electronic design

THE AUTHORITY ON
EMERGING TECHNOLOGIES
FOR DESIGN SOLUTIONS

August 2016 electronicdesign.com

TOP 50 EMPLOYERS IN ELECTRONIC DESIGN

BROADCOM LIMITED • CADENCE DESIGN SYSTEMS IN... APPLE INC.
LOGIC INC. • ANALOG DEV... TRONIC INC. • FORD MO...
COMCAST CORP. • ... INC. • ... AYTHEON
HONEYWEL... ARCH CORP
INSTRUMENTS ... CKHEED MAR
THE BOEING CO... INC. • LEAR C
TEXTRON INC. • V... • HARMAN INTE
INDUSTRIES INC. ... • ROCKWELL C
LINEAR TECHNOL... E TECHNOLOGY
PUBLIC LIMITED CO... RAMBUS INC. •
CYPRESS SEMICON... GENERAL DYNAMIC
HARRIS CORP. • E... MICROSOFT CO
TECHNOLOGY INC. • ... ON SCIENTIFIC CORP
AUTOMATION INC. • THE ... TIFIC INC. • QUALCOM
GENERAL MOTORS CO. • 3M ... • NATIONAL INSTRUMENTS
MICRON TECHNOLOGY INC. • G... AL ELECTRIC CO. • JUNIPER N
WESTERN DIGITAL CORP. • XILINX INC. • NORTHROP GRUMMAN

$10.00 Powered by Penton

electronic design

In This Issue

FEATURES

14

24

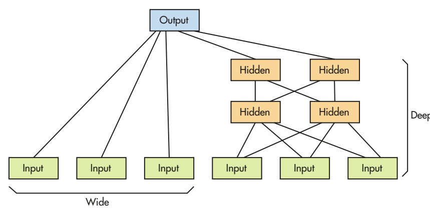

28

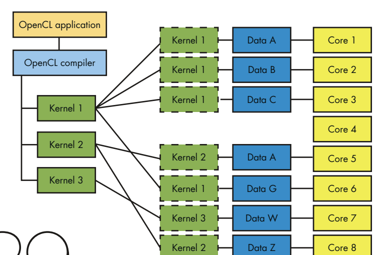

29

NEWS & ANALYSIS

IDEAS FOR DESIGN

COLUMNS & DEPARTMENTS

EDITORIAL MISSION:
To provide the most current, accurate, and in-depth technical coverage of the key emerging technologies that engineers need to design tomorrow's products today.

ELECTRONIC DESIGN (ISSN 0013-4872) is published monthly by Penton Media Inc., 9800 Metcalf Ave., Overland Park, KS 66212-2216. Paid rates for a one-year subscription are as follows: $120 U.S., $180 Canada, $240 International. Periodicals postage paid at Kansas City, MO, and additional mailing offices. Editorial and advertising addresses: ELECTRONIC DESIGN, 1166 Avenue of the Americas, New York, NY 10036. Telephone (212) 204-4200. Printed in U.S.A. Title registered in U.S. Patent Office. For subscriber services or to order single copies, write to Electronic Design, PO Box 2100, Skokie, IL 60076. POSTMASTER: Send change of address to Electronic Design, PO Box 2100, Skokie, IL 60076. Canadian Post Publications Mail agreement No. 40612608. Canada return address: IMEX Global Solutions, P.O. Box 25542, London, ON N6C 6B2.

on electronicdesign.com

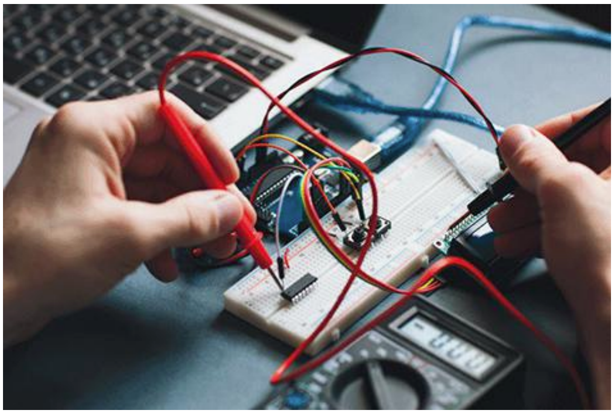

HOW TO CHOOSE THE BEST MICROCONTROLLER

http://electronicdesign.com/blog/how-choose-best-microcontroller

Contributing Tech Editor Lou Frenzel is in the midst of updating his book *Electronics Explained*, and has turned to the readers of *Electronic Design* for advice choosing the best microcontroller as an example. Read the blog online and add your thoughts to the lively collection of comments.

NI'S 2ND-GEN VST SIMPLIFIES SCALABLE RF TEST

http://electronicdesign.com/blog/ni-s-2nd-gen-vst-simplifies-scalable-rf-test

The second-generation vector-signal transceiver from National Instruments, the PXIe-5840, features a bandwidth of 1 GHz (4x the 200-MHz of the previous generation), occupies only two PXI slots instead of three, and swaps out the Virtex-6 for a much more powerful Virtex-7 FPGA.

join us online

Follow us on Facebook: http://tinyurl.com/odzo6hc
and Twitter: http://tinyurl.com/k5pum39

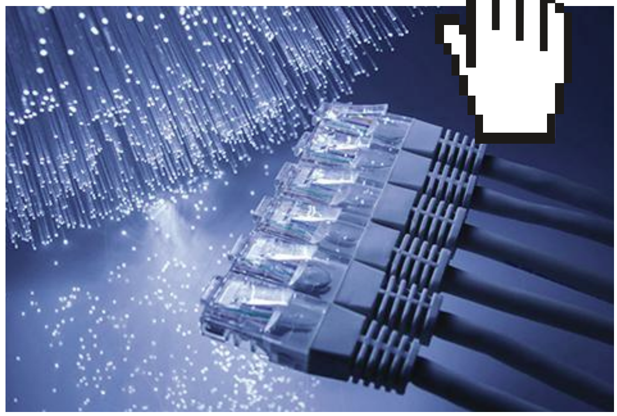

PAM4 SINGLE LAMBDA ADDRESSES TODAY'S ETHERNET CHALLENGES

http://electronicdesign.com/communications/qa-pam4-single-lambda-addresses-today-s-ethernet-challenges

Electronic Design talks to AppliedMicro's Omar Hassen about the world's first 100G PAM4 single-wavelength solution for 100G and 400G Ethernet.

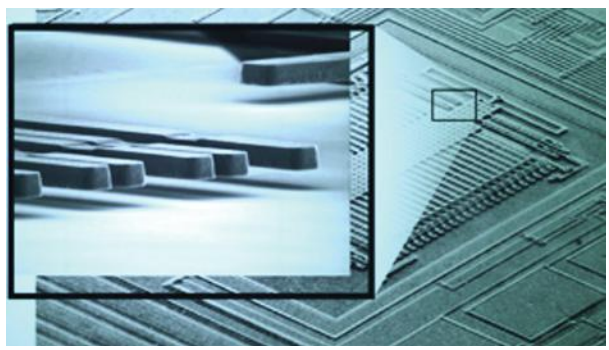

BENEFITS OF MEMS ACCELEROMETERS FOR CONDITION MONITORING

http://electronicdesign.com/analog/bring-benefits-mems-accelerometers-condition-monitoring

A growing number of condition-monitoring products now use a microelectromechanical-system (MEMS) accelerometer as the core sensor. These economical, highly integrated solutions help reduce the cost of deployment and ownership. As a result, they make it possible for more facilities and equipment to benefit from a condition-monitoring program.


GETTING
YOU THERE
IN MORE
WAYS
THAN YOU
KNOW
MPD DEVICES CAN BE FOUND IN NUMEROUS
TRANSPORTATION VEHICLES INCLUDING MOTORCYCLES AND SCOOTERS.
MPD
MEMORY PROTECTION DEVICES
MPD IS A GLOBAL MANUFACTURER OF BATTERY HOLDERS AND OTHER ELECTRONIC COMPONENTS. WE BELIEVE THAT OUR COMPONENTS SHOULD FIT EASILY INTO YOUR DESIGNS, WHICH IS WHY WE ARE ALWAYS CREATING INNOVATIVE NEW
BATTERYHOLDERS.COM

electronic design
Powered by Penton

AUGUST 2016

EDITORIAL

CONTENT DIRECTOR: **NANCY K. FRIEDRICH** nancy.friedrich@penton.com
CONTENT PRODUCTION DIRECTOR: **MICHAEL BROWNE** michael.browne@penton.com
CONTENT PRODUCTION SPECIALIST: **ROGER ENGELKE** roger.engelke@penton.com
PRODUCTION EDITOR: **JEREMY COHEN** jeremy.cohen@penton.com
DISTRIBUTION: **VICTORIA FRAZA KICKHAM** SourceESBeditor@penton.com
EMBEDDED/SYSTEMS/SOFTWARE: **WILLIAM WONG** bill.wong@penton.com
ANALOG/POWER: **MARIA GUERRA** maria.guerra@penton.com
CONTENT OPTIMIZATION SPECIALIST: **WES SHOCKLEY** wes.shockley@penton.com
ASSOCIATE CONTENT PRODUCER: **LEAH SCULLY** leah.scully@penton.com
ASSOCIATE CONTENT PRODUCER: **JAMES MORRA** james.morra@penton.com
CONTRIBUTING EDITOR: **LOUIS E. FRENZEL** lou.frenzel@penton.com

ART DEPARTMENT

GROUP DESIGN DIRECTOR: **ANTHONY VITOLO** tony.vitolo@penton.com
SENIOR ARTIST: **JIM MILLER** jim.miller@penton.com
CONTRIBUTING ART DIRECTOR: **RANDALL L. RUBENKING** randall.rubenking@penton.com
CONTENT DESIGN SPECIALIST: **JOCELYN HARTZOG** jocelyn.hartzog@penton.com
CONTENT & DESIGN PRODUCTION MANAGER: **JULIE JANTZER-WARD** julie.jantzer-ward@penton.com

PRODUCTION

GROUP PRODUCTION MANAGER: **CAREY SWEETEN** carey.sweeten@penton.com
PRODUCTION MANAGER: **FRAN VAUGHN** fran.vaughn@penton.com

AUDIENCE MARKETING

USER MARKETING DIRECTOR: **BRENDA ROODE** brenda.roode@penton.com
USER MARKETING MANAGER: **DEBBIE BRADY** debbie.brady@penton.com
FREE SUBSCRIPTION/STATUS OF SUBSCRIPTION/ADDRESS CHANGE/MISSING BACK ISSUES
T | 866.505.7173 **F** | 847.763.9673 electronicdesign@halldata.com

SALES & MARKETING

MANAGING DIRECTOR: **TRACY SMITH** **T** | 913.967.1324 **F** | 913.514.6881 tracy.smith@penton.com

REGIONAL SALES REPRESENTATIVES
AZ, NM, TX: **GREGORY MONTGOMERY** **T** | 480.254.5540 gregory.montgomery@penton.com
AK, CA, CO, ID, MT, ND, NV, OR, SD, UT, WA, WI, WY, W/CANADA: **JAMIE ALLEN** **T** | 415.608.1959 **F** | 913.514.3667 jamie.allen@penton.com
AL, AR, IA, IL, IN, KS, KY, LA, MI, MN, MO, MS, NE, OH, OK, TN: **PAUL MILNAMOW** **T** | 312.840.8462 paul.milnamow@penton.com
CT, DE, FL, GA, MA, MD, ME, NC, NH, NJ, NY, RI, PA, SC, VA, VT, WV, EASTERN CANADA:
SHANNON ALO-MENDOSA **T** | 978.501.7303 shannon.alo-mendosa@penton.com

INTERNATIONAL SALES
GERMANY, AUSTRIA, SWITZERLAND: **CHRISTIAN HOELSCHER** **T** | 011.49.89.95002778 christian.hoelscher@husonmedia.com
BELGIUM, NETHERLANDS, LUXEMBURG UNITED KINGDOM, SCANDINAVIA, FRANCE, SPAIN, PORTUGAL:
JAMES RHOADES-BROWN **T** | +011 44 1932 564999 **M** | +011 44 1932 564998 james.rhoadesbrown@husonmedia.com
PAN-ASIA: **HELEN LAI** **T** | 886 2-2727 7799 helen@twoway-com.com
PLEASE SEND INSERTION ORDERS TO: orders@penton.com
PENTON REPRINTS: **WRIGHT'S MEDIA** **T** | 877.652.5295 penton@wrightsmedia.com
CIRCULATION: **CUSTOMER SERVICE** **T** | 866.505.7173 **F** | 847.763.9673 electronicdesign@halldata.com
LIST RENTALS:
SMARTREACH CLIENT SERVICES MANAGER: **DAVID SICKLES** **T** | (212) 204 4379 david.sickles@penton.com

ONLINE

PRODUCT DEVELOPMENT DIRECTOR **RYAN MALEC** ryan.malec@penton.com

DESIGN ENGINEERING & SOURCING GROUP

EXECUTIVE DIRECTOR OF CONTENT AND USER ENGAGEMENT: **NANCY K. FRIEDRICH**
GROUP DIRECTOR OF OPERATIONS: **CHRISTINA CAVANO**
GROUP DIRECTOR OF MARKETING: **JANE COOPER**

PENTON

CHIEF EXECUTIVE OFFICER: **DAVID KIESELSTEIN** david.kieselstein@penton.com
CHIEF FINANCIAL OFFICER: **NICOLA ALLAIS** nicola.allais@penton.com
INDUSTRY GROUP, PRESIDENT: **PAUL MILLER** paul.miller@penton.com
1166 AVENUE OF THE AMERICAS, 10TH FLOOR
NEW YORK, NY 10036 **T** | 212.204.4200

Penton

Electronic Design | Machine Design | Microwaves & RF | Source ESB | Hydraulics & Pneumatics | Global Purchasing | Distribution Resource | Power Electronics | Defense Electronics

BILL WONG | Embedded/Systems/Software Editor
bill.wong@penton.com

Is It Time to Learn About Deep Learning?

Deep learning is the latest artificial intelligence craze, but it's one that can deliver significant advantages for many applications.

Artificial intelligence (AI) is one of those hyped topics that often comes and goes with little understanding of what's involved or how it can be monetized. It's not unique in that respect; robotics and a host of other topics come to mind. They typically peak as research turned into products that ride the wave and then disappear, or at least it seems that way.

In actuality, though, they rarely become products. Also, at this juncture, the practical applications hide from the spotlight, but continue to grow in acceptance. This is true even for AI, where techniques drive everything from financial transactions to Roombas.

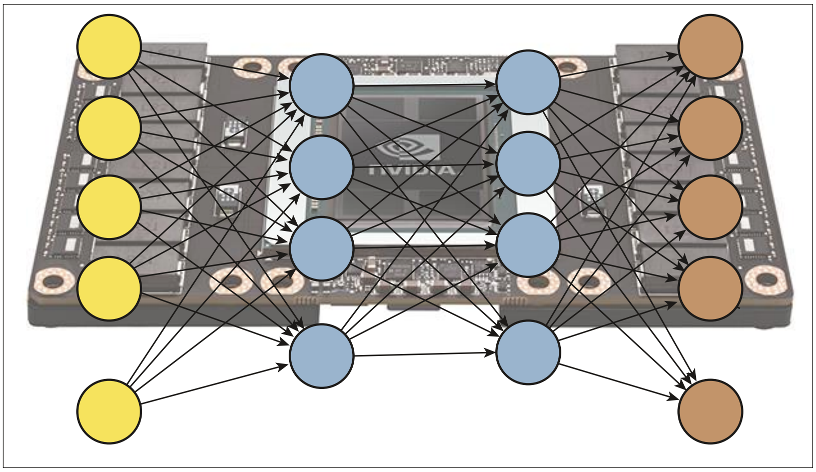

The latest idea to bring artificial intelligence to the forefront is "deep learning" technology, which is based on deep neural nets (DNNs). Deep learning has moved AI into the spotlight again because of improved performance enabled by platforms like GPUs.

The latest idea to bring AI to the forefront is "deep learning" technology, which is based on deep neural nets (DNNs). Deep learning has moved AI into the spotlight again because of improved performance enabled by platforms like GPUs (see *"GPU Targets Deep Learning Applications" on electronicdesign.com*). This has turned many applications of DNN into practical concepts. Keep in mind that some applications may use GPU clusters for training a DNN, but that deployment may only require a microcontroller.

So what is DNN, and how does it apply to your application area?

I would say yes, at least to get enough of an understanding to decide whether it's worth it to learn more. Take a look at some of the applications other than facial recognition or the more-common DNN applications to see how they might impact your application area. Recursive neural networks take time into account, allowing them to be used in applications like speech recognition.

Most applications will not benefit from DNNs or other AI techniques, such as rule-based or behavior-based systems. However, you won't know until you at least understand the basics. From there, it might warrant the inclusion of a library or two, or something more extensive. The advantages can often be significant, providing an edge over the competition that may not know about this secret sauce. ed

News

Infineon Buys Cree's Wolfspeed to EXPAND POWER AND WIRELESS CHIPS

In a surprise restructuring of the market for advanced semiconductors, Infineon Technologies announced that it struck a deal to buy Wolfspeed, a division of Cree making power and radio frequency chips.

Infineon is paying $850 million in cash for Wolfspeed, which fabricates chips for power management and wireless infrastructure. The deal fits the German chipmaker's mantra of energy efficiency, connectivity, and mobility. In recent years, Infineon has placed bets on the growing markets for renewable energy, autonomous cars, and the Internet of Things.

The deal bolsters Infineon's power management products, adding to a wide array of sensors, radars, and security chips. While the company supplies chips for things like mobile phones, its largest buyer is the automotive industry, which needs new products for managing car batteries and protecting against hackers.

Seeking growth in the market for power semiconductors, Reinhardt Ploss, Infineon's chief executive, has steered the chipmaker toward silicon chip alternatives. With the deal, Infineon gains access to new compound semiconductors,

An Infineon Technologies engineer holds a 200 millimeter wafer in a semiconductor cleanroom in Dresden, Germany. *(Image courtesy of Infineon)* which are smaller and more energy-efficient than silicon devices. After years of confinement to military electronics, these chips are finding their way into devices like light bulbs and power inverters.

"Joining forces with Wolfspeed represents a unique growth opportunity," Ploss said in a statement. "This will enable us to create additional value for our customers with the broadest and deepest portfolio of innovative technologies and products in compound semiconductors."

Wolfspeed is one of the biggest makers of wide bandgap semiconductors, which are capable of handling higher voltages and switching frequencies than silicon chips. Their products are smaller and thinner than conventional chips, cutting energy losses. Its major product is silicon-carbide for power management devices. It has developed a process for layering gallium-nitride material on SiC substrates, forming integrated circuits for wireless applications.

FOCUS ON SiC AND GaN

With the deal, Infineon will take control of Wolfspeed's SiC and GaN devices, in addition to the related chip manufacturing. Those assets include around 550 Wolfspeed employees and 2000 patents and patent applications. Cree will keep manufacturing substrates used in its main lighting business.

The business that Infineon is buying generated revenues of $173 million over the fiscal year ending in March. In 2014, the business that would become Wolfspeed recorded revenues of $120 million. Infineon plans to close the deal by the end of 2016.

GaN-on-SiC substrates are proving extremely useful in satellite, telecommunications, and radar systems. Wolfspeed's manufacturing process can handle high-electron-mobility transistors (HEMTs) on massive microwave integrated circuits, which are used in RF power amplifiers. GaN-on-SiC are extremely efficient at frequency bands up to 80 gigahertz, which could be tapped for 5G wireless networks.

The transaction was announced only months after Cree rebranded the business as Wolfspeed and began working out the details of an initial public offering. The spinout was an attempt for Cree to focus more closely on its main business of LED lighting, including chips and bulbs. Chuck Swoboda, CEO of Cree, said that selling Wolfspeed was also meant to refocus its business operations.

After Cree revealed that Wolfspeed would hold an initial public offering, several companies made offers to buy the business directly, Swoboda said. "After much consideration and due diligence over the past year, we concluded that selling Wolfspeed to Infineon was the best decision for our shareholders, employees and customers," he said. ■

SUNPOWER MAKES QUICK WORK
of Efficiency Record

SunPower's prototype solar panel converts 24.1% of the sunlight that hits it into electricity, making it vastly more efficient than conventional solar cells, which are usually between 15% and 18% efficient. *(Image courtesy of coniferconifer, Flickr)*

SUNPOWER, THE SECOND largest solar company in the United States, said that it has struck a new benchmark for solar panel efficiency, beating the record it set only four months ago using similar technology.

The prototype solar panel converts 24.1% of the sunlight that hits it into electricity, a significant step above conventional panels, most of which are between 15% and 18% efficient. According to the National Renewable Energy Laboratory, which tested the module and keeps the record charts, it is the most efficient ever made from monocrystalline silicon—the most common material used to make solar cells.

The new module is composed of individual solar cells that are slightly more efficient than the entire panel. Most solar research has focused on the cell level, but attempts to make them into panels are fraught with high costs and design challenges. Scientists have made individual solar cells almost twice as efficient as SunPower's module, but most are gathering dust on laboratory shelves.

Boosting efficiency has always been on the solar industry's checklist, but recent efforts have taken a backseat to expanding operations and building out infrastructure. SunPower, though, takes pains to highlight efficiency in its marketing. The company says that solar panels with 15% efficiency generate only half the electricity of its X-series panels, which are 21.5% efficient.

"With greater efficiency, we can fit more watts on the roof," Peter Cousins, SunPower's senior vice president of research and development, said in a statement. With higher efficiency, solar panels can produce more electricity in a smaller area, driving down production costs.

Other companies are taking an aggressive stance on efficiency, too. SolarCity, the largest solar provider in the United States, has built modules with 22% efficiency and revealed plans to start making them at its flagship facility in Buffalo, N.Y., which is still under construction.

Typical solar cells consist of a thin layer of silicon treated with chemicals to produce electricity when sunlight shines on it. Narrow wires embedded on the front of the cell, usually made of silver, collect the energy. The circuit is completed with a back plane of aluminum.

SunPower's modules are based on special Maxeon solar cells, which contain slight tweaks to the traditional design. The basic cell is the same, but the silver wiring is replaced with an electrode material, which has poorer electrical characteristics but covers the entire front of the panel. Without the wiring, more sunlight filters into the cell, while mirrors underneath reflect diffused light into the electrode, further improving efficiency.

Maxeon has fueled several SunPower records. In February, solar panels from the company's X series reached 22.8% efficiency in tests confirmed by the National Renewable Energy Laboratory. Those panels, which SunPower said are in production, held the record until the latest prototype.

In the laboratory, more complex designs on the cell level can be vastly more efficient. Earlier this year, NREL scientists built a multi-junction cell, which combines two layers of silicon crystals, that converts 29.8% of sunlight into electricity. In 2014, before it ended research on the design, France's Soitec tested a multi-junction device with 46% efficiency, using mirrors to concentrate the sunlight into intense beams.

But efficiency is not always everything, and certain solar cells have struggled to translate into full-blown panels. Their complex structures and high manufacturing costs keep multi-junction cells off the average person's home. Crystalline silicon is far more affordable, though SunPower uses more expensive materials and construction methods to make its cells.

SunPower is a slightly unusual solar company in other ways. It has expanded more slowly than its competitors and in previous years generated profits in a struggling industry. In the first quarter of 2016, however, its losses swelled to $85.4 million, up from $9.6 million a year earlier—at least in part linked to huge investments in California and Nevada power plants.

Rivals including SolarCity have been pouring money into new infrastructure, building losses in the pursuit of fast growth. Tesla's recent buyout offer for SolarCity is widely speculated to be an attempted bailout by Elon Musk, the primary shareholder in both companies. In recent months, other solar companies, including SunEdison and Spain's Abengoa, have declared bankruptcy.

SunPower stressed that its latest module will eventually leave the laboratory. Cousins declined to reveal when the new panels might be available, but he said that in the meantime the company would continue testing. In a blog post, he said that "new cell designs will keep pushing laboratory results higher, even above 26%." ■

SELF-DRIVING BUS ARRIVES in Maryland, with IBM Watson as Concierge

LOCAL MOTORS IS a strange animal, even for an automotive industry in transition. The automaker has become known for using 3-D printers to make certain parts and relying on a community of developers to sketch out vehicle features. Now it is trying to stick another odd-colored feather in its cap: a fully autonomous bus called Olli.

The company said that it will begin testing the electric bus on public roads in National Harbor, Md., where it recently opened a new manufacturing and research facility. Local Motors plans to give free public rides in the vehicle toward the end of the summer. The bus seats 12.

Olli was intended for driving around cities, providing a service that falls somewhere in between public transportation and ride-hailing apps from the likes of Uber Technologies and Lyft. Once the bus is cleared to operate autonomously, passengers will be able to call it with an app.

"Olli offers a smart, safe, and sustainable transportation solution that is long overdue," John B. Rogers, chief executive of Local Motors, said in a statement. Supporters of autonomous driving have underlined its potential to alleviate gridlock traffic and make the roads safer not only for other drivers but also for pedestrians and bicyclists.

Olli will also incorporate IBM's Watson artificial intelligence platform. Watson, which is located in the cloud, will not actually navigate the vehicle, but will collect data about Olli's driving patterns and talk with passengers through a voice interface.

Only a single Olli model is currently located at the new facility. But Local Motors, which recently received an undisclosed amount of funding from Airbus Ventures, is building more vehicles at its Phoenix, Ariz., headquarters. ■

Olli and Local Motors chief executive John B. Rogers. *(Image courtesy of IBM)*

LOU SOSA | Contributing Editor

TOP 50 EMPLOYERS in Electronic Design

Though sprinkled with some high spots during 2015, the electronics industry as a whole mirrored 2014 in terms of relatively flat revenue and growth.

Almost 10 years out from the 2007-09 subprime mortgage financial crisis, the U.S. economy has yet to show a full recovery. For 2015, the economy expanded at a rate of 2.4%, below the average 1947-2016 annual rate of 3.23% according to the U.S. Bureau of Economic Analysis.

China's slowdown plus weakness in commodity-dependent emerging market economies caused commodity prices to fall. Export shipments weakened due to the strength of the dollar, and after six years of continuous annual gains, the U.S. stock market pretty much flattened out.

Oil prices, which began dropping in 2014, hit an 11-year low in December 2015. While this appeared to imply more disposable income for consumers, the end result was that consumers did not spend but rather increased their savings rates.

In addition, the Fed raised interest rates in December 2015 for the first time in a decade after delaying it from September, if only by a quarter of a point, to fight inflation. The actual effect may have been to slow housing and investment down by increasing the cost of borrowing and further reduce exports by further strengthening the dollar.

INDUSTRY GAINS IN KEY AREAS

Category	Fiscal 2015 vs. 2014	Fiscal 2014 vs. 2013
Employee growth	0.4%	0.4%
Sales growth	0.6%	3.5%
Pretax income growth	8.4%	–7.9%
Pretax margin improvement	1.0	–1.5 pts.
Long-term debt-to-equity ratio improvement	17.4	–12.3 pts.
Research & development expense	2.7%	5.1%
ED Reader Profile Survey number of respondents	395	467

It appears that fiscal 2015 brought more of the same since the financial crisis of 20-2009. Some ups, some downs, but no real traction for sustained high growth. We're at the point where continued productivity gains are becoming elusive. Has the spirit of Moore's law finally come to an end?

Any gains in corporate profitability are not being passed down to the average worker, with wage increases continuing to be minimal. Employment growth has been encouraging at times, such as the June 2016 growth of 287,000 jobs, but keep in mind that May 2016 saw a jump of only 11,000. And despite the fact that unemployment now sits at 4.9%, a number of the workers who rejoined the full-time work force had to accept substantially lower-paying opportunities and many others have only been able to find part-time work.

The Top 50 list is based on a formula using public financial data from a "pool" of 92 public companies, with bonus points awarded using the results of our annual Electronic Design Reader Profile Study *(see the full methodology at electronicdesign.com)*.

PERCENTAGE OF COMPANIES THAT SAW GROWTH IN KEY AREAS

Category	Fiscal 2015 vs. 2014	Fiscal 2014 vs. 2013
Sales growth	45%	70%
Pretax income	42%	57%
Employee growth	49%	48%
R&D	54%	67%

2015 SHOWS LITTLE TOP-LINE MOVEMENT FROM 2014

The "Industry Gains in Key Areas" table compares collective data between fiscal 2015 versus 2014 for our current group of 89 companies.

Both sales growth and employment stayed basically flat year over year, and while pretax income and pretax margin showed marked improvement, their gains basically made up for last year's declines.

Collective debt-to-equity ratio improvement was significant, too, but it also offset the erosion from last year. However, at 43.9% for fiscal 2015, debt repayment and refinancing should not be an issue as long as the low-interest-rate environment continues and long-term debt can continue to be used to fund growth. If interest rates go up, though, trouble looms.

While R&D investment grew at almost a 3% rate and higher than sales, we don't like to see the decrease versus last year's rate.

While Brexit has brought global interest rates down even lower, the "feel" remains the same: Companies still appear reluctant to invest and the average worker feels like wage increases are almost nonexistent.

THE REST OF 2016—NOTHING REALLY NEW HERE

According to the Conference Board, real GDP is expected to grow at 1.9% for the next two quarters and at 1.7% annually for 2016, versus 2.4% in 2015. Real consumer spending at 2.2% for the rest of 2016 implies an annual growth rate of 2.5%, below 2015's 3.1%. Housing starts should show a 1.7% increase versus 2015 levels, maintaining their positive trend. Net exports continue to decline another 4.3% versus 2015 due to the strong dollar.

Real capital spending will decline by 0.9% versus a 2.8% increase in 2015, largely due to the 6.2% first-quarter decrease. Slow productivity gains are reducing profitability,

Fiscal 2015 Rank	Company Name	Total Company Line Score
1	BROADCOM LIMITED	78
2	CADENCE DESIGN SYSTEMS INC.	77
3	APPLE INC.	75
4	CIRRUS LOGIC INC.	74
4	ANALOG DEVICES INC.	74
4	MEDTRONIC INC.	74
7	FORD MOTOR COMPANY	73
7	COMCAST CORP.	73
9	VERIZON COMMUNICATIONS INC.	72
9	RAYTHEON COMPANY	72
9	HONEYWELL INTERNATIONAL INC.	72
12	AT&T INC.	71
12	LAM RESEARCH CORP.	71
14	TEXAS INSTRUMENTS INC.	70
15	NORTHROP GRUMMAN CORP.	68
15	LOCKHEED MARTIN CORP.	68
15	THE BOEING COMPANY	68
15	SYNOPSYS INC.	68
19	TERADYNE INC.	67
19	LEAR CORP.	67
19	TEXTRON INC.	67
22	VISTEON CORP.	66
22	CISCO SYSTEMS INC.	66
22	HARMAN INTERNATIONAL INDUSTRIES INC.	66
22	ROCKWELL COLLINS INC.	66
22	LINEAR TECHNOLOGY CORP.	66
22	INTEL CORP.	66
28	SEAGATE TECHNOLOGY PUBLIC LIMITED CO.	65
28	APPLIED MATERIALS INC.	65
30	RAMBUS INC.	64
31	CYPRESS SEMICONDUCTOR CORP.	63
32	GENERAL DYNAMICS CORP.	62
32	HARRIS CORP.	62
32	EMC CORP.	62
35	ITT CORP.	61
35	MICROSOFT CORP.	61
37	MICROCHIP TECHNOLOGY INC.	60
37	WHIRLPOOL CORP.	60
39	BOSTON SCIENTIFIC CORP.	58
39	ROCKWELL AUTOMATION INC.	58
41	THERMO FISHER SCIENTIFIC INC.	57
42	XILINX INC.	56
42	QUALCOMM INC.	56
44	GENERAL MOTORS CO.	55
44	GENERAL ELECTRIC CO.	55
44	3M COMPANY	55
44	NATIONAL INSTRUMENTS CORP.	55
44	MICRON TECHNOLOGY INC.	55
49	JUNIPER NETWORKS INC.	54
49	WESTERN DIGITAL CORP.	54

which is impacting capital investment and wage increases. With sales flat, China's growth slowing down, Brazil in recession, the uncertainty of Brexit, and more expensive exports, investment has stalled other than for replacement purposes.

Kiplinger's economic forecast calls for slower job growth for the rest of 2016 despite the strong addition of 287,000 jobs in June, which represented a huge bounce back from just 11,000 new jobs added a month earlier. Monthly job growth is likely to range from 150,000 to 200,000—less than the average 229,000 jobs added per month in 2015.

According to the U.S. Census Bureau's Retail Trade report, retail sales will increase around 4%, down from 2015's 4.8% due to lower new car sales. The increase is mainly from online stores (+12.2% in May), as brick-and-mortar department stores continue to struggle (-5.8% in May). Consumers are optimistic but cautious.

None of the above suggests a favorable environment for substantially increased capital spending or consumer spending, which accounts for around 67% of total U.S. economic output, at least not enough to spur significant GDP growth. Thus, 2016 may very well be more of the same.

Following is a closer look at the top three companies in our 2016 report.

WIRED PRODUCTS PUSH BROADCOM TO NO. 1

Sitting atop our Top 50 is Broadcom Limited (AVGO), the successor to Avago Technologies. Following Avago's acquisition of Broadcom Inc. on Feb. 1, 2016, Singapore-based Broadcom Limited became the ultimate parent company of Avago and BRCM.

To achieve cost savings by divesting non-core segments, Broadcom and Cypress Semiconductor announced on April 28 the signing of a definitive agreement under which Cypress will acquire Broadcom's wireless Internet of Things (IoT) business and related assets. The all-cash transaction was valued at $550 million.

Broadcom Limited's products cover a wide range of semiconductors. They include chips for wireless and wired communications, as well as optoelectronics, radio-frequency and microwave components, power amplifiers, and application-

specific integrated circuits (ASICs, i.e., custom chips). The company's thousands of products are used in myriad applications, including mobile phones, data-networking and telecommunications equipment, consumer appliances, displays, printers, servers and storage networking gear, and factory automation.

The company has four reportable segments: wired infrastructure, wireless communications, enterprise storage, and industrial & other, which align with its principal target markets. Wired infrastructure is by far the largest segment (58% of revenues), followed by wireless communications (22%) and enterprise storage (15%). Positive market trends are driving long-term growth in each of these markets:

Wired infrastructure
- Cloud, social media, and video streaming
- Big data and data analytics
- Rapid evolution of the connected home

Wireless communications
- LTE Advanced/carrier aggregation
- Increasing RF bands per phone
- Mobile connectivity and high-speed Wi-Fi

Enterprise storage
- Data center—higher reach/bandwidth
- Massive growth in cloud storage
- Exponential digital universe growth

Industrial & other
- Increased factory automation
- Energy efficiency/energy conversion
- Emerging markets

Broadcom's net revenue was $3.541 billion, an increase of 100% from $1.77 billion in the previous quarter and an increase of 119% from $1.6 billion in the same quarter last year. Hock Tan, president and CEO of Broadcom, says, "We delivered solid second-quarter revenue, while exceeding EPS expectations for our first-quarter operating as a combined company. Our increased scale and diversity is already proving very resilient, with strong product cycles in our now largest segment, wired, offsetting weaker demand in our enterprise storage and wireless segments. We are expecting a robust third quarter, led by strong growth in wireless revenue, and continued strength in wired networking, and remain confident in our ability to leverage earnings growth as we work toward full integration and achievement of our operating model."

Wired Infrastructure experienced strong demand in standard switching and routing products, including very

> "Our increased scale and diversity is already proving very resilient, with strong product cycles in our now largest segment, wired, offsetting weaker demand in our enterprise storage and wireless segments." —Hock Tan, president and CEO, Broadcom

strong traction for the new Tomahawk switching and Jericho routing platforms, especially from cloud and service provider end customers. The company also saw strong demand for broadband products from set-top box refreshes driven by the startup adoption of 4K video. Service providers continue to invest in broadband access infrastructure, including ongoing fiber-to-the-home build-out in China, as well as DSL and cable modem build-out in Europe.

Growth is expected to resume in its custom ASIC business, driven by increasing shipments to wireless base stations and a product ramp-up into new data-center switches. AVGO's standard ASSP (application-specific standard parts) switching, routing, and physical-layer products are getting a boost from enterprise demand.

The wireless-communications segment experienced a noticeable slowdown in demand from large North American smartphone customers due to seasonal product lifecycle-related reduction in shipments. This was partially offset by an increasing shipment to a large Asian handset OEM.

The rest of the year should show strong growth in this segment as Apple ramps up for the iPhone 7. Broadcom is already producing RF filters for the phone, which it will start shipping in September with at least 20% content growth in the new model. According to *BARRON'S*, Taiwan's *Economic Daily* said on May 23 that Apple had asked its suppliers to produce 72 to 78 million new iPhones by the end of the year, the highest production target in about two years. The expectation had been for 65 million iPhone 7's to be produced this year, so this is good news for Apple suppliers such as Broadcom.

Continued innovations in mobile Wi-Fi and Bluetooth technologies are also a growth contributor. This is important as around 80% of mobile data now moves through Wi-Fi as opposed to LTE.

The company's annual average sales growth over the last five years through both acquisitions and organic growth has been very strong at 27%, and the average EPS growth has also been robust at 24%. Furthermore, average annual estimated EPS growth for the next five years continues to be strong at 18%.

The company achieved 60% gross margins for five consecutive quarters, an impressive accomplishment. In addition, the company's margins, growth rates, and return on capital have been above the industry median, its sector median, and the S&P 500 median.

IoT SPURS CADENCE'S RISE UP THE CHARTS

Our second-ranked company, Cadence Design Systems (CDNS), helps engineers pick up the development tempo to support the newly connected and application-rich world we live in. It's a market leader in electronic-design-automation (EDA) software and semiconductor intellectual property (IP), offering custom/analog tools that help engineers design the transistors, standard cells, and IP blocks within systems-on-a-chip (SoCs).

Cadence's digital tools automate the design and verification of gigascale, gigahertz SoCs at the latest semiconductor processing nodes. Its IC packaging and PCB tools permit the design of complete boards and subsystems. And a growing portfolio of design IP and verification IP is available for memories, interface protocols, analog/mixed-signal components, and specialized processors.

Customers have included Pegatron, Silicon Labs, and Texas Instruments. Based in San Jose, Calif., Cadence gets about 55% of its sales from customers outside the U.S.

The company is well-positioned to benefit from the IoT, whose technology involves complex circuits. Cadence provides the software that allows its customers to design and produce them.

The *2016 Ericsson Mobility Report* projects that IoT sensors and devices will exceed mobile phones as the largest category of connected devices in 2018, growing at a 23% compound annual growth rate (CAGR) from 2015 to 2021. Ericsson predicts a total of approximately 28 billion connected devices worldwide by 2021, with nearly 16 billion related to IoT. Around 400 million IoT devices with cellular subscriptions were active at the end of 2015, and cellular IoT is expected to have the highest growth among the different categories of connected devices, reaching 1.5 billion connections in 2021.

The value in Cadence's products is that they allow customers to reduce time-to-market for electronic systems and lowers design, development, and manufacturing costs. As an example, Cadence is collaborating with ARM on IoT and wearable devices targeting TSMC's (Taiwan Semiconductor) ultra-low-power technology platform. The collaboration hopes to accelerate development of IoT and wearable devices by optimizing the system integration of ARM IP and Cadence's integrated flow for mixed-signal design and verification, along with its leading low-power design and verification flow.

The partnership will deliver reference designs and physical design knowledge to integrate ARM Cortex processors, ARM CoreLink system IP, and ARM Artisan physical IP, in addition to RF/analog/mixed-signal IP and embedded flash in the Virtuoso-VDI Mixed-Signal Open Access integrated flow for TSMC's new 55ULP, 40ULP, and 28ULP process technologies. The ULP technology platform is an important development in addressing the IoT's low-power requirements.

Another example would be the Cadence PSpice Analog/Digital with system simulation and modeling technology that enables a unified design environment for mixed-signal design. A customer can design all three blocks of IoT devices (sensor, controller, and actuator) and simulate the complete system.

TOP 10 OEM EMPLOYERS

Company	Fiscal 2015 OEM rank	Fiscal 2015 overall rank	Category
CADENCE DESIGN SYSTEMS INC.	1	2	Test equipment
CIRRUS LOGIC INC.	2	4	Components & subassemblies
ANALOG DEVICES INC.	2	4	Components & subassemblies
SYNOPSYS INC.	4	14	Components & subassemblies
TEXAS INSTRUMENTS INC.	5	15	Test equipment
ROCKWELL COLLINS INC.	6	22	Avionics & space; government & military electronics
LINEAR TECHNOLOGY CORP.	6	22	Components & subassemblies
INTEL CORP.	6	22	Components & subassemblies
RAMBUS INC.	9	30	Components & subassemblies
CYPRESS SEMICONDUCTOR CORP.	10	31	Components & subassemblies

As for fiscal Q1 2016 results, total revenues were up 9% year over year in a challenging environment, while pretax income was up 36%. The company's portfolio of solutions across chip, package, board, systems and software, and IP, guided by their System Design Enablement strategy, position it to drive new business in verticals including automotive, aerospace, medical, and IoT applications.

Strong, broad-based demand for the new Palladium Z1 contributed to Cadence's best-ever hardware revenue quarter. Rapidly growing complexity and time-to-market requirements make emulation more critical than ever for customers designing chips and systems for mobile, cloud, automotive, and other verticals.

The company's innovative new Virtuoso platform strengthens and solidifies its position in custom, analog, and mixed-signal design. At its CDNLive Silicon Valley user conference in June, Cadence announced the next-generation Virtuoso platform, including the Virtuoso Analog Design Environment Suite and the Virtuoso Layout Suite. The new Virtuoso offers designers an average 10X improvement in performance and capacity across the platform. The platform includes new technologies to address requirements of automotive-safety, medical-device, and IoT applications.

In addition, the acquisition of Rocketick Technologies will significantly increase the performance of Cadence's incisive enterprise simulator using parallel computing on standard multicore servers.

In IP, the company had a key design win for 5G baseband digital signal processors (DSPs) with a leading mobile handset company. And Spreadtrum licensed the Tensilica HiFi Audio/Voice DSP because of its ultra-low power.

Digital and signoff solutions are also growing, especially with customers in the mobile, consumer, automotive, and IoT segments. In Q1, a leading mobile chip company adopted Cadence's digital and signoff flow for its most demanding 10-nm projects. The Innovus implementation system added more than 15 new customers in Q1, while the Genus RTL synthesis solution added more than 25 new customers. Taiwan Semi certified its digital and signoff tools for 7-nm design and 10-nm production, and Samsung Foundry certified its tools for its 14LPP process.

Cadence is well-positioned to outgrow its competition and has a very strong, free-cash-flow position.

APPLE STILL RANKS THIRD DESPITE TOUGH YEAR

Apple Inc. engages in the design, manufacture, and marketing of mobile communication, media devices, personal computers, and portable digital music players. The firm offers products and services under the iPhone, iPad, Mac, iPod, Apple Watch, and Apple TV brands; consumer and professional software applications under the iOS, OS, X, and watchOS brands; and operating systems under the iCloud and Apple Pay brands. It operates through the following segments: Americas, Europe, Greater China, Japan, and Rest of Asia Pacific. Apple, based in Cupertino, Calif., generates nearly two-thirds of sales outside the U.S.

Fiscal Q2 2016 results show that net sales dropped 13% year over year with pretax income decreasing by 24%. Without currency effects, net sales would have decreased 9%.

The most significant factor influencing Apple's financial results is the iPhone—unit sales of the iPhone were down 16%. For the quarter, iPhone sales represented 65% of net sales, marking the first year-over-year decline for iPhone sales.

While both iPad and Mac unit sales were down 19% and 12%, respectively, for the quarter, these product categories only make up 9% and 10% of net sales, respectively. Services (including revenue from Internet Services, AppleCare, Apple Pay, licensing, and other services), which make up around 12% of net sales, grew at 20%. Other products (including sales of Apple TV, Apple Watch, Beats products, iPod and

MOST IMPROVED COMPANIES, 2014-2015

Company	Rise in the ranks
CIRRUS LOGIC INC.	71
VISTEON CORP.	56
VERIZON COMMUNICATIONS INC.	52
NORTHROP GRUMMAN CORP.	50
FORD MOTOR COMPANY	49
TERADYNE INC.	46
AT&T INC.	44
RAYTHEON COMPANY	42
JUNIPER NETWORKS INC.	39
SEAGATE TECHNOLOGY PUBLIC LIMITED CO.	37

Apple-branded and third-party accessories) grew at 30% and make up only 4% of net sales.

Without a doubt, the fact that the iPhone 6 and 6 Plus offered larger screen sizes for the first time is a factor in weakening the growth rate of the iPhone 6s. However, according to the latest Consumer Intelligence Research Partners (CIRP) survey, the iPhone refresh cycle appears to be lengthening. In mid-2013, around 33% of iPhones were more than two years old. Currently, that figure is around 50%.

The end of the subsidized phone era by U.S. cellular providers has had an impact on iPhone renewals. In addition, Fortune reported at the end of May that, according to Nikkei Japan, Apple is switching to a three-year new iPhone lifecycle. That would mean only a slight upgrade for the upcoming iPhone 7, thus not making it a "must-have."

Rumors are also out that Apple would skip the 7s and would go directly to the iPhone 8 in 2017, with purported major form-factor changes such as OLED, no home button, wireless charging, and all glass design. Also, reports from Korea-based news outlets the *Korea Herald* and *Hankyung* state that a rumored $2.6 billion deal between Apple and Samsung will have Samsung supply Apple with 100 million OLED panels for the company's 2017 iPhone version.

As Android phones came into their own around 2013, Apple was able to maintain global market share. In an April report from market-research firm Kantar Worldpanel ComTech, for the three months ended February 2016 versus the three months ended February 2015, Apple has basically been able to maintain its global market share with the exception of a 3-pt. dip in China and a 2-pt. dip in Europe. While Android has increased its market share, for the most part it has come at the expense of Microsoft and other operating systems.

In another report, Kantar reported a 95% loyalty rate for U.S. iPhone owners, the highest ever measured for any smartphone brand. Based on that combination of loyalty and market share, Apple simply dominates industry profitability as the genius is in its marketing. The company's products are very expensive compared to the competition, with arguably inferior specs, but people still buy them.

A February 2016 Kantar report reveals that Apple accounts for 91% of smartphone profits on around 17% global market share, while Samsung accounts for 14% on around 24% global market share. The data is imperfect, as Chinese manufacturers tend not to disclose profitability, but virtually everyone else lost money, with Microsoft losing the most.

Both the entry-level iPhone SE and the 9.7-in. iPad Pro results were not reflected in fiscal Q2 results, but Apple reports both are selling well. Although the iPhone SE may reduce overall margins, it's introducing a new customer to the Apple ecosystem.

Apple surpassed an installed device base of over 1 billion in fiscal Q2 2016 between the iPhone, iPad, Mac, Apple Watch, and Apple TV communicating with the App Store and the iCloud. The iOS ecosystem enables the iPhone to protect market share and maintain premium iPhone pricing and strong margins. Once a consumer has spent years in the ecosystem and purchased hundreds of dollars of content, he or she isn't likely to make the switch to Android and have to rebuild a content library from scratch.

This is the main reason why Apple's consumer loyalty is so high. Samsung's Android phones have had better specs than the iPhone for a while now. However, as long as Apple continues to offer a positive user experience, the costs of switching make it too inconvenient and expensive to leave the ecosystem.

If they can also continue to maintain market share in developing markets to take advantage of both population and economic growth, the future still bodes well; if not later this year, then certainly in 2017. ed

MARIA GUERRA | Technical Editor

Bit-Error-Rate Testers Face Ethernet Speed Challenges

With the demands on BERTs growing, industry groups and standards are considering the use of multilevel signals.

The bit error ratio (BER) is defined as the probability that a bit is received in error when, between transmission and reception, one or more bits are changed (e.g., a binary 0 is transmitted and a binary 1 is received). This measured ratio is affected by many factors, including noise, interference, distortion, attenuation, and jitter. Depending on the telecommunication protocol, BER requirements might vary from $1x10^{-6}$ to $1x10^{-12}$.

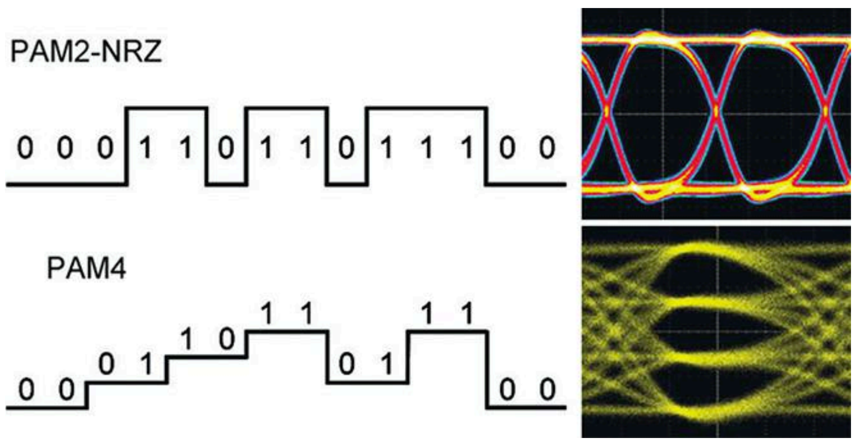

1. An NRZ eye diagram contains a single eye, while a PAM-4 eye diagram has three vertical eyes. (Courtesy of Tektronix)

With the move to higher data rates, bit-error-rate testing will adopt new multilevel-signal transmissions techniques for better performance, such as PAM-4. It will coexist with currently active techniques like non-return-to-zero (NRZ). Let's take a closer look at the principles behind BER testers (BERTs) and the challenges that the key players in the market are facing—along with the products they are offering to cover the need for data transmission using different forms of signal modulation.

BER TEST

BERTs are electronic systems used to detect the BER over any communication link (fiber-optic, Ethernet, radio communication, etc.). A BERT typically consists of a test pattern generator and a receiver to test that pattern. The pattern generator sends a bit stream (stimulus) to the device under test (DUT), which then responds back with another bit stream. The receiver compares the actual response from the DUT with the expected response, which is provided by the user.

A test pattern is usually chosen to emulate the type of data expected to occur during normal operation, but it is also chosen to verify a receiver's tolerance margins. The most common types of BERT patterns are:

- *Pseudorandom binary sequence (PRBS):* This binary sequence pseudo-randomly generates a stream of repeated or alternating bits. PRBS is difficult to predict. It can be generated using linear feedback shift registers, where the most common used linear function of single bits is an exclusive OR (XOR). The PAM-4 version of PRBS is referred to as a quaternary PRBS (QPRBS) pattern.

• *Quasi-random signal source (QRSS):* This pseudorandom binary sequence is based on a combination of a 20-bit word. It repeats every 1,048,575 bits and suppresses consecutive zeros to no more than 14.

The quality of the BER estimation increases as the number of transmitted bits approaches infinity. But in practice, how many bits do we need to transmit in order to achieve a practical BER performance close to an infinite BER performance? How long does that test take?

2. Keysight's M8040A 64-Gbaud high-performance BERT for PAM-4 and NRZ allows for simplified and accurate receiver characterization. (Courtesy of Keysight Technologies)

The concept of BER confidence level can help answer these questions because, in most cases, we only need to test that the BER is less than a predefined threshold. The confidence level is the percentage of tests that the system's true BER is less than the specified BER.

"The confidence level is an important factor when doing bit-error-ratio measurements," explains Ellen Spindler, product manager for BERTs at Keysight Technologies' Digital & Photonic Test Division. "Typically, digital communication links have to achieve a target BER of less than 10^{-12}. If you just compare 10^{12} bits with no error, the achieved confidence level is just 63.21%, which is mostly insufficient.

"If you want to achieve a higher confidence level of 95%, for example, you need to compare at least 2.996×10^{12} bits without a single error. The BER test time is directly related to the desired confidence level."

Mathematically, the confidence level can be expressed as:

CL = PROB [BERT < R] given E and N

where:

CL = BER confidence level

PROB [ ] = "Probability that"

BERT = True BER

R= Specified ratio

Confidence level is a probability whose values vary between 0 and 100%. Once the BER confidence level has been calculated, we may say that we have a CL percent confidence that the true BER is less than R. The quality of the estimate improves as the test time increases. This quality can be quantified using statistical confidence level methods.

TEST CHALLENGES

NRZ is a digital data transmission that has been used for several decades. With higher data rates of 400 Gb/s, however, it is better to use other digital data-transmission methods, such as pulse amplitude modulation (PAM-N). Such methods might perform better at such a high rate regardless of the challenges that come with it. IEEE 802.3bs has declared that 400-Gb/s links will use eight 56-Gb/s channels and will allow for the use of either 56-Gb/s NRZ or 28-Gb/s PAM-4.

An eye diagram is a common indicator of the quality of signals in high-speed digital transmissions. An NRZ eye diagram contains a single eye, while a PAM-4 eye diagram contains three vertical eyes *(Fig. 1)*. When using PAM-4, there is a tradeoff bandwidth for signal-to-noise-ratio (SNR); the smaller the vertical eye opening (differential amplitude), the more difficult it becomes to maintain a SNR that allows us to interpret the signal at the receiving end of the link.

Some of the challenges are:

• *Clock recovery:* It identifies the crossover point, or the place at which a signal crosses the threshold; however, multilevel signals have both symmetric and asymmetric zero crossings, resulting in a complicated task.

• *Noise:* There are three vertical eyes in a PAM-4 system, and each of them is independently affected by noise; therefore, there is only 33% of amplitude due to level spacing. Due to different rise and fall times between crossings, intersymbol interference (ISI) can be more significant.

• *Linearity:* PAM-4's three eye diagrams present a challenge, as each eye should contribute equally to the BER. In other words, the amplitude levels must be evenly spaced, with maximum opening of the three eyes.

• *Equalization:* PAM-4 systems can use equalization to open closed eye diagrams caused by ISI. Channel loss requires enhanced equalization techniques to identify symbols at the receiver or transmitter using techniques such as continuous-time linear equalization (CTLE) and decision feedback equalization (DFE) to correct.

3. Anritsu's MP1800A performs bit-error-rate tests from 0.1 Gbit/s to 32.1 Gbit/s; 64.2 Gbit/s with an external multiplexer/demultiplexer. (Courtesy of Anritsu)

Current options in the market include:

1. Keysight Technologies' M8040A high-performance BERT

This modular solution can be used to test both PAM-4 and NRZ devices that operate up to 64 Gbaud. The M8040A *(Fig. 2)* can be used for receiver (input) testing for many popular interconnect standards that implement PAM-4 and NRZ data formats, such as 400 GbE; 50/100/200 GbE; OIF CEI-56G and CEI-112G; 64G/112G Fibre Channel; Infiniband-HDR; and proprietary interfaces for chip-to-chip, chip-to-module, backplanes, repeaters, and active optical cables.

The M8040A is controlled via the M8070A system software, which can be used to control digital applications where non-NRZ data formats are required, such as PAM. The analyzer module provides true PAM-4 error analysis in real-time for long PRBS and QPRBS patterns. In addition, "the remote head concept allows short connections to the device under test, which is beneficial for data-center testing," says Spindler.

2. Anritsu's MP1800A Signal Quality Analyzers

The MP1800A *(Fig. 3)* is an expandable plug-in modular BERT with built-in pulse pattern generator (PPG). It supports the output of high-quality, high-amplitude signals.

In addition, an error detector (ED) with high input sensitivity supports signal analysis, such as bathtub jitter and eye-diagram measurements. The BERT also includes a jitter-modulation source for generating various jitters, such as SJ/RJ/BUJ/SSC, and supporting jitter-tolerance tests.

It can be configured to generate the high-quality, low-S/N, PAM-4 and PAM-8 data signals required for the characterization of high-speed backplanes and 400-GbE interfaces using the PAM-4/PAM-8 Converter (MZ1834A/MZ1838A). The bit error rates of three PAM-4 eye patterns can be measured simultaneously. The MP1800A performs BER measurements of PAM-4 signals using the long-memory programmable pattern function and error-mask function for filtering out unwanted errors. It offers data patterns for various applications—e.g., PAM-4 PRBS.

Additional modules, the MP1861A multiplexer and MP1862A demultiplexer, can be added to the MP1800A. When used in conjunction with the MP1800A, the two modules support a generation of NRZ data and BER measurements at data rates up to 64.2 Gb/s.

3. SHF's 12104 A and SHF 11104

The bit pattern generators (BPGs) and error analyzers (EAs) from SHF Communication Technologies AG (SHF 12104 A and SHF 11104 A, respectively) deliver a speed of more than 60 Gb/s per channel. The BPGs and EAs are available in two formats: plug-in and benchtop. The BPG generates digital bit sequences, such as standard PRBS or user-defined bit patterns. The 60/64-Gb/s outputs/inputs can operate at both full clock and half clock; for example, a 20- or 40-GHz signal is required for 40-Gb/s operation. The BPG is controlled by the intuitive graphical-user-interface BERT Control Center (BCC). The EA can be used with multilevel applications—e.g., 100GbE (using 2x 28 Gbaud PAM-4) and 400 GbE (using 8x 28 Gbaud PAM-4).

CONCLUSIONS

The bit error ratio is a vital figure of merit in the digital communication world. 400G Ethernet will play a key role across digital data applications, especially in applications where PAM-4 and NRZ will need to coexist. Hence, considerations for new designs and tests are essential to verify PAM-4 compliance and guarantee interoperability. ed

REFERENCES

1. "PAM4 signaling in high speed serial technology: Test, analysis, and debug" by Tektronix
2. "PAM-4 design challenges and the implication test" by Keysight Technologies
3. "Physical Layer Performance: Testing the Bit Error Ratio (BER)" by Maxim Integrated

WILLIAM WONG | Technical Editor

A Deeper Look at Deep-Learning Frameworks

In artificial intelligence, deep learning continues to gain ground, thanks to multicore hardware such as GPGPUs, with tools and frameworks also providing more accessibility to the technology.

Deep learning is one aspect of artificial intelligence (AI) that continues to advance, owing to performance improvements in multicore hardware such as general-purpose computation on graphics processing units (GPGPUs). Tools and frameworks have also made deep learning more accessible to developers, but as of yet, no dominant platform like C has emerged. The plethora of choices can be confusing, and not all platforms are created equal. It's also an area where cutting-edge development perpetually makes it more difficult to create new applications on top of a solid base.

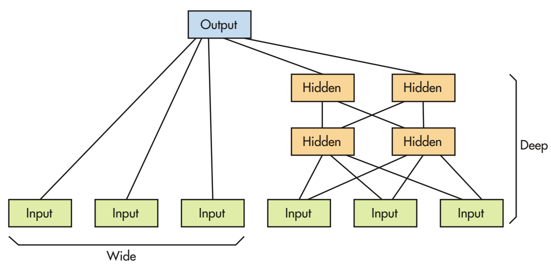

Some frameworks allow wide (left) and deep-learning (right) neural networks to be combined within a single implementation.

Deep learning is another name for deep neural networks (DNNs). This type of neural network has many layers, which affects computation requirements: As the size of a layer and number of layers increases, so do those requirements. In addition, wide neural networks, which are shallow in nature, can be useful for many applications. In fact, it's possible to mix them (*see figure*) using some frameworks. Neural networks are also in play with recurrent neural networks, convolutional neural networks, and logistic regression.

Some of the more popular, open-source, deep-learning frameworks include Caffe, CNTK, TensorFlow, Torch, and DeepLearning4J. Caffe, developed at the Berkley Vision and Learning Center (BVLC), probably has the greatest following and support. Microsoft's Computational Network Toolkit (CNTK) is an active open-source project. Torch and Theano are Python libraries that provide deep-learning support. MatConvNet is a toolbox designed for Mathworks' MATLAB.

Google started TensorFlow. The TensorFlow Playground is a website where you can experiment with predefined networks to see how changes affect the recognition process and its accuracy. DeepLearning4J, developed by Skymind, is a deep-learning framework written in Java that's designed to run on a Java Virtual Machine (JVM). The Skymind Intelligence Layer (SKIL) is based on Deeplearning4J, ND4J, and LibND4J (an n-dimensional array library).

DNN UNDERPINNINGS

Most of the DNN platforms often utilize new or existing computational frameworks to do the heavy lifting required by applications. Two well-known computational frameworks are OpenCL and Nvidia's cuDNN (CUDA DNN). OpenCL has the advantage of running on a range of hardware from multicore CPUs to GPGPU arrays. Nvidia's solution targets its own GPUs, including the latest Pascal architecture, Tesla P100 (*see "GPU Targets Deep Learning Applications" on electronicdesign.com*).

DNN applications often require significant amounts of training on large computational clusters to determine the weights associated with the nodes or neurons within a neural net. The plus side is that the resulting network can be implemented on much simpler hardware that may include microcontrollers.

The biggest challenge for developers is to become familiar with DNN. The frameworks typically have a number of preconfigured networks for sample applications, such as image recognition. The tools can be used for much more, but it often takes an expert to develop and tune new configurations.

Unfortunately, commercial support is only available for some frameworks. Select companies like Nvidia have an active support program with tools like the Deep Learning GPU Training System (DIGITS), which is designed to handle image classification and object detection tasks. It can help with a range of functions, such as training multiple networks in parallel.

DNNs and their associated tools are not applicable to all applications. However, they can make a significant difference in terms of capability and performance for many application areas, from cars to face recognition on smartphones.

11 MYTHS About OpenCL

OpenCL is a computational framework designed to take advantage of multicore platforms. Unsurprisingly, there are a few myths surrounding it.

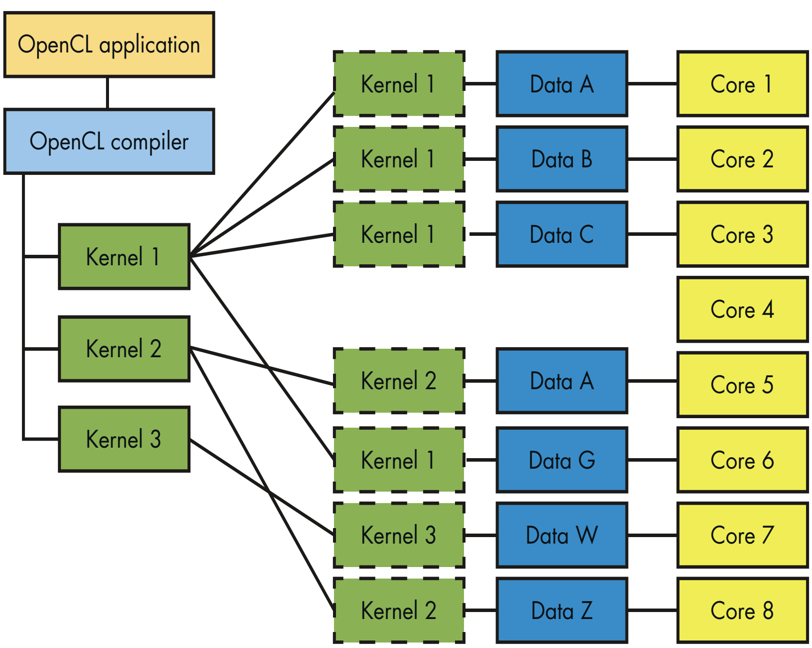

An OpenCL program is compiled into kernels or nuggets of code, which can be applied to data. The resulting data can be passed to the next kernel that requires it. A multicore system is able to run multiple kernels at the same time. They may be the same or different, depending on the application and its current state.

OpenCL is an open, cross-platform, parallel-computing standard from the Kronos Group. Applications are written in a variant of C, and the latest OpenCL 2.2 brings a static subset of C++14 *(see "C++14 Adds Embedded Features" on electronicdesign.com)* into the mix.

OpenCL has wide support on both the hardware and software side. It is often a checkbox item for hardware feature sets, and developers can utilize OpenCL-based software without knowing what is involved by either the software or the hardware that it is running on. Those already writing OpenCL already know how it works and will probably not find anything new in the following myths. Likewise, those simply using OpenCL-based software will not care other than to have their software run properly on OpenCL-supported hardware. For the rest of you, please read on.

1. I JUST NEED TO WRITE C CODE TO USE OpenCL.

True, sort of. OpenCL is a language specification, a runtime API, and a software framework that includes an OpenCL compiler and matching runtime. Application code is written in the form of small kernels that are then scheduled to execute by the runtime *(see figure)*. OpenCL kernels are written in C or C++ (as of OpenCL 2.2).

Part of the challenge in writing OpenCL code is to take advantage of the extensions and constructs for parallelizing computation, since this is now OpenCL speeds up execution of an application. This does not come free simply by using an OpenCL compiler. Even determining how functions will be implemented in a kernel will affect performance.

The runtime schedules kernels and their matching data to run on the hardware that normally consists of multiple cores, thereby potentially providing a speed-up of the application. A kernel runs until it terminates, and the resulting data can be used in future computations.

The code for a kernel may be shared among cores in a shared memory system, or it may be copied to a node as in a cluster system. Kernels are not portable across different hardware, but the source code is.

2. OpenCL ONLY RUNS ON GPUs.

Not true. OpenCL is a specification that's implemented by an OpenCL compiler. It can generate code for target hardware that can include CPUs, GPUs, or a mix of the two. The compiler dictates what targets it will support. An OpenCL application can run on a single-core CPU, but normally multiple-core systems are the target to get more overall performance from a system.

OpenCL can also run on FPGAs. This approach is a bit more static because the kernels are implemented in a FPGA, and its configuration doesn't normally change over time. Altera's SDK for OpenCL *(see "How To Put OpenCL Into An FPGA" on electronicdesign.com)* has an OpenCL compiler that generates an FPGA configuration; that includes supporting the configuration for what is essentially the OpenCL runtime. There's also software support to match that moves data between a host and the FPGA and to initiate kernel execution in the FPGA. Essentially, data is placed into the FPGA's memory where a kernel has access to it. The results can then be extracted or made available to another kernel in the same fashion, since OpenCL operates on a CPU or GPU. FPGAs have the advantage of performing many operations in parallel.

3. CUDA IS JUST NVIDIA'S VERSION OF OpenCL.

False. Nvidia's CUDA is similar to OpenCL, but they are distinct. Both use the kernel approach to partitioning code, and both support C and C++ for kernel code.

CUDA is an NVidia architecture that specifically targets Nvidia GPU hardware, but it can also generate code for some CPUs, such as x86 platforms providing similar functionality as OpenCL. In addition, the CUDA Toolkit includes additional libraries like cuDNN for deep learning, cuFFT, cuBLAS, and NPP (NVIDIA Performance Primitives for imaging). Nvidia device drivers for Nvidia hardware support CUDA, as well as OpenCL. Typically, developers targeting Nvidia GPUs will choose CUDA or OpenCL for their projects.

4. OpenCL WILL RUN BETTER ON A GPU THAN A CPU.

Yes and no. Typically, a GPU will have many more cores than even a multicore CPU. Having lots of cores can help in some applications, but any speed-up tends to be application-specific. Some applications will do better on a multicore CPU, while others will do better with GPUs. A lot depends on how various kernels are written and executed, as well as the type of operations being performed. Operations that take advantage of GPU functionality will usually run better on a GPU. Sometimes a mix of hardware may be the best alternative. It is possible to have an OpenCL system that spans both.

5. I HAVE TO PROGRAM IN OpenCL TO TAKE ADVANTAGE OF IT.

False. It's possible to use applications that were written and compiled for an OpenCL platform. In this case, the platform needs to have the OpenCL runtime installed (usually as a device driver), and the application needs to be compiled for that platform. This is comparable to conventional native-code applications designed for an operating system. The applications require the matching operating system and hardware to execute.

Another way to use OpenCL is to use an OpenCL-based library and call functions in the library from an application that runs on the CPU. The OpenCL APIs are defined for a number of programming languages, including Python, Java, C, and C++.

Writing an OpenCL program is only necessary if neither of the prior scenarios is suitable.

6. OpenCL CAN ONLY BE USED BY C APPLICATIONS.

False. Kernel code is compiled from C or C++ source, but an OpenCL application doesn't just consist of the kernel code. It's possible to use OpenCL-based libraries and applications in conjunction with application code running on the CPU side that's written in almost any programming language.

OpenCL specifies C and C++ APIs for interfacing with OpenCL applications. These APIs have been translated into other programming languages, including popular ones such as Java and Python.

7. OpenCL IS HARD TO LEARN AND PROGRAM.

True and false. OpenCL uses C and C++, making it easy to start with if one has a background in C or C++. The trick is that OpenCL has its own conventions, techniques, and debugging methodologies that differ from standard C and C++ development.

There's a large amount of information, documentation, and examples for OpenCL on the internet, as well as numerous books written on the subject. Still, there are quite a few things to learn in order to create efficient and bug-free OpenCL applications. Debuggers such as AMD's CodeXL have similar but different characteristics than conventional

CPU debuggers. OpenCL also uses different tracing and profiling tools.

8. OpenCL CANNOT BE USED FOR STREAM PROGRAMMING.

False. Stream programming is normally associated with a data stream such as an audio or video stream. Managing the data is a little more complex than a typical OpenCL application, but it is possible.

9. OpenCL ONLY RUNS ON AMD AND NVIDIA GPUs.

False. OpenCL will run on most GPGPUs, including GPUs from ARM, Imagination Technologies, Intel, and other vendors. It will not run on all GPUs, though, and it requires a matching runtime/driver and OpenCL compiler. The GPGPUs can be integrated with a CPU, or else they can be standalone GPUs that interface with a host (usually via PCI Express).

10. OpenCL REQUIRES LOTS OF HARDWARE TO RUN.

False, although more hardware tends to improve performance. OpenCL can run on single-core microcontrollers, but it makes more sense on devices with multiple cores or GPUs that support OpenCL. OpenCL systems can scale to very large clusters that contain multiple nodes with multiple CPUs or GPUs. Distributing code and data within a cluster can be a more complex chore, but this allows developers to concentrate on the code for the kernels.

11. OpenCL IS NOT GOOD FOR EMBEDDED APPLICATIONS.

False. OpenCL can be very useful in embedded applications. The embedded system will need support for OpenCL, but there are a number of microcontrollers and SoCs that have sufficient resources, along with OpenCL support, to make their use worthwhile.

Embedded designers actually have an advantage, since the scope of performance and requirements are known entities. They can often adjust the chips used to scale up or down to match those requirements. OpenCL can also help with power-sensitive applications, since it's often more efficient.

Still, OpenCL is not a free lunch. The necessary hardware costs money and uses power. It can allow implementation of features that would not be possible without utilizing OpenCL. Deciding when and how to employ OpenCL is how developers make a difference. ed

STEVE SHEARD | Strategic Program Manager
ON Semiconductor

CAPACITIVE TOUCH Broadens Switching Options for Equipment Apps

Touch controls can enhance the high-tech appearance of smart appliances, but some caveats should be kept in mind when considering the sensing technology, touchpad design, and control strategy.

Today's homes and workplaces are increasingly pervaded by smart appliances, which are designed to be attractive and reliable, and ideally should support interactions that are as easy and rewarding as those delivered by the smartphone sitting in the user's pocket.

As new smart appliances like televisions, set-top boxes, audio equipment, white goods, PC peripherals, security panels, and industrial controls enter the market, the design of the user interface and the quality of the user experience delivered have a critical influence on market appeal and, ultimately, sales and revenue. Creating a good impression through the user interface is vital, and designers can choose from a range of technologies, such as conventional pushbutton switches or various types of touch sensors.

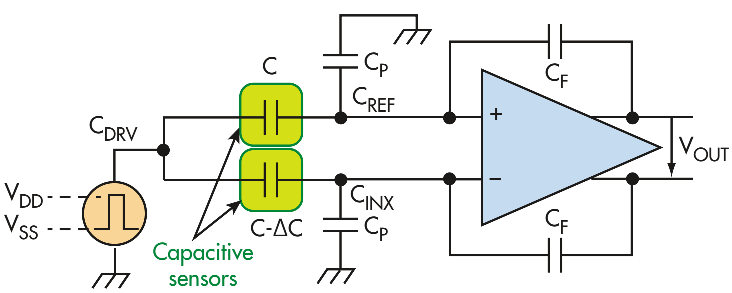

1. Differential sensing cancels the effects of parasitic capacitances (C_P) between pads and ground.

SWITCH OR TOUCH—PROS AND CONS

Mechanical switches like membrane or tactile switches are a mature technology that offers advantages including low cost, easy availability, and straightforward integration requiring minimal electronic design or software engineering. These types of switches also provide reassurance for the user by giving tactile feedback to confirm that the button-press has been detected. On the other hand, some disadvantages of mechanical switches include relatively slow response times, poor reliability owing to reliance on moving parts, and the possibility that additional noise-suppression or de-bounce circuitry may be required.

On the other hand, a touch sensor can create a more modern impression, and it simplifies construction and assembly. It does away with the panel cutouts typically needed for installing mechanical switches, and helps simplify the mechanical design of the control panel. The panel also can be sealed more easily against ingress of liquids, such as water. This is a valuable feature for many types of equipment, such as kitchen appliances like coffee machines, and simultaneously enhances reliability as well as facilitating cleaning.

Among touch-sensing technologies in use today, which include capacitive, resistive, or piezoelectric sensing, capacitive sensing is widely used. Capacitive touch sensors can be deposited easily on glass or plastic using a process such

DESIGN NOTES

Electrolytic Capacitor-Based Data Backup Power Solution for a 12V System Regulated from 5V to 36V Input

Design Note 553

Victor Khasiev

Introduction

Data loss is a concern in telecom, industrial and automotive applications where embedded systems depend on a consistent supply of power. Sudden power interruptions can corrupt data during read and write operations for hard drives and flash memory. Often, embedded systems need just 10ms to 50ms to backup volatile data to prevent loss.

Data backup is used in embedded systems for maintenance, troubleshooting and repair work. In complex industrial metal machining equipment, it's important to store the position and state of multiple tools after power disconnect to prevent equipment failure when power is later restored. These applications require a stable power supply and data retention, but unreliable power sources make it difficult to accomplish. Long supply lines, discharged batteries, unregulated AC adapters, load dumps and switching high power electrical motors result in widely fallible input supplies. As a result, developers of embedded systems prefer to design with the widest possible input voltage range, enabling use in a variety of applications and environments.

Circuit Description

Figure 1 shows a system that delivers reliable primary power plus holdup power for data backup. This solution is centered on the LTC®3643 bidirectional power backup supply. When the input voltage is present, the LTC3643 charges the storage capacitor, $C_{STORAGE}$, up to 40V in boost mode. When the input voltage is interrupted, the LTC3643 discharges the storage capacitor into the load in buck mode, keeping the nominal voltage at the load (V_{SYS}) in the range of 3V to 17V.

The relatively high voltage of the backup storage rail increases stored energy of this solution ($E = CV^2/2$) and enables the use of electrolytic capacitors as a backup storage component. Electrolytic capacitors are inexpensive and widely available, significantly reducing the cost of the backup solution. Another advantage of the LTC3643 is its ability to support 12V systems, the default standard voltage rail in many automotive and industrial applications.

In Figure 1, the LTM®4607 μModule® buck-boost converter acts as the front end regulator, producing 12V at up to 5A from a 5V to 36V input, such as a vehicle battery. The buck-boost regulator maintains a steady 12V output so long as the input voltage

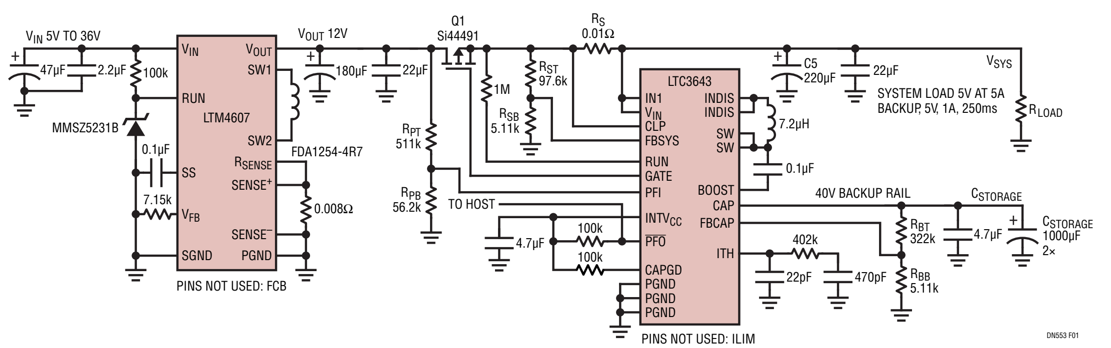

Figure 1. LTC3643 Backup Supply, Electrical Schematic

stays within the specified range, allowing V_{SYS} to ride through brownout and overvoltage conditions such as automotive cold crank and load dump. When the input voltage is interrupted or moves out of this range, the LTC3643 based backup power solution maintains the V_{SYS} system voltage to allow for short-term data backup.

Circuit Functionality

In normal operation, when the P-channel MOSFET Q1 is on, the flag PFO is low and the electrolytic capacitor array $C_{STORAGE}$ is charged to 40V. When the input voltage is interrupted, the LTC3643 turns Q1 off, sets the flag PFO high and starts to discharge the $C_{STORAGE}$ capacitor array, maintaining 12V to the load. When Q1 is in the off state, the body diode of this transistor effectively isolates the load from the input lines. The PFO flag identifies the fault and signals the host computer to disconnect the noncritical loads and supply circuitry. Here it is assumed that the critical circuitry related to data retention consumes 1A for up to 100ms.

Figure 2 illustrates the entire switchover process. At the start, the system load is supplied by the LTM4607, as the input voltage is present. When the input voltage is interrupted, the LTC3643 supports the system load by discharging the storage capacitor. Figure 3 shows the timing of the switchover in more detail. The load voltage falls to 10V, a value set by the resistor divider R_{PT}/R_{PB} and then recovers to the nominal 12V, set by the resistor divider R_{ST}/R_{SB}.

The formulas for an estimation of the required storage capacitance and holdup time are below. If a more detailed analysis is needed, the necessary information can be found in vendor's documentation.

1. Energy Stored

$$E_{CAP} = \frac{C_{STORAGE}}{2} \bullet (V_{CAP}^2 - V_{SYS}^2)$$

2. Energy Needed to Supply Load for Time T_H

$$E_{LOAD} = I_{SYS} \bullet V_{SYS} \bullet T_H$$

3. Holdup Time

$$T_H = \frac{C_{STORAGE} \bullet \left(V_{CAP}^2 - V_{SYS}^2\right) \bullet \eta}{2 \bullet I_{SYS} \bullet V_{SYS}}$$

η = efficiency

4. Storage Capacitance

$$C_{STORAGE} = \frac{2 \bullet V_{SYS} \bullet I_{SYS} \bullet T_H}{V_{CAP}^2 - V_{SYS}^2}$$

Conclusion

The LTC3643 is a highly integrated, high performance backup regulator. The design shown in this Design Note combines the advantages of this IC with a high efficiency buck-boost LTM4607 μModule regulator. Together, these devices enable a small footprint, efficient and cost effective solution for data retention and backup in automotive and industrial applications.

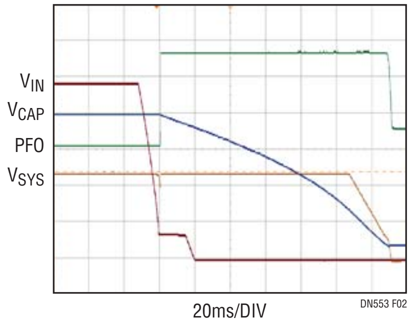

Figure 2. Switchover Waveforms, V_{SYS} = Load Voltage, V_{IN} = Input Voltage, PFO = Flag Status, V_{CAP} = $C_{STORAGE}$ Voltage (V_{SYS} and V_{IN} = 5V/DIV, V_{CAP} = 10V/DIV, PFO 1V/DIV)

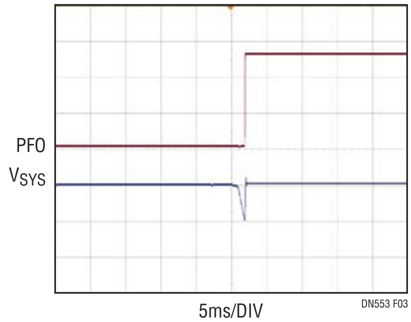

Figure 3. Detailed View of Switching Waveforms (PFO 1V/DIV, V_{SYS} 2V/DIV)

Data Sheet Download

www.linear.com/LTC3643

For applications help,
call (408) 432-1900, Ext. 3161

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

DN553 LT/AP 0816 71K • PRINTED IN THE USA

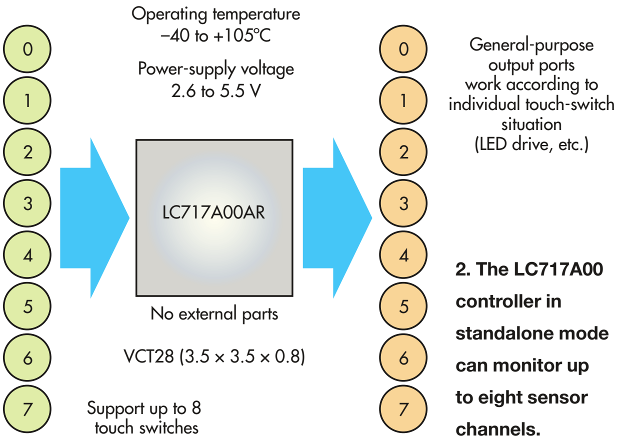

2. The LC717A00 controller in standalone mode can monitor up to eight sensor channels.

as screen printing, making it possible to create a variety of shapes such as buttons, sliders, or rotary selectors. Backlighting is also easy to arrange.

In the absence of any mechanical response confirming actuation, as provided by a tactile switch, the designer may use light to confirm detection of touch. This can be exploited to create stylish effects, such as keeping the panel dark until touched and using color-change effects or high-tech icons. Discreet audible beeps or jingles may also be introduced to complement the touch controls. All of these effects are in keeping with the sophisticated look and feel enabled by a touch-sensitive user interface.

SENSING PRINCIPLES

Generically, capacitive sensing operates by detecting a change in the basic capacitance of the sensor pad upon the approach of the user's fingertip. In the self-sensing type of detector, the approaching finger effectively adds to the sensor's capacitance. The change increases the time constant of the touch-pad circuit, which is detected by charge-time measurement circuitry implemented in the touch-sensor controller.

An alternative technique is mutual differential sensing. It detects a reduction in the capacitance formed between the sensor pad and a separate excitation pad as the user's fingertip comes into proximity with the touch panel. The change in capacitance (ΔC) is detected and converted into a voltage (V_{OUT}). *Figure 1* illustrates a differential-amplifier circuit used to generate V_{OUT}. This has a linear relationship with ΔC, and is compared with a threshold voltage (V_T). Touch detection is indicated when V_{OUT} exceeds V_T.

Mutual differential sensing delivers an advantage through improved sensitivity for a large dynamic range. The practical benefits of this include extended flexibility in terms of sensor-pattern design, and tolerance of long sensor trace lengths due to superior cancellation of parasitic capacitances (C_P) between pad and ground. In practice, trace lengths up to 500 mm have been demonstrated successfully.

TOUCH-CONTROLLER IMPLEMENTATION

The touch sensor must have an electronic controller with enough channels to support the required number of touch buttons. Historically, the controller has been implemented either as a standalone application-specific IC, or integrated in a microcontroller. A microcontroller-based solution may use a combination of analog peripherals and software, while others provide dedicated on-chip touch-control functionality such as a charge-time measurement unit (CTMU) for use with self-sensing detectors. The microcontroller vendor may provide free capacitive-touch software IP, which can help simplify code development.

Implementing a standalone controller brings a number of advantages. An IC such as ON Semiconductor's LC717A00 doesn't require a host microcontroller, and has several built-in features including noise cancellation and water detection, to prevent spurious responses. In addition, patented automated noise and environmental-change compensation technology is also built in to prevent noise-related malfunctions and compensate effects that can alter sensitivity, such as changes in humidity or accumulation of dirt on the sensor area.

The controller's internal circuitry features a unique C/V-conversion amplifier, which converts the monitored differential capacitance to an output voltage. No need for programming helps shorten development time, and because no additional external components are required, it minimizes bill-of-materials costs.

Figure 2 shows the LC717A00 used in standalone mode for controlling up to eight capacitive touch sensors. The device features a communication port, either I^2C or SPI, allowing for convenient connection to a host microcontroller if required.

The controller uses the differential signaling method, which delivers high sensitivity and extends freedom for designers to create control panels in a wide variety of shapes with long connection lengths. *Figure 3* shows a sample touch-panel design with extremely long trace lengths. The data rate is configurable to over 200 Hz, thus enabling its use in equipment that requires rapid touch response.

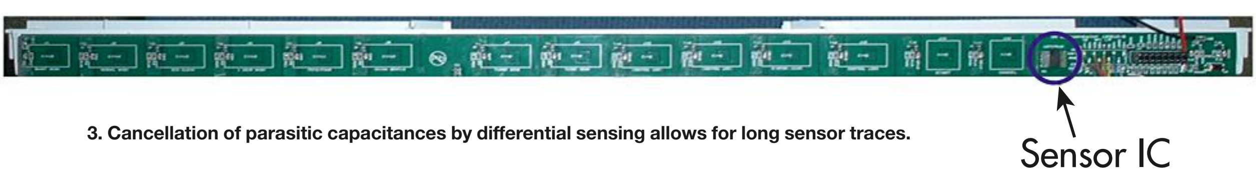

3. Cancellation of parasitic capacitances by differential sensing allows for long sensor traces.

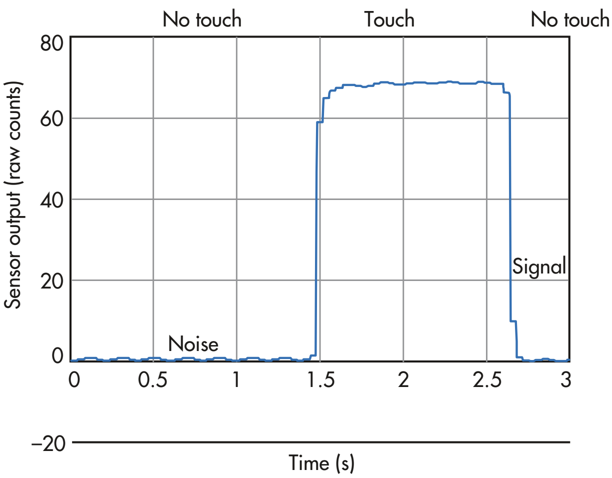

4. A superior signal-to-noise ratio ensures reliable touch detection.

The controller's high signal-to-noise ratio (SNR) increases touch-sensing reliability, even while wearing gloves. Oftentimes, this isn't possible with self-sensing type detectors. Moreover, the design rules governing the shape and size of senor electrodes, and the effects of air gaps as well as overlay thickness and material, can be relaxed. *Figure 4* shows an example of signal and noise levels using real sensor data. In this example, the signal is a 17,000-count difference and the noise is a 160-count difference; therefore, the SNR is 106:1.

CONCLUSION

Capacitive touch has many strengths as a user-interface technology for modern smart appliances. Mutual differential sensing ensures reliable touch detection and allows for greater design freedom when laying out the control panel and selecting materials. A standalone application-specific touch-control IC supports convenient and fast implementation, based on this mutual-differential-sensing technique, and can deliver advantages such as elimination of software design, saving on external components, and acceleration toward project completion. ed

STEVE SHEARD is a program development manager at ON Semiconductor, supporting the System Solutions Group. He has over 35 years in the electronics industry. He previously worked at Motorola as an applications and marketing engineer, supporting products from image sensors and virtual displays to GPS receivers and navigation systems. Sheard holds an MBA from the University of Phoenix and a BSEE from the City & Guilds of London.

AARON YARNELL | Field Applications Engineering Manager
CUI

Peeling Back the Onion on Point-of-Load Modules

Looking beyond typical marketing specifications, here are some of the key differences in point-of-load (POL) dc-dc converter performance and how they relate to your specific system design

The endlessly creative art of marketing is a powerful tool that helps companies highlight their differentiations and advantages over industry competitors. However, do these supposed advantages make a difference in your particular application? Does it really matter if the product of interest is smaller, faster, stronger? As it turns out, like most things in life, the answer is "it depends."

This article will attempt to look beyond typical marketing specifications to truly understand some of the key differences in point-of-load (POL) dc-dc converter performance and how they relate to your specific system design. Specifically, we will look at efficiency, output capacitance, compensation schemes, and cooling requirements.

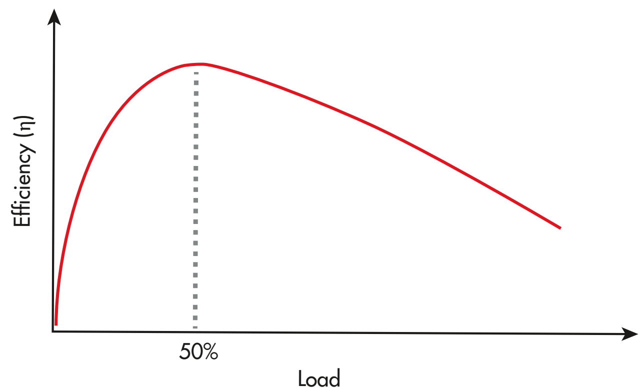

1. Above is an illustration of a typical efficiency curve.

PEAK EFFICIENCY VS. EFFICIENCY UNDER REAL-WORLD LOADING CONDITIONS

The efficiency of power converters is typically denoted by the lowercase Greek letter eta (η), and is expressed as the ratio of power delivered to the output relative to the power consumed by the input ($\eta = P_{OUT}/P_{IN}$). The ideal ratio, or efficiency, of any converter is 1. This indicates that 100% of the power coming into the converter is delivered to the load with zero loss. In real-world applications, though, there will always be some loss/inefficiency associated with converting energy from one form to another, which will reduce the η from 1 down to something less.

Knowing that 100% efficiency is ideal, marketing teams often tout their highest achievable conversion efficiency possible in an attempt to stand out as the "best" for your application. This is often referred to as the "peak efficiency." The challenge is that efficiency isn't just a single number, but rather

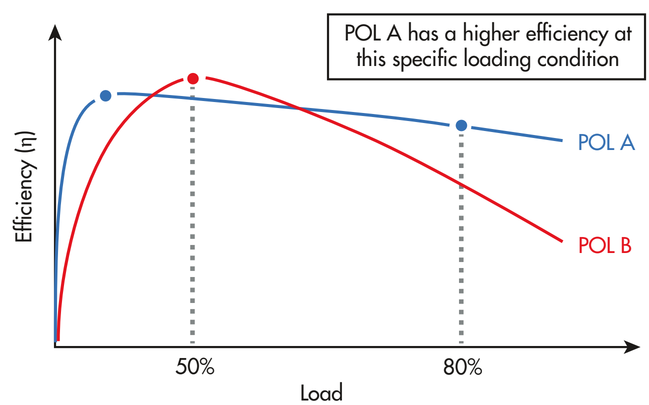

2. Here is a comparison of efficiency curves and an application's loading conditions.

	POL A	POL B
Current Rating	60 A	40 A
ΔV_{out}	10 mV	10 mV
ΔI_{out}	30 A	20 A
Ceramic Capacitors	3x10 µF = 30 µF	4x47 µF = 188 µF
Polymer Capacitors	9x330 µF = 2970 µF	27x330 µF = 8910 µF
Total Capacitors	~ 3000 µF	~ 9000 µF

3. Output capacitance is compared between two POLs.

a multi-variant function that's typically expressed as a function of output current/power delivered to the load. To illustrate how efficiency is affected by output load, *Figure 1* shows a hypothetical example of a point-of-load efficiency curve.

In this hypothetical example, the peak of the efficiency curve occurs when the output load is at 50% of full load. At light loads, the efficiency is much lower; at loads beyond the peak, the efficiency gradually declines. It's important to understand these curves when designing a power-delivery system because operating at any load above or below the peak efficiency point will result in wasted power and unwanted heat in your system. *Figure 2* shows that although POL B has the higher peak efficiency, POL A is the preferred choice for this application (in terms of efficiency) due to the amount of power demanded by the load.

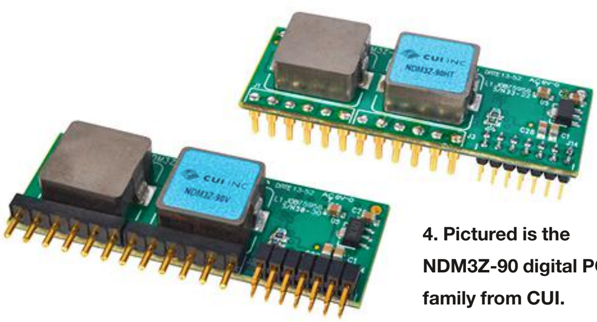

4. Pictured is the NDM3Z-90 digital POL family from CUI.

ACHIEVING DESIRED RIPPLE/TRANSIENT PERFORMANCE

Another metric of interest with POL converters is the amount of added system-level capacitance required to arrive at the desired ripple and transient performance. Details surrounding the theory of external capacitor quantity and type are beyond the scope of this article, but it should be noted that not all POL modules are created equal when it comes to performance—even if the datasheets present similar numbers.

On the surface, it may appear that different POLs have similar ripple and transient performance. However, by digging deeper into the test conditions, you will often see big differences that can affect the overall cost and size of your power-delivery solution.

Figure 3 is a comparison of two competing POL modules. From the datasheets' high-level marketing bullet-point numbers, these two potential solutions look to be fairly identical in terms of ripple and noise.

However, when analyzing the fine print, we see that one of the modules (POL B) requires 300% more external capacitance to achieve the same voltage deviation performance as the other. This represents a significant amount of added cost and under-utilized board space.

Fortunately, more advanced POL modules now offer full digital implementations. They bring significant improvement over traditional analog modules in terms of ripple/transient performance relative to the size of the overall solution.

The CUI NDM3Z-90 series *(Fig. 4)* is an example of one such solution. It delivers up to 90 A of current to the load with good ripple/transient performance, often with a significant reduction in output capacitance.

COMPENSATION SCHEMES

A point-of-load module provides a stable and regulated output in an attempt to produce a clean voltage rail to its load. This implies that the POL inherently contains a negative feedback loop. Therefore, whenever a deviation from the ideal output occurs, the POL's feedback network will compensate and attempt to bring the output back into ideal regulation.

Many different nuanced compensation schemes are available in the market. What follows is a review of the high-level strengths and weaknesses surrounding common analog- and digital-compensation schemes.

• ***Analog compensation:*** In an analog-compensation network, the output of the module is sensed, filtered, and compared to a reference voltage to generate an error signal. This error signal is used to compensate the output and correct any deviations that may have occurred *(Fig. 5)*.

The advantage with analog-compensation schemes is that they have been around for a long time and can be implemented using standard off-the-shelf components. The drawback with such schemes is that it can be quite challenging to "tune the loop" to be stable across all operating conditions while maintaining a wide bandwidth for a fast transient response. This typically requires many hours in the lab to do soldering, testing, re-soldering, re-testing, etc. Analog-compensation schemes are also susceptible to picking up external noise, which could inadvertently be coupled to the output.

Although the analog compensation scheme and its many variants have been the standard for quite some time, there are newer digital compensation schemes that have come along in the last decade or so, which offer some significant advantages.

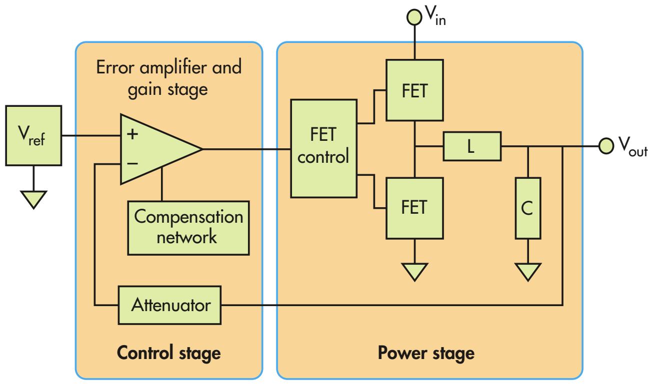

5. Above is a schematic of a typical analog switching voltage regulator.

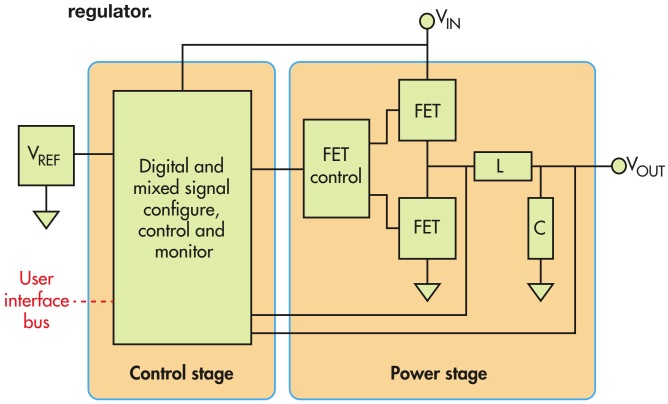

6. Above is a schematic of a typical digital switching voltage regulator.

• ***Digital compensation:*** Similar to analog schemes, a digitally compensated implementation senses the output, filters it, compares it to a reference, generates an error, and ultimately compensates the output to correct any deviations that may have occurred *(Fig. 6)*.

The major difference is that all of this is done in the digital domain with 1s and 0s. The "sensing" of the output is accomplished with an analog-to-digital converter, and then all of the comparing, error-generation, and compensation is done digitally inside an integrated circuit (IC). Operating within the digital domain also significantly improves noise rejection, which helps prevent inadvertent coupling of external noise sources to the output.

Utilizing a digital-compensation scheme means you no longer have to spend hours in the lab soldering different components to tweak the feedback loop. Instead, you can simply modify a few digital parameters within the IC and change the POL's behavior to meet the needs of your application.

The more advanced digital POLs on the market today take this luxury one step further to offer "compensation-free" designs. In these designs, the POL makes all of the necessary measurements and adjustments for you in-system to continuously provide a fast-reacting and stable output-voltage rail.

COOLING REQUIREMENTS

One of the biggest limiting factors in point-of-load modules is heat dissipation. Inefficiencies in the module's design results in unwanted internal heat generation, causing critical components (e.g. FETs, inductors, capacitors, etc.) to approach their maximum rated operating temperature. Operation at or above these components' thermal limits can reduce reliability and result in hardware failure.

To combat the damaging effects of internal heat generation, POL vendors often recommend a minimum amount of airflow to draw heat away from the module. This prevents heat from building up inside the components and raising temperatures beyond their rated limits. Using airflow to remove heat from the module can often increase the amount of power that can be delivered to the load, as well as improve the ambient operating-temperature range. *Figure 7* shows a POL module's ability to operate in different airflow environments ranging from natural convection (still air) to 3 m/s.

We see that under natural convection (still air) conditions (denoted by the lowest solid line in *Fig. 7*), the module can deliver 43 A to the load up to 60°C. Adding just 2 m/s of airflow increases both the current capacity and ambient operating-temperature range up to the full 50 A at 64°C ambient (denoted by the dash-dot line in *Fig. 7*). However, forced-air cooling does have its drawbacks—it consumes power, which can negate some of the efficiency gains, and can generate unacceptable levels of noise. Designers must carefully weigh the thermal requirements of the power module with the cooling capabilities of their system when selecting a POL.

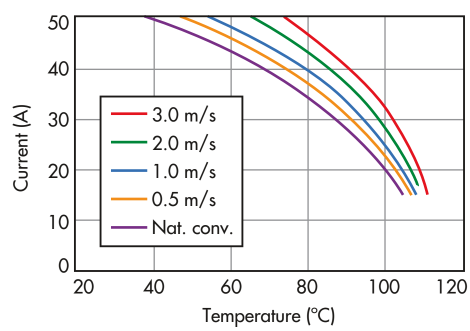

7. This illustration shows a typical derating curve under different airflow conditions.

CONCLUSION

Every application is different and values different performance metrics. For some, a fast transient response may be the most important consideration, while others may require the smallest size, highest efficiency, or widest operating-temperature range. No one POL can meet all of these requirements for each application, no matter what the marketing teams may tell you. It's crucial to first understand your application's needs under its specific operating conditions. Only then can you compare and select the optimal POL for your design.

For further information regarding modules, visit www.cui.com.

Modeling Magnetic Core Offers Insight into Behavior, Operating Range, Saturation

GREGORY MIRSKY | mirskiy@usa.net

IT'S OFTEN NECESSARY to understand the behavior of magnetic components during operation, primarily when applied to chokes or inductors being used for magnetic energy storage and subsequent release. Since most of these components use the magnetic properties of their cores, it's difficult to build an efficient and reliable electronic device without knowing how the magnetic materials behave in the field created by the inductor's winding.

Any magnetic core may saturate, and since the saturation border is not sharp, a first step is to agree on a saturation point. Some devices can afford a higher level of the core saturation, others have a very low one, and some devices use saturation. An analytical tool for assessing the core saturation mode will therefore save considerable design time.

This analysis uses an analytical expression that models the B-H curve diagram by closely describing the process of magnetic-core saturation. It's based on general principles of a magnetic core's behavior in magnetizing and de-magnetizing processes, but cannot take into account all specificities of magnetic materials being used in the electronics industry.

Efforts to find a function that could closely represent a B-H curve resulted in Equation 1. This part of analysis is intended to model the behavior of materials having negligible coercive force, such as soft ferrites:

$$B(H_x) = B_{sat}\cdot\left[\frac{2}{1 + e^{-\left(\frac{H_x}{H_0}\right)}} - 1\right] \qquad (1)$$

where $B(H_x)$ is the function that simulates the B-H magnetization curve; B_{sat} is the magnetic material saturation flux density; H_x is the current value of the magnetic field strength, and H_0 is a reference magnetic field strength, which is still to be found. It defines the slope of the magnetizing curve.

In the corresponding B-H curve *(Fig. 1)*, the slope of the B-H magnetic curve (which is the full magnetic permeability of the material), is found by differentiating Eq. 1 over H_x. The slope clearly reveals the value for H_0:

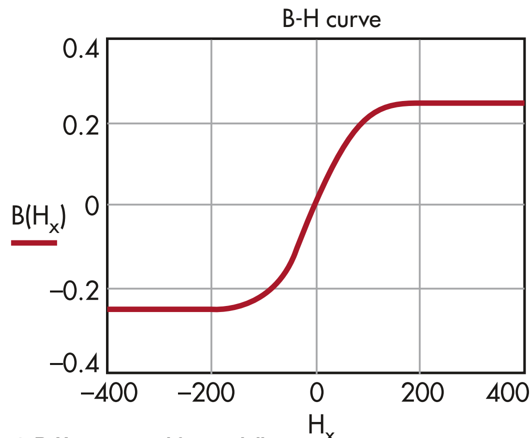

1. B-H curve used for modeling.

$$\frac{d}{dH_x}\left[B_{sat}\cdot\left[\frac{2}{1 + e^{-\left(\frac{H_x}{H_0}\right)}} - 1\right]\right] \rightarrow \frac{2\cdot B_{sat}\cdot e^{-\frac{H_x}{H_0}}}{H_0\cdot\left(e^{-\frac{H_x}{H_0}} + 1\right)^2} \qquad (2)$$

Designating the derivative as:

$$\text{Slope}(H) = \frac{2\cdot B_{sat}\cdot e^{-\frac{H_x}{H_0}}}{H_0\cdot\left(e^{-\frac{H_x}{H_0}} + 1\right)^2} \qquad (3)$$

and defining the slope of the B-H curve at H = 0, we can obtain the missing value for H_0:

$$\text{Slope}(0) = \frac{2\cdot B_{sat}\cdot e^{-\frac{0}{H_0}}}{H_0\cdot\left(e^{-\frac{0}{H_0}} + 1\right)^2} \rightarrow \text{Slope}(0) = \frac{B_{sat}}{2\cdot H_0} \qquad (4)$$

At the same time, we have:

$$\mathrm{Slope}(0) = \frac{B_{sat}}{H_{sat}}$$

Therefore, using Eq. 4, we get:

$$\frac{B_{sat}}{2 \cdot H_0} = \frac{B_{sat}}{H_{sat}}$$

As a result:

$$H_0 = \frac{H_{sat}}{2}$$

However:

$$H_{sat} = \frac{B_{sat}}{\mu_0 \cdot \mu_r}$$

therefore:

$$H_0 = \frac{B_{sat}}{2\mu_0 \cdot \mu_r}$$

Using this in Eq. 1, we have:

$$B(H_x) = B_{sat} \cdot \left[\frac{2}{1 + e^{-\left(\frac{H_x}{\frac{B_{sat}}{2\mu_0 \cdot \mu_r}}\right)}} - 1 \right] \quad (5)$$

Unfortunately, this formula is barely usable, as designers are primarily interested in the dependence of the inductance on the magnetizing current. Instead, move from the classical B-H curve, which we will also plot later to the required L(I) function. Since:

$$L = \frac{N_t^{\,2} \cdot S_{mag}}{l_{mag}} \cdot \frac{dB}{dH} \quad (6)$$

where L is the inductance; S_{mag} is the magnetic core cross-sectional area; l_{mag} is the core's central magnetic line length; and N_t is the inductor's number of turns.

The first fraction is the parameter:

$$\frac{A_L \cdot N^2}{\mu_0 \cdot \mu_r}$$

where A_L is the inductance factor:

$$A_L = \frac{\mu_0 \cdot \mu_r \cdot S_{mag}}{l_{mag}}$$

Therefore:

$$\frac{N_t^{\,2} \cdot S_{mag}}{l_{mag}} = \frac{A_L \cdot N^2}{\mu_0 \cdot \mu_r}$$

The value for A_L can be found in the magnetic-core datasheet. The second fraction is the slope:

$$\frac{dB}{dH} = \frac{2 \cdot B_{sat} \cdot e^{-\frac{H_x}{\frac{B_{sat}}{2\mu_0 \cdot \mu_r}}}}{\frac{B_{sat}}{2\mu_0 \cdot \mu_r} \cdot \left(e^{-\frac{H_x}{\frac{B_{sat}}{2\mu_0 \cdot \mu_r}}} + 1 \right)^2} \quad (7)$$

Simplifying, we obtain:

$$\frac{dB}{dH} = \frac{4 \cdot \mu_0 \cdot \mu_r \cdot e^{-\frac{2 \cdot H_x \cdot \mu_0 \cdot \mu_r}{B_{sat}}}}{\left(e^{-\frac{2 \cdot H_x \cdot \mu_0 \cdot \mu_r}{B_{sat}}} + 1 \right)^2}$$

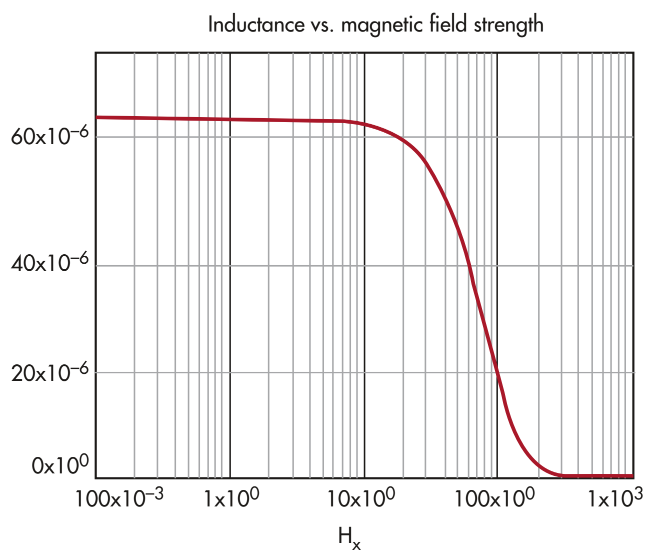

2. Plot indicating inductance vs. magnetic field.

Therefore:

$$L(H_x) = \frac{A_L \cdot N_t^2}{\mu_0 \cdot \mu_r} \cdot \frac{4 \cdot \mu_0 \cdot \mu_r \cdot e^{-\frac{2 \cdot H_x \cdot \mu_0 \cdot \mu_r}{B_{sat}}}}{\left(e^{-\frac{2 \cdot H_x \cdot \mu_0 \cdot \mu_r}{B_{sat}}} + 1\right)^2} \qquad (8)$$

By Ampere's Law, we have:

$$H_x = \frac{I_m \cdot N_t}{l_{mag}}$$

where I_m is the magnetizing current, Using this in Eq. 8, we find:

$$L(I_m) = \frac{4 \cdot A_L \cdot N_t^2 \cdot e^{-\frac{2 \cdot I_m \cdot N_t \cdot \mu_0 \cdot \mu_r}{B_{sat} \cdot l_{mag}}}}{\left(e^{-\frac{2 \cdot I_m \cdot N_t \cdot \mu_0 \cdot \mu_r}{B_{sat} \cdot l_{mag}}} + 1\right)^2} \qquad (9)$$

This expression describes the dependence of inductance on the magnetizing current. Simplifying it yields:

$$L(I_m) = \frac{A_L \cdot N_t^2}{\cosh\left(\frac{I_m \cdot N_t \cdot \mu_0 \cdot \mu_r}{B_{sat} \cdot l_{mag}}\right)^2} \qquad (10)$$

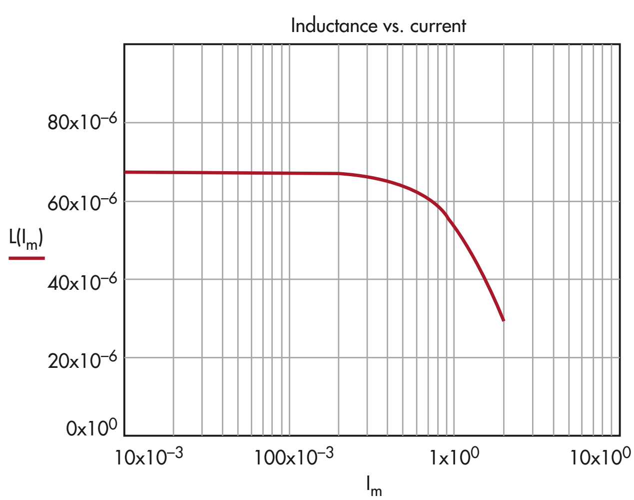

3. This illustrates inductor behavior in the magnetic field.

Equation 10 allows us to find the magnetizing current and the value at which saturation begins:

$$I_m = \frac{B_{sat} \cdot l_{mag} \cdot \operatorname{acosh}\left(\frac{\sqrt{A_L} \cdot N_t}{\sqrt{L}}\right)}{N_t \cdot \mu_0 \cdot \mu_r}$$

Some examples and figures show how these equations can be put to use.

Example 1: Basic B-H curve

Using mT := 0.001T; nH := 10^{-9}H; μ_r := 2500; and B_{sat} := 0.25T, plus the equation:

$$B(H_x) := B_{sat} \cdot \left[\frac{2}{1 + e^{-\left(\frac{H_x}{\frac{B_{sat}}{2\mu_0 \cdot \mu_r}}\right)}} - 1\right]$$

produces the B-H curve of *Figure 1*.

Example 2: Inductance vs. magnetic field strength

Assuming A_L := 4204 nH and N_t := 4, and:

$$L(H_x) := \frac{A_L \cdot N_t^2}{\cosh\left(\frac{H_x \cdot \mu_0 \cdot \mu_r}{B_{sat}}\right)^2}$$

yields the inductance versus magnetic field strength of *Figure 2*.

Example 3: Inductance versus current

Assume l_{mag} := 0.103 m, and

$$L(I_m) := \frac{A_L \cdot N_t^2}{\cosh\left(\frac{I_m \cdot N_t \cdot \mu_0 \cdot \mu_r}{B_{sat} \cdot l_{mag}}\right)^2}$$

The plot in *Figure 3* shows inductor behavior in the magnetic field in terms of the field strength and applied current. It's reasonable to set up the borderline for saturation at –3 dB of the maximum value of the core permeability or inductance.

Example 4: Permeability is equal to the slope of the B-H curve

The expression:

$$\delta_s := \frac{\sqrt{2}}{2}$$

defines the position of the borderline at –3 dB *(Fig. 4)*:

$$\mathrm{Slope}(H_x) := \frac{2 \cdot B_{sat} \cdot e^{-\frac{H_x}{\frac{B_{sat}}{2\mu_0 \cdot \mu_r}}}}{\frac{B_{sat}}{2\mu_0 \cdot \mu_r} \cdot \left(e^{-\frac{H_x}{\frac{B_{sat}}{2\mu_0 \cdot \mu_r}}} + 1 \right)^2}$$

$$\mathrm{Slope}(0) = 3.142 \times 10^{-3} \frac{m \cdot kg}{A^2 \cdot s^2}$$

$$\mathrm{Border} := \mathrm{Slope}(0) \cdot \delta_s = 2.221 \times 10^{-3} \frac{m \cdot kg}{A^2 \cdot s^2}$$

$\mathrm{Border}(H_x) := \mathrm{Border}$

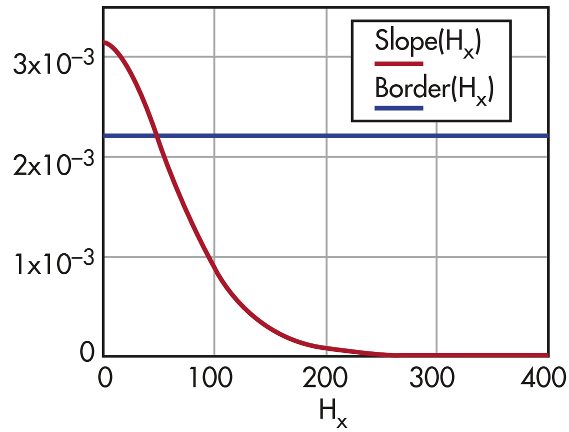

4. In this example, permeability equals the slope of the B-H curve.

Example 5: Core saturation plot

For inductance, the borderline that defines core saturation is shown in *Figure 5*, using $\mu_r := 2500$; $l_{mag} := 103$ mm; $S := 138$ mm^2; $N_t := 4$; $B_{sat} := 0.25$T; $L1 := N_t^2 \bullet A_{L1} = 6.735 \times 10^{-5}$H; $\mathrm{Border}(I_m) := \mathrm{Border}_L$; and:

$$A_{L1} := \frac{\mu_r \cdot \mu_0 \cdot S}{l_{mag}} = 4.209 \times 10^{-6} H$$

$$L(I_m) := \frac{A_{L1} \cdot N_t^2}{\cosh\left(\frac{I_m \cdot N_t \cdot \mu_0 \cdot \mu_r}{B_{sat} \cdot l_{mag}} \right)^2}$$

$$\mathrm{Border}_L := \frac{\sqrt{2}}{2} \cdot L(0) = 4.762 \times 10^{-5} H$$

Solving with respect to the saturation current, we obtain:

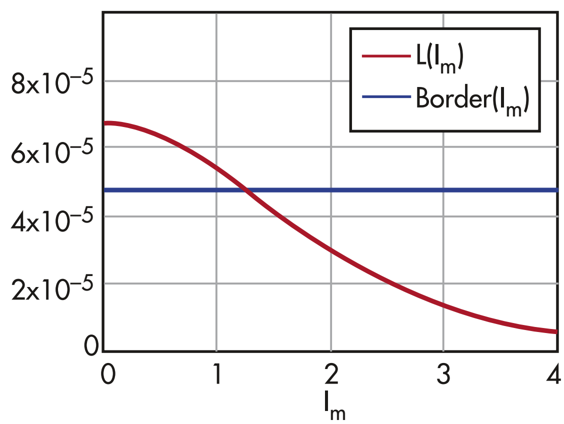

5. The borderline that defines core saturation is shown in this plot.

$$I_{sat} := \frac{B_{sat} \cdot l_{mag} \cdot \mathrm{acosh}\left(\frac{\sqrt{A_{L1}} \cdot N_t}{\sqrt{\mathrm{Border}_L}} \right)}{N_t \cdot \mu_0 \cdot \mu_r}$$

In conclusion, this modeling method makes it possible to estimate a magnetic core's operating range and helps avoid unwanted behavior.

GREGORY MIRSKY holds an MS from the St. Petersburg Baltic Technical University, Russia, and a PhD from the Moscow State Pedagogical University. He can be reached at mirskiy@usa.net.

Simple, Novel Switch Exploits Triboelectric Effect

DEV GUALTIERI | Tikalon LLC

TRIBOELECTRICITY, A FORM of electricity known for millennia, is the static electricity that pulls a spark from your finger to a switch-plate on a winter day when the air is dry and the humidity is low. It's also the reason why most integrated circuits are stored in anti-static foams and bags.

It's generated when certain combinations of materials are rubbed together. The quantity of electric charge that can be generated from rubbing depends on how well the materials can generate or accept electrons. The *table* ranks the triboelectricity of common materials. The ones which are widely separated in this table, such as acrylic and Teflon, will generate the most electrical charge when rubbed.

TRIBOELECTRIC MATERIALS

Material	Relative triboelectricity
Paper (uncoated)	0.5
Silicone II	0.3
Acrylic (polymethyl methacrylate)	–0.5
Epoxy	–1.7
PET (mylar)	–2.1
Polystyrene	–3.7
Polyimide (Kapton)	–3.7
Vinyl (flexible)	–3.9
Cellulose nitrate	–4.9
PVC (rigid vinyl)	–5.3
Latex rubber	–5.5
Teflon	–10.0

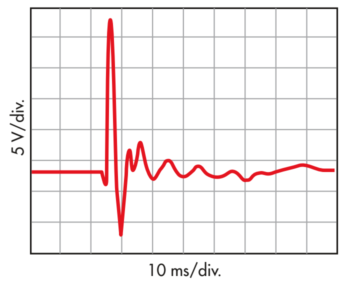

2. The peak voltage of the pulse obtained by tapping the switch depends on the force and speed of the finger tap, and the loading resistance (the probe impedance was 10 MΩ).

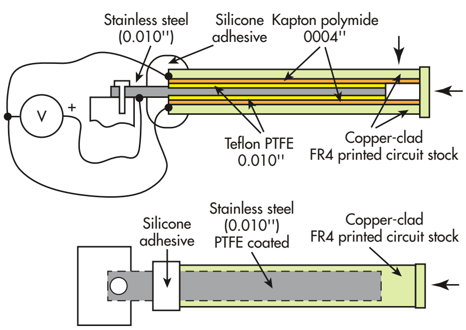

1. Construction of the triboelectric switch shows the layered implementation, which allows for rubbing and thus generation of high voltages.

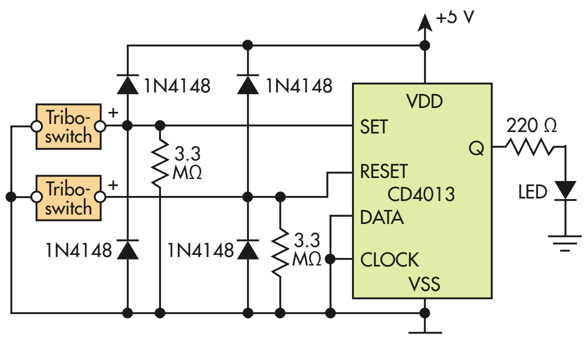

3. The demonstration circuit uses the switch to activate a simple LED; standard 1N4148 diodes are used as clamps to prevent high-voltage spikes from damaging the flip-flop IC.

The triboelectric effect can be used in a simple switch *(Fig. 1)*. A metal tine wrapped with Teflon tape is sandwiched between pieces of copper-clad epoxy circuit board stock and covered with Kapton tape (a polyimide film developed by DuPont in the late 1960s that remains stable across an extremely wide temperature range, from –270 to +400°C). The surfaces are in tight contact with each other, and are joined by a blob of silicone adhesive. In this case, Teflon and Kapton are used as the triboelectric couple, but other materials can be used.

The silicone adhesive allows a rubbing notion when the assembly is tapped from above or on end, as shown by the arrows. The generated charges are collected by the metal tine and the copper surfaces of the circuit board. When discharged into a large resistance, impressively high voltages can be generated *(Fig. 2)*. (The author's demonstration switch was four inches long and ¾-inch wide, but smaller switches are possible.)

You can use such switches to activate CMOS digital circuitry *(Fig. 3)*, where two triboelectric switches control a set-reset flip-flop and control an LED. The diodes prevent high voltages from damaging the input transistors of the CD4013 integrated circuit.

DEV GUALTIERI received his Ph.D. in Solid State Science from Syracuse University in 1974. After many years doing research for a major aerospace company, he now does various computer, electronic and embedded systems projects at his consulting company, Tikalon LLC (http://www.tikalon.com), Ledgewood, New Jersey. He is the author of three science fiction novels and a children's book on cryptography, available on Amazon. He can be reached at gualtieri@ieee.org.

Advanced Debug Tech Available for Cortex-M3 MCUs

SEGGER IS NOW OFFERING J-Link debug and flash programming support for Zilog's ZNEO32! Cortex-M3 based family of MCUs. Production line programming is also available via SEGGER's Flasherfamily. Boasting up to 3 MB/s download speed to RAM and ground-breaking flashloaders, as well as the ability to set an unlimited number of breakpoints in flash memory of MCUs, the J-Link debug probes optimize debugging and flash programming.

J-Trace is a full superset of J-Link with the added benefit of trace capabilities. J-Trace models are designed for one particular family of cores. Available J-Link and J-Trace models vary in price, speed and other hardware features such as built-in trace memory and on-board Ethernet, as well as advanced software features, which include: high-speed production grade flash programming software (J-Flash); an unlimited number of breakpoints while debugging in flash memory (J-Link Unlimited Flash Breakpoints); and support for RDI/RDDI compatible debuggers (J-Link RDI/RDDI).

J-Link debug probes support ARM 7/9/11, Cortex, Microchip PIC32, Renesas RX CPUs and are supported by all major IDEs such as IAR EWARM, Keil MDK-ARM, Rowley Crossworks, SEGGER's Embedded Studio and other Eclipse/GDB based offerings.

SEGGER

https://www.segger.com

Module Offers Multi-GNSS in Ultra-Small Footprint

THE MULTI MICRO SPIDER (ORG4033) from OriginGPS, a fully integrated and highly sensitive multi-GNSS module within an ultra small footprint, is positioned for applications that require quick movement, minimal power consumption and ultra-small form factors. Like its predecessor, the Multi Micro Hornet (ORG1510-MK), the new module utilizes MediaTek's MT3333 chip and its onboard flash memory to achieve a rapid update rate and positioning speed of up to 10 Hz.

The multi-GNSS module's key features include sub-one second TTFF and a sensitivity of -165 dBm for two simultaneous constellations packed in a 5.6 x 5.6 x 2.65 mm footprint, while consuming less than 9 mW of power. The company's proprietary Noise Free Zone technology provides continued noise immunity and razor-sharp sensitivity even in poor signal conditions and a design that allows for seamless migration from GPS to GNSS pin-to-pin compatibility. Utilizing MediaTek's MT3333 chip, the Multi Micro Spider multi-GNSS module extends the functionality of GPS and GNSS solutions in wearables, drones, and IoT devices.

ORIGINGPS

www.origingps.com

Ruggedized Power Connectors Snap-On, Push-to-Release

SNAP-LUG RUGGEDIZED POWER CONNECTORS from TE Connectivity boast a dependable and rapid method for interconnecting heavy gauge power cables to bus bars, relays, batteries, power cables, bulkhead power feedthroughs and grounding devices. The self-locking power connectors, featuring snap-on/push-to-release installation and a sealed interface preventing contamination, are color coded and keyed for polarity to help prevent damage from accidental polarity reversal. A shrouded, isolative design, with no exposed metal, provides a touch-proof connection that helps prevent accidental shock incidents when under power. With a -65°C to +150°C operating temperature range and vibration and mechanical shock compliant to MIL-STD-1344 and MIL-STD-810G, the connectors can withstand extreme temperatures and high vibrations to enable continuous power in harsh environments. Snap-Lug ruggedized power connectors are used in military ground vehicles, commercial aviation and industrial commercial vehicles.

TE CONNECTIVITY

http://www.te.com/usa-en/home.html

ADAS DevKit Features Reference Design with AVP

D3 ENGINEERING is now offering its DesignCore Development Kit for Advanced Driver Assistance Systems. Based on D3's DesignCore Rugged Vision Platform Reference Design along with TI's TDA3x Automotive Vision Processor, the new development kit allows product development teams to evaluate ADAS technology under realistic on-vehicle conditions.

The DesignCore RVP-TDA3x Reference Design enables synchronous acquisition of four HD video streams with real-time vision processing and analytics. The ADAS development kit includes the reference design module, four FPD-Link III rugged camera modules, HD video display, cables, software and calibration tools. It features multiple connectivity interfaces and automotive 12 Vdc power. Delivered on a wheeled test platform, the components can be readily de-mounted and installed on a test vehicle after initial verification. Available now to qualified design integrators, the DesignCore ADAS Development Kit is priced at $4,750.

D3 ENGINEERING

www.d3engineering.com

Synchronous Rectification ICs Provide Efficiency in SMPS

INFINEON TECHNOLOGIES' latest SSR family of controller ICs, the IR1161L and IR11688S, meet the 2016 standards set by the U.S. DOE and the ECCEEDC which demand a 1% to 3% efficiency improvement over previous requirements. Together with Infineon's energy saving OptiMOS and StrongIRFET MOSFETs, the family of SSR ICs seek to help improve overall efficiency for SMPS. The IR1161L is targeted for flyback SMPS in chargers and adapters, while the IR11688S is targeted for LLC SMPS in televisions, desktops, silver boxes and micro-servers.

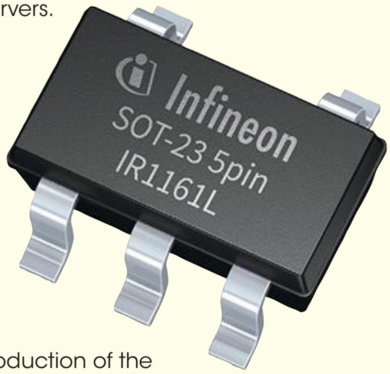

Both, the IR1161L and IR11688S offer 200 V direct sensing, which eliminates the need for external voltage dividers. The family also features programmable MOT (minimum on time), ensuring reliable operation from no-load to full-load. The low quiescent current capability meets standby requirements of the 2016 efficiency standard in the industry. Volume production of the IR1161L and IR11688S families of controller ICs has started.

INFINEON TECHNOLOGIES

https://www.infineon.com/

Mini SMT Gas Discharge Tubes Boost Surge Ratings

LITTELFUSE, TWO NEW SERIES of miniature SMT two-terminal Gas Discharge Tubes designed to protect sensitive electronics from low-to-medium level lightning-induced surges and other voltage transients. The 3.5 mm diameter sized CG6 Series is a two-electrode single chamber SMT GDT that features two square terminals and offers 3 kA (8/20 µs), surge withstand capability (10 shots). The CG7 Series provides surge handling capability up to 1kA @ 8/20µs (10 shots) in a 2.8 mm SMT footprint with one square and one round terminal.

Both series provide rapid response to fast-rising transients and have ultra-low capacitance, which helps to minimize insertion and return losses to ensure signal integrity in high-speed applications. The GDTs are suited for a range of applications, including: broadband, satellite and CATV equipment; dataline and Ethernet (up to 10bE) equipment; ADSL, xDSL and general telecom equipment; G.Fast modems; EMTAs; set-top boxes; RF connectors; MoCA equipment; base station RF antenna transmitters; and IADs. The CG6 Series with two square terminals is being offered in quantities of 2,000. The CG7 Series with one square terminal is offered in quantities of 2,500.

LITTELFUSE

www.littelfuse.com

Movable EV DC Chargers Meet IEC 61851 Standards

PHIHONG'S FAMILY OF MOVABLE EV DC CHARGERS provide highly efficient charging solutions for personal and commercial electric vehicles with 200 V-700 V battery systems. Featuring a moveable design, the off-board charger family is currently comprised of 20, 40 and 60 kW models, all of which are compliant with the latest GB/T standards for EV interface and conductive charging system requirements issued by the CNC of the ISO and IEC. Implemented January 2016, these latest GB/T standards directly correspond to IEC 61851 requirements regarding electrical safety, charging function and performance.

Each charger features a compact, digital, 47~63 Hz power module that delivers 93%+ efficiency with a power factor of >0.99 at full load. Equipped with intelligent RFID card readers and 7" LCD touchscreens, each model is outfitted with tactile start, stop and emergency stop push-buttons, and features an Ethernet-based connection that enables remote assistance, troubleshooting, repairs, integration and upgrades. Each model also features a charging gun with an electronic lock function. The Movable EV DC Chargers have an input voltage rating of 380 Vac ±15% or 480 Vac ±10%, and are suited for highway gas/service stations, commercial fleet operators, EV dealer workshops, parking garages, EV infrastructure operators and household applications.

PHIHONG

https://phihong.com

Arduino Primo Features Native BLE Technology

NORDIC SEMICONDUCTOR'S NRF52832 Bluetooth low energy SoC is specified at the heart of Arduino's latest product: an IoT-targeted programmable SBC called the Arduino PRIMO. In addition to being able to wirelessly connect to an array of BLE sensors, the Arduino Primo uses the chip's integrated NFC for secure authentication and "Touch-to-Pair." The BLE SoC features a 64-MHz, 32-bit ARM Cortex-M4F processor that boosts generic processing power while increasing Floating Point and DSP performance. The SoC's 2.4 GHz multiprotocol radio is fully compatible with Bluetooth 4.2 specification, and features -96 dB RX sensitivity, and 5.5 mA peak RX/TX currents. The chip also features 512 kB Flash memory and 64 kB RAM, plus a fully-automatic power management system.

In addition to controlling the BLE RF protocol stack and application code, the SoC's ARM processor has ample computational overhead to supervise and control the Arduino Primo's on-board accelerometer, temperature, humidity, and pressure sensors. For common projects, programming is done via the Arduino IDE programming interface. Developers who want access to Primo's full potential will also be able to use any Nordic nRF52 Series-compatible SDK or programming tools. For example, the nRF5 SDK for IoT facilitates development of IPv6 over BLE applications on the SoC, enabling the Arduino Primo to communicate with other Internet-connected things without requiring the resources of a Wi-Fi router or smartphone.

ARDUINO

https://www.arduino.cc

Rugged Fanless PC for Embedded Applications

PREMIO'S THREE latest families of rugged fanless PC are designed for embedded applications where operating environments are harsh, such as machine automation, security surveillance, IOT gateway and outdoor signage player. The hardware offerings include the RCO-1000 Series, RCO-3000 Series, and RCO-6000 Series Fanless PC series, which incorporate Intel Bay Trail, Broadwell-U and Skylake processor support, respectively. All models support extended operating temperature (-25°C to 70°C), 9~48 Vdc input with over/under-voltage and reverse polarity protection, 5 Grms vibration and 50 G/11ms shock protection, and cable-less design.

All of the PCs are built with one-piece extruded aluminum, and each has several form factors to support different I/O and expansion requirements. Custom I/O cards are available to bring additional I/O out without internal cabling. For example, COM port, DisplayPort, Digital I/O and 4-port GbE (with or without 802.3at PoE) in RJ-45 or M12 connectors expansion can be added. Additional expansion options include mini-PCIe with universal I/O slots and PCIe, or PCI slot(s) that allow further I/O expansion. All models also come with the option of dual SIM card slots to support cellular network connection. In addition to the barebones systems, Premio also offers full system integration services from system build, revision control, inventory and distribution, and reverse logistics.

PREMIO INC.

www.premioinc.com

BILL WONG | Embedded/Systems/Software Editor
bill.wong@penton.com

Don't Forget Those Emulators and Simulators

When developing applications, designers tend to overlook emulators and simulators, which can play key roles in the debug stage.

It has been a long time since I worked on minicomputers like the Data General Nova (*see figure*). I remember entering boot sequences using the front-panel switches to load a paper tape-based loader that finally loaded an operating system. The process was not too involved, and it kept the OS running (which was a good idea, because rebooting was not fun).

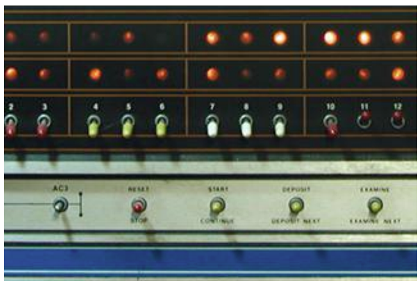

The Data General Nova was a 16-bit minicomputer that first appeared in 1969.

I also dabbled with a number of minis, including DEC LSI-11s and Burroughs B1700s. I recently decided to see if there was a simulator or emulator for the Nova. Sure enough, I found a couple of them, along with ones for most of the other platforms I remember using. Some are limited simulations of the processor, while others emulate peripherals.

Most of these were put together for nostalgic purposes. It continues with latter-day platforms like the 6502 that drove the original Apple II and Atari game consoles. There's actually a Javascript emulator that probably runs faster than any 6502 ever made on most browsers.

Emulators and simulators are often overlooked when developing applications. They're heavily used for high-end system-on-chip (SoC) platforms these days because of the cost and timing. Misconceptions still persist about this type of emulation (*see "11 Myths About Hardware Emulation" on electronicdesign.com*), but it has many advantages.

For instance, such emulators allow software developers to get started before hardware is available, simulating and debugging using hooks that can only come from software. A wide range of options exists between these high-end, high-resolution simulators and the nostalgic implementations that are often overlooked by developers, who typically assume they have to develop on their target hardware.

My favorite implementations are emulators for low- to mid-range microcontrollers. The emulators often run applications faster on a PC than on the actual hardware. They can be ideal for software regression or unit testing. The systems are often available from the hardware vendors and integrated with the development tools. I have used Microchip's in the past to good effect.

Another option is the QEMU system emulator. It's integrated with most Linux distributions and runs on a range of hardware to emulate a similar range of hardware. For example, it's possible to run an ARM Cortex-A environment on top of an x86 version of Linux. Hardware virtualization can be used when emulating similar hardware such as x86 on an x86 system.

Emulating or simulating architectures that differ from the host can be done in a number of ways. Two methods are interpretation and dynamic translation of machine code. The latter is essentially just-in-time (JIT) compilation of one machine code to another. The requirements placed on simulation often determine how these tasks are performed. For example, does execution have to replicate a chip's timing and, if so, to what level of accuracy? At the other end, just the functionality may need replication.

Emulating multiple systems can be very useful in debugging communication scenarios for the Internet of Things (IoT). Often, the debug hooks are more sophisticated and less expensive to implement than using real hardware for tracing and replication of test scenarios. Using real hardware will still be required, but doing more with only software can reduce development and debugging time.

I'm always amazed by developers who don't know about debuggers or trace tools. I should probably add emulators and simulators to that list. Make sure you don't overlook any of these tools. ed

20-Bit SAR ADC Driver

Fast, Accurate Data Acquisition

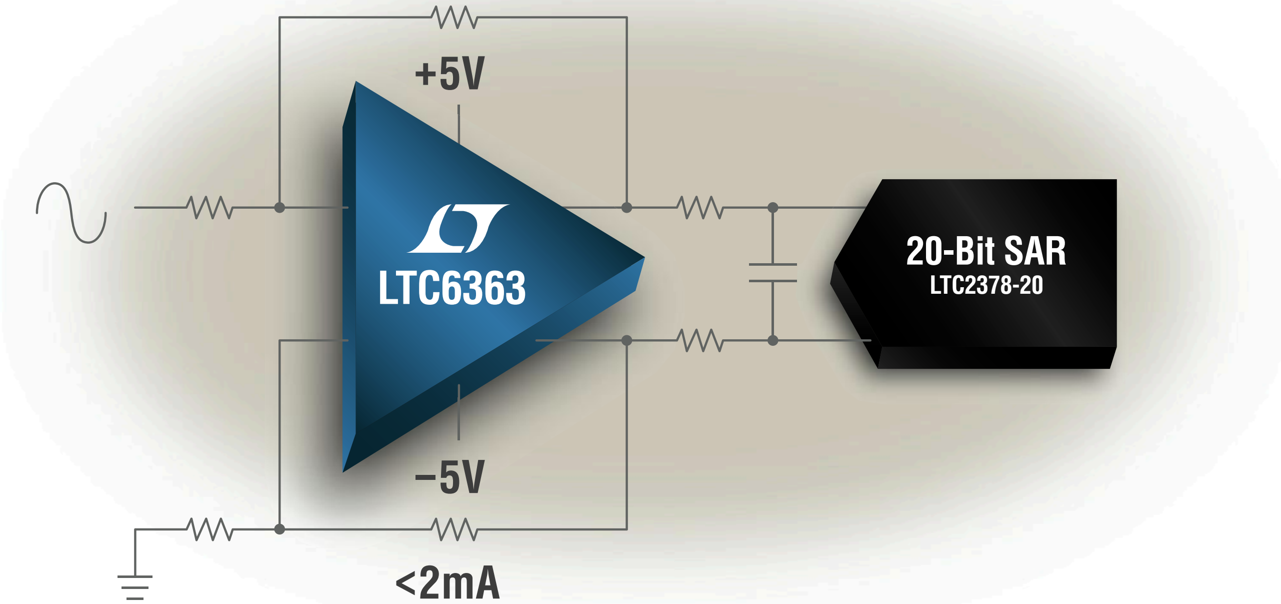

Drive your precision ADC with the power efficient LTC®6363. Drawing just 1.9mA, this fully differential amplifier achieves 100μV max input offset voltage and 2.9nV/√Hz broadband noise. With a 2.8V to 11V supply range and rail-to-rail outputs, the LTC6363 is well matched to the ADC's full range, to achieve maximum signal path performance.

Features

- 100μV Max Offset
- 50nA Max Input Offset Current
- Fast Settling: 780ns to 18-Bit, $8V_{P-P}$ Output
- 1.9mA Supply Current
- 2.9nV/√Hz Input-Referred Noise
- 35MHz -3dB Bandwidth
- 2.8V to 11V Supply Range

Recommended Signal Chain Products
ADCs
LTC2380-24: 1.5Msps, 24-Bit, Digital Filter
LTC2378-20: 1Msps, 20-Bit, Low Power
LTC2373-18: 1Msps, 18-Bit, 8-Channel
LTC2387-18: 15Msps, 18-Bit
Voltage References
LT6657: 1.5ppm/°C Drift
LTC6655: $0.25ppm_{P-P}$ Noise
LT6654: SOT-23
Matched Resistors
LT5400: Quad, 0.01% Matching

Info & Free Samples

www.linear.com/product/LTC6363

1-800-4-LINEAR

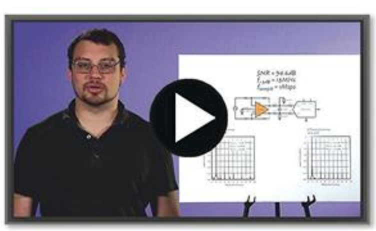

video.linear.com/5950