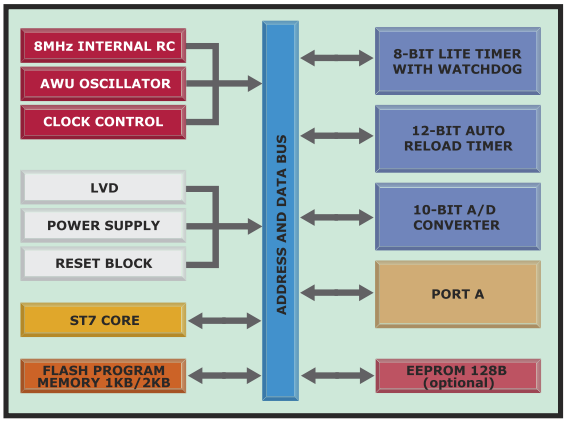

For further information, datasheets and application notes, visit www.st.com/st7

Solutions @

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