UBM Electronics

EDN®

VOICE OF THE ENGINEER

FEB 3
Issue 3/2011
www.edn.com

UBM Electronics

EDN®

VOICE OF THE ENGINEER

FEB 3
Issue 3/2011
www.edn.com

ADVANCES IN ENERGY-STORAGE TECHNOLOGY POWER WIRELESS DEVICES

Page 26

UNDERSTANDING EMBEDDED-SYSTEM-BOOT TECHNIQUES

Page 18

LOWERING THE COST OF MEDICAL-IMAGING R&D

Page 36

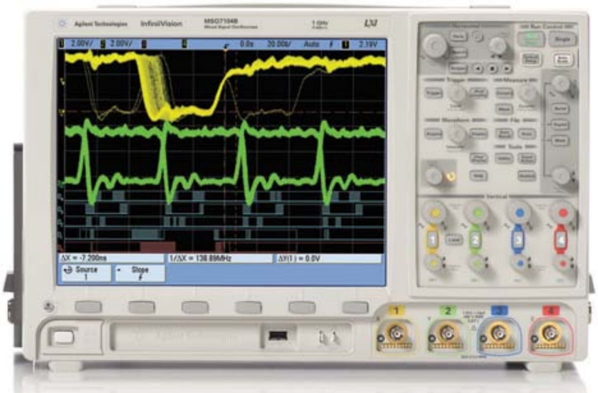

Agilent InfiniiVision 7000B Series (MSO and DSO models) Bandwidth: 100 MHz - 1 GHz

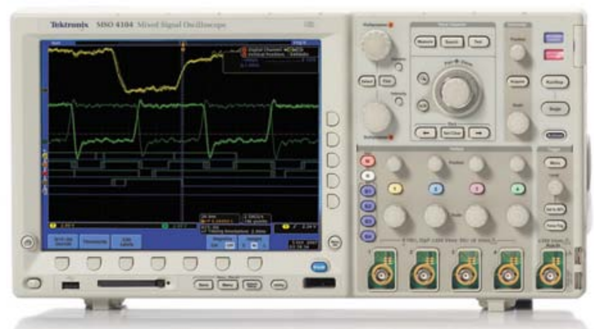

Tektronix 4000 Series

40% bigger screen.

475,000 times faster...Yes, really.

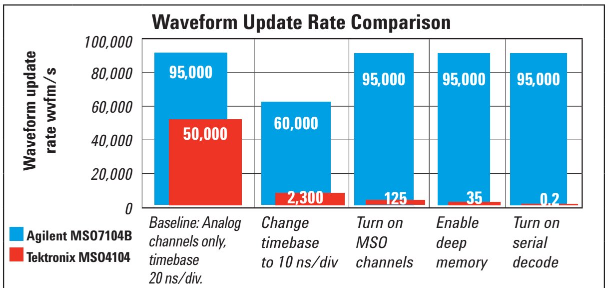

Notes:
- Measurements taken on same signal using Agilent MSO7104B and Tektronix MSO4104 .
- Memory depth = display window X sample rate with up to 8 Mpts on Agilent.
- Tektronix measurements taken with version 2.13 firmware.
- Screen images are actual screen captures, and scopes are shown to scale.

A scope's real world performance depends on more than just bandwidth and sample rates. It takes an ultra fast update rate to capture infrequent events and display the most accurate signal detail. That's where Agilent 7000B Series scopes really stand out. Thanks to a proprietary chipset, you get faster update rates not only in analog mode, but in every mode. That's why you see greater signal detail. That's Agilent.

Agilent and our Distributor Network

Right Instrument.
Right Expertise.
Delivered Right Now.

METRICTEST™
6,000 instruments. One source.

866-436-0887
www.metrictest.com/agilent

Are you using the best scope?
Take the 5-minute Scope Challenge.
Visit www.metrictest.com/agilent/scope.jsp

Agilent Technologies

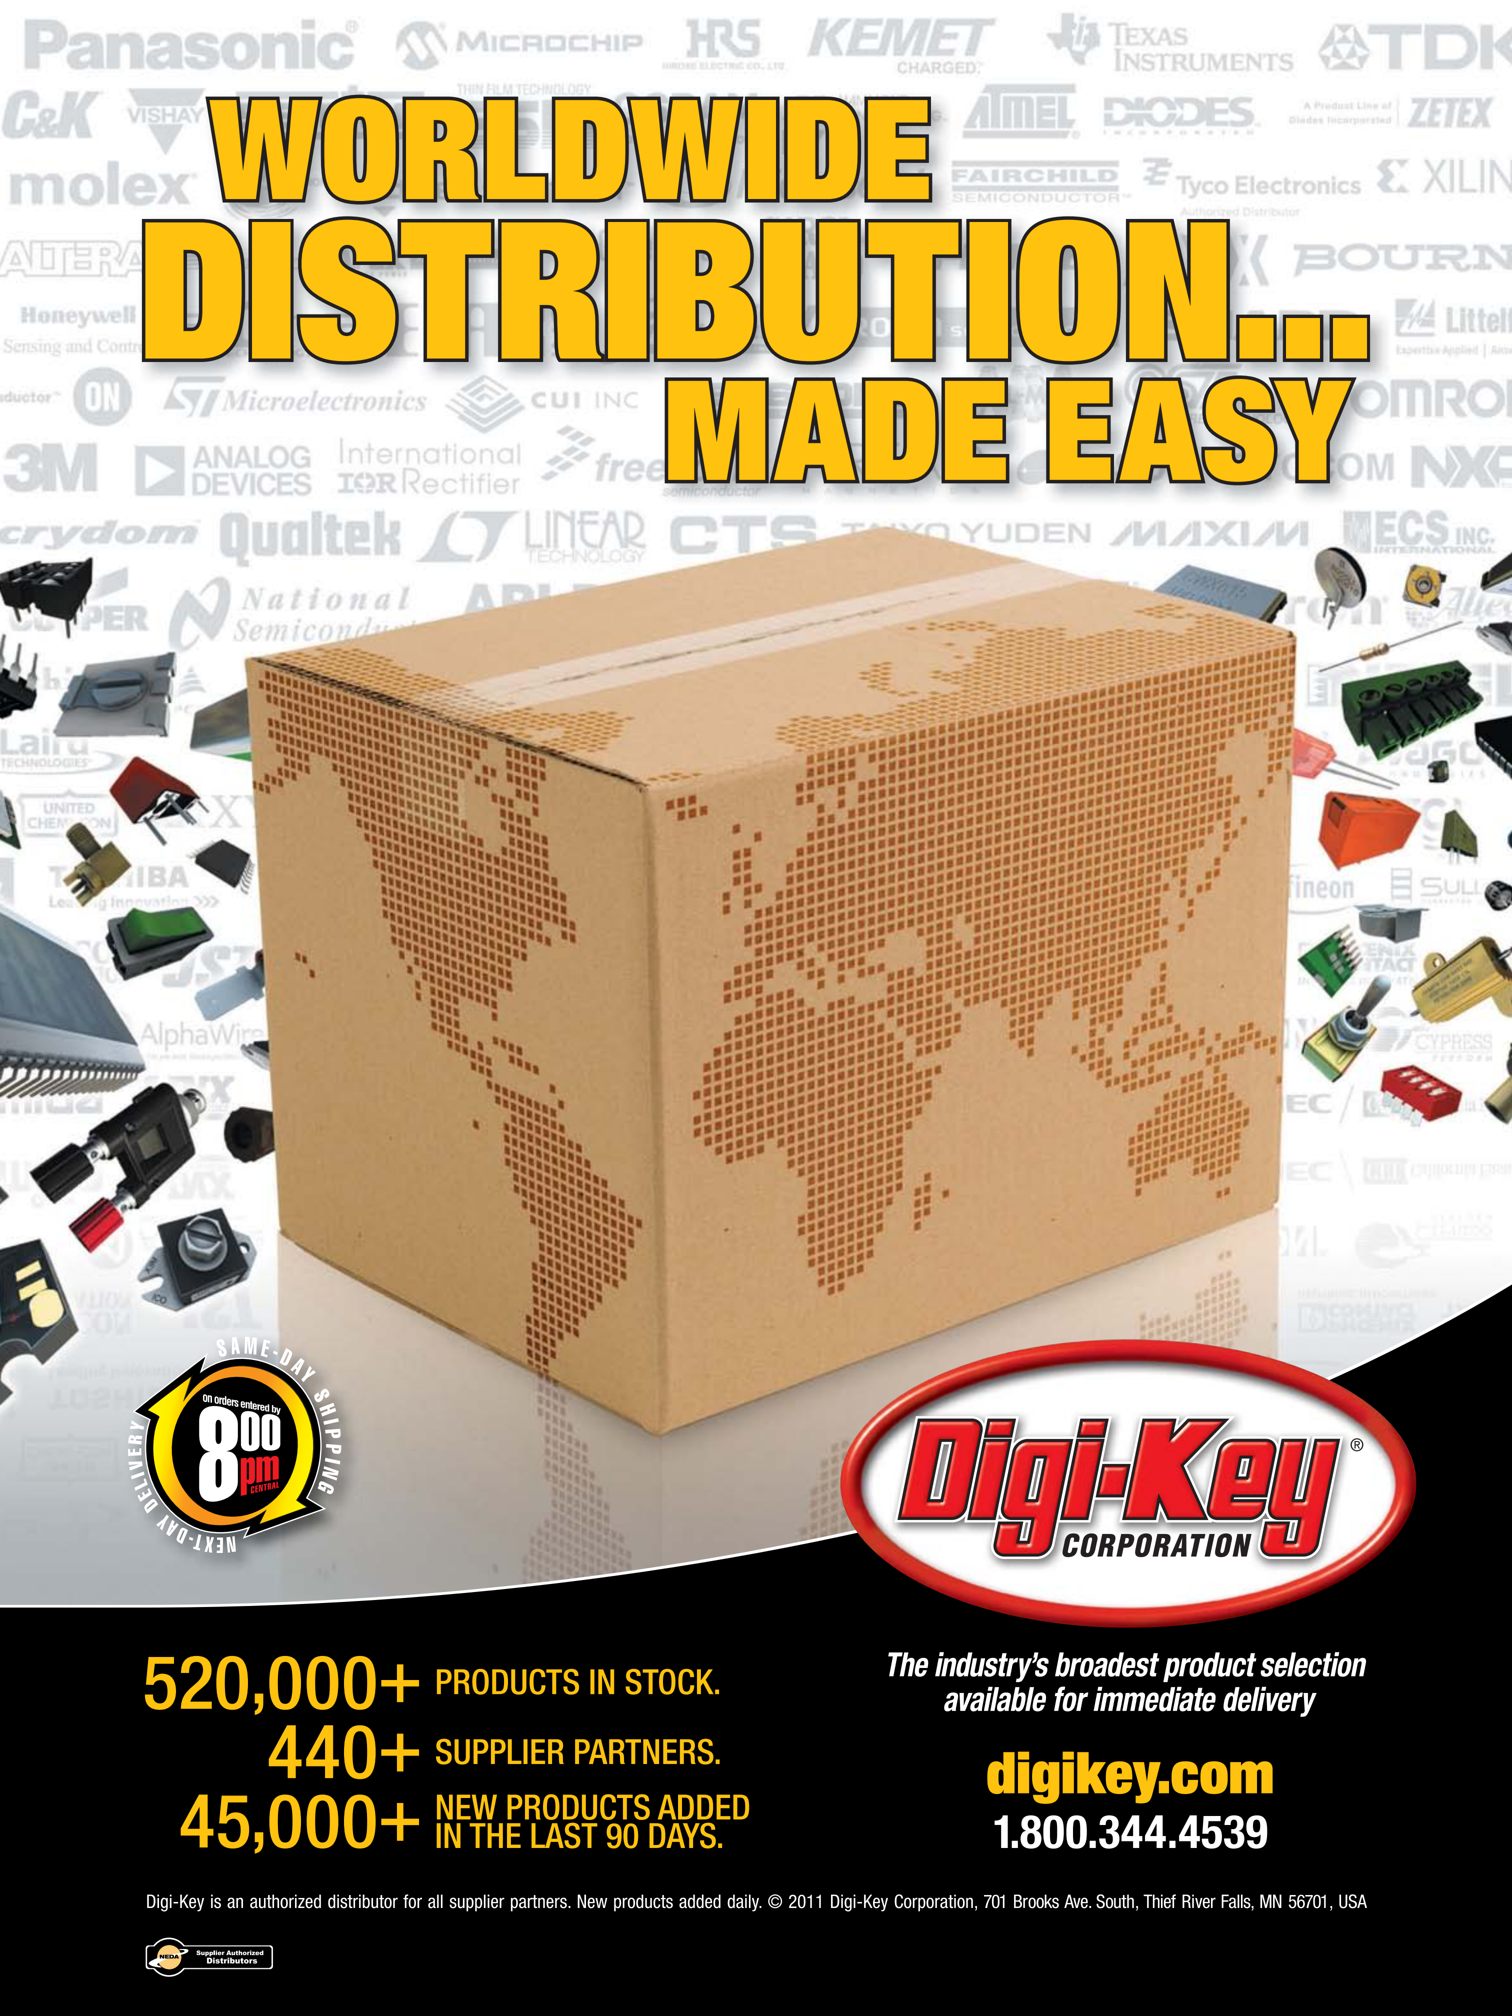
WORLDWIDE
DISTRIBUTION...
MADE EASY
SAME-DAY SHIPPING
on orders entered by
8:00 pm CENTRAL
NEXT-DAY DELIVERY
Digi-Key®
CORPORATION
520,000+ PRODUCTS IN STOCK.
440+ SUPPLIER PARTNERS.
45,000+ NEW PRODUCTS ADDED IN THE LAST 90 DAYS.
The industry's broadest product selection available for immediate delivery
digikey.com
1.800.344.4539
Digi-Key is an authorized distributor for all supplier partners. New products added daily. © 2011 Digi-Key Corporation, 701 Brooks Ave. South, Thief River Falls, MN 56701, USA
NEDA Supplier Authorized Distributors

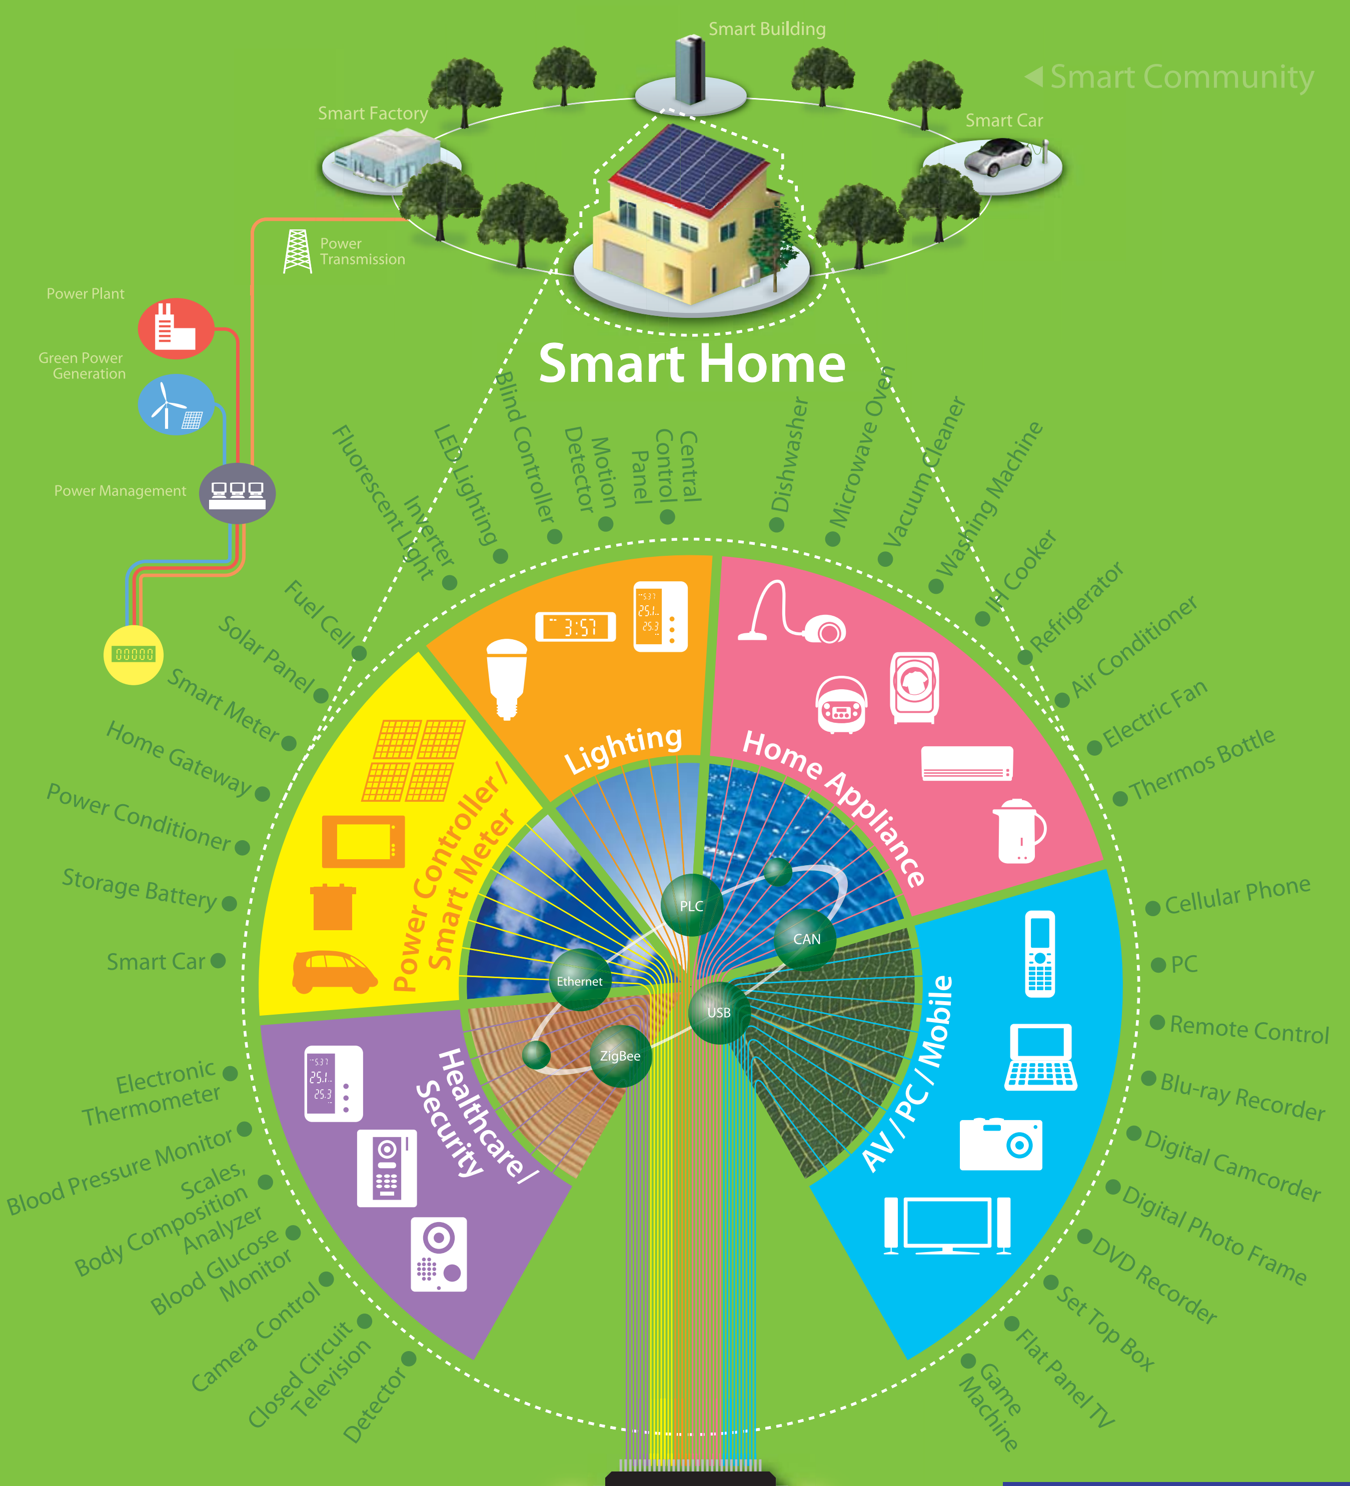

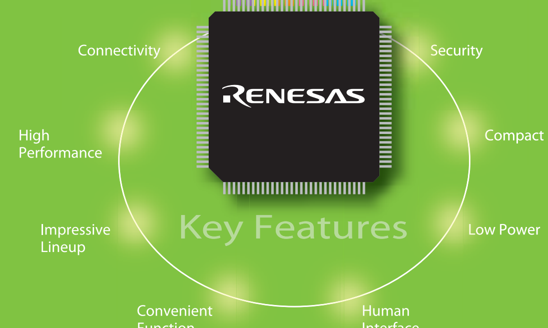

Get started now
for smart development at
www.renesas.com/smart

Renesas Electronics Corporation

Green Solutions ready for you

RENESAS

EDN contents 2.3.11

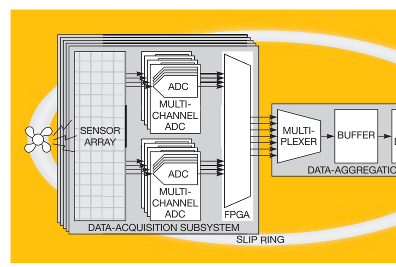

pulse

COVER: COMPOSITE ILLUSTRATION BY TIM BURNS.
VAULT: MARK EVANS/ISTOCKPHOTO.COM;
CIRCUIT PATTERN: MON5TER/ISTOCKPHOTO.COM

DESIGNIDEAS

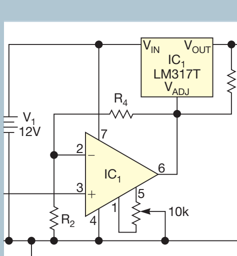

Improving RF Transmission Efficiency

Tiny RF PA Power Management IC

RF Power amplifiers used in cell phone applications are typically powered directly from the battery. This creates a challenge for the designers: as RF devices, the power amplifier adjusts the voltage and current based on the transmission. Inefficiencies are created when the PA supply voltage and the output voltage are different.

The MIC2808's DAC controlled adjustable voltage output feature minimizes the input to output voltage differential, thus improving efficiency and battery life. The voltage is adjusted based on real-time needs of the PA to achieve the most efficient operation at any given transmit power level.

For more information, contact your local Micrel sales representative or visit Micrel at: www.micrel.com/ad/mic2808.

Ideal for use in:

- CDMA2000 mobile phones
- UMTS/WCDMA mobile phones
- Power amplifier modules (PAMs) with linear PAs
- WiMAX/Wibro modules
- WiFi modules
- DECT handsets

MICREL®
Innovation Through Technology®
www.micrel.com

contents 2.3.11

52

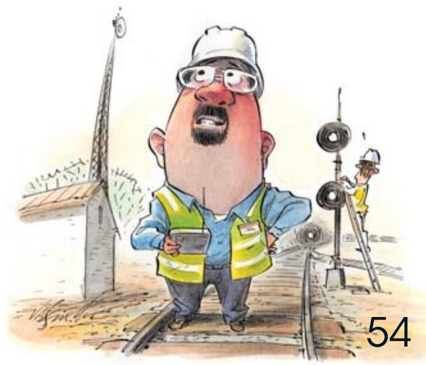

54

DEPARTMENTS & COLUMNS

EDN online contents www.edn.com

ONLINE ONLY

Check out these Web-exclusive articles:

Photovoltaic market to see sunny 2011
Photovoltaic solar installations are expected to rise in 2011, but rollbacks in other countries add clouds to iSuppli's positive outlook.
→www.edn.com/110203toca

Energy Star turning into black hole, technology companies fear
Although IT manufacturers have long supported and would like to continue supporting Energy Star, additional time and money requirements may drive some companies out of the program.
→www.edn.com/110203tocb

EDN HOT 100

EDN proudly presents its list of the Hot 100 products that in 2010 heated up the electronics world and grabbed the attention of our editors and our readers.
→www.edn.com/110203tocc

WANTED: Design Ideas

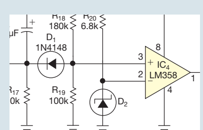

Want to see your work featured in *EDN*? Submit your own Design Idea—a short, compact article that helps solve problems or shows innovative ways to accomplish design tasks. We review articles that fall into virtually any technology area: analog or digital circuits, programmable logic, hardware-design languages, systems, programming tips, useful utilities, test techniques, and so on. The idea should be useful, innovative, or tricky. And if we accept your Design Idea, we pay you $150.

To learn what makes a good Design Idea entry, spend some time browsing our archive of Design Ideas at www.edn.com/designideas. For information on how to submit a Design Idea, see the Design Ideas Writers' Guide: www.edn.com/110203tocd.

EDN® (ISSN#0012-7515) is published semimonthly, 24 times per year, by UBM Electronics, 11444 W. Olympic Blvd., Los Angeles, CA 90064-1549; 310/445-4200; FAX 310/445-4299. Periodicals postage paid at Los Angeles, CA, and at additional mailing offices. SUBSCRIPTIONS—Free to qualified subscribers as defined on the subscription card. Rates for nonqualified subscriptions, including all issues: US, $150 one year; $250 two years; $300 three years. Except for special issues where price changes are indicated, single copies are available for $10 US and $15 foreign. For telephone inquiries regarding subscriptions, call 763/746-2792. E-mail: EDN@kmpsgroup.com. CHANGE OF ADDRESS—Notices should be sent promptly to PO Box 47461, Plymouth, MN 55447. Please provide old mailing label as well as new address. Allow two months for change. NOTICE—Every precaution is taken to ensure accuracy of content; however, the publisher cannot accept responsibility for the correctness of the information supplied or advertised or for any opinion expressed herein. POSTMASTER—Send address changes to EDN, PO Box 47461, Plymouth, MN 55447. Canada Post: Publications Mail Agreement 40612608. Return undeliverable Canadian addresses to Pitney Bowes Inc, PO Box 25542, London, ON N6C 6B2.

Volume 56, Number 3 (Printed in USA).

BY MIKE DEMLER, TECHNICAL EDITOR

The design-to-cost imperative and customer value

Engineers are constantly challenged to meet aggressive performance specifications and to produce designs that don't exceed manufacturing cost targets. That approach may be insufficient to ensure product success, however. I hold this opinion because, early in my career as an IC designer for Texas Instruments, the company adopted a "design-to-cost" philosophy. J Fred Bucy, then TI's president, was clearly intent on driving design-to-cost consciousness into the culture of the corporation. "The application of the design-to-cost concept must become second nature not only to those who design but also to those who manufacture and market," he said (**Reference 1**).

Looking back at TI's product strategy in 1977 (**Reference 2**), you'll see that things really haven't changed much. As a novice IC designer, the design-to-cost imperative struck home for me during what appeared to be a simple design project: a thermal print-head driver. The driver that I designed (**photo**) required three chips with a fixed print head. With this approach, we could simultaneously burn the dot matrix row by row as the paper scrolled past.

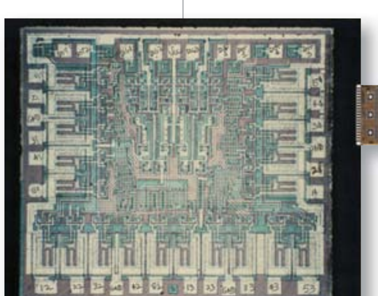

The only problem was that the design-to-cost targets had determined the chip size before I even started the design. With the best layout draftsmen that we had working on it, we produced a beautifully symmetrical die. Unfortunately, the constrained chip area limited the size of the output drivers, and the on-resistance was too high to get the print head hot.

After I convinced my boss that there was nothing more to squeeze out of the layout, he let me in on a simple fix: Shrink it! The new fabrication process was not yet in volume production, but it would meet the design-to-cost goals.

All of these memories came back to me during the 2011 Consumer Electronics Show in Las Vegas last month. Multicore processors for smartphones are among the hottest new IC-design and -application topics, and a battle for leadership is under way. Nvidia has laid claim to "the world's most advanced mobile processor," the Tegra 2 (**Reference 3**), which it built from the ground up as a heterogeneous multicore SOC (system on chip). Meanwhile, TI (**Reference 4**), Samsung, and Qualcomm are also competing for next-generation design wins.

Broadcom recently announced its second-generation dual-core BCM2157 SOC (**Reference 5**). It integrates two ARM11 cores with an HSDPA (high-speed-downlink-packet-access) baseband processor and an Android application processor. With the dual-core design, Broadcom is targeting the low-end market for $100 Android smartphones. More than 30 years ago, TI pursued the same strategy for market penetration.

Let's not get carried away with this design-to-cost concept, though. It is also critical to keep in mind the customer side of the cost equation—that is, the value proposition: Value equals benefit divided by cost. My second-generation thermal printer design wasn't just less expensive; it also provided added benefit in higher reliability. Both Broadcom and Nvidia use the free Android operating system but design to deliver value to consumers—from low-end phones to high-end "superphones." Through each generation of Moore's Law, the semiconductor industry has provided more customer value, lowering cost but also increasing benefits.

Too many companies in "mature" competitive industries blame customers or the tough economy for driving down prices. The lesson from the most successful companies is to continually deliver greater value. Companies should focus more on how their engineers can design for value rather than obsess over balance-sheet-driven cost-cutting strategies, in which layoffs and outsourcing are all too prevalent. **EDN**

REFERENCES

1. White, John A, *Production Handbook*, Wiley, 1987.
2. *Texas Monthly*, July 1977, http://bit.ly/feXMDW.
3. Demler, Mike, "CES Day-2: Nvidia, Motorola, and you say you want to join the press corps?" *EDN*, Jan 6, 2011, http://bit.ly/eXUH31.
4. Demler, Mike "More from CES: Blackberry PlayBook enters the battle of the 4G tablets," *EDN*, Jan 11, 2011, http://bit.ly/f2eNgy.
5. "Broadcom Announces New Android Platform to Enable Mass Market Smartphones," Broadcom, Dec 14, 2010, http://bit.ly/fG8SZr.

Contact me at mike.demler@ubm.com.

UBM Electronics

EDN. 33 Hayden Avenue, Lexington, MA 02421. **www.edn.com**. **Subscription inquiries:** 1-763-746-2792; EDN@kmpsgroup.com. **Address changes:** Send notice promptly to PO Box 47461, Plymouth, MN 55447. Please provide an old mailing label as well as your new address. Allow two months for the change.

UBM Electronics, 11444 W. Olympic Blvd., Los Angeles, CA 90064-1549; 1-310-445-4200; fax: 1-310-445-4299.

DSP-based analyzers boast 85-MHz bandwidth

With the introduction of the RSA5000 series signal analyzers, Tektronix has significantly expanded its spectrum- and vector-signal-analysis portfolio. The series targets use in design and operations applications, including spectrum management, radar, electronic warfare, radio communications, and EMC (electromagnetic compatibility).

As RF signals become more complex and the use of the ISM (industrial/scientific/medical) band becomes more common, designers, engineers, and operators must reliably and efficiently discover transient spectral phenomena that digital-RF circuits can create over wider bandwidths. Traditional signal analyzers cannot trigger on transient problems, and midrange units' maximum acquisition bandwidth until now has topped out at 40 MHz. Featuring advanced time, amplitude, and DPX (digital-phosphor-technology)-triggering functions, as well as a swept-DPX mode, the RSA5000 series enables you to discover and capture intermittent and rapidly changing signals at bandwidths as high as 85 MHz. This bandwidth covers the entire ISM band, which is home to such common technologies as Bluetooth, Zigbee, RF identification, and wireless LANs.

The series' DPX-based live RF-spectrum display enables the analyzers to quickly detect previously unseen signal behavior and improves test confidence by catching transients as brief as 5.8 μsec. The swept-DPX engine can also collect as many as 292,000 spectrum updates/second over bandwidths as high as 85 MHz and can sweep the DPX display across instrument inputs that range to 6.2 GHz.

To capture transients for analysis, the analyzers offer frequency-mask, frequency-edge, density, time-qualified, and runt triggering. You can also use cross-domain triggering among multiple instruments to isolate hard-to-find hardware and software anomalies and capture in deep memory seamless 7-second time records of 85-MHz-bandwidth RF signals. The series lets you analyze captured data in any domain at any time with correlated markers. Automatic pulse measurement and detection support multiple measurements on the same set of captured data, which simplifies testing; reduces test time; and, by providing one instrument that replaces multiple test sets, reduces test cost.

The analyzers offer 17-dBm third-order intercept and –154-dBm/Hz DANL (displayed average-noise level) at 1 GHz and carrier-referred phase noise of –131 dBc/Hz at a 10-kHz offset from the carrier and –150-dBm DANL at a 10-MHz carrier frequency. Prices range from $34,900 for a 1-Hz to 3-GHz unit with 25-MHz bandwidth to $68,800 for a 1-Hz to 6.2-GHz unit with 85-MHz bandwidth.

—by Dan Strassberg

▷**Tektronix**, www.tektronix.com.

The RSA5000 signal analyzers' dynamic range rivals that of comparably priced swept-frequency spectrum analyzers but offers double the RF bandwidth.

TALKBACK

"While open-loop uses of optocouplers are fun, useful, and profitable—meaning they work well for many applications—if lots of stability over long periods of time is required, then it's hard to beat a closed-loop feedback system, with the optocoupler inside the loop."

—Engineer Steve Hageman, in *EDN's* Talkback section, at http://bit.ly/egx2kt. Add your comments.

Wireless-power technology, development kit target Chevy Volt

General Motors (www.gm.com) hopes to soon establish the plug-in hybrid Chevy Volt as a technology leader in the eyes of the "green"-car-buying public. Toward that end, GM Ventures (www.gm.com/ventures), the company's venture-capital subsidiary, has invested $5 million in Israeli wireless-power-technology start-up Powermat. The first GM deployment of the Powermat technology will be a wireless charging station in the 2012 Volt, and the technology will eventually roll out to the Chevy, Buick, GMC, and Cadillac brands. GM announced the development at last month's Consumer Electronics Show in Las Vegas, but no demos took place.

Wireless power charging is an apt application for consumer devices in a car's cabin because, in space-constrained applications, you'd like to avoid myriad chargers and dangling wires. On the other hand, the Powermat technology requires a receiver for each device, so you're basically replacing your charger and dangling wire with a $40 receiver, which looks like a rigid cell-phone cover.

Powermat currently supports all iPhone models; the second- and third-generation iPod Touch; all docking iPods; HTC Evo 4G and HD2; Motorola Droid X; Nintendo DS Lite and DSi; and Blackberry Tour 9630, Bold 9000/9650/9700 series, Curve 8300/8500/8900/9300 series, and Pearl.

Powermat is not the only wireless-charging technology available. The Wireless Power Consortium (www.wirelesspowerconsortium.com), of which Texas Instruments is a founding member, last summer released Qi, its own version of an industry standard. The consortium based its approach on Fulton Technology's (www.fultontechnology.com) eCoupled technology.

TI last month released what it calls the industry's first Qi-certified development tools and chip set for wireless power. The $499 bqTesla development kit includes the bq500110 wireless-power-transmitter manager; the single-input, 5V bq25046 power-supply IC; and the MSP430bq1010 wireless-power-control and -communications microcontroller.

—**by Margery Conner**

▷**Powermat**, www.powermat.com.

▷**Texas Instruments**, www.ti.com.

The Qi-certified bqTesla development kit (right) helps developers design wireless power for plug-in hybrid vehicles, such as the Chevy Volt (below).

DILBERT By Scott Adams

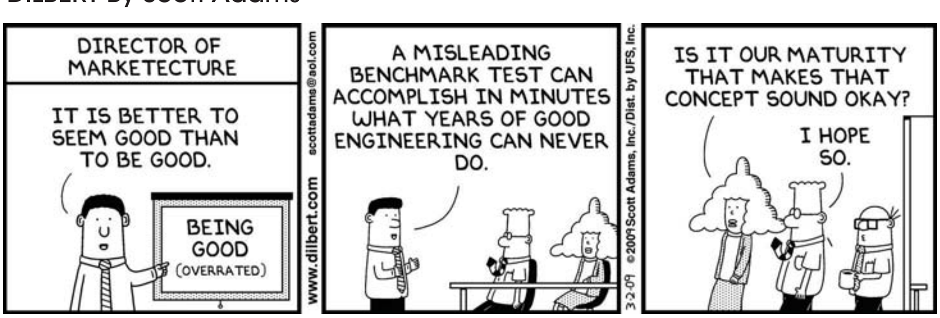

02.03.11

HIGH-VOLTAGE GATE-DRIVER ICs IMPROVE NOISE IMMUNITY

Fairchild Semiconductor recently introduced a series of high-voltage gate-driver ICs, including the two-input, two-output, high- and low-side FAN7392 with shutdown; the one-input, two-output FAN7393 half-bridge with shutdown and controllable dead time; the one-input, two-output FAN73932 half-bridge with shutdown and fixed dead-time control; and the two-input, two-output FAN73933 half-bridge with controllable dead time. The parts suit use in industrial applications, such as motor-drive inverters, distributed power supplies, and telecom-system power supplies.

The devices feature a common-mode dV/dt noise-canceling circuit that enables stable operation of the high-voltage gate driver under high-noise circumstances, as well as a level-shift circuit that offers high-side gate-driver operation with negative-floating-supply-return-voltage swings as high as –9.8V at a high-side floating-supply voltage of 15V. The FAN739x series provides stable operation over a temperature range of –40 to +125°C. The series also features floating channels for bootstrap operation to 600V. The devices sell for $1.52 (1000) each.

—by Paul Rako

▶Fairchild Semiconductor, www.fairchildsemi.com.

Class D amplifier suppresses EMI

Silicon Laboratories recently introduced a 5W-stereo Class D amplifier that uses multilayer technology to suppress EMI (electromagnetic interference), common in Class D devices, at its source. The Si270x amplifier finds use in a range of price- and noise-sensitive consumer-audio products, including smartphone-docking stations, tabletop radios, TV sound bars and monitors, boom boxes, and battery-powered radios.

Until now, two issues have impeded the adoption of Class D amplifiers: high EMI emissions, which interfere with AM/FM radio and smartphone operation, and the high cost of adding filtering and shielding for EMI-regulatory compliance. The Si270x addresses these issues by having 10 times less radiated interference in the EMI-compliance band, 100 times less in the FM-radio band, and 1000 times less across the AM band than do other Class D products.

The amplifier also features 2.5-times more play time than Class AB-based systems and uses half the number of batteries. A consumer-audio system employing the Si270x amplifier can provide as much as 8.4 hours of play time using four AA alkaline batteries.

You can combine the Si270x amp with the Si473x AM/FM-radio tuner. The latest Si473x devices offer a stereo analog input and internal ADCs multiplexed with the radio-tuner front end to support auxiliary analog-system inputs without additional external ADCs.

> Until now, high EMI emissions and costly filtering have impeded the adoption of Class D amps.

Samples and preproduction quantities of the Si270x amplifiers are now available in 24-pin QFN packages. Prices begin at $1.17 (10,000).

To help accelerate application development, Silicon Labs offers audio engineers the Si270x-A-EVB (evaluation board). The $325 board comprises a baseboard and a daughtercard. It includes a graphical user interface that runs on a standard PC and connects to the EVB through a USB (Universal Serial Bus) interface.—**by Paul Rako**

▷**Silicon Laboratories**, www.silabs.com.

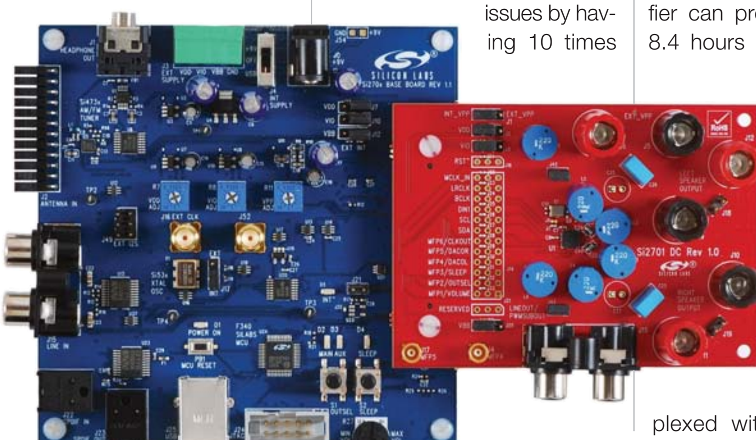

The Si270x-A evaluation board accelerates application development of the Si270x Class D amplifier.

Freescale unveils Cortex-A9-based devices

Cortex-A9 is currently the cream of the ARM (www.arm.com)-core crop, so it's the next logical step in Freescale's i.MX product progression. Freescale plans to begin making the first few members of its new i.MX6 family available for sampling this year, with volume quantities in 2012. The family comprises the four-core i.MX6Quad, the two-core i.MX6Dual, and the single-core i.MX6Solo. The i.MX6Quad and i.MX6Dual will appear first. Freescale is, at least initially, fabricating them on a common sliver of silicon, with two of the four CPU cores disabled on the i.MX6Dual devices.

Freescale isn't yet releasing details on whose graphics technology it's licensing for i.MX6; likely candidates include Imagination Technologies' (www.imgtec.com) popular PowerVP cores or ARM's Mali multimedia engine. The devices include as many as four ARM Cortex-A9 cores running at greater than 1 GHz per core; as much as 1 Mbyte of Level 2 system cache; and ARM Version 7, Neon, VFP Version 3, and Trustzone support. A multistream-capable high-definition video engine delivers 1080p60 decode, 1080p30 encode, and 3-D video playback in high definition. The devices also include 2-D and Vertex acceleration engines and interface to stereoscopic image sensors for 3-D imaging.

Interconnect is through HDMI (high-definition-multimedia interface) Version 1.4 with integrated PHY (physical layer), SD (secure digital) 3.0, multiple USB (Universal Serial Bus) 2.0 ports with integrated PHY, GbE (gigabit Ethernet) with integrated PHY, SATA (serial-advanced-technology attachment)-II with integrated PHY, PCIe (Peripheral Component Interconnect Express) with integrated PHY, MIPI (mobile-industry-processor-interface) CSI (camera serial interface), MIPI DSI (device intellectual-property interface), MIPI HSI (high-speed serial synchronous interface), and FlexCAN (controller-area network) for automotive applications. The devices also support the VP8 codec, along with optional integration of an e-paper display controller for e-reader applications.—by Brian Dipert

▶Freescale Semiconductor, www.freescale.com.

Step-down dc/dc controller features fast transient response

Linear Technology Corp's new LTC3810H-5 synchronous step-down dc/dc controller guarantees junction temperatures as high as 150°C. The device employs a constant on-time valley-current-mode-control architecture, delivering low duty cycles and fast transient response with accurate cycle-by-cycle current limit without requiring a sense resistor. The controller can directly step down voltages from 60V to an output voltage ranging from 0.8V to 93% of the input voltage. The 1Ω onboard gate drivers minimize switching losses associated with driving N-channel MOSFETs at high frequency and high voltage for output currents as high as 25A. Applications include 48V power conversion in telecom and datacom, as well as automotive and industrial systems susceptible to high surge voltages.

The LTC3810H-5 features an adjustable minimum on-time that you can set as low as 100 nsec, enabling a high step-down ratio at high frequency. The operating frequency is selectable from 100 kHz to 1 MHz, or you can synchronize it to an external clock over the same range. You can configure pulse-skipping operation to maintain high efficiency at light loads. Other features include programmable soft start or tracking, adjustable cycle-by-cycle current limit, output overvoltage protection, a power-good output signal, and programmable undervoltage lockout. The LTC3810H-5 comes in a 5×5-mm QFN-32 package, and prices start at $3.76 (1000) each.—**by Fran Granville**

▷**Linear Technology**, www.linear.com/3810.

Solid-state drives are one-eighth the size of their siblings

Intel recently announced the 310 Series of m-SATA (mini serial-advanced-technology-attachment) solid-state drives and is currently shipping them in $99 (1000), 40-Gbyte and $179, 80-Gbyte versions. The drives deliver performance equivalent to that of the company's comparable 34-nm, MLC (multilevel-cell)-flash-memory-based X25 solid-state drives but at one-eighth the size. They measure 50.8×29.85 mm, are 4.85 mm thick, and weigh less than 10g. Their PCIe (Peripheral Component Interconnect Express) Mini Card-based mechanical interfaces are fully compatible with 1.5- and 3-Mbps SATA protocols, including NCQ (native-command queuing).

Other features include sustained sequential reads as high as 170 and 200 Mbytes/sec and sustained sequential writes as high as 35 and 70 Mbytes/sec for the 40- and 80-Gbyte versions, respectively. Read latency is 65 µsec for both versions, and write latency is 110 and 75 µsec for the 40- and 80-Gbyte versions, respectively. Random 4-kbyte reads are, respectively, as high as 25,000 and 35,000 IOPS (input/output operations per second); respective random 4-kbyte writes are 2500 and 6600 IOPS. The devices have a life expectancy of 1.2 million MTBF (mean time between failures), a typical power consumption of 150 mW in active mode and 75 mW in idle mode, and an operating-temperature range of 0 to 70°C.

—**by Brian Dipert**

▷**Intel Corp**, www.intel.com.

Intel's m-SATA drives are one-eighth the size of the company's 34-nm, MLC-flash-memory-based X25 solid-state drives.

Tool performs FPGA-power designs

National Semiconductor recently introduced the Webench FPGA-power-architect-design tool to model power supplies for FPGAs. The tool incorporates the detailed supply requirements of more than 130 FPGA devices from Altera (www.altera.com) and Xilinx (www.xilinx.com). Modern power-supply systems, including advanced FPGAs, are complex to design, often incorporating multiple unique loads to drive precisely specified voltages. In addition to the required voltages and currents, each load may have limitations for ripple, noise filtering, synchronization, and separation of supplies, as well as start-up definitions.

Webench delivers power-supply designs that incorporate the comprehensive power requirements that FPGA manufacturers publish, giving designers the confidence that their power supplies will meet these constraints and save time. To begin a design, a designer selects an advanced FPGA from Altera or Xilinx, and the tool automatically populates the unique power requirements, identifying core and I/O options for every potential load of the array. Once the designer tunes the design, Webench collects all the loads and creates multiple power-supply architectures as options for driving the FPGA and the total system.

The power supply may involve one or more intermediate voltage rails between the input supply and the point-of-load regulators. The designer can tune the recommended power supply with the turn of a dial, in seconds reducing size, increasing efficiency, or lowering complete system cost. The designer can also order components for prototyping; share the complete system with others; or easily print a complete project report, including schematics, BOM (bill of materials), and performance characteristics.—by Paul Rako

▶National Semiconductor, www.national.com.

02.03.11

BY HOWARD JOHNSON, PhD

Take the fifth

Fourier analysis represents a perfect square wave as the infinite sum of harmonically related sine waves. The requisite sine-wave components comprise all the odd harmonics of the fundamental square-wave frequency with their respective amplitudes set according to this pattern: 1, 1/3, 1/5, 1/7, and so forth. Even harmonics are not required. The Fourier theory calls out an amplitude pattern that guarantees a sum, which—in the limit—approximates a square-wave shape with amplitude π/4, or 0.78537316.

Figure 1 illustrates the first few terms of the harmonic series. The first waveform is a simple sine wave at the fundamental frequency. The next waveform shows the fundamental plus its third harmonic, the next takes harmonics to the fifth, and so on. The more terms you add, the more "squarish" the waveform becomes, but it still looks wiggly.

Many engineers wonder how many harmonic terms they must take to adequately represent a good square wave. In other words: What bandwidth does a digital transmission system need? Let's explore that question by first examining the amplitudes that the harmonic-sum model uses. The Fourier theory applies only to an infinite sum, which never occurs in practice. If you have a finite sum, with only a finite number of sine-wave components, why not try other amplitudes different from the ideal infinite-sum values? Instead of chopping off the infinite sum after the seventh harmonic, with subsequent amplitudes dropping abruptly to zero, the right side of **Figure 1** smoothly tapers the harmonic amplitudes so that, by the time you get to the last one, its amplitude is small. That fact mitigates the discontinuity at the end of the harmonic sequence, making a better waveform.

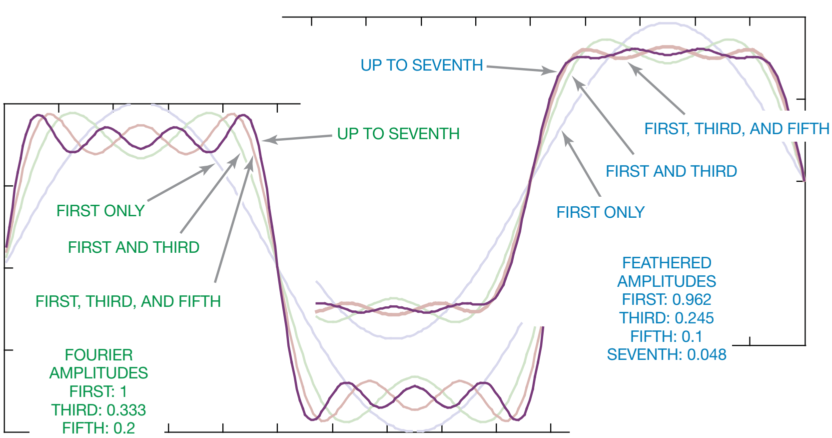

Figure 1 Additional harmonics sharpen rising and falling edges but do not improve the worst-case ripples. In contrast, feathering the harmonic amplitudes improves the ripple amplitude, but at the expense of some loss in edge sharpness.

Both the number of harmonics and the precise schedule of their proportions affect the quality of the result. Stated differently, when considering the adequacy of bandwidth, not only the −3-dB frequency of your system but also the precise shape of its entire transfer function may affect your result. Harry Nyquist fleshed out that notion in his theorem of data transmission, concluding that, given infinite complexity in the receiver and perfect control over the exact system-transfer function, the bandwidth must equal or exceed half the system data rate to achieve reliable binary communication.

Suppose that a system transmits data with a baud frequency of *B*. The fastest alternating pattern you can send is the sequence 1, 0, 1, 0, … . That pattern is based on a repeating cell two bits long, 1 and 0, so the pattern has a fundamental frequency equal to only half of *B*. According to Nyquist, if you pass that much of the signal, you can in theory recover all the data.

However, you don't have access to a receiver of infinite complexity or perfect control over the system-transfer function. You have a digital system that needs a wide-open eye with plenty of margin for setup and hold, so you must design your system to accommodate the rise and fall time, not just the baud rate.

I maintain that the bandwidth required to preserve a rise and fall time of *T* as it propagates through a system equals approximately one-half divided by *T*. If your rise and fall times equal 10% of the data-baud interval, so that *T* equals 0.1/*B*, then you can express the required system bandwidth of 0.5/*T*, after suitable algebraic manipulation, as five times *B*.

All of these words attempt to explain the habit among old digital-system designers, when someone asks them how many harmonics of a data sequence are necessary for data transmission, to say, "I take the fifth." **EDN**

Howard Johnson, PhD, frequently conducts technical workshops for digital engineers.

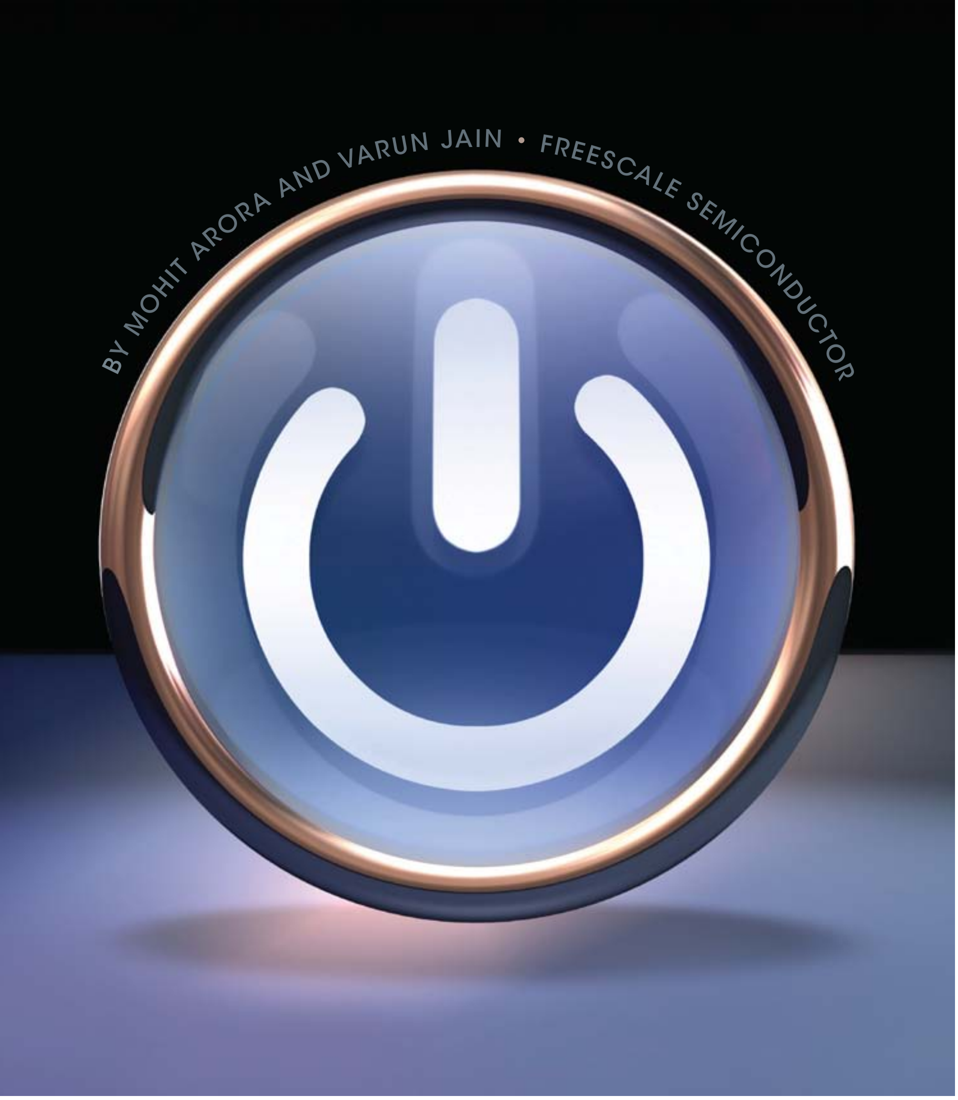

BY MOHIT ARORA AND VARUN JAIN • FREESCALE SEMICONDUCTOR

UNDERSTANDING EMBEDDED-SYSTEM-BOOT TECHNIQUES

BOOT-UP, THE SEQUENCE OF STEPS THAT A SYSTEM PERFORMS BETWEEN WHEN YOU SWITCH ON POWER AND LOAD APPLICATIONS, IS SIMPLE IN THEORY BUT OFTEN COMPLEX IN REALITY. THE MAIN JOB OF A BOOT LOADER IS TO LOAD THE OPERATING SYSTEM, BUT SOFTWARE AND HARDWARE ENGINEERS VIEW THIS PROCESS IN DIFFERENT WAYS.

To load a program into memory, you must first load a program into memory. The boot-up process, often a complex multistep sequence involving numerous substeps, solves this problem. Any boot-up process, including booting up Windows, Linux, or an embedded RTOS (real-time operating system), begins with the application of power to the system and the subsequent removal of system reset. During POR (power-on-reset) assertion, you may have to reconfigure hardware peripherals if operational values differ from those of default settings. Embedded microcontrollers, for example, often offer various hardware-reset-configuration schemes.

Over the last several decades, booting up has evolved from a simple DOS-based step to more complicated multiple-operating-system choices or even peripheral-based boot-up techniques. A USB (Universal Serial Bus) interface, for example, allows you to boot up a disk image from an external storage device; this approach is increasingly popular in industrial- and embedded-system applications because it provides abundant flexibility. In the case of software corruption, for example, in which a system requires the reloading of new firmware, the USB technique allows a service engineer to simply copy new software onto a flash drive and boot from it. The service department therefore saves the thousands of dollars in expenses that it would otherwise incur in transporting the equipment to the manufacturer for repair.

To enable system-boot flexibility both from USB, PCIe (Peripheral Component Interconnect Express), and SDHC (secure-digital-high-capacity) interfaces and from conventional memory devices requires in-depth hardware and software capabilities. Open-source firmware—specifically, the U-Boot (Universal Boot Loader) utility, which finds wide use in embedded-system platforms—may also be of value. The Linux-based boot loader can automatically boot up the operating system; alternatively, it allows a user to manually run explicit commands to start the operating system, and it supports booting from a variety of interfaces (see **sidebar** "The U-Boot").

WINDOWS XP SYSTEM BOOT

A simple x86 boot sequence is fairly self-explanatory (**Figure 1**). Windows XP follows comparable steps, albeit

AT A GLANCE

- The system-boot process, although potentially simple in concept, becomes complex when you consider various implementation options.
- Windows XP provides a popular case study of a classic boot sequence.
- Numerous hardware and software techniques provide different means of getting postreset-configuration data to the processor.
- Primary and secondary boot options comprehend different start-up and kernel-code sizes, read- and write-performance expectations, and other variables.
- U-Boot (Universal Boot Loader) is a powerful open-source tool that deserves consideration in Linux-based designs.

with more sophistication. The boot sequence begins with the application of system power; the processor remains in reset. When all voltages and current levels are acceptable, the power supply sends a power-good signal to the processor. The next step, POR negation, takes place when the availability of the good-power signal negates the processor reset to allow the CPU to begin operation. The CPU points to the ROM boot address and begins executing the BIOS (basic-input/output-system) code. The ROM BIOS then performs POST (power-on self-test), a basic test of core hardware to verify basic performance. The boot-up process then reports any errors that occur at this point using beep codes because the video subsystem has not yet initialized.

The next step is video-card initialization, during which the BIOS searches for a video-card adapter by scanning memory addresses C000:0000 through C780:0000 to find a video ROM. The video test initializes and tests any discovered video adapter, potentially along with its video memory, and displays configuration information. If a "cold" start is taking place, the ROM BIOS executes a full POST. If it is a "warm" start, the boot-up process skips the memory-test portion of the POST.

The CMOS now reads from the BIOS. During this step, the BIOS locates and reads configuration information stored in the CMOS. A small coin battery cell on the motherboard typically maintains the CMOS, a small—typically, 64-byte area of memory—on the motherboard. The CMOS memory stores information such as date, time, and peripheral-boot order. If the first bootable disk is a fixed disk, the BIOS examines its first sector for an MBR (master boot record). The MBR comprises a partition table, which describes the layout of the fixed disk, and a partition-loader code, which includes instructions for continuing the boot process. The boot, or partition, loader then examines the partition table for an active partition. The partition loader searches the first sector of that partition for a boot record. The boot process than checks the active partition's boot record for a valid boot signature. If it finds one, it executes the boot-sector code as a program.

The NTLDR (New Technology Loader) for Windows, a hidden system file that resides in the root directory of the system partition, controls the loading of Windows XP. During NTLDR's initial phase, it moves the processor from real mode to protected mode, enabling 32-bit memory accesses and turning on memory paging. It then loads the appropriate minidrivers to allow NTLDR to load files from a partition formatted with any of the file systems that Windows XP supports, including FAT (file-allocation table)-16, FAT-32, and NTFS (New Technology File System).

If the boot-initialization file resides in the root directory, NTLDR reads its contents into memory. If the file contains entries for more than one operating system, NTLDR suspends the boot sequence, displays a menu of choices, and waits for the user to make a selection. NTLDR then continues the boot-up process by locating and loading the DOS-based New Technology executable file to perform hardware detection. After selecting a hardware configuration, NTLDR begins loading the Windows XP New Technology kernel file. During this process, the screen clears, and a series of white rectangles subsequently progresses across it. NTLDR now loads device drivers that are marked as boot devices. Before performing this load, NTLDR relinquishes control of the computer.

At this point, the system displays a graphics screen with a status bar indicating the load status. During later initialization phases, the system cannot accept device interrupts. The I/O manager also begins to load the remainder of the system drivers, picking up where NTLDR left off. The last task for this initialization phase is to launch the session-manager subsystem, which is responsible for creating the user-mode environment. The session-manager subsystem then loads the Windows device driver, which implements the graphics subsystem. Windows XP boot is not complete until a user has successfully logged onto the system. The Windows log-in file begins the log-in proc-

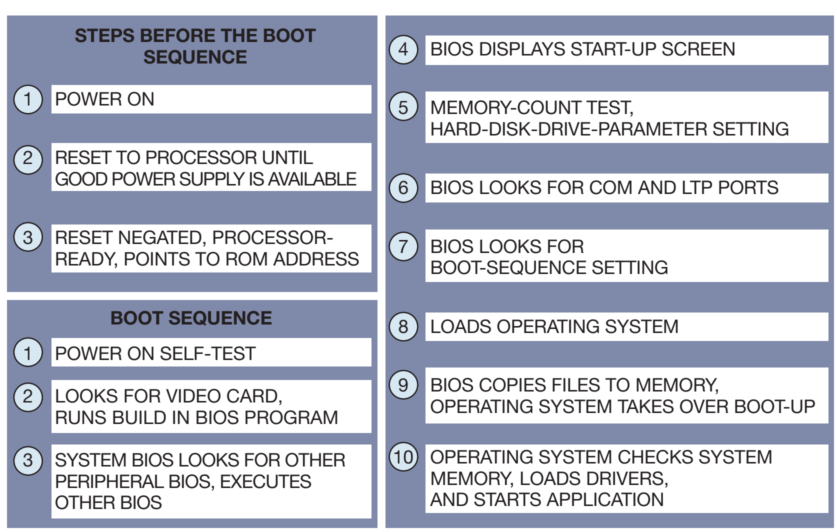

Figure 1 Microsoft's Windows XP moves through a suite of steps as it boots a system.

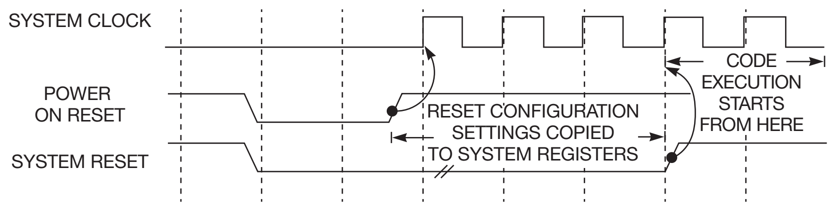

Figure 2 Loading default values during boot-up is the most common reset configuration.

ess, which the kernel loads as a service, and displays the log-on dialogue box.

EMBEDDED-SYSTEM RESET

Older microcontrollers had a single fixed configuration state for the entire register suite after reset deassertion. This situation translated into fixed values for parameters, such as clock-speed configuration, start-address location, pad slew, drive strength, external-memory-port size, and peripheral enable/disable state. It enforced restrictions on the way users employed chips after reset deassertion. For example, if the system disabled an on-chip oscillator providing a clock to off-chip peripherals at reset, those peripherals would be able to work only after software re-enabled the oscillator.

In simpler systems, such behavior might be acceptable, but it may not meet requirements if that same chip is finding use in more complex applications that require different configurations during reset. To provide flexibility to a configuration of registers after reset, you can implement various microcontroller-based design schemes. Some microcontrollers also support multiconfiguration schemes, in which the system selects any configuration by reading the state of pins during reset. Here we discuss four reset configurations: default, fuse programming, external pins, and serial interface.

Loading default values during boot-up is the most common reset configuration, and it requires no special onboard setup (**Figure 2**). Conversely, it provides no flexibility or options to configure any register. All the registers initialize to fixed values; thus, the chip exits reset in only one fixed state. This approach is the fastest for initializing the system before the boot process, but it is the least powerful mode in terms of capabilities for controlling the system state. It might work for some applications but is less preferable if you use the same microcontroller in varying applications with varying boot requirements.

A system POR asserts internal chip reset, with both signals being active low; when deasserted, it restarts the clock and loads the reset configuration in the system registers. Depending on the system design, other tasks might gate the boot-up process. For example, the system might wait for a PLL (phase-locked loop) to achieve clock locking before the internal reset deasserts and the system begins executing the boot code.

Fuse programming involves a reset configuration that is the result of programming through a chip's special test mode, fuses, or on-chip nonvolatile-flash registers (**Figure 3**). This mode implements special bits and registers through either on-chip fuses or an array of nonvolatile-flash-memory registers for configuring reset-control-word information. These registers require write-once capability; therefore, you can program them only once in a chip's lifetime. This mode usually requires a hardware-setup or software sequence to program these special registers or fuses. Once fuses are programmed and reset is asserted, reset-control-word information is read from the fuses and copied to the desired system registers. The system then internally deasserts the reset and begins executing code.

This scheme provides the flexibility to configure different system-register options but requires the implementation of special fuse registers in the design. Because the fuses are one-time programmable and secure, they can effectively enable and disable functions within the chip, thereby creating "phantom" parts with lower prices. This strategy is common with semiconductor vendors, which sell the same sliver of silicon with different costs and features—achieved by blowing the fuses to enable or disable functions.

You can also reset the configuration through external pins. This scheme uses a group of microcontroller pins to control the reset configuration. These pins are externally pulled high or low during reset to define a configuration option. Once the system reset deasserts, the microcontroller internally latches these values and decodes them to configure the system registers. This scheme provides limited flexibility in selecting control-word configurations. The available number of configurations is directly proportional to the number of pins for this purpose.

Engineers often implement the design by means of an external buffer or a line driver, such as the 74LVC125, which drives either a logic one or a logic zero to the pins for reset configuration (**Figure 4**). The reset signal usually connects to the enable pin of a tristate buffer—that is, an external line driver—so any change from one to zero or vice versa to the input of the tristate buffer appears as an input to the pins that eventually handle the reset configuration. This approach provides ad-

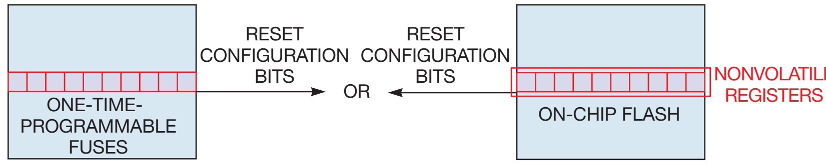

Figure 3 Obtaining reset-configuration bits from on-chip fuses or flash memory provides device customization during manufacturing.

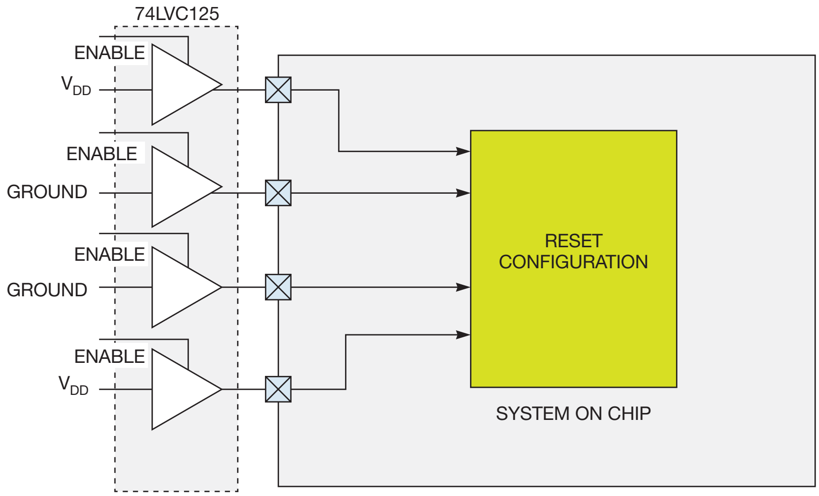

Figure 4 Accessing the reset configuration through an external bus provides flexibility that only the number of pins for this function limits.

ditional flexibility and control. For example, if one of the buffers controls whether PLL is enabled or disabled during boot-up, a buffer output can allow users to enable the PLL when the buffer connects to the drain-to-drain voltage and conversely disable it when the buffer connects to the source-to-source voltage. The number of pins available for this purpose usually constrains this approach. Note, too, that most external line drivers, such as that of the 74LVC125, integrate groups of four or eight buffers. To limit the implementation's cost, define the buffer numbers in multiples of four.

You can also load the reset configuration through an external serial interface. For a highly integrated complex microprocessor, it is impractical to either dedicate or share pins for the numerous available power-up options. This scheme conversely involves loading the chip-reset-configuration data from an external serial memory (**Figure 5**). In a typical implementation flow, when system reset asserts, the chip establishes communication with the serial memory, subsequently transferring reset-configuration information to the microcontroller. Upon serial-data reception, the microcontroller configures system registers based on the received data and deasserts reset. This method provides the maximum flexibility in configuring options in system registers because serial memories can store a large number of data bytes.

In advanced serial-configuration schemes, the serial memory can even hold software code. In such cases, the system reads both reset-configuration data and boot code from external seri-

THE U-BOOT

Many standard boot loaders, such as DINK32, Open Firmware, and x86 bios, find use in embedded-system applications. They facilitate the loading of an operating system and bring the system to a safe operating state. U-Boot (Universal Boot Loader) is an open-source program that is popular in Linux-based embedded-system applications. U-Boot provides an automated interactive environment, which offers substantial flexibility and diverse options for various boot schemes and interfaces. It provides a platform to the end application-development user who doesn't want or need to delve into the low-level specifics of chip hardware.

After basic initialization of the system, U-Boot starts an interactive program, which allows the user to provide input through a serial-communication interface-console utility, such as Windows' HyperTerminal. The user can also choose to run U-Boot in an automated fashion without any intervention. U-Boot can reside in an internal ROM or a flash memory. After basic initialization of the CPU, local memories, and buses, U-Boot can relocate itself to a RAM location and then continue executing from there. A splash screen appears on the serial console when U-Boot is running.

U-Boot supports a powerful set of commands that you can execute through the interactive command window. Other than loading the operating system, these commands provide abundant functions, such as memory load and dump; serial-interface access; and read, erase, and program functions for external memories, such as NAND- and NOR-flash memory, serial-flash devices, and EEPROMs. U-Boot also allows the system to boot from a variety of interfaces, such as USB (Universal Serial Bus), SD (secure digital), PCIe (Peripheral Component Interconnect Express), or SATA (serial advanced-technology attachment).

To load an operating system in a typical scheme over a network such as Ethernet, for example, U-Boot first initializes the network environment's variables, copies the operating-system kernel's image to the target board, and then moves execution to the operating-system kernel. After this point, U-Boot plays no role in the system.

U-Boot's image size generally depends on the boot-peripheral support you select during source-code compilation or building. You can customize U-Boot's features to appropriately suit the application. U-Boot follows a standardized directory structure, which allows high scalability and portability to platforms. When porting U-Boot to a new platform, most of the files, other than CPU-, peripheral-, and board-specific files, remain the same.

All of these features make U-Boot a desired choice for embedded-system developers. Because U-Boot is an open-source boot loader, embedded-system developers are heavily contributing to the U-Boot environment, adding device support and otherwise keeping U-Boot rich and up to date. For more details on U-Boot-development resources, go to http://bit.ly/eZsILO.

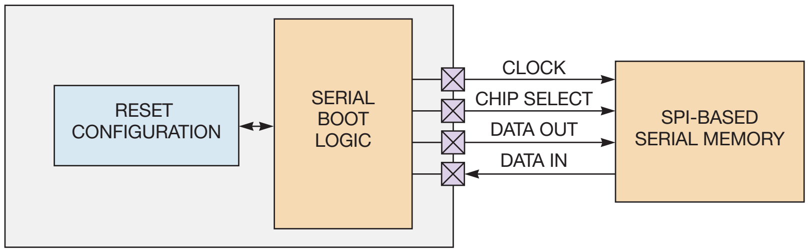

Figure 5 An external memory can house not only reset-configuration data but also boot code.

al memory during the microprocessor-reset sequence, thereby requiring few I/O pins. By reading data stored in, for example, external SPI (serial-peripheral-interface) memory, the system would also need to configure the SPI memory's clock frequency along with the power-up options for the microprocessor and optionally load code into the microprocessor's memory. The system would have to accomplish all these tasks before the negation of device reset, thereby ensuring that the chip is properly configured when exiting the reset state. Serial memories' low cost, simple implementation, high flexibility, and optional software-boot code often make this option the preferred one for booting or loading the reset configuration.

SYSTEM-BOOT COMPONENTS

Boot components are primary or secondary devices, depending on their ability to support boot immediately after reset deassertion. The primary option provides a direct-boot capability and facilitates the processor's first instruction fetch from the memory location in which software-initialization code resides. You should enable the primary interfaces immediately after reset and can sometimes configure them through the reset option. You can place boot loaders or small initialization-code sequences in these locations. You initialize and configure secondary interfaces in these primary boot-code sequences and boot loaders. They mainly hold the large operating-system kernel's code, which loads after the boot loader performs basic initialization. You can execute the operating-system kernel directly from these interfaces or copy it in processor-accessible memory (**Figure 6**). For example, an on-chip ROM component can act as the primary boot device, including initial initialization code, which subsequently switches execution to a secondary option, such as external DDR SDRAM, which in turn holds the complete operating-system kernel.

Common hardware-boot components

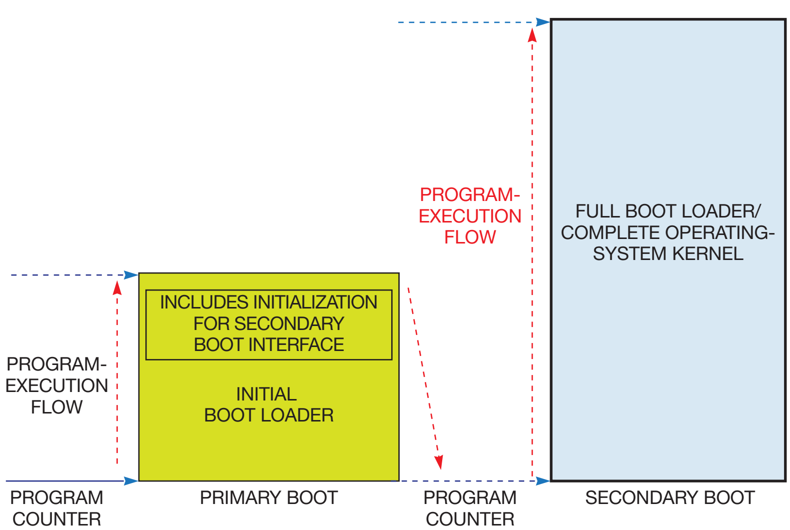

Figure 6 Primary and secondary boot options combine to comprehend large operating-system-kernel images.

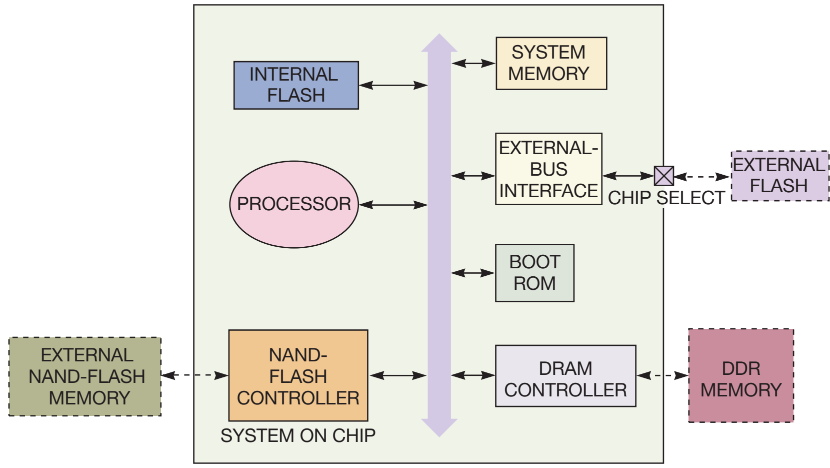

Figure 7 Common hardware-boot components allow a system to boot from a variety of interfaces.

allow a system to boot from a variety of interfaces (**Figure 7**). Booting from internal flash memory is among the most common and simplest methods of configuring an embedded microcontroller that includes the necessary on-chip resources. This method reduces dependencies on external interfaces because the boot loader resides in on-chip memory. After system-reset deassertion, the processor points to the flash memory's starting address and loads the necessary initialization code and operating systems. Operating systems with small footprints are compatible with this approach because a practical limit exists on the amount of on-chip flash memory that can be available. This approach is also one of most secure ways to boot a processor because modifying code residing in on-chip memory requires fewer changes than do off-chip boot options.

As with Windows XP, some microcontrollers integrate ROM as the primary boot option. The boot ROM includes a basic boot loader so that the microcontroller can subsequently perform a more sophisticated boot sequence on its own, loading programs from various sources, such as Ethernet, NAND-flash memory, an SD (secure-digital) card, an MMC (multimedia card), or a USB interface. Boot-ROM usage enables more flexible boot sequences than does hard-wired logic, and it allows users the choice to boot up from various peripherals. Users often employ the boot-ROM feature for system recovery when someone inadvertently erases the usual boot software in nonvolatile memory other than ROM. You cannot reprogram boot ROM, so applications that require a secure boot may include security checks so that the boot-up stops if one or multiple security checks fails.

AS WITH WINDOWS XP, SOME MICROCONTROLLERS INTEGRATE ROM AS THE PRIMARY BOOT OPTION.

An external bus interface allows the system to boot directly from external NOR flash or other parallel memories. It is one of the fastest ways to boot the system because the interface to external memory can be 32 bits or more with a reasonable frequency of operation. For a full-fledged operating system such as Linux, or Windows, it can take several milliseconds to seconds to boot the system due to the size of the operating system, which can be annoying to the user. Keeping the boot loader/operating system in external parallel memory reduces the boot-up time drastically for systems in which boot time is critical.

NAND-flash memories are gaining popularity in numerous applications due to their higher read throughput, although this throughput is lower than that of NOR-flash memory. NAND flash also offers faster erases and lower cost per byte than does NOR flash. The primary use

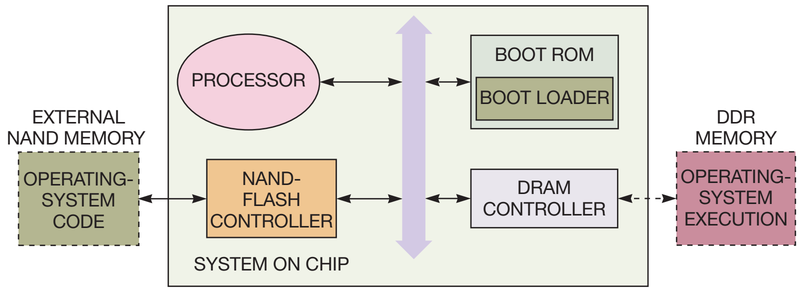

Figure 8 Secondary boot using high-performance DRAM requires a corresponding nonvolatile primary-boot device, such as a ROM or a flash memory.

for NAND-flash memory—for example, in a USB flash drive—is to store large quantities of data and code. However, in recent years, an increasing number of embedded-system applications also support NAND-flash memory as a primary boot option. In these cases, the microcontroller must include a NAND-flash-memory controller to handle read and write accesses. The NAND-flash-memory interface also requires more than 20 pins, so this option may not be cost-effective and otherwise practical unless you use a package with more pins.

Booting from internal volatile memory is always a secondary boot technique because a primary boot device must load the RAM before it can execute code. Commonly, a primary boot loads the operating system and drivers, copying them to RAM. At this point, the RAM-based code takes control of the system. Executing code from RAM is faster and consumes less power than do other memory technologies, including either external or internal flash memory. However, because the internal RAM is volatile, some system designs enable the RAM to switch to battery power in the event of the failure of the primary supply, thereby eliminating any further need to copy the code from external or internal memory when the primary power returns and thus reducing subsequent initialization and boot time.

Booting from a DRAM is also a secondary boot technique. DRAM use is common in high-end applications that must handle abundant multimedia content and that require high throughput. You can view DRAM as a bigger and faster extended RAM buffer for managing complicated applications. In one example, ROM includes the boot loader, and NAND-flash memory contains the operating system and the application code (**Figure 8**). The boot process begins with system initialization through the ROM-boot loader, which also includes the reset vector. The main operating-system code copies from NAND-flash memory to DDR SDRAM, subsequently switching execution control to the external volatile memory. This scheme is efficient because executing from DDR SDRAM is faster than reading directly from NAND-flash memory. You can use similar schemes to copy code from other interfaces, such as Ethernet.

You may want to boot from various interfaces, such as SDHC, SPI, I^2C (inter-integrated circuit), USB, SATA (serial advanced-technology attachment), PCIe, and Ethernet. These interfaces all represent secondary boots. A primary boot interface, such as ROM, initializes a secondary boot interface, such as USB, before code execution switches to the secondary boot device. Storing boot code in external nonvolatile serial memories, such as SPI flash memory or I^2C EEPROM, can be useful for microcontrollers that have low pin counts and can afford to have longer boot-ups. Such schemes first copy the boot code from the external memory to the on-chip RAM and code execution switches to the RAM after reset, so the scheme can immediately fetch boot code after reset deassertion. EDN

AUTHORS' BIOGRAPHIES

Mohit Arora is a systems engineer at Freescale Semiconductor. He has a bachelor's degree in electronics and communications engineering from Netaji Subhas Institute of Technology (New Delhi, India).

Varun Jain is a senior design engineer at Freescale Semiconductor. He has a bachelor's degree in engineering from Delhi College of Engineering (Delhi, India).

ADVANCES IN ENERGY-STORAGE TECHNOLOGY POWER WIRELESS DEVICES

BY MARGERY CONNER • TECHNICAL EDITOR

COMPOSITE ILLUSTRATION BY TIM BURNS. VAULT: MARK EVANS/ISTOCKPHOTO.COM; CIRCUIT PATTERN: MONSTER/ISTOCKPHOTO.COM

ENERGY HARVESTING RELIES ON INTERMITTENT ENERGY SOURCES AND REQUIRES ENERGY STORAGE, SUCH AS A CAPACITOR OR A BATTERY. NEW IMPROVEMENTS IN THESE COMPONENTS POINT TOWARD SMALLER, LONGER-LIVED WIRELESS SENSOR NODES, BUT A LONG-LIVED PRIMARY BATTERY MAY BE YOUR SIMPLEST ENERGY-SOURCE CHOICE.

Going wireless in applications such as sensor nodes or PC peripherals requires not only wireless communication but also an alternative energy source to an ac wall plug. It makes little sense to eliminate wired communication if your design still requires a power cord. Harvesting the ambient energy, which can be solar or artificial lighting or mechanical energy, such as equipment vibration, provides an energy source for ultra-low-power devices. Devices such as wireless sensor nodes conform to their ultra-low-power budget by spending most of their time asleep, waking up briefly to take a sensor reading and transmit it in a burst to their node network (see **sidebar** "Characterizing average harvested energy").

However, the burst of power necessary for communication, although brief, requires more energy than ambient sources can provide, so some form of energy reservoir must provide this short, intense burst, making energy storage the other side of the coin for energy harvesting. Capacitors, including supercapacitors, and batteries are the main forms of electrical energy storage for energy harvesting. Both device types have constraints that energy-harvesting circuits for wireless sensor networks must accommodate to make wireless nodes practical. Those constraints are high leakage current for supercapacitors and charge/discharge-cycle lifetimes for batteries. New lithium-ion technologies promise to expand the scope of both supercapacitors and batteries.

For applications in which your wireless design can operate from the energy that one long-life battery stores, the simplest source of energy is a large-capacity battery with minimum self-discharge, or leakage, current (see **sidebar** "Power constraints of wireless sensor nodes"). You may worry about the maintenance costs for replacing spent batteries. Depending on your device's power budget, however, some primary nonrechargeable batteries can provide power for 10 to 20 years, a longer lifetime than many electronic devices. Ask yourself how many electronic devices you have been using for more than 10 years.

Not just any battery chemistry is a good candidate. The common alkaline battery, for example, has only half the energy-storage capacity of an equivalent-sized lithium battery. On the other hand, alkaline batteries are inexpensive and have low leakage current, an important characteristic for use in an application that must go many years without service. An advantage of alkaline batteries in "mostly asleep" wireless networks is that the keep-alive power during the sleep phase can be at the alkaline battery's lower nominal voltage (**Reference 1**). Combining an alkaline battery with a dynamic-voltage regulator allows the voltage to drop to an alkaline battery's nominal voltage of 1.5V and supply 600-nA keep-alive current. When it's time to take a reading or talk to the network, the regulator can boost the battery voltage to the 3V that many electronic circuits require. Using a primary lithium battery with a nominal voltage of 3V, the keep-alive power is twice as large as that of an alkaline

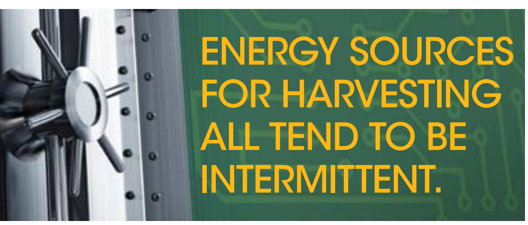

battery because of the two-times-larger nominal voltage (**Reference 2**).

Lithium-thionyl battery chemistry, on the other hand, can easily last for 10 to 20 years with little leakage. Lithium-thionyl batteries can provide a simple approach to your energy needs if your application has a finite life of a maximum of 20 years and can bear the higher cost of a lithium-thionyl battery, which is approximately $2 compared with 50 cents for an alkaline battery. In addition to low self-discharge rates, lithium thionyl also benefits from a flat discharge profile over time so that the

AT A GLANCE

- Energy storage is a key component of any power circuit because harvested energy is intermittent.
- Determining your energy budget tells you how much energy you need to get from either a primary battery or an energy-harvesting circuit.
- Supercapacitors have high leakage current but virtually unlimited life cycles.
- Batteries have low leakage current and long shelf life, but rechargeable batteries have limited charge/recharge cycles.

terminal voltage stays relatively constant over the battery's service life. This battery chemistry is more costly than other lithium chemistries and finds use in applications demanding long battery life, such as water and gas meters and other industrial- and military-electronics applications. The cells have nominal voltages of 3.6V and discharge voltages of 2.2V. Industrial- and military-grade lithium-thionyl-chloride primary batteries supply approximately 80 mAhr/year with a failure rate of only one per 1 million cells. As a comparison, a microgenerator that generates power from machine vibrations might cost $100 (**Reference 3**).

Consider how many electronic devices are still in use in industrial or infrastructure settings after more than 20 years. Especially for wireless sensor networks, which are still in their infancy, hardware will likely change over the next couple of years, let alone decades. A primary battery may be the answer for your energy-storage needs. However, some applications lend themselves well to energy harvesting. Take Logitech's K750 solar-power-harvesting wireless keyboard. Wireless keyboards can go for a long time on a battery, but the battery at some point becomes exhausted, and the keyboard dies, most likely when it's least convenient. Logitech has taken a page from Texas Instruments' solar-powered calculators and incorporated solar power in its K750 batteries (**Figure 1** and **Reference 4**). The approach allows a user to store a keyboard, for example, in a dark closet for three months before its battery power leaches away, and the keyboard can run for two months on the battery power alone.

In the world of mostly asleep sensor networks, the mostly asleep mode enables the long life of an energy-miserly system. At some point, however, the sensor node will wake up, sense its surroundings, and transmit its status and data to some receiving node. This burst of activity requires more power than the ambient amount of energy available in that one instant from the environment. In this case, it comes in handy to be able to store the trickle of available ambient energy and spend it all in one burst before going back to trickle-charging the storage device. In this way, it will be ready for the next time the node awakes to sense and communicate.

Advanced Linear Devices makes modules that harvest energy from solar cells and transducers (**Figure 2**). "Energy storage is key to energy harvesting because energy sources for harvesting all tend to be intermittent," says John Skurla, director of marketing and sales at the company. "It's not steady-state input energy. A microgenerator can act as a transducer to transform mechanical energy into electrical energy, and, from there, the electrical energy goes into some form of storage." Advanced Linear Devices' modules

Figure 1 The Logitech K750 wireless solar keyboard can run for two months on battery power alone. The solar cells top off the lithium-ion button batteries.

Figure 2 Advanced Linear Devices' EM300 energy-harvesting modules provide power management for a variety of energy sources, including piezoelectric transducers.

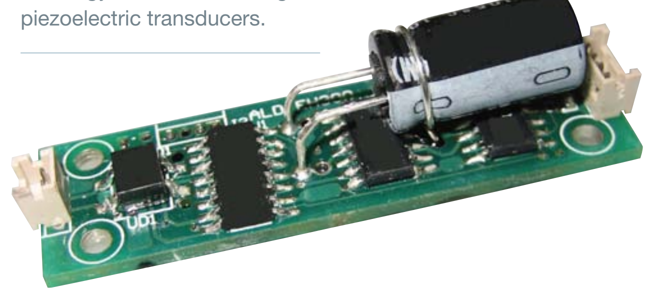

CHARACTERIZING AVERAGE HARVESTED ENERGY

By Tony Armstrong, Linear Technology Corp

Ambient-energy sources include light, heat differentials, mechanical vibration, transmitted RF signals, or any other source that can produce an electrical charge through a transducer. These energy sources are all around us, and you can convert them into electrical energy by using a suitable transducer, such as a thermoelectric generator for temperature differential, a piezoelectric element for vibration, a photovoltaic cell for sunlight or indoor lighting, and even galvanic energy from moisture. You can use these so-called free energy sources to autonomously power electronic components and systems.

The applications in which energy-harvesting benefits are compelling are parts of a long and growing list, and all these applications have two things in common: They all have a low average-power requirement, and they reside in situations in which access to ac-line power is costly or unavailable, batteries are costly or difficult to replace, or both. Examples include wireless sensors in shipping containers, asset-tracking devices, structural monitors in aircraft and bridges, environmental monitors in buildings or factories, and remote outdoor locations.

A good starting point for the feasibility of implementing an energy-harvesting system or wireless sensor node is to consider how much energy is necessary to power it. The good news is that only 48 μW of continuous power is necessary to run a sensor node for an indefinite amount of time. Furthermore, lower duty cycles produce even lower average-power levels. Another piece of good news is that a transducer's output power is generally more than adequate when you properly match it to the application, thereby enabling you to charge an external storage element for use when the ambient-energy source is unavailable.

Because energy harvesting is generally subject to low, variable, and unpredictable levels of available power, it is advisable to implement a hybrid structure that can interface to both the harvester and a secondary power reservoir. The harvester, because of its unlimited energy supply and deficiency in power, is the energy source of the system. The secondary power reservoir, either a battery or a capacitor, yields higher output power but stores less energy, supplying power when necessary but otherwise regularly receiving charge from the harvester. Thus, applications lacking ambient energy from which to harvest power must use the secondary power reservoir to power the sensor node.

From system designers' perspectives, this requirement adds a further degree of complexity because they must now consider how much energy the secondary reservoir must store to compensate for the lack of an ambient-energy source. Just how much energy is necessary depends on several factors, including the length of time the ambient-energy source is absent and the duty cycle of the sensor node—that is, how often a data reading and transmission must take place. Other factors include the size and type of a secondary reservoir—that is, whether it is a capacitor, a supercapacitor, or a battery—and whether enough ambient energy is available to act as the primary energy source and have sufficient leftover energy to charge up a secondary reservoir when it is unavailable for a certain time.

The need for these factors varies from application to application; however, consideration of these factors will assist system designers in choosing an energy-harvesting-system topology. Moreover, with analog switch-mode-power-supply-design expertise in short supply worldwide, it has been difficult to design an effective energy-harvesting system. The primary hurdle is the power-management aspects of remote wireless sensing.

Fortunately, recent product introductions in this area include devices that can extract energy from almost any source of light, heat, or mechanical vibration. Furthermore, their comprehensive features greatly simplify the difficult power-conversion-design aspects of an energy-harvesting chain. As a result, system designers and planners must prioritize the needs of their power management from the outset to ensure efficient designs and successful long-term deployments.

Author's biography

Tony Armstrong is director of product marketing, power products, at Linear Technology Corp.

currently use capacitors, but the company is monitoring advances in rechargeable batteries and supercapacitors.

Rechargeable lithium-ion coin batteries, such as the ML2032 that the Logitech keyboard uses, have been available for only about a year. Because of their lack of maturity, Skurla has a cautious approach to evaluating them. "[A lithium coin battery] tends to be very finicky in that you can't overcharge it, and you can't undercharge it, either," he says. "If you don't manage it just right, the battery will die." Chemical reactions occur during overcharging and undercharging. "If you overcharge, you trigger a different chemical reaction," says Skurla. "Another chemical reaction occurs during undercharging, and the cell becomes permanently disabled." The cells don't explode from undercharging, he adds. They just die. "That's one of the characteristics of these postage-stamp-sized rechargeable coin batteries," he adds. "You never used to have to deal with undercharge."

Hitachi Maxell publishes the most complete information on ML-series lithium-manganese-dioxide rechargeable batteries. The company sells them only to equipment manufacturers as built-in parts and does not supply them directly to users of equipment with these batteries (**Figure 3**). Hitachi Maxell specifies the parts for 1000 charge and discharge cycles.

Advanced Linear Devices has also considered supercapacitors for storing harvested energy. "Supercapacitors are

One great instrument deserves another. And another. And another…

The legendary performance and ease-of-use you expect from our oscilloscopes are now available for your entire bench. Introducing the first fully integrated set of test bench instruments from Tektronix designed to streamline testing and offer unparalleled connectivity, accuracy and usability. From shortcut keys for common functions to USB ports for quick data export, every Tektronix instrument is built to maximize measurements and minimize inefficiency. Plus, all instruments come with National Instrument's LabVIEW SignalExpress™ software, so you can automate measurements, analyze data across multiple instruments and control your entire bench from your PC. Performance, usability and connectivity. Now available for your entire test bench.

Oscilloscopes

- 2 or 4 analog channels
- Up to 25 automated measurements
- Battery operation available

Mixed Signal Oscilloscopes

- Up to 4 analog and 16 digital channels
- Parallel and serial bus automated analysis
- Power and video automated analysis
- Wave Inspector® for fast navigation of data

Arbitrary/Function Generators

- 1 or 2 channels
- 12 standard waveforms
- Arbitrary waveform capability

Digital Multimeters

- Up to 11 automated measurements
- Statistics, histogram and trending analysis modes

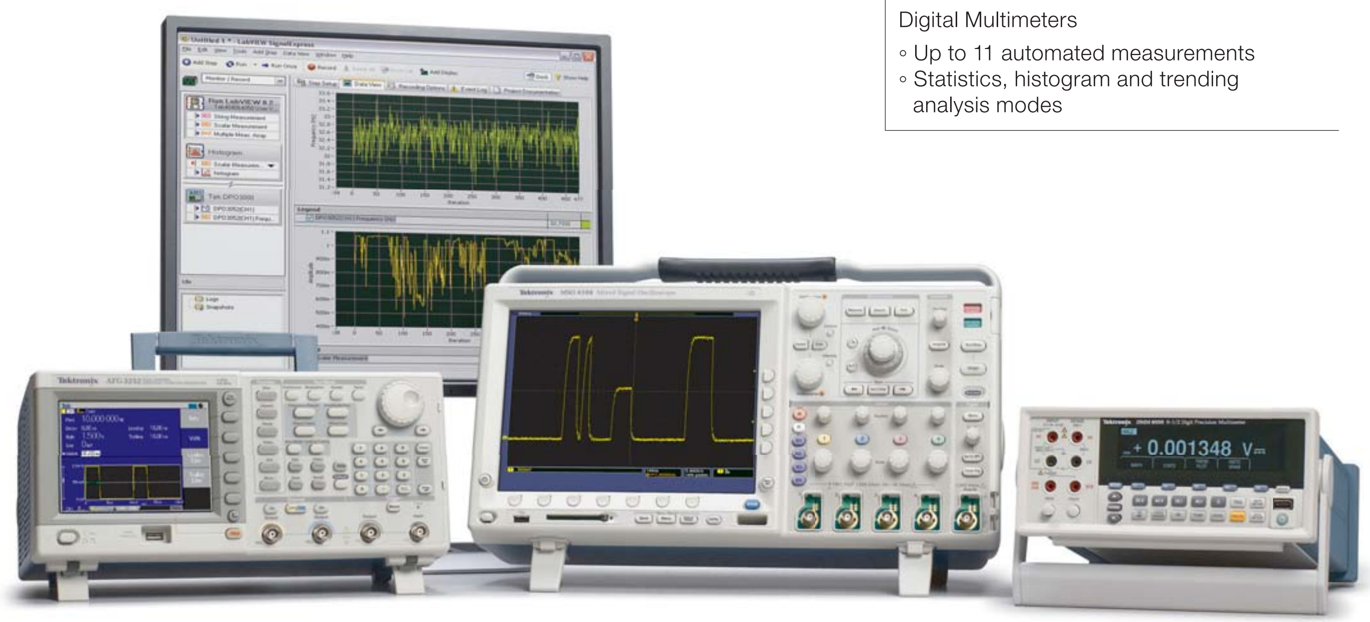

Take the new bench for a virtual spin at Tektronix.com/NewBench.

Tektronix®

POWER CONSTRAINTS OF WIRELESS SENSOR NODES

By Dave Freeman and Sriram Narayanan, Texas Instruments

Recent developments in hardware and software designs for wireless-sensor-network nodes motivate the need for system-level design methods. Figure A shows a network, along with the subsystems of each node. Considerations for ease of deployment and cost of installation require that the nodes be able to communicate wirelessly. To reduce communication overhead and improve response time, the nodes should be able to locally process sensor data and control actuators. Performing routine maintenance, such as replacing batteries on a large number of nodes, may be prohibitively expensive. Sensor nodes should also last for several years by relying only on stored or harvested energy.

The choice of sensors, radio, and microcontroller depends on the nature of the application. A sensor network in an office environment addresses applications such as energy management, security, or resource planning.

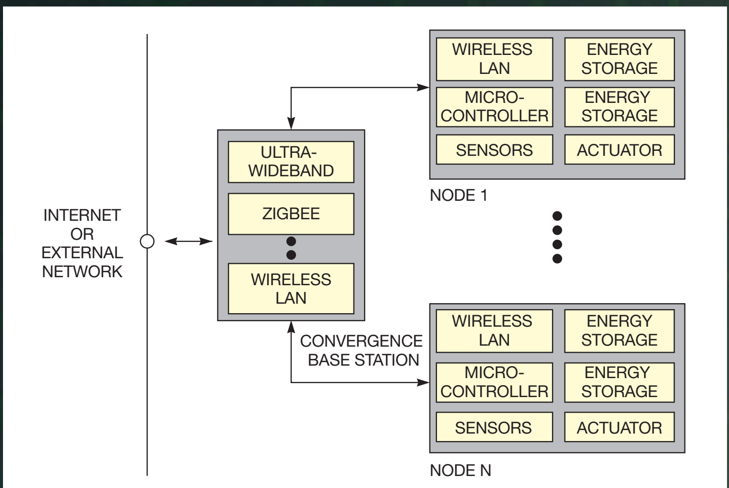

Figure A Harvested energy powers sensor nodes that make autonomous decisions about their environment and can communicate using multiple protocols.

Light energy is usually the most abundant form of ambient energy in indoor environments. Modern solar cells comprising amorphous silicon generate approximately 5 μW/cm^2 when illuminated by a 200-lux fluorescent-light source. Table A lists estimates of energy-accrual rates and suggests that a 10-cm^2 solar cell may generate 70 to 120 μW.

Microthermoelectric generators use a gradient in temperature to produce electrical energy. To generate a power density of 15 μW/cm^3, however, thermal harvesters may need a thermal gradient of approximately 10°C. Many applications, particularly indoor environments, do not experience large swings in temperature. Thus, thermal harvesters have limited applicability. Today's vibrational-energy harvesters need acceleration of approximately 1.75 to 2g—magnitudes usually absent in indoor environments—to produce 60 μW of power.

With limited onboard capacity to store energy and limited opportunities to harvest ambient energy, the sensor node must be frugal in energy use. Consider, for example, a 100-mAhr battery and a solar cell that accrues 70 μW for 50% of a 10-year node lifetime. This node must operate its various subsystems and consume no more than 39 μW of average power.

The microcontroller, radio, sensors, and actuators exhibit dramatically different power and performance characteristics. Meeting the system's power budget requires the sensor node to optimally manage its subsystems. Modern low-power microcontrollers consume approximately 345 μW of peak power to operate at a

TABLE A ACCRUAL RATES

Setting	**Intensity** (lux)	**Power density** (μW/cm^2)
Conference room	300	7.5
Office	500	12.5
Hallway	400	10

TABLE B LOW-POWER ARCHITECTURES

	Frequency (GHz)	**Data rate** (Mbps)	**Nominal range** (m)	**Peak power** (mW)	**Energy per bit** (nJ/bit)
Bluetooth	2.4	1	1 to 10	92.4	92
IEEE 802.15.4a ultrawideband	3.1 to 10.6	1 to 27	1 to 30	Less than 10	Approximately 1
Zigbee	2.4	0.25	10 to 100	29.7	119
Wireless LAN	2.4	54	100	851	15

super in capacity, but they also leak like sieves," says Skurla. It's also difficult to get precise specifications for leakage and shelf-life characteristics. He says that the company bought and tested a bunch of supercapacitors that were available on the market, running the tests over days and weeks. The leakage current is proportional to the capacity: As the capacitance increases, the leakage current also increases. "What you really want is a 10F supercap with the leakage-current characteristic of a 0.1F supercap," says Skurla. "They still have about 11-times more leakage than we like to see for an energy-harvesting application."

SUPERCAPACITORS ARE SUPER IN CAPACITY, BUT THEY ALSO LEAK LIKE SIEVES.

clock speed of about 1 MHz. Given that the requirements for sensor processing are often modest, you can use a microcontroller with an aggressive duty cycle—less than 1%, for example—to reduce average power consumption.

Sensor nodes usually communicate information about physical phenomena and related control messages at relatively low rates. Table B summarizes the salient features of some key low-power wireless-communication technologies. The power-consumption quantities in the table serve only as general guidelines for system design. As transceiver designs evolve, they consume less power. When choosing a transceiver architecture, it is important to consider all aspects of the design. A wireless-LAN transceiver consumes lower energy per bit than does a Zigbee transceiver, but the LAN transceiver has a higher data rate and consumes more peak power.

Sensors that may find use in indoor applications include thermometers, humidity sensors, microphones, and passive infrared sensors. Today's temperature and humidity sensors and microphones consume approximately 70 to 80 μW of peak power. Passive infrared sensors that can detect human activity typically consume peak power of 100 to 500 μW. Temperature and humidity sensors monitor slowly varying phenomena and can operate on low duty cycles, but you cannot power off other sensors that detect motion without compromising detection performance. In many applications, the sensor demands more energy than the processing of the data or wireless communication. Therefore, meeting the system power budget requires innovation in the way you manage the sensors.

The lack of an adequate power supply and energy remains a formidable challenge to implementing wireless sensor networks, despite advancements in computation, communication, and sensing. Technological advancements in energy harvesting and storage continue to ease power constraints, but the demands of the end application continue to push their limits. Bridging this persistent power gap requires a system-level-design approach that optimally trades off performance for energy savings and that guarantees a minimum quality of service. Future wireless sensor nodes will autonomously adapt to time-varying-application demands and energy availability.

See the Web version of this article at www.edn.com/110203cs for a list of references used in writing this sidebar.

FOR MORE INFORMATION

Advanced Linear Devices
www.aldinc.com

Analog Devices
www.analog.com

Cap-XX
www.cap-xx.com

Hitachi Maxell
www.maxell.co.jp

Ioxus
www.ioxus.com

JM Energy
www.jmenergy.co.jp

Linear Technology
www.linear.com

On Semiconductor
www.onsemi.com

Texas Instruments
www.ti.com

Figure 3 Maxell sells rechargeable lithium-manganese-dioxide ML-series rechargeable batteries. The company offers these batteries only to equipment manufacturers as built-in parts, not directly to end users.

Another advancement in the technology of supercapacitors is the emergence of lithium-ion supercapacitors. For example, JM Energy two years ago announced its first lithium-ion supercapacitors, the 1000F/2000F series. The devices remain in limited production, however. In December 2010, supercapacitor-manufacturer Ioxus announced its lithium-ion series with an energy density 115% higher than that of standard electric double-layer capacitors.

Ioxus last year received a federal grant to work with Binghamton University (www.binghamton.edu) on increasing energy density in supercapacitors. The company's resulting work yielded a "hybrid" capacitor, which comes in 220, 300, 800, and 1000F versions. It uses a positive carbon electrode and a negative lithium-ion electrode and has a 20,000-cycle lifetime compared with approximately 1000 cycles for a battery. Leakage current is 0.6 mA at 72 hours. The devices' height and diameter measurements, respectively, range from 45×22 mm for the 220F version to 88×35 mm for the 1000F version. Prices for the 300F version range from $4 to $6 each. EDN

REFERENCES

1. Conner, Margery, "Run for your life: ultralow-power systems designed for the long haul," *EDN*, May 26, 2005, pg 46, http://bit.ly/f6q6rP.
2. Odland, Keith, "Selecting the best battery for embedded-system applications," *EDN*, Nov 18, 2010, pg 36, http://bit.ly/evzniq.
3. Conner, Margery, "Energy-harvesting conference covers both cutting-edge and established technologies," *EDN*, June 13, 2007, http://bit.ly/gCHicK.
4. Conner, Margery, "Energy harvesters extract power from light, vibrations," *EDN*, Oct 27, 2005, pg 45, http://bit.ly/hx8Xqw.

You can reach Technical Editor Margery Conner at 1-805-461-8242 and margery.conner@ubm.com.

BY ALLAN EVANS • SAMPLIFY SYSTEMS INC

Lowering the cost of medical-imaging R&D

MEDICAL-EQUIPMENT COMPANIES ARE SHIFTING FROM A VERTICAL-INTEGRATION MODEL TO A SYSTEM-INTEGRATION MODEL TO DO MORE WITH LESS RESEARCH AND DEVELOPMENT.

A look inside a typical piece of medical-imaging equipment reveals myriad technologies from many engineering disciplines that must collaborate to form a system. For example, developing X-ray scintillators in CT (computerized-tomography) machines to convert X-ray-particle energy into photon energy requires materials-science research into rare-earth materials. Capturing this photon energy over wide dynamic ranges requires custom analog-front-end electronics, and acquiring the data from tens or even hundreds of thousands of such X-ray detectors requires customized storage subsystems and interface technologies. Complex software algorithms then form images from this substantial volume of raw data. Radiologists, sonographers, and other medical-imaging professionals examine these images to make their diagnoses. This vertically integrated, multidisciplinary engineering from materials-science research to custom silicon and complex software design has historically required hefty R&D budgets, which only large companies could afford.

The current economic environment has caused an increased focus on price and performance and forced medical-equipment OEMs to re-evaluate their R&D budgets to focus on technologies that are truly keys to their competitive advantage, outsourcing other nondifferentiated imaging functions to outside semiconductor and subsystem suppliers. These merchant suppliers of technologies can, in turn, enable cost reductions by amortizing their R&D costs over the entire industry, not just across the market share of one equipment manufacturer.

The evolution of the ultrasound-equipment market provides an instructive example of how this trend will play out. **Figure 1** depicts a 10-year-old ultrasound machine whose probe contains 64 to 256 piezoelectric transducers, each individually connected through a microcoaxial cable to the console. The probe required a specialized manufacturing process for matching gains and delays across elements. As manufacturers introduced imaging techniques, analog front ends primarily comprised discrete components. Top-tier imaging-equipment manufacturers held this "black art" close to the vest. Processing the tens or hundreds of gigabits per second of data could take place only in custom ASICs. When medical applications began to use analog CRTs for ultrasound dis-

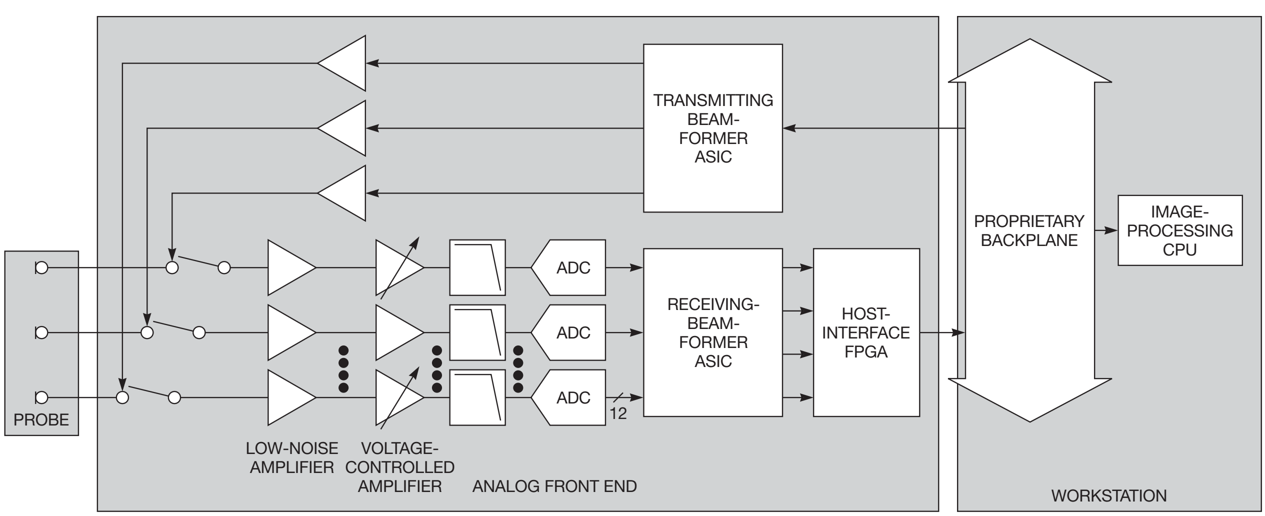

Figure 1 This 10-year-old ultrasound machine required a specialized manufacturing process for matching gains and delays across elements.

LINEAR TECHNOLOGY

DESIGN NOTES

Supercapacitor-Based Power Backup Prevents Data Loss in RAID Systems – Design Note 487

Jim Drew

Introduction

Redundant arrays of independent disks, or RAID, systems, by nature are designed to preserve data in the face of adverse circumstances. One example is power failure, thereby threatening data that is temporarily stored in volatile memory. To protect this data, many systems incorporate a battery-based power backup that supplies short-term power—enough watt-seconds for the RAID controller to write volatile data to nonvolatile memory. However, advances in flash memory performance such as DRAM density, lower power consumption and faster write time, in addition to technology improvements in supercapacitors such as lower ESR and higher capacitance per unit volume, have made it possible to replace the batteries in these systems with longer lasting, higher performance and "greener" supercapacitors. Figure 1 shows a supercapacitor-based power backup system using the LTC®3625 supercapacitor charger, an automatic power crossover switch using the LTC4412 PowerPath™ controller and an LTM®4616 dual output µModule® DC/DC converter.

The LTC3625 is a high efficiency supercapacitor charger ideal for small profile backup in RAID applications. It comes in a 3mm × 4mm × 0.75mm 12-lead DFN package and requires few external components. It features a programmable average charge current up to 1A, automatic cell voltage balancing of two series-connected supercapacitors and a low current state that draws less than 1µA from the supercapacitors.

Backup Power Applications

An effective power backup system incorporates a supercapacitor stack that has the capacity to support a complete data transfer. A DC/DC converter takes the output of the supercapacitor stack and provides a constant voltage to the data recovery electronics. Data transfer must be completed before the voltage across the supercapacitor stack drops to the minimum input operating voltage (V_{UV}) of the DC/DC converter.

To estimate the minimum capacitance of the supercapacitor stack, the effective circuit resistance (R_T) needs to be determined. R_T is the sum of the ESR of the supercapacitors, distribution losses (R_{DIST}) and the $R_{DS(ON)}$ of the automatic crossover's MOSFETs:

$$R_T = ESR + R_{DIST} + R_{DS(ON)}$$

Allowing 10% of the input power to be lost in R_T at V_{UV}, $R_{T(MAX)}$ may be determined:

$$R_{T(MAX)} = \frac{0.1 \bullet V_{UV}^2}{P_{IN}}$$

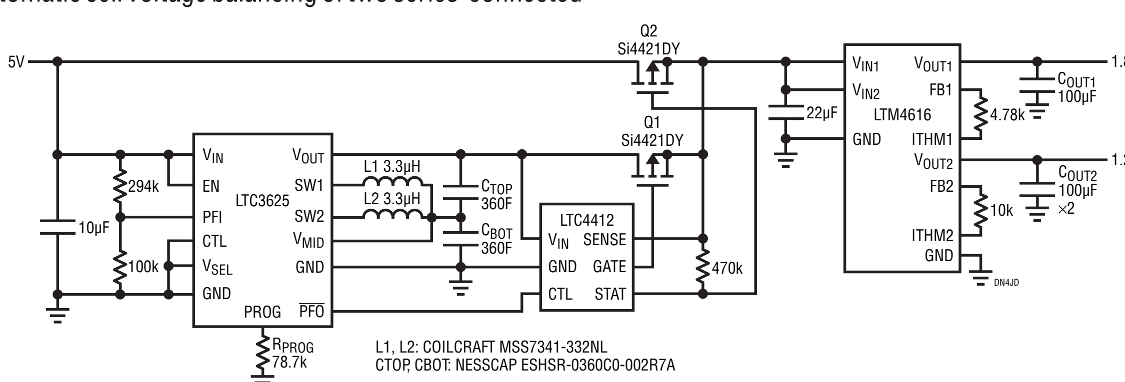

Figure 1. Supercapacitor Energy Storage System for Data Backup

The voltage required across the supercapacitor stack ($V_{C(UV)}$) at V_{UV}:

$$V_{C(UV)} = \frac{V_{UV}^2 + P_{IN} \bullet R_T}{V_{UV}}$$

The minimum capacitance (C_{MIN}) requirement can now be calculated based on the required backup time (t_{BU}) to transfer data into the flash memory, the initial stack voltage ($V_{C(0)}$) and ($V_{C(UV)}$):

$$C_{MIN} = \frac{2 \bullet P_{IN} \bullet t_{BU}}{V_{C(0)}^2 - V_{C(UV)}^2}$$

C_{MIN} is half the capacitance of one supercapacitor. The ESR used in the expression for calculating R_T is twice the end-of-life ESR. End of life is defined as when the capacitance drops to 70% of its initial value or the ESR doubles.

The Charge Profile into Matched SuperCaps graph in the LTC3625 data sheet shows the charge profile for two configurations of the LTC3625 charging a stack of two 10F supercapacitors to 5.3V with R_{PROG} set to 143k. This graph, combined with the following equation, is used to determine the value of R_{PROG} that would produce the desired charge time for the actual supercapacitors in the target application:

$$R_{PROG} = 143k \bullet \frac{10F}{C_{ACTUAL}} \bullet \frac{5.3V - V_{C(UV)}}{V_{OUT} - V_{C(UV)}} \bullet \frac{t_{RECHARGE}}{t_{ESTIMATE}}$$

$V_{C(UV)}$ is the minimum voltage of the supercapacitors at which the DC/DC converter can produce the required output. V_{OUT} is the output voltage of the LTC3625 in the target application (set by V_{SEL} pin). $t_{ESTIMATE}$ is the time required to charge from $V_{C(UV)}$ to the 5.3V, as extrapolated from the charge profile curves. $t_{RECHARGE}$ is the desired recharge time in the target application.

Design Example

For example, say it takes 45 seconds to store the data in flash memory where the input power to the DC/DC converter is 20W, and the V_{UV} of the DC/DC converter is 2.7V. A $t_{RECHARGE}$ of ten minutes is desired. The full charge voltage of the stack is set to 4.8V—a good compromise between extending the life of the supercapacitor and utilizing as much of the storage capacity as possible. The components of R_T are estimated: R_{DIST} = 10mΩ, ESR = 20mΩ and $R_{DS(ON)}$ = 10mΩ.

The resulting estimated values of $R_{T(MAX)}$ = 36mΩ and R_T = 40mΩ are close enough for this stage of the design.

Data Sheet Download

www.linear.com

$V_{C(UV)}$ is estimated at 3V. C_{MIN} is 128F. Two 360F capacitors provide an end-of-life capacitance of 126F and ESR of 6.4mΩ. The crossover switch consists of the LTC4412 and two P-channel MOSFETs. The $R_{DS(ON)}$, with a gate voltage of 2.5V, is 10.75mΩ (max). An R_T of 26.15mΩ is well within $R_{T(MAX)}$. The value for R_{PROG} is estimated at 79.3k. The nearest standard 1% resistor is 78.7k. The data sheet suggests a 3.3µH value for both the buck and boost inductors.

The LTC3625 contains a power-fail comparator, which is used to monitor the input power to enable the LTC4412. A voltage divider connected to the PFI pin sets the power fail trigger point (V_{PF}) to 4.75V.

Figure 2 shows the actual backup time of the system with a 20W load. The desired backup time is 45 seconds, whereas this system yields 76.6 seconds. The difference is due to a lower R_T than estimated and an actual V_{UV} of 2.44V. Figure 3 shows the actual recharge time of 685 seconds compared to the 600 seconds used in the calculation, a difference due to the lower actual V_{UV}.

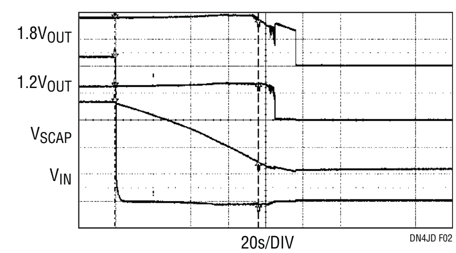

Figure 2. Supercapacitor Backup Time Supporting a 20W Load

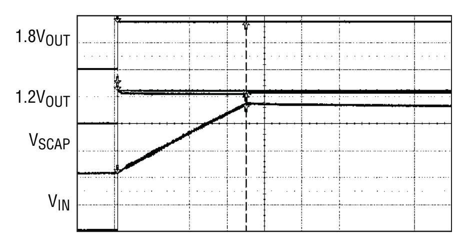

Figure 3. Recharge Time After Backup

Conclusion

Supercapacitors are replacing batteries to satisfy green initiative mandates for data centers. The LTC3625 is an efficient 1A supercapacitor charger with automatic cell balancing that can be combined with the LTC4412 low loss PowerPath controller to produce a backup power system that protects data in storage applications.

For applications help,
call (978) 656-3768

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

dn487f LT/AP 0211 226k • PRINTED IN THE USA

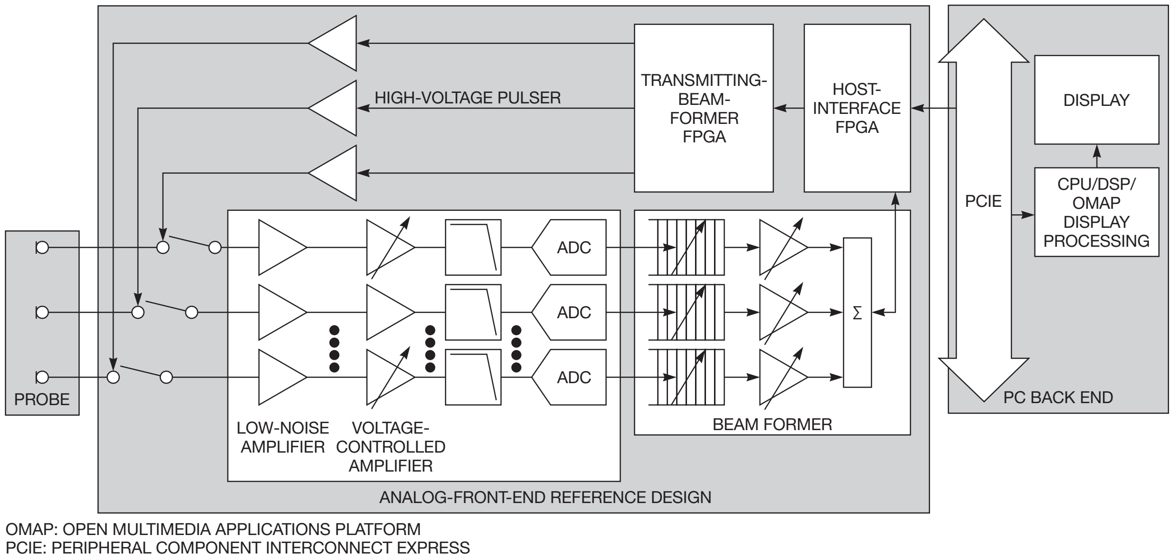

Figure 2 Semiconductor manufacturers have introduced highly integrated analog-front-end devices and modules for ultrasound applications and replaced ASICs with lower-cost FPGAs for digital processing.

plays, manufacturers began to develop proprietary interfaces and backplanes for image processing.

Now fast-forward 10 years (**Figure 2**). The high-level picture looks the same, but the components are different. Manufacturers now typically outsource the labor-intensive manufacturing of the probe to ODMs (original-design manufacturers). As the imaging modes matured, semiconductor companies began to understand the specifications for analog performance. Consequently, they began to introduce highly integrated analog front ends and modules with as many as 32 channels for the ultrasound-equipment market. Since mid-2008, companies have introduced more than a dozen analog- and mixed-signal semiconductor devices for this market.

Digital signal processing for functions such as beam forming is also moving to the merchant semiconductor suppliers. Previously, at the 180-nm node, top-tier ultrasound OEMs could justify the NRE (nonrecurring-engineering) costs for custom ASIC designs for digital processing. Many of these designs also featured mixed-signal integration with ADCs at 8 or 10 bits/sample. However, with new imaging technologies, such as tissue harmonic imaging and pulse-mode color Doppler, leading-edge ASIC designs have become cost-prohibitive, and the integration of high-performance 12-bit ADCs has become too complex. Today's mask costs, even at the 65-nm node, are too expensive even for market leaders. As a result, ultrasound OEMs have turned to FPGAs for their DSP functions. Manufacturers are building FPGAs at the 45- or 40-nm node and are announcing 28-nm parts, and they feature gate densities many times larger than ASIC densities at the 180-nm node. Now ultrasound OEMs can purchase an entire analog-front-end subsystem complete with beam forming, which can generate ultrasound images out of the box, rather than taking six months and many hundreds of thousands of dollars to custom-design their own.

OEMs ARE ENTERING THE ULTRASOUND MARKET WITH A FOCUS ON NICHES.

The range of applications for ultrasound equipment has increased greatly by enabling OEMs to focus their human and capital resources on their core features. Specialized ultrasound machines are now available for mammography, surgical guidance, and cardiography, and new imaging form factors are emerging for emergency-room and ambulance applications. The availability of these products has lowered R&D costs, allowing OEMs to enter the ultrasound-equipment market with a focus on niches. Top-tier OEMs also have benefited from these subsystems and silicon by allowing these OEMs to rapidly broaden their product portfolios rather than rely on a one-size-fits-all machine. The health-care industry has benefited through an improved price-to-performance ratio across a range of applications.

These approaches can provide the same benefits to imaging technologies besides ultrasound. For example, a CT machine is the result of vertically integrated development, much the same as ultrasound was 10 years ago. Again, OEMs' research into materials science results in customized X-ray scintillators that turn X-ray-particle energy into photons. These scintillators require customized analog-front-end silicon. Customized algorithms and protocols amplify, digitize, encode, and aggregate these scintillators' data and then transmit it to a workstation across a slip-ring device. In the CT-image-reconstruction workstation, customized network-interface cards terminate custom protocols, and custom RAID (redundant-array-of-independent-disk) controllers buffer the raw data at high rates until the workstation can reconstruct the image at a lower rate (**Figure 3**). These workstations can often reconstruct a 30-second CT scan in two to five minutes.

The high data rates for real-time CT data acquisition drive much of the need for custom hardware in the workstation in which the data is stored. Consider a 64-slice CT machine with 1000 detectors per slice with a rate of 10k samples/

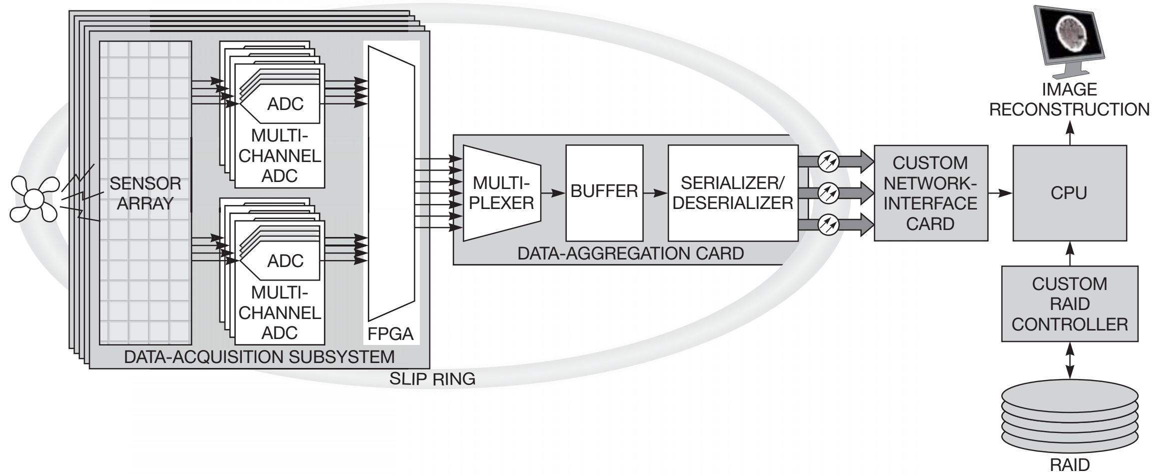

Figure 3 In a CT-image-reconstruction workstation, customized network-interface cards terminate custom protocols, and custom RAID controllers buffer the raw data at high rates until the system can reconstruct the image.

sec and 16 bits/sample. The aggregate throughput requirement is 1280 Mbytes/sec, a rate that exceeds the throughput of commercially available RAID controllers, which top out at 800 Mbytes/sec. CT-signal compression can reduce the throughput requirement of RAID controllers. For example, GE Healthcare (www.gehealthcare.com), Samplify Systems (www.samplify.com), and Stanford University (www.stanford.edu) have demonstrated the efficacy of signal compression in both lossless and nearly lossless modes to achieve compression ratios that provide compelling reductions of 3-to-1 or even 4-to-1 in data-acquisition rates (**Reference 1**).

Figure 4 shows an original image and a sample image of 3-to-1 compressed data. For more than 400 images, a Stanford radiologist could not distinguish the image of compressed samples from the image of noncompressed samples. With signal compression on the rotor side of the slip ring and decompression in software on the workstation, the CT system realizes the benefit of bit-rate reduction across the entire signal chain, including the slip ring, the network-interface card, and the RAID controller. Integration of signal compression into CT slip rings makes the compression transparent to the rest of the CT system so that such compressing slip rings can be drop-in replacements for older slip rings.

In CT machines, image reconstruction typically occurs three to 10 times slower than does CT data acquisition. This asymmetry creates the need for RAIDs in the system to buffer the raw data until the CT machine can reconstruct it into an image. Because the RAID is the last link in the data-acquisition chain, CT-machine designers must build the workstation on the CT gantry by considering data-acquisition rates instead of image-reconstruction rates. A new architecture helps designers build CT workstations based on the lower image-reconstruction rates rather than the higher data-acquisition rates (**Figure 5**).

With compression of the X-ray-detector data in the slip ring, you can now replace a custom multilane fiber-optic interface with a standard storage interface, such as FibreChannel or InfiniBand. This approach allows you to use off-the-shelf RAIDs outside the workstation. In this architecture, an off-the-shelf FibreChannel or InfiniBand switch provides the connectivity between the gantry or the data-acquisition subsystem, storage subsystem, and workstation. The workstation can now use standard low-cost network-interface cards rather than the customized interfaces that today's gantry-to-console connectors use.

Furthermore, because image-reconstruction rates are several times lower than data-acquisition rates, designers can build decompression into the RAID, enabling the integration of compression into CT systems in a manner that is transparent to image-reconstruction software. With the elimination of the need for a RAID or its transition into the gantry, the designer can now base the workstation on commodity PC-server hardware.

With R&D budgets tightening, medical-imaging OEMs must deliver next-generation technologies across a wider range of market segments with lower development costs. This move requires them to rationalize technologies that target the competitive advantages of their machines. EDN

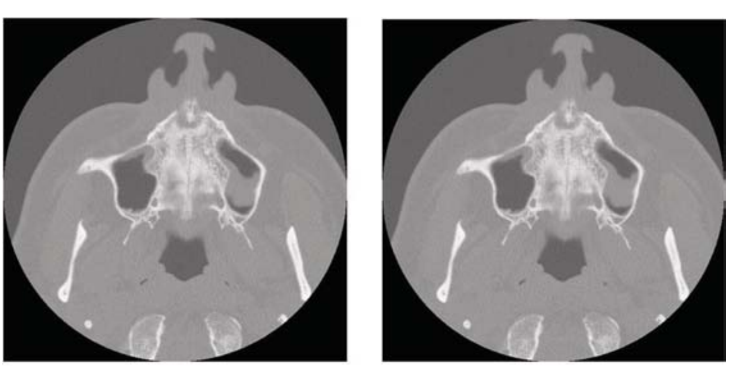

Figure 4 A cranial image (a) is indistinguishable from the reconstructed image (b) after 3-to-1 data compression.

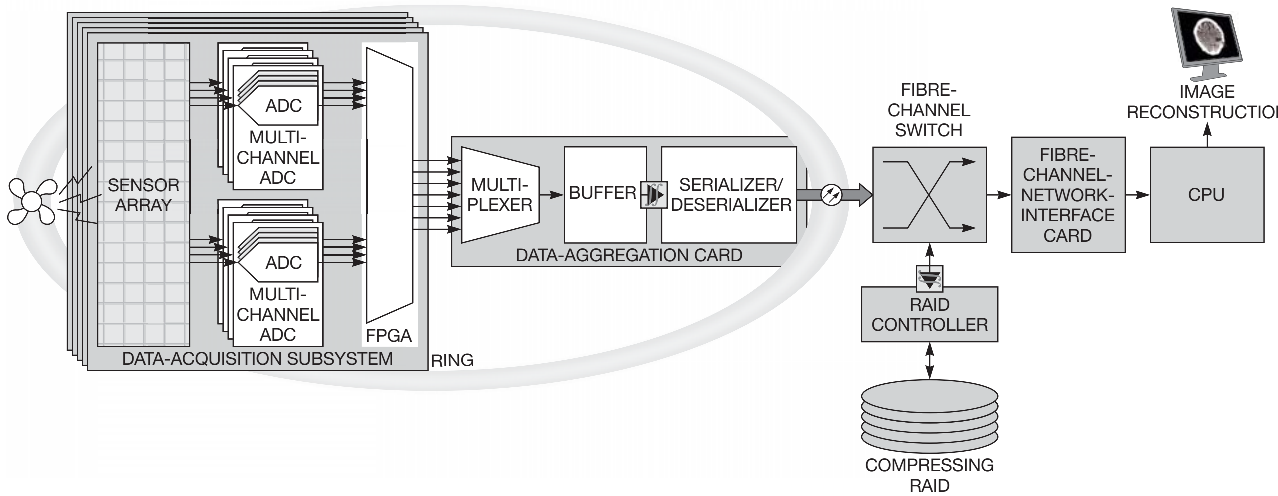

Figure 5 A new architecture for CT workstations puts compression of the X-ray-detector data into the slip ring, so that you can replace custom multilane fiber-optic interfaces with a standard storage interface, allowing the use of off-the-shelf RAIDs.

REFERENCE

1 Wegener, Albert; Naveen Chandra; Yi Ling; Robert Senzig; and Robert Herfkens, "Real-time compression of raw computed tomography data: technology, architecture, and benefits," *Proceedings of the SPIE Medical Imaging Conference*, Volume 7258, pg 7258H, March 2009, www.samplify.com/ct.

AUTHOR'S BIOGRAPHY

Allan Evans is vice president of marketing at Samplify Systems Inc (Santa Clara, CA), a provider of signal-compression and beam-forming technology for the medical-imaging market. He holds a master's degree in electrical engineering from the University of California—San Diego and a master's degree in business administration from Santa Clara University (Santa Clara, CA).

5A LDO for Digital ICs

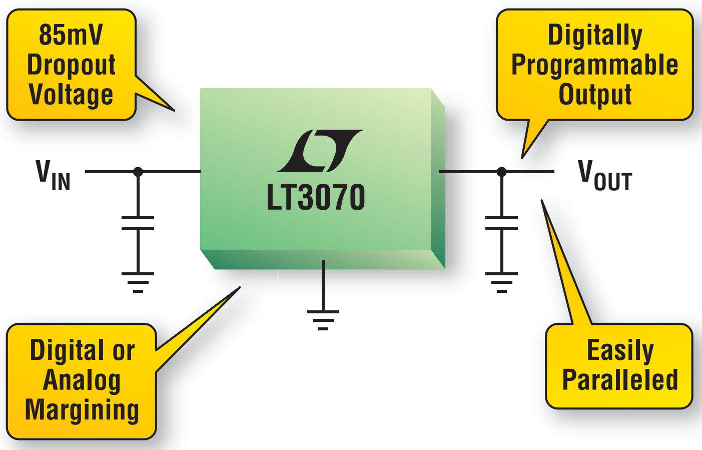

Transient Load Response

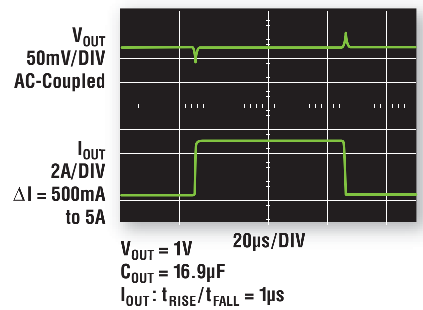

Ultrafast Transient Response Minimizes Output Capacitance

Our new generation linear regulator, the LT®3070, provides all the necessary performance and features to power high current and low voltage digital ICs such as FPGAs and DSP cores. It delivers 5A of output current with voltages as low as 0.8V, noise performance of only 25μV_{RMS} and very fast transient response with minimal output capacitance. Its MOS-based architecture enables ultralow dropout of only 85mV at full load with no variation in ground pin current even as the input or output voltage changes. The device's high PSRR makes it excellent for post-regulating a switching regulator. Devices can also be paralleled for 10A or more output current.

Features

- Ultralow Dropout Voltage: 85mV
- Ultralow Noise: 25μV_{RMS} (10Hz to 100kHz)
- Fast Transient Response Over Line & Load
- Input Voltage Range: 0.95V to 3V
- Digitally Programmable Output: 0.8V to 1.8V in 50mV Increments
- Digital Margining (LT3070) Analog Margining (LT3071)
- Stable with only 15μF of Ceramic Output Capacitance

Ripple Rejection vs Frequency

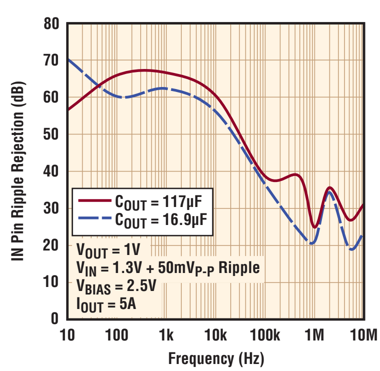

Info & Free Samples

www.linear.com/3070

1-800-4-LINEAR

Free Telecom, Datacom and Industrial Power Products Brochure

www.linear.com/48vsolutions

LINEAR TECHNOLOGY

Compute a histogram in an FPGA with one clock

Mohit Kumar, Texas Instruments, Bangalore, India

Histograms are often useful tools for analyzing digital data. To get reliable results from a histogram, though, you must collect large amounts of data, often with 100,000 to 1 million points. If you need to collect an ADC's digital outputs for analysis, you can use an FPGA (**Figure 1**).

The **figure** shows the histogram, RAM, and pulse-generator blocks, which let you capture and display the histogram computation based on 14-bit data. The RAM block is the FPGA's built-in RAM, and the histogram block is the VHDL (very-high-level-design-language) code to compute the histogram. You can also download the VHDL code for this application from the online version of this Design Idea at www.edn.com/110203dia.

The 14-bit parallel data, Device_Data[13..0], from an ADC goes to the histogram block and to the RAM Rd_Addr input. The RAM provides the data at its address location, RAMDataOut[15..0]. This data loops back to the histogram block, which increments it by one and sends it to output pin DataOut[15..0], a 16-bit data output. When the WREN (write-enable) pin is at logic level one, the data is written at the address at pin Wr_Addr[13..0]. That approach is the same as if the data were coming from Device_Data[13..0].

The RAM has a fixed delay from its input to its output. That is, when the input is at Rd_Addr, data becomes available at its output pin, RAMDataOut, after a fixed delay. This delay might vary across the FPGA. To take care of this delay, allow a two-clock delay to Device_Data before calculating the histogram. The delay in RAM should be less than two clock periods; otherwise, there could be data loss. This constraint limits the maximum frequency of Device_Clk.

The Cntr_Value gives the number of input data for which the histogram block computes the histogram. To reset the

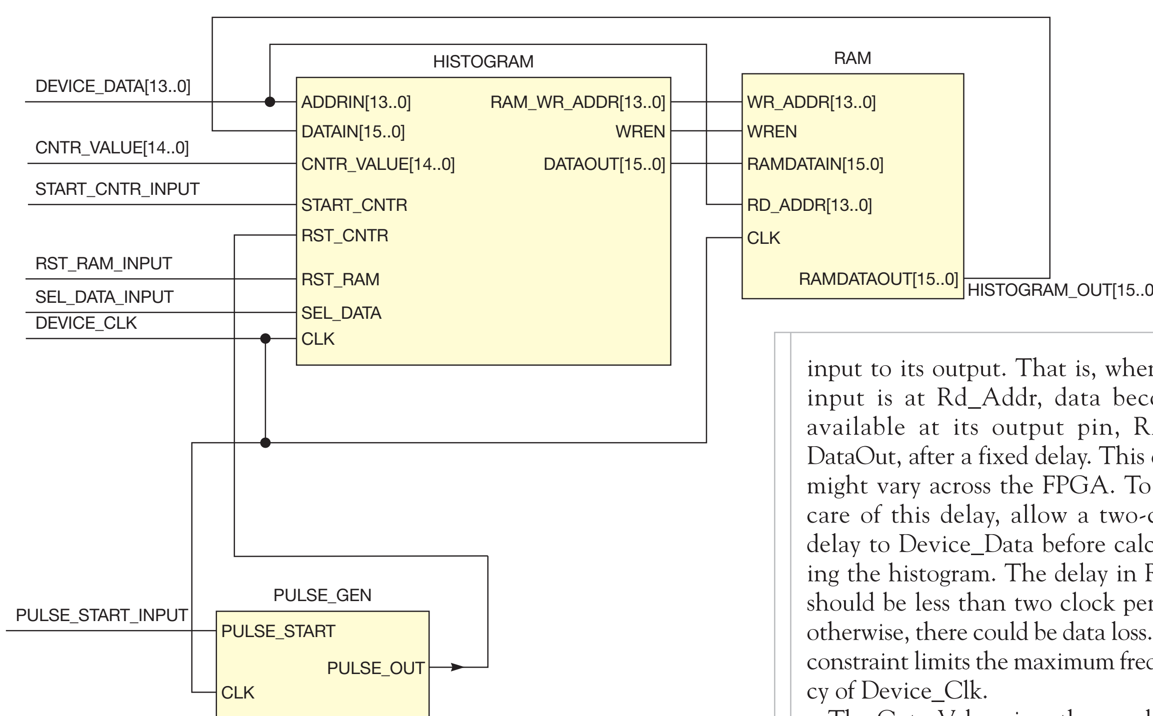

Figure 1 A histogram computational circuit retrieves data from an FPGA's RAM block.

counter, the Pulse_Gen block generates a pulse, which enters at input Rst_Cntr. At this point, the histogram again computes the histogram for the next set of input data from Cntr_Value. The Cntr_Value is 15 bits, but you can increase it to collect more histogram data.

The signals Sel_Data and Rst_RAM reset the data stored in the FPGA's RAM. Whenever the high signal is at the Rst_RAM pin, the DataOut pin of the histogram block gives all bits as 0. When the high signal is at the Sel_Data input pin of the histogram block, the output from RAM_Wr_Addr is not the Device_Data but an internally generated ramp that ramps up from 0 to 16,384. The histogram block does no computation because doing so would reset the address of the RAM.

When the FPGA completes the histogram computation, the RAM can read the histogram data by selecting Sel_Data as logic high and keeping Rst_RAM as logic low. The data in the RAM address sequentially exits the output pin, and you can transfer the data to a PC. Because all the blocks run on a single clock, Device_Clk, the design is simple and helps you meet timing constraints. You can easily modify the design to accommodate 16- or 12-bit data histograms.EDN

Protect MOSFETs in heavy-duty inductive switch-mode circuits

Marián Štofka, Slovak University of Technology, Bratislava, Slovakia

The MOSFET power switch is commonly the most vulnerable part of a new switched-mode high-power circuit. One threat for this device is exceeding the value of the maximum allowed pulse current. You cannot exceed this limit, even for pulse durations as short as 10 nsec. You could still thermally damage the MOSFET with a high duty cycle even when the drain-to-source current has a value between the peak and the dc ratings. The FET might eventually enter self-oscillations at a frequency, which might be an order of magnitude higher than your planned operating repetition rate. To protect the FET, you can limit the duty cycle by ac coupling the FET-driver circuit. If you further limit the repetition rate to tens of kilohertz, you needn't worry about thermal considerations.

To limit the duty cycle of the pulses, use the Schmitt-trigger input of IC_1 (**Figure 1**). You pass the input-voltage waveform through a derivative circuit comprising C_D, R_D, and R_S. The low-to-high transition of the clock causes an abrupt rise of voltage at resistor R_D. The output of the noninverting driver therefore goes high. Immediately after this transition, the voltage on R_D starts to decrease exponentially. When it falls below V_{TL}, the lower threshold of input IN_A, output OUT_A abruptly falls to 0V. The time constant $(R_D+R_S+R_{GEN})\times C_D$ yields the rate of exponential decrease. R_{GEN} is the output resistance of the generator of the input clock. You can calculate the value of capacitor C_D using the desired pulse width, T_P:

$$C_D=\frac{T_P}{R_D+R_S+R_{GEN}}\times\frac{1}{\ln\left[\frac{V_{DD}}{V_{TL}}\times\frac{R_D}{R_D+R_S+R_{GEN}}\right]}.$$

The **equation** employs an estimate of

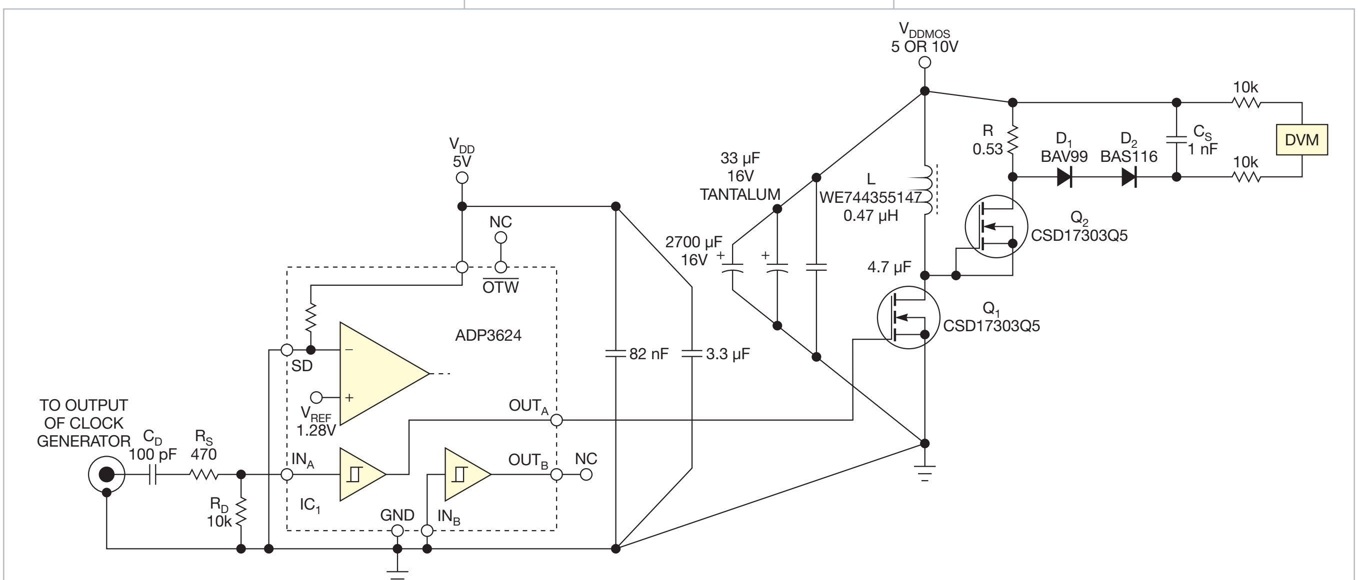

Figure 1 The test setup allows you to operate a MOSFET power switch at a fixed turn-on time. The power circuit remains cool with a 10-kHz repetition rate, even at peak inductor currents of tens of amperes.

the value of V_{TL}:

$$V_{TL} \simeq \sqrt{V_{INL} \times V_{INH}} - 0.1V = 1.165V.$$

The IC's data sheet gives the values of 0.8 and 2V as the limits of the low and high input voltages, respectively. The high-to-low transition of the clock has no effect. This transition causes a sharp negative exponential pulse, which an internal Schottky diode at input IN_A suppresses. The anode of this internal diode connects to ground, and its cathode connects to input IN_A. Resistor R_S limits the peak current flowing through the protective diode to about 10 mA.

The IC has an output current of ±4A. The typical on-resistance of Q_1 is 2 mΩ. You interconnect Q_2's gate and source pins to create a freewheeling diode. This diode has a typical reverse-recovery time of 33 nsec at a 25A forward current. When Q_1 turns off, the peak inductor current flows through Q_2. Voltage V_R occurs on power resistor R and is superimposed onto the supply voltage, V_{DDMOS}. The sum of these voltages must be lower than or equal to the manufacturer's specified value of the drain-to-source voltage of transistors Q_1 and Q_2.

When testing the circuit, you should monitor the dc-supply current. You can calculate the ideal-case supply current as a function of supply voltage on the power section and the pulse period as follows:

$$I_{SID} = \frac{1}{2} \times \frac{V_{DDMOS} T_{PON}^2}{L} \times f_{REP} = \frac{1}{2} I_{LPEAK} T_{PON} f_{REP}.$$

You calculate the pulse width of a single interval when the channel of Q_1 is conductive as an approximation relating the rise, fall, on, and off times of the FET:

$$T_{PON} \simeq T_P + t_{DOFF} - t_{DON} + t_{DMOSOFF} - t_{DMOSON} + (t_R + t_F) \times \left(\frac{1}{2} - \frac{V_T}{V_{DD}} \right).$$

The sum of differences in the propagation delays of IC_1 and Q_1 is positive and totals 32.1 nsec. V_T is the gate-to-source threshold voltage of Q_1. The data sheet gives a typical V_T of 1.1V, and the supply voltage, V_{DD}, has a value of 5V. These values yield 9.8 nsec for the last term of the preceding **equation**. T_{PON} is thus larger by 41.9 nsec. For a good design, an ammeter will indicates a supply current one to 1.5 times the ideal value of the current.

You can check the peak voltage at load resistor R. D_1 and D_2 and storage capacitor C_S function as a peak detector. The peak-voltage pulses at resistor R cause a dc voltage at C_S of roughly the same value as the peak voltage. You can determine the peak current flowing through the inductor from the voltage at the peak detector using the following **equation**:

$$I_{LPEAK} \simeq \frac{V_{RPEAK}}{R}.$$

Set the auxiliary supply voltage at 5.078V, the supply voltage at 10V, and the clock-pulse repetition frequency to 11,387 Hz. This approach causes the supply's current to be 0.327A and the peak voltage to be 16.4V. The peak current of the inductor reaches 30.94A. The experimentally determined turn-on time is approximately 1.502 μsec.

The IC driver contributes to the protection of the MOSFETs with an undervoltage lockout. If the supply voltage is on, but the auxiliary supply voltage is off, a voltage could get from the IN_A pin through internal protective diodes to the V_{DD} pin. The undervoltage lockout disables the control outputs until the auxiliary supply's voltage reaches a typical value of at least 4.2V.

The 0.65-mΩ dc resistance of ferrite-core inductor L might seem to be overrated for the circuit. However, the slope of current pulses in the inductor represents a megahertz-range equivalent frequency. The effective resistance at these slopes increases due to the skin effect and the proximity effect. This effective resistance can be many times the dc value. EDN

Control an LM317T with a PWM signal

Aruna Prabath Rubasinghe, University of Moratuwa, Moratuwa, Sri Lanka

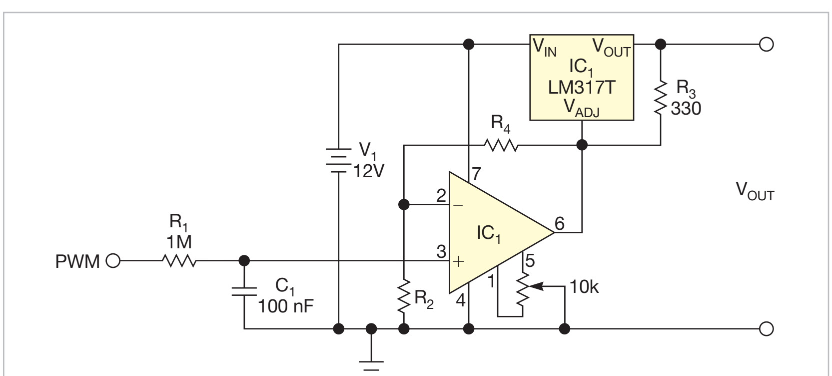

Figure 1 This circuit replaces a potentiometer with an analog voltage that you can control from a PWM signal.

The LM317T from National Semiconductor (www.national.com) is a popular adjustable-voltage regulator that provides output voltages of 1.25 to 37V with maximum 1.5A current. You can adjust the output voltage with a potentiometer. The circuit in **Figure 1** replaces the potentiometer with an analog voltage that you can control from a PWM (pulse-width-modulation) signal. You control this signal with a microcontroller or any other digital circuit. You can use the same microcontroller to dynamically monitor the output and adjust the LM317T.

Using an RC lowpass filter and an op amp, you can convert the PWM signal to a dc level that can adjust the

LM317T's voltage output. Varying the pulse width of the input signal lets you generate an analog voltage of 0 to 5V at the output of the lowpass filter. The op amp multiplies the voltage to achieve the desired voltage range.

For scenarios in which you must multiply the input voltage by two, the LM317T's adjustment pin receives 0 to 10V. Its output-voltage range is 1.25 to 11.25V. The **equation** $V_{OUT}=V_{ADJ}+1.25V$ governs the LM3175T's output voltage. You can change the op amp's gain by choosing proper values for R_4 and R_2. You must be able to remove offset voltages from the op amp. Use an op amp, such as a National Semiconductor LM741, with null adjustment. The selection of values for the capacitor and resistor for the RC lowpass filter depends on the PWM signal's frequency. This circuit uses values for a 1-kHz PWM signal.

YOU CAN IMPROVE THE CIRCUIT BY REPLACING THE RC LOWPASS FILTER WITH AN ACTIVE FILTER.

You can improve the circuit by replacing the RC lowpass filter with an active filter and then feeding a feedback signal from the circuit's output into the microcontroller for dynamic adjustments. EDN

High-speed buffer comprises discrete transistors

Lyle Russell Williams, St Charles, MO

Circuits sometimes need a gain-of-one buffer to lower output impedance and prevent the load from interfering with the previous stage. For an application involving a 1.5-MHz, low-power transmitter and antenna, a Burr Brown (www.ti.com) BUF634 buffer IC would work, but a discrete transistor buffer may be more convenient and less expensive than the IC.

Figure 1 shows the classic design of such a buffer. This circuit can drive a load as low as 200Ω with a peak output voltage of 2V. The maximum collector current of the transistors limits the output. You can use larger output transistors if your application requires more output current. Trimmer resistor R_4 across the diodes is, however, a relatively expensive part, and you must adjust it to produce the correct bias current for Class

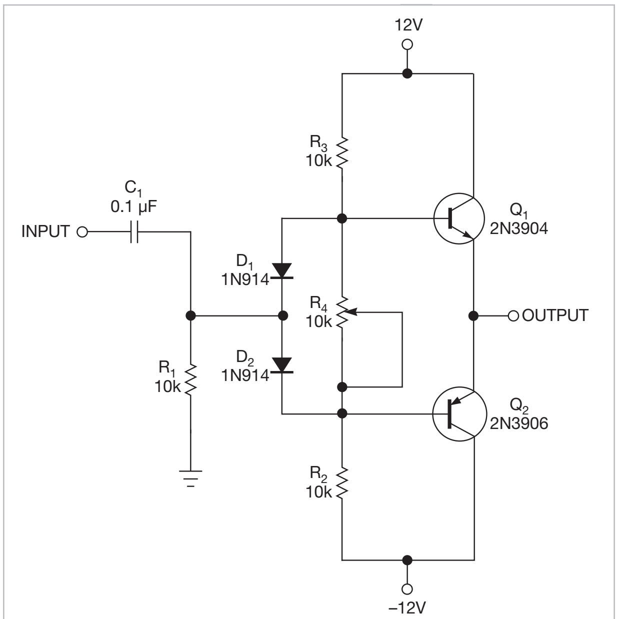

Figure 1 A typical buffer can drive loads as low as 200Ω.

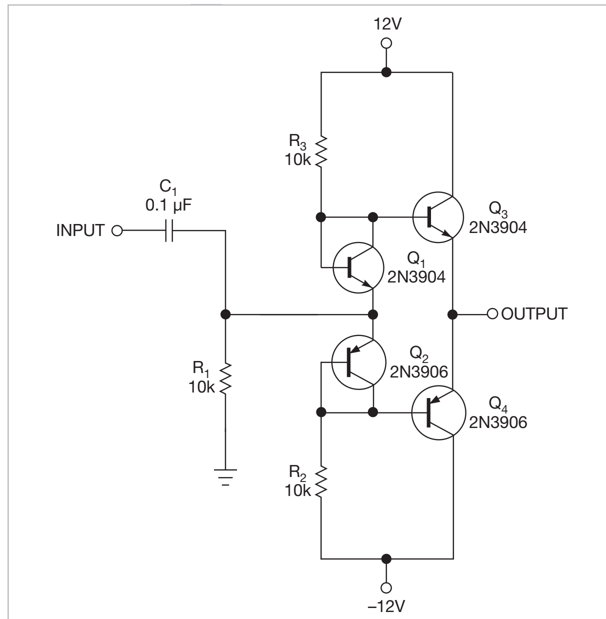

Figure 2 Current-mirror transistors replace the diodes in Figure 1.

AB operation. The adjustment is likely to drift over time.

A simpler circuit, such as the one in **Figure 2**, uses current-mirror transistors Q_1 and Q_2 instead of diodes. Resistors R_2 and R_3 set the zero-signal bias current in the bias-transistor circuits. The current-mirror effect causes the current in the output transistors to be nearly equal to the current in the bias transistors—approximately 1.2 mA in this case.

Because the current-gain-bandwidth product of the 2N3904 and 2N3906 transistors is 300 MHz, this circuit should work at 100 MHz or higher frequencies. At these frequencies, however, the circuit layout may be critical, and the slew rate, which is unknown, may limit usefulness. The offset of the circuit is approximately 0.1V, which is not a problem for this application because the circuit uses capacitive coupling through C_1. If you use the buffer in the feedback loop of an op amp, the op amp can null the offset.

You may want to monitor the current in the output transistors, so the circuit in **Figure 3** adds 100Ω resistors R_3 and R_4 and 10-μF bypass capacitors C_2 and C_3 in the collectors of output transistors Q_3 and Q_4. The voltage across these resistors reveals the collector currents, which are nearly equal in the two output transistors, and is close to the value that the values of R_1 and R_2 predicted. EDN

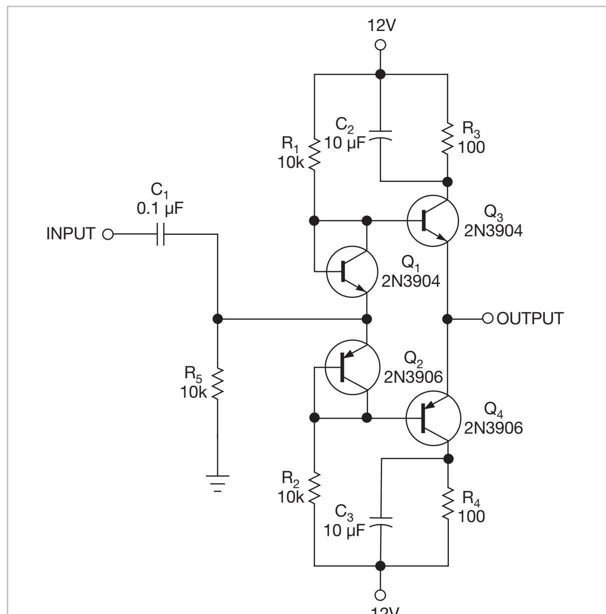

Figure 3 Resistors R_3 and R_4 limit the output current.

Limit inrush current in high-power applications

JB Castro-Miguens, Cesinel, Madrid, Spain, and
C Castro-Miguens, University of Vigo, Vigo, Spain

A high-power offline supply is nothing more than a half- or full-bridge dc/dc converter. Rectifying the ac line yields a dc voltage that feeds the converter. At power-supply turn-on, the bulk capacitor of the uncontrolled rectifier is completely discharged. It results in a huge charging current for a high instantaneous line voltage because the discharged bulk capacitor temporarily short-circuits the diodes of the rectifier stage. The high inrush current can trigger a mains circuit breaker, burn a fuse, or even destroy a power supply's rectifier diodes unless you take precautions. The circuit in **Figure 1** limits the inrush current.

At turn-on, if the instantaneous rectified ac-line voltage, V_{ACR}, is greater than approximately 10V, Point A in **Figure 2**, MOSFET Q_2 turns on, forcing thyristor Q_1 off. In this situation, a little current flows through R_1 and Q_2, injecting a small charge into bulk capacitor C_O, Path A to B in **Figure 2**.

When $V_{ACR}-V_O \leq 8V$ or so, where V_O is the output voltage, Q_2 is off, letting Q_1 conduct. In this situation, the bulk capacitor receives the necessary charge through Q_1, Path B to C in **Figure 2**, to match V_O to V_{ACR}. After this point, V_{ACR} falls below V_O, and the bulk capacitor alone must support any power the dc/dc converter demands until $V_{ACR}-V_O \geq 5V$ or so, Path C to D in **Figure 2**. At Point D, $V_{ACR}-V_O \approx 5V$ and thyristor Q_1 triggers, which conducts the capacitor's charge current and the current the dc/dc converter demands until V_{ACR} matches the sinusoidal peak at Point E.

When V_{ACR} falls, thyristor Q_1 cuts off, and the bulk capacitor alone feeds the dc/dc converter. The thyristor conducts again when V_{ACR} matches V_O to the sinusoidal peak. This process then repeats. Use a nonsensitive gate thyristor with a breakdown voltage of at least 400V for an ac voltage of 220V rms (root mean square) and with twice the rms-current rating of the rectifier diodes.

This circuit uses a TYN610 thyristor. You can calculate the value of R_1 using $R_1=(6.8-V_{GT})/I_{GT-20°}$, where V_{GT} is the minimum gate-cathode voltage necessary to produce the gate-trigger current for Q_1 and I_{GT} is the minimum gate current to trigger Q_1 down to –20°C. The NTD4815NHG MOSFET is suitable for this circuit. A MOSFET with a different threshold voltage may require different values for R_2 and R_3.EDN

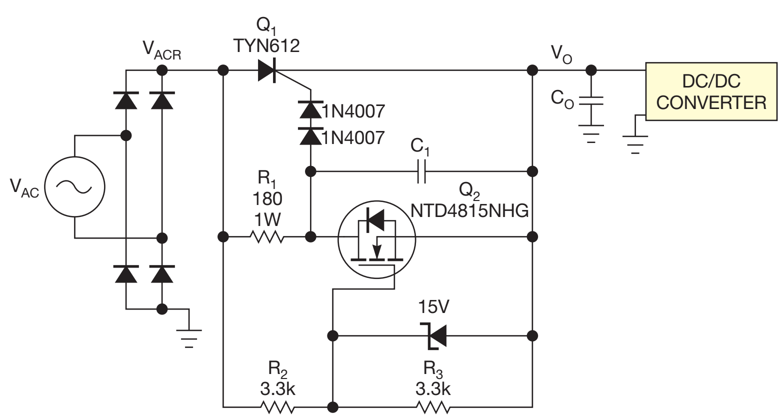

NOTE: 15 nF IS AN OPTIONAL VALUE FOR C_1.

Figure 1 A thyristor and a MOSFET control current to bulk capacitor C_O. This circuit limits the inrush current.

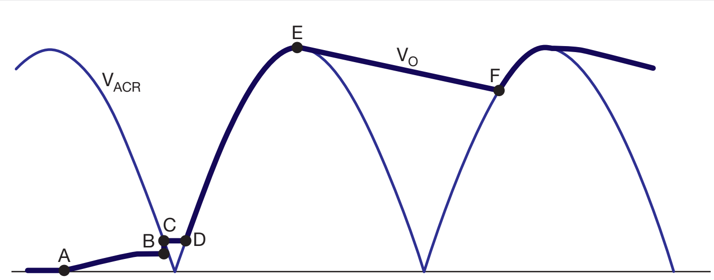

Figure 2 if V_{ACR} is greater than approximately 10V, MOSFET Q_2 turns on; current flows through R_1 and Q_2, injecting a small charge into bulk capacitor C_O.

A new twist on design chain: EBV collaborates to launch chips

Design-chain strategy usually involves meeting design with suppliers and encouraging product development with support from in-house engineers or educational materials. Distributor EBV Elektronik (www.ebv.com), an Avnet company, is taking that strategy one step further by launching chips it developed in collaboration with its line-card suppliers.

In doing so, EBV stresses that it will not become a chip supplier or producer. Instead, the company will accent component design and distribution as parts of its business strategy, aiming to create opportunities by matching customer needs for focused designs with appropriate suppliers for access to specially customized products. EBV will combine requests from customers in a segment and then collect inputs to design a rough specification. Using that input, it will look to suppliers that provide the desired technology for production.

"We collect inputs, and we will make design proposals, but one of our suppliers does the production," says Bernd Schlemmer (**photo**, left), director of communications at the company. The suppliers' logos are on the parts, and the suppliers provide the warranty, he adds.

The effort, which EBV announced in January 2009 and elaborated on at 2010's Electronica show, aims to provide access to specially customized products to small and midsized customers, beyond larger key-account customers.

"The major advantage for our customers is access to the latest technologies and access to our suppliers," says Klaus Schlund (**photo**, right), technical director at EBV. "You can imagine when there is a midsize environment in Italy or France; usually, the big players in the semiconductor market are not very open to go into negotiations for an ASIC or for any kind of IC. EBV can bundle those designs and talk to the right guys at our suppliers."

> "The basic principle behind this approach is a tailored solution for certain applications.

Because it will not produce the semiconductors, EBV says that its suppliers have no concerns about competition. These collaborations will generate chips that will not compete with any products, Schlemmer explains. He notes that an EBV chip is something that is not available yet and that comes as additional business to suppliers that want to provide access to such specially customized products.

"Some of our customers are happy to share their IP [intellectual property] and knowledge with other customers," Schlund says. "The basic principle behind this approach is a tailored, dedicated solution for certain applications."

EBV chips will target EBV's vertical automotive, lighting, medical, RFID (radio-frequency-identification), consumer, and renewable-energy segments. The company has also partnered with Texas Instruments (www.ti.com) on industrial-interface design and is working with STMicroelectronics (www.st.com) and Vishay (www.vishay.com). Products, which EBV and Avnet will distribute, should become available in 2011.

"We are the first distributor to start a program like this," Schlemmer says.

OUTLOOK

CHINA'S RFID MARKET TO DOUBLE BY 2014

China's RFID (radio-frequency-identification) market has grown significantly, with revenue set to more than double from 2009 to 2014, according to iSuppli (www.isuppli.com), now part of IHS Inc. The market comprises tags, readers, and software/middleware. In early January, iSuppli predicted a 22% increase—from $1.1 billion in 2009 to $1.4 billion in 2010. By 2014, iSuppli expects the Chinese RFID market to reach $2.4 billion.

"The rise of China's RFID market is the result of rising demand from applications in transportation, warehouse logistics, electronic payment, medical-equipment tracking, food-security systems, asset management, and more," says Vincent Gu, senior analyst for China research at iSuppli. "At the same time, technology integration in the RFID market has been advancing. Many enterprises have begun to study technical innovations to broaden the applications for RFID devices, including mobile payment."

The company notes that, although costs are declining, they remain a barrier to RFID's acceptance. As the technology matures, iSuppli predicts, so will manufacturing of the devices, which will allow the devices to reach economies of scale.

productroundup

PASSIVES

1.2- and 1.5-mm inductors come in 1212 cases

The IHLP-1212AB-11 and IHLP-1212AE-11 inductors have 1.2- and 1.5-mm profiles, respectively, and both have 3×3.6-mm footprints. They target use in voltage-regulator-module and dc/dc-converter applications, including next-generation mobile devices, notebooks, desktop computers, and graphics cards. The IHLP-1212AB-11 offers an inductance of 0.22 to 0.56 μH, a saturation current of 6.7 to 9.3A, and a typical dc resistance of 9.5 to 18.7 mΩ. The IHLP-1212AE-11 offers an inductance of 0.22 to 1 μH, a saturation current of 5.3 to 9A, a typical dc resistance of 9.5 to 29.5 mΩ, and a maximum dc resistance of 11.4 to 33 mΩ. The devices sell for 35 cents (10,000).

Vishay Intertechnology Inc, www.vishay.com

Precision potentiometer offers slip-clutch option

The 3547, 3548, and 3549 multiturn potentiometers now include a slip-clutch option, a safety feature that helps prevent damage to the potentiometer in applications using full travel. The option thus increases reliability and reduces potential service costs and system downtime. The feature targets use in mechanized or machine-to-machine applications, such as linear actuators, cable transducers, and mechanical position sensors in which the mechanical actuator may occasionally rotate the shaft beyond the mechanical stops or beyond a set number of turns. The internal design option does not alter the physical dimensions of the component and allows the device to idle at the extreme clockwise or extreme counterclockwise ends without inducing damage to the stops on the potentiometer. The three-turn model 3547 and five-turn model 3548 potentiometers with the slip-clutch option sell for $16.99 to $18.31 (1000). The 10-turn model 3549 with the option sells for $14.44 to $15.49 (1000).

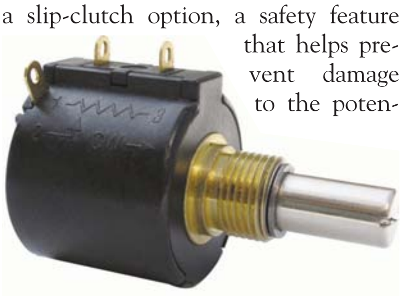

Bourns Inc, www.bourns.com

Hybrid capacitor increases energy density over competitors

This new hybrid capacitor combines an ultracapacitor and a lithium-ion battery, providing a 115%-higher energy density than that of standard electric double-layer capacitors. The devices provide more than 20,000 charge/discharge cycles. The hybrid capacitor finds use in flashlights, LEDs, memory backup, portable hand tools, solar chargers, off-grid lighting, and automotive power windows and door locks. The capacitor operates over a range of −25 to +60°C.

Ioxus Inc, www.ioxus.com

Low-profile power inductors target use in mobile devices

The BRL2515 and BRFL2518 wire-wound power inductors measure 2.5×1.5 and 2.5×1.8 mm, respectively, with maximum heights of 1.2 and 1 mm, respectively, to offer dc

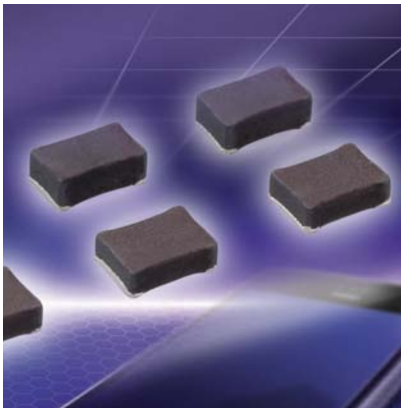

bias characteristics in a thinner, smaller size. The BRL2515 offers a saturation current of 1000 mA and inductance of 2.2 μH and has a dc resistance of 0.07 to 0.265Ω. The BRFL2518 offers a dc resistance of 0.9 to 0.33Ω. Both ROHS-compliant inductors are available for sampling at 20 cents each.
Taiyo Yuden, www.t-yuden.com

EDN ADVERTISER INDEX

EDN provides this index as an additional service. The publisher assumes no liability for errors or omissions.

Tower of babble

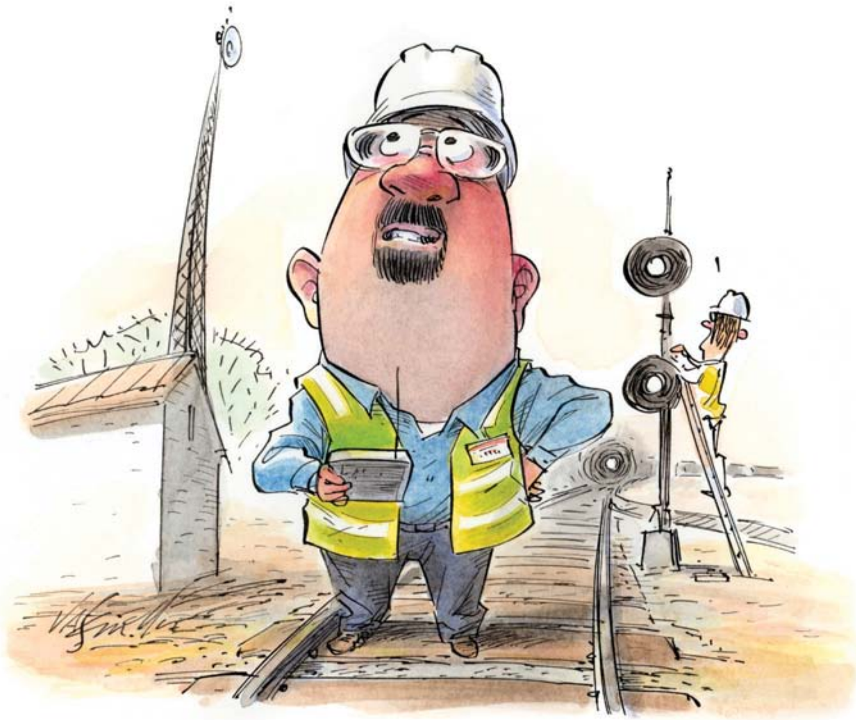

While working for a railroad, I once had to troubleshoot a remote microwave site that had failed while communicating with a track-side switch controller in southern Missouri. To throw a switch on one of the tracks, a dispatcher in Fort Worth, TX, would send a signal over the microwave link to a VHF radio that connected to the microwave link. While investigating this situation, I discovered that the VHF radio was receiving a tone that should not have been there, and this tone was of a sufficient level that it masked the data pulses that the dispatcher was sending. On my spectrum analyzer, I saw and heard a steady tone imposed on the signal that I was looking at. Someone else must have been transmitting on our assigned VHF channel.

This location was a heavy-traffic area for the railroad, and the pressure was on to correct the situation as soon as possible. This investigation was costing about $50,000 an hour. Expenses included train delays and paying personnel overtime rates to manually throw this critical switch. How could I find the source of this interfering signal? The antenna on the tower of the microwave hut was 350 feet high.

The receiving equipment at the track-side location had an antenna that was just 8 feet above the tracks, so I took the spectrum analyzer to the track-side location, where I couldn't see the interfering signal. This step confirmed my theory that the interference was not local and that the 350-foot-tall antenna was picking it up.

Other staff members were checking the Federal Communications Commission license base to see which transmitters were near this location. I drove around in the darkness for several hours until I found a radio facility with an attached VHF antenna. We managed to identify the owner of this facility, and the owner's maintenance personnel met me there a few hours later. We had to pay a premium for calling out the maintenance workers after normal working hours. Nevertheless, when they checked their equipment, they found that it wasn't the source of the interference.

The chief engineer at my company decided to send a "tower man" to my location with a new antenna to replace the one on the tower. As I waited—for five hours—for his arrival, I mulled over my options because I knew a replacement antenna would not fix our problem.

When the tower man arrived to climb this tower, I told him to wait because I had a little test that I would like to make first. He agreed. After all, climbing a 350-foot tower is no fun. I had him place the new antenna close to the ground, supported by the security fence that surrounded the microwave hut. We then ran coaxial cable into the hut from this ground-level antenna and connected it to the VHF radio in the microwave site. Bingo! The radios once again began to communicate, and normal operations resumed. The microwave site was on a hill overlooking the track-side antenna, resulting in a good path between the two antennas.

Replacing the antenna on the tower would not have solved our problem. The signal remained on the original antenna as I called every licensee I could find in a large radius of the microwave site. No one ever acknowledged having any trouble with their radio systems, but the interfering signal vanished three days later. My best guess was that a transmitter somewhere had a defective antenna, resulting in a high standing-wave ratio, which in turn caused the radiation of spurious emissions. One of those emissions fell into our frequency, and our antenna on its tall tower picked up this spur. Although you would normally place an antenna as high as possible, lower was better in this case. EDN

Earl Schlenk is a retired engineer for Burlington Northern Railroad. He resides in Saint Louis, MO.

DANIEL VASCONCELLOS

ANSYS® 13.0

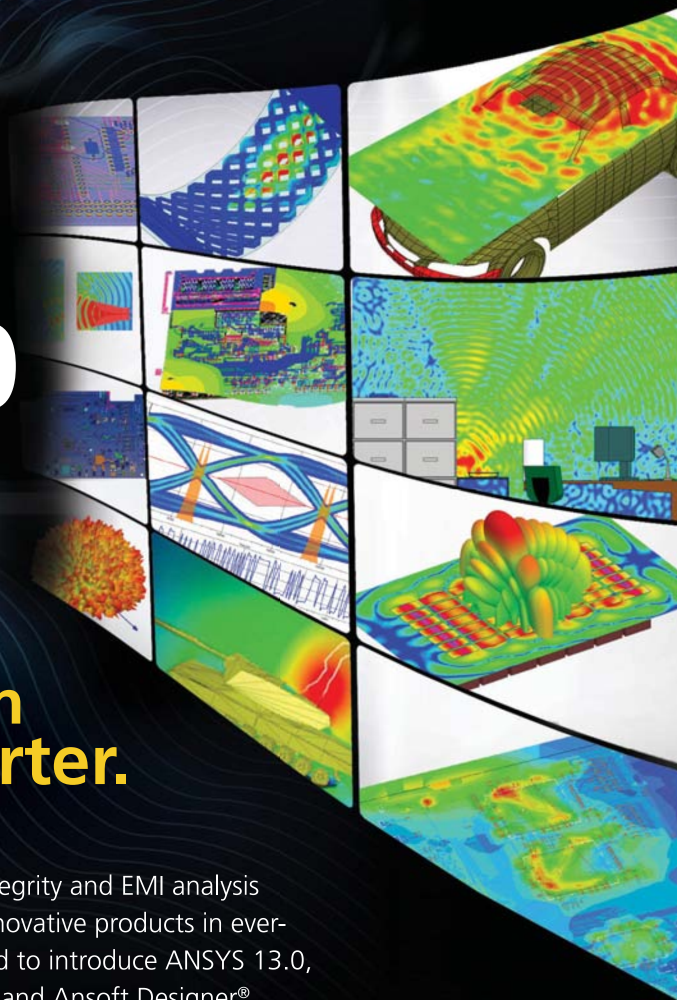

Smart engineering simulation is now even smarter.

In today's world of smart electronics, signal integrity, power integrity and EMI analysis are mission critical. Engineers must deliver both reliable and innovative products in ever-shrinking development cycles. For this reason, ANSYS is pleased to introduce ANSYS 13.0, our latest release which includes tools such as HFSS™, SIwave™ and Ansoft Designer®, allowing engineers to identify and correct potential design or performance problems before they reach the marketplace.

ANSYS 13.0 can help your design team achieve an even shorter total time to solution with several new and advanced features, including:

- Greater accuracy and solution fidelity for very large electromagnetic structures through the use of new solvers
- Higher productivity through automated 3D electromagnetic model creation, resulting in more time to design and innovate
- Increased computational efficiency, through the use of HPC, resulting in a dramatic speedup of wide-frequency band or parametric simulations

For speed, reliability and accuracy, more product development leaders worldwide trust ANSYS for simulation solutions to help create smarter products in less time.

ANSYS 13.0 **Advanced Fidelity.** **Higher Productivity.** **Leading Performance.**

ANSYS®

For more information, visit www.ansys.com/electronics13
or call us at 1.866.267.9724

Ansoft is now part of the ANSYS platform of simulation solutions.