electronic design

THE AUTHORITY ON EMERGING TECHNOLOGIES FOR DESIGN SOLUTIONS

July 2015 electronicdesign.com

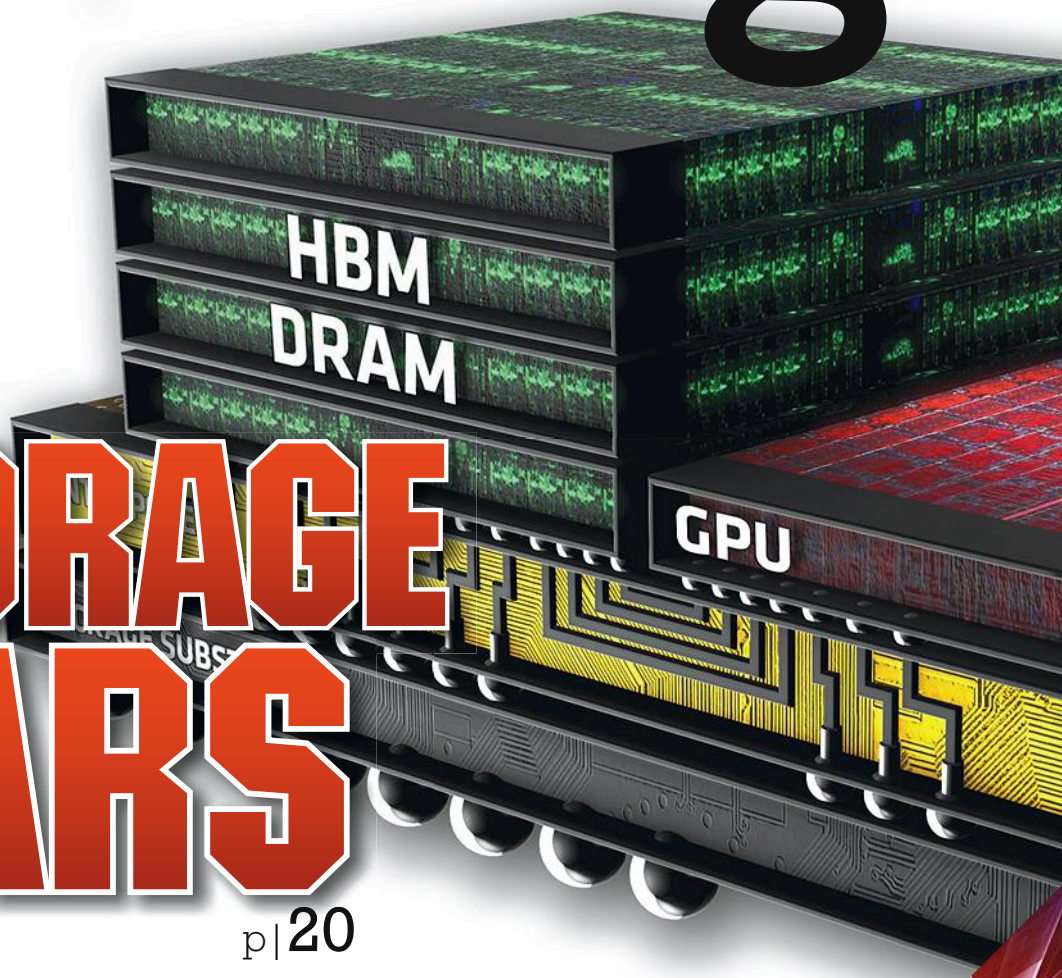

STORAGE WARS

p|20

$10.00 Powered by Penton

Periodicals Postage Paid • USPS 100 Approved Poly

From Search to Samples in Seconds

Our web tools help you quickly pick the perfect inductor. Then get free samples overnight!

We make picking inductors fast, easy, and almost fun!

Start with a quick parametric search for parts that meet your specs. (We even show you volume pricing.) Click to instantly compare core and winding losses for multiple inductors.

Want to see actual L and I curves at a specific temperature? There's a tool for that!

When you've picked your inductors, just click to get free samples. Do it by 5:00 p.m. CST and they'll ship to you the same day!

See how fast *you* can go from search to samples. Visit **coilcraft.com/tools**

Coilcraft®

WWW.COILCRAFT.COM

coilcraftdirect.com
No min order. Next day delivery.

PAM-4 insights don't schedule meetings.

They come when they're good and ready.

Some call them Eureka moments. Others call them epiphanies. We call them insights, the precise moments when you know you've found great answers. As the networking industry considers transitioning to more complex signaling, we can help you achieve insights to meet the technical challenges of PAM-4 that lie ahead. From simulating new designs to characterizing inputs, outputs and connectors, we have the software, hardware and measurement expertise you need to succeed.

Keysight Advanced Design System bundle for signal integrity
Simulation-measurement correlation and workflow for Ethernet PAM-4 and NRZ compliance with the Ethernet Compliance Test Bench also available

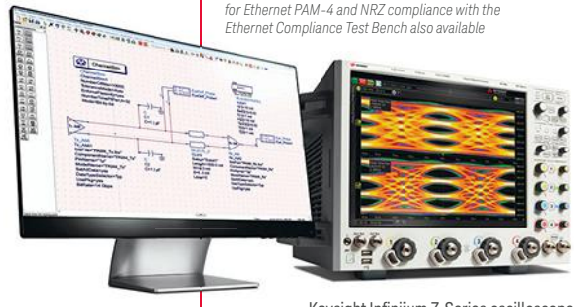

Keysight Infiniium Z-Series oscilloscopes
Compliance solutions available for current and emerging PAM-4/Ethernet standards

HARDWARE + SOFTWARE + PEOPLE = PAM-4 INSIGHTS

PEOPLE

- *Member representatives in test working groups including IEEE, OIF-CEI, and Fibre Channel Industry Association*
- *Applications engineers in more than 100 countries around the world*
- *Nearly 1,000 patents granted or pending*

Download our app note *PAM-4 Design Challenges and the Implications on Test* at www.keysight.com/find/PAM-4-insight

USA: 800 829 4444 CAN: 877 894 4414

HARDWARE + SOFTWARE

- *Instruments designed for testing PAM-4 from simulation to compliance*
- *Advanced Design System software for simulation-measurement correlation and workflow*
- *More than 4,000 electronic measurement tools*

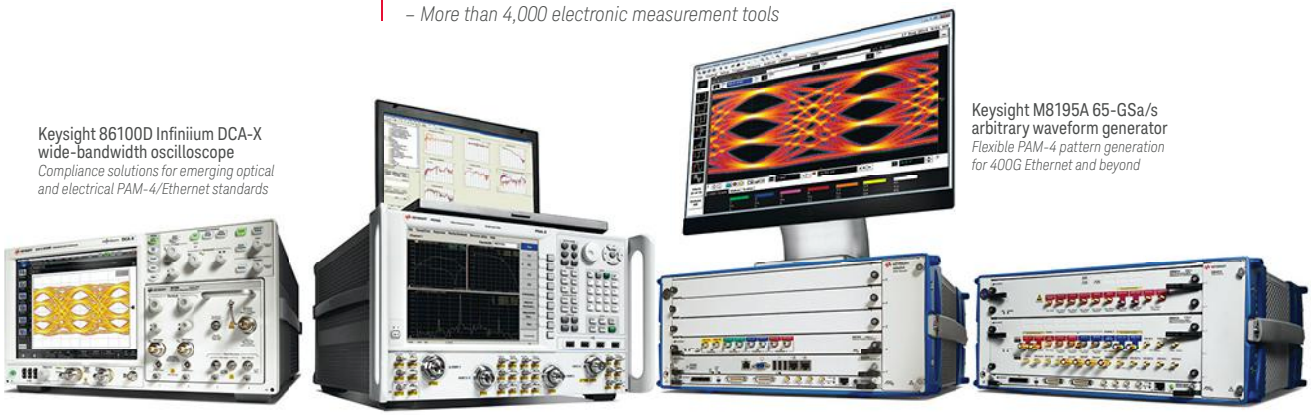

Keysight 86100D Infiniium DCA-X wide-bandwidth oscilloscope
Compliance solutions for emerging optical and electrical PAM-4/Ethernet standards

Keysight M8195A 65-GSa/s arbitrary waveform generator
Flexible PAM-4 pattern generation for 400G Ethernet and beyond

Keysight N5245A PNA-X microwave network analyzer with N1930B physical-layer test system software
Gigabit Ethernet interconnect and channel test solutions

Keysight J-BERT M8020A high-performance BERT
The most integrated solution for 100G Ethernet input testing

KEYSIGHT TECHNOLOGIES

Unlocking Measurement Insights

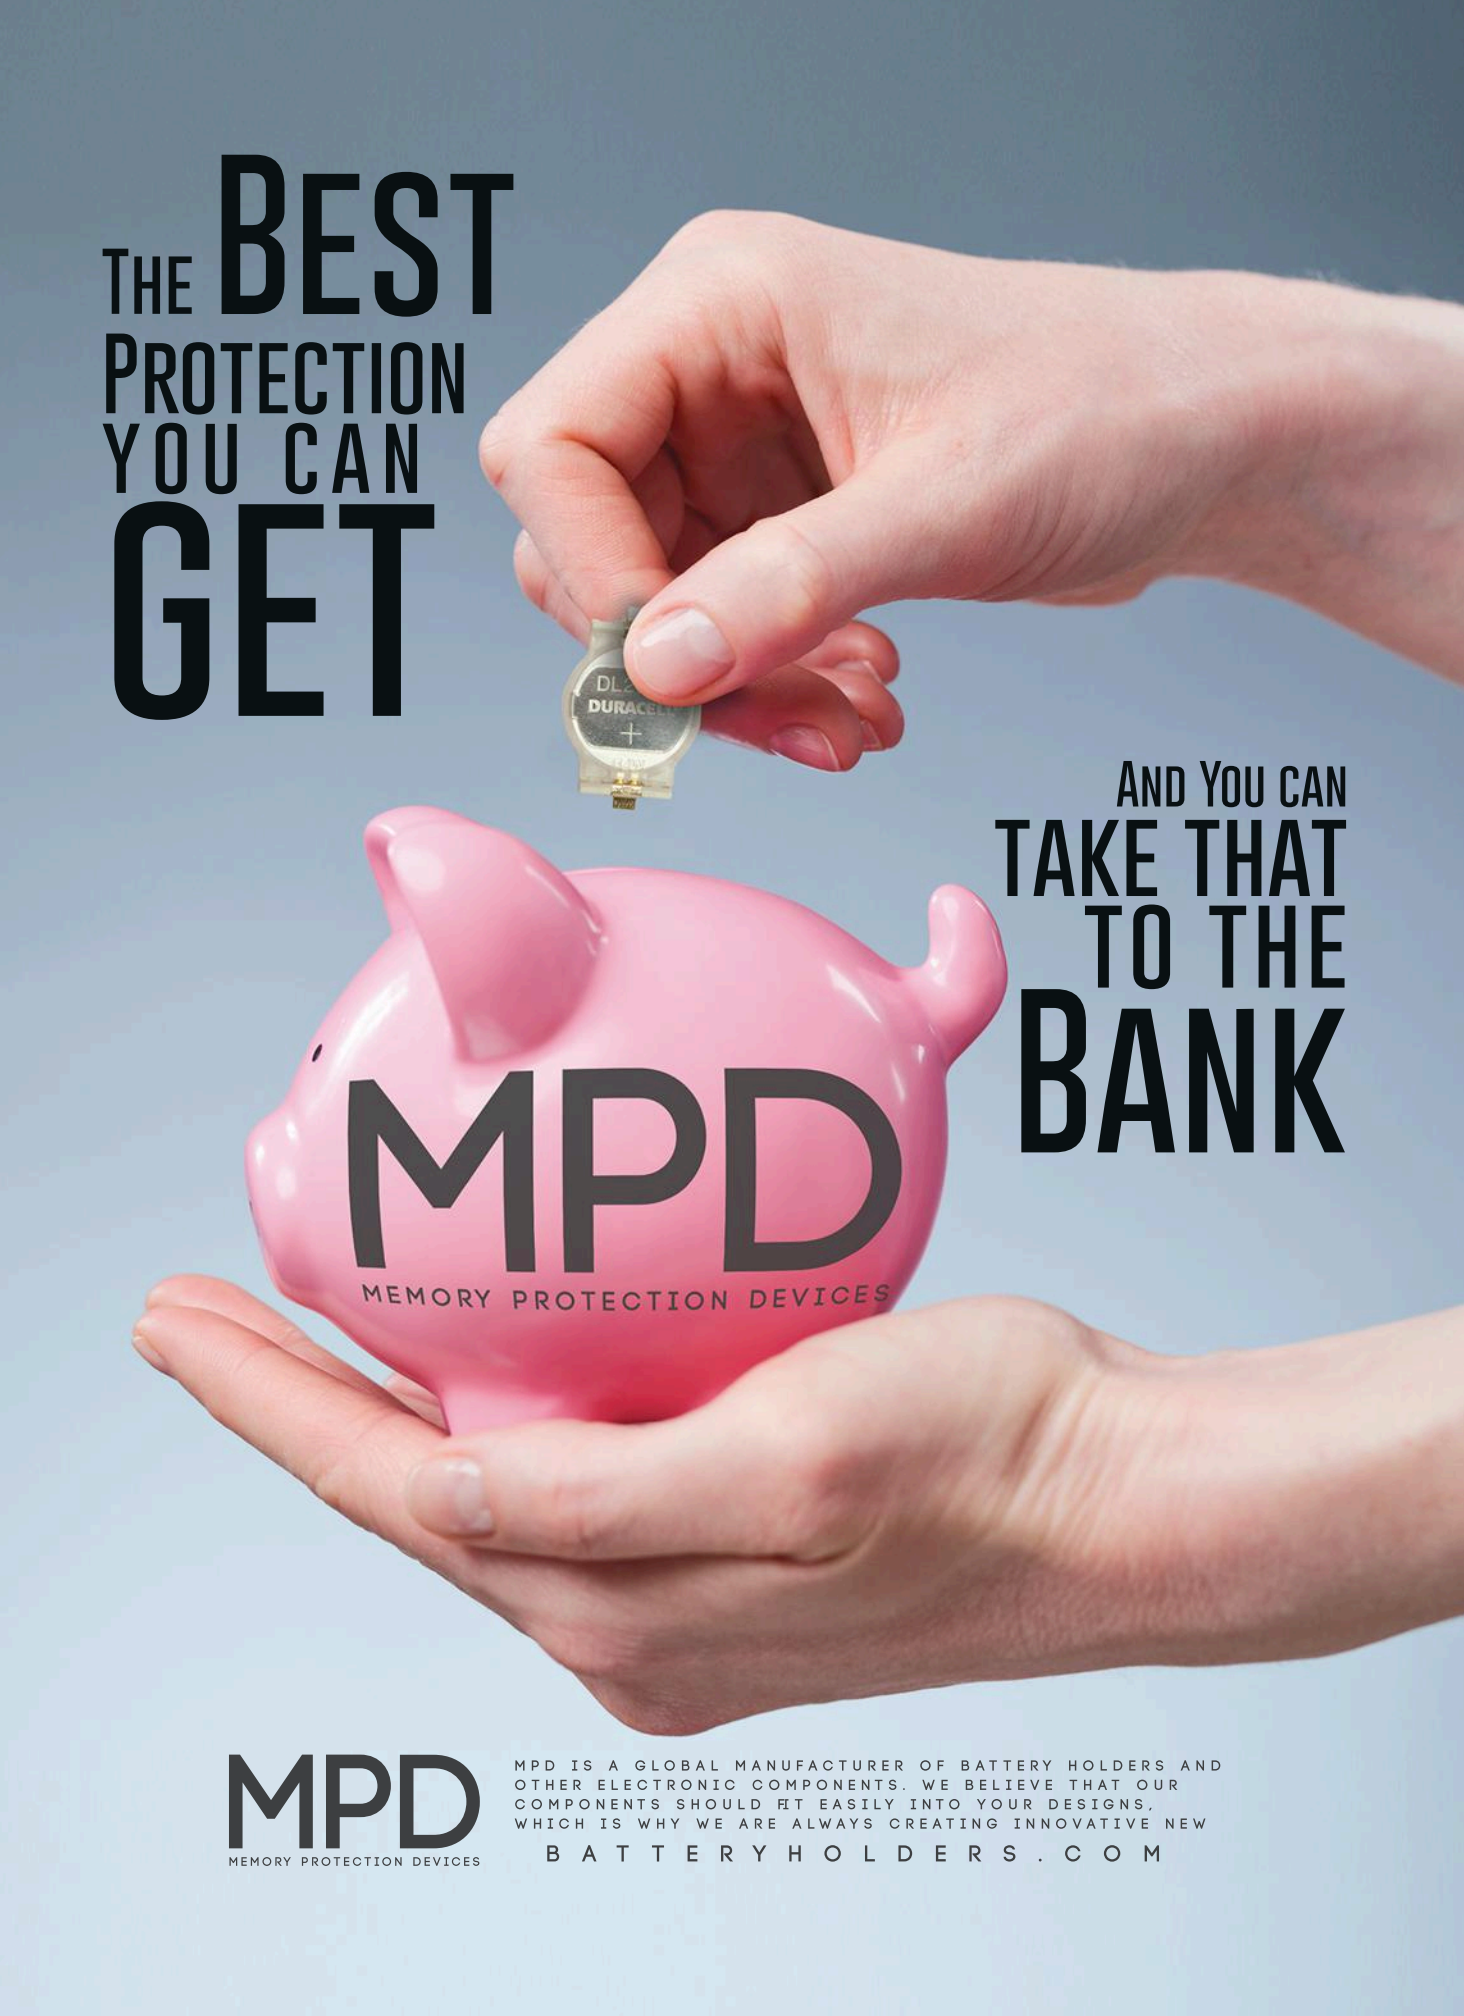
THE BEST PROTECTION YOU CAN GET
AND YOU CAN TAKE THAT TO THE BANK
MPD
MEMORY PROTECTION DEVICES
MPD
MEMORY PROTECTION DEVICES
MPD IS A GLOBAL MANUFACTURER OF BATTERY HOLDERS AND OTHER ELECTRONIC COMPONENTS. WE BELIEVE THAT OUR COMPONENTS SHOULD FIT EASILY INTO YOUR DESIGNS, WHICH IS WHY WE ARE ALWAYS CREATING INNOVATIVE NEW
BATTERYHOLDERS.COM

In This Issue

FEATURES

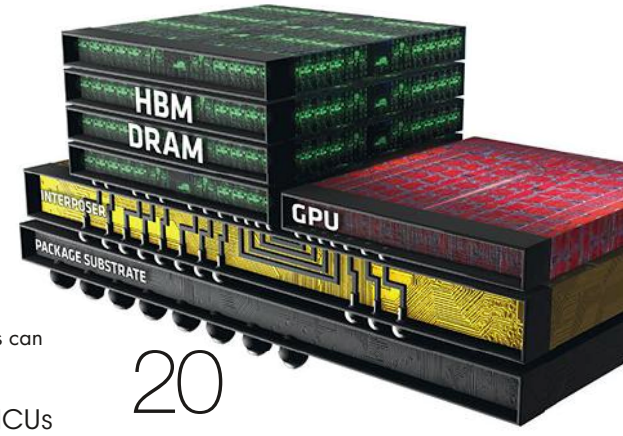

20

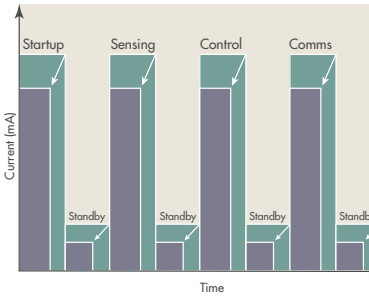

24

33

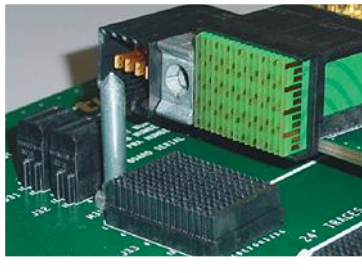

34

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. 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

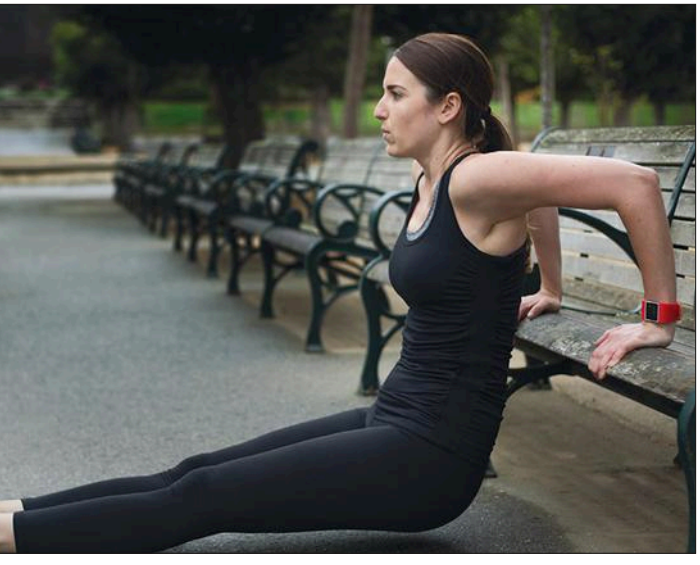

STAY UP-TO-DATE WITH WIRELESS POWER

New standards have emerged that allow for interoperability of charging surfaces and portable devices across different manufacturers, as well as higher efficiency and higher power levels up to several watts.

http://electronicdesign.com/power/stay-date-wireless-power

WHAT'S THE DIFFERENCE BETWEEN NRZ AND PAM?

PAM-4 may be to used to reach 400-Gb/s Ethernet, as well as being a good fit for other high-speed serial interfaces like Fiber Channel. So what's the difference between the usual non-return to zero (NRZ) and pulse amplitude modulation (PAM)?

http://electronicdesign.com/communications/what-s-difference-between-nrz-and-pam

blogs

LOUIS FRENZEL
COMMUNICATIONS

- Justifying 2.5G/5G Ethernet

BILL WONG
EMBEDDED/SYSTEMS/SOFTWARE

- Safety Document Package Targets ADAS

LIGHTWEIGHT VR GLASSES USE BLUETOOTH

As wearables continue to evolve, one of the newest, more stylish advances are Dlodlo VR Glasses, weighing only 4 oz. (120 g) and using Bluetooth to connect to a host device.

http://electronicdesign.com/displays/lightweight-virtual-reality-glasses-use-bluetooth

SIMULATING A BATTERY REAPS BENEFITS

When proving battery-powered device designs, a battery emulator reduces test setup time, creates a safer test environment, and provides more repeatable results versus using a real battery during test.

http://electronicdesign.com/test-measurement/simulating-battery-power-supply-reaps-benefits

join us online

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

Characterizing multiport devices?

Time to speed up your production

Characterizing multiport devices in production and development has never been easier. Rohde & Schwarz offers a large portfolio of network analyzers for parallel measurement and control of one DUT with many ports or many DUTs with a few ports.

- Up to 24 test ports for ultimate speed and parallel testing with the new R&S®ZNBT8 vector network analyzer
- Up to 48 test ports with switch matrices controlled directly by the R&S®ZNB vector network analyzer

More to explore: www.rohde-schwarz.com/ad/multiport

ROHDE&SCHWARZ

electronic design
A Penton® Publication
JULY 2015

Higher frequency signals, meet True*form* signal integrity.

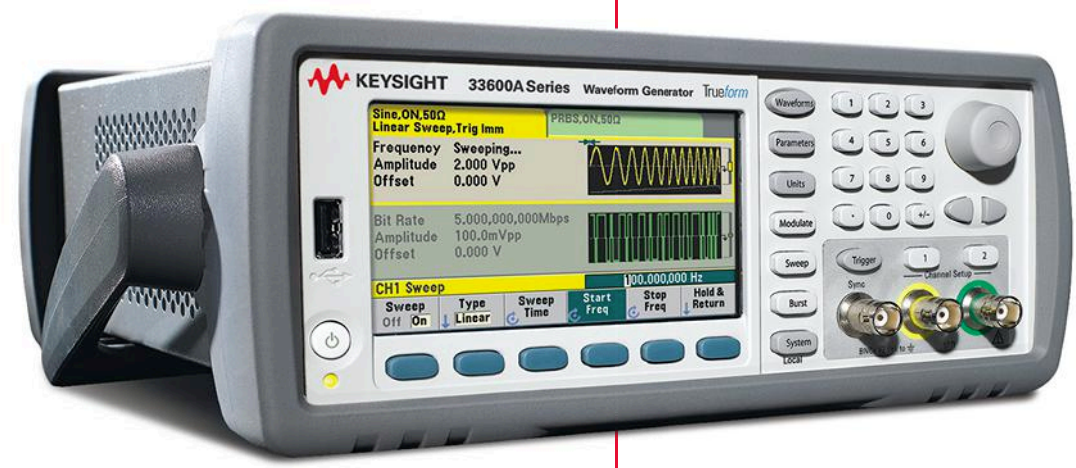

Higher frequency Keysight 33600A Series waveform generators with True*form* technology deliver unparalleled signal fidelity. Choose from four models to accelerate even your most complex tasks.

33600A Series Waveform Generators

- 1 GSa/s sampling rate and up to 120 MHz bandwidth
- Arbs with sequencing and optional 64 MSa memory
- 1 ps jitter, 200 times better than DDS generators
- 5x lower harmonic distortion than DDS generators
- Compatible with Keysight BenchVue software

Newark
element14

View the 33600A Series library of measurement briefs, videos, and more.
www.newark.com/Keysight_Function_Generator
Buy from an Authorized Distributor 800 463 9275

KEYSIGHT TECHNOLOGIES
Unlocking Measurement Insights

Redefining Automated Test

with open software and modular hardware

How we interact with devices is changing. As the world becomes more software oriented, what we can accomplish increases exponentially. This shift should apply to our test equipment, too. Unlike traditional instruments with predefined functionality, the NI automated test platform provides the latest technologies to build complex systems while reducing development time and cost.

Through an intuitive graphical programming approach, NI LabVIEW reduces test development time and provides a single environment that simplifies hardware integration and reduces execution time.

>> Accelerate your productivity at ni.com/automated-test-platform

800 891 8841

NATIONAL INSTRUMENTS

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

Home Security Defense in Depth

How many wireless IoT gizmos do you have in your home? I have a few, including a couple of Nest thermostats. This does not include the plethora of smartphones, tablets, e-readers, and network printers that also use my local Wi-Fi network. I also have four servers running. I like lots of backup. The system is complex compared to the typical home network, since most of the servers are running a collection of virtual machines. It is probably overkill for even a small business.

My attempts to lock things down include very long root passwords and only using root access to manage the system versus regular access. The Linux servers run SELinux, and even SSHD does not allow root login. Automatic updates are turned off, but I keep a calendar reminder for manual updates on a regular basis. I have even tinkered with TripWire and AIDE, two intrusion detection systems (IDSs) for servers. I also run IDS software on my network gateway. I was paranoid even before Edward Snowden's revelations.

I have been trying to follow best practices for an in-depth defense of the network as well as the individual devices I have control over. That does not really include the IoT devices other than controlling how they connect to the Internet and the local network.

My network gateway plugs into my Verizon router, which has its own firewall and wireless access point but is located in our ranch house where it can provide limited support to the rest of the house. I have three other access points (APs) on my local network. Unfortunately, most of the IoT devices in the house need to use those APs. The APs run the open-source DD-WRT software that actually support VLANs and virtual access point support, making it possible to have an SSID for the wireless IoT devices to gain Internet access through a VLAN to my gateway. This at least isolates the remaining LAN from the IoT devices.

Does it sound complicated? It is even possible to get more complicated by having multiple AP/VLANs for different collections of devices. It is not something I would recommend to anyone without a lot of time, a network certification, or a lot of network experience.

What I am hopefully highlighting here is the need for security in depth within IoT devices and their frameworks. An enterprise may have the ability to lock down, partition, and track the security of its network, but it will be impossible for the average consumer or user to even come close.

It should be standard practice today to have different security checks for actions like over-the-wire updates, Web-based management, and command authentication. Passwords or keys should never be stored as cleartext, and security should not be an add-on or a simple firewall.

So consider why someone might want VPN or VLAN support on their IoT or network device. A hacked IoT device or application can be a gateway into a network of devices if those devices assume they are secured by a third-party firewall.

News & Analysis

IO-LINK REFERENCE
Design Kit Simplifies Relay of Diagnostic Information from Sensors

In order to meet the growing demand for larger automated electrical networks, Linear Technology Corp. has introduced a new IO-Link reference design kit, which includes the DC2228A 8-Port IO-Link master and DC2227A IO-Link device reference designs. This kit implements an IO-Link v1.1 physical interface, which sends diagnostic information on hardware, cables, and software directly to the control level of a system.

The IO-Link is a point-to-point serial communication standard that allows for the simplified integration of sensors and actuators to an automated control mechanism. Used primarily in complex automated systems, such as sorting and packaging technologies, an IO-Link design typically uses monitoring relays for preventative maintenance.

The IO-Link reference design kit implements an IO-Link v1.1 physical interface, which sends diagnostic information on hardware, cables, and software directly to the control level of a system. *(Image courtesy of Linear Technology Corp.)*

The DC2228A is a complete master reference design that can integrate eight IO-Link devices simultaneously with individual port isolation and fault reporting support. This technology reduces the number of integrated circuits (ICs) within the system, resulting in a more direct exchange of parameter data from sensors to switching devices. Operating on a pair of LTC2874 quad master integrated circuits, it features controller-protected L+ outputs with 200mA supply and 400mA start-up pulse capability and a 39V transient voltage suppression diode. It functions with an input range of 20 to 30 V and supports a Power over Ethernet (PoE) connection. With the master reference design implemented, the system can be monitored and controlled directly from software on a Windows PC or remotely through Ethernet for more accurate factory-floor simulation. The DC2228A supports all IO-Link devices, including the DC2227A, and can be modified based on operational requirements.

Compatible with all IO-Link master reference designs, the DC2227A is a complete IO-Link device reference design equipped with an onboard temperature sensor, light sensor, and incandescent lamp. It uses an LT3669-2 device transceiver integrated circuit that features ±60V line protection and high-current driving capabilities, a 300mA step-down regulator, and 150mA low-dropout regulator. Operating within an input range of 18 to 36 V supplied via an IO-Link cable, the DC2227A builds IO-Link device description files from service data and transfers them directly to the control software managing the system. These files also support both IO-Link and SIO device operation. ■

IF PERFORMANCE COUNTS, IF RELIABILITY COUNTS,

THEN YOUR CHOICE OF INTERCONNECTS COUNTS A LOT.

Maintaining a seamless, uninterrupted signal is the objective for any electronic system. Regardless of how elegantly designed or sophisticated the individual components might be, continuity and reliability will only be as good as the interconnects that link them. Machined pin receptacles from Mill-Max can make the difference between a product that only performs adequately and an exceptional product that lasts.

To learn more, download our new white paper at www.mill-max.com/design.

MILL-MAX®

Maximum Reliability.

MillMaxMfg MillMaxMfg Mill-Max Mfg. Corp. MILL-MAX MFG. CORP. • 190 PINE HOLLOW ROAD • OYSTER BAY, NY 11771 • (516) 922-6000

CHIPSET CONTROLS DEAD TIME to Reduce MOSFET Failure Risk

ACCURATELY CONTROLLING DEAD time in synchronous rectifiers helps minimize power loss and risk of MOSFET failure. With that in mind, Texas Instruments incorporated intelligent digital control and a body-diode sensing feature into its latest chipset to optimize secondary-side synchronous rectification in ac-dc and isolated dc-dc power supplies. The chipset comprises the UCD3138A digital controller and UCD7138 low-side gate driver.

The controller/driver combo uses body-diode voltage information with digital control algorithms to optimize dead time and compensate for power-stage component variations without calibration or screening during mass production. Intelligent sensing adjusts timing for minimal diode conduction to boost efficiency and reliability while eliminating the traditional signal-to-noise ratio channels of MOSFET VDS_ON sensing devices.

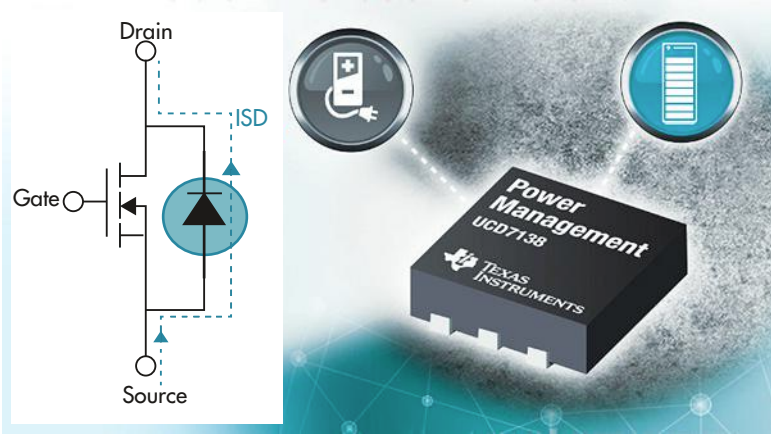

A new power-management chipset from Texas Instruments exploits intelligent digital control and a body-diode sensing feature that optimizes secondary-side synchronous rectification in ac-dc and isolated dc-dc power supplies.
(Image courtesy of Texas Instruments)

The gate driver features an asymmetrical, rail-to-rail, 4-A source and 6-A sink peak-current drive to support load range of a few hundred watts to a kilowatt (for multiple paralleled FETs). Operation maintains high efficiency at frequencies up to 2 MHz, thanks to the controller's hardware peripherals in addition to 14-ns propagation delays and fast rise/fall teams. A 3- by 3-mm QFN package helps conserve board space. ■

TELEMATICS WILL CONTINUE Driving M2M Market Forward

DRAWING FROM NEW research on the global technology industry, Juniper Research is predicting that the telematics sector will lead the machine-to-machine (M2M) market, in terms of revenue, until the end of the decade. Its report, "M2M & Embedded Devices: Strategic Analysis & Vertical Market Forecasts 2015-2019," suggests that the automotive industry will continue to expand its telematics capabilities over the next five years. At the same time, other industries will follow with increased efficiency and connectivity.

In the next five years, the M2M market is estimated to generate global service revenues of over $40 billion—more than double the amount generated in today's market. The automotive industry is expected to contribute significantly to that growth. Over the last several years, most telematics advancements have involved linking smartphones to the vehicle interface. But in addition to smartphone capabilities, Juniper forecasts that one in five vehicles will support Internet access and wireless local area networks (WLANs) by 2019.

The utilities industry has followed a similar path toward increased efficiency while forging connections between service providers and users. Juniper points to smart-metering technology as one of the most promising initiatives within the M2M market over the next five years. As governments begin to implement digital technologies and high-efficiency systems on a large scale, it will be necessary to reduce costs by matching resource consumption with the levels of production.

Smart meters are able to measure the consumption of electric energy in short intervals and communicate that data to a central system for analysis. This information will allow companies to adjust the amount of electricity distributed at certain times of the year and even to certain households, using the smart meter as a remote control mechanism.

According to Juniper Research, smart metering initiatives underline the increased role of Big Data analytics in the M2M market. The insight gained through analytics will help to predict areas of growth within emerging forces in the market, such as agricultural and environmental applications for digital technology. ■

Touch. Discover. Solve.

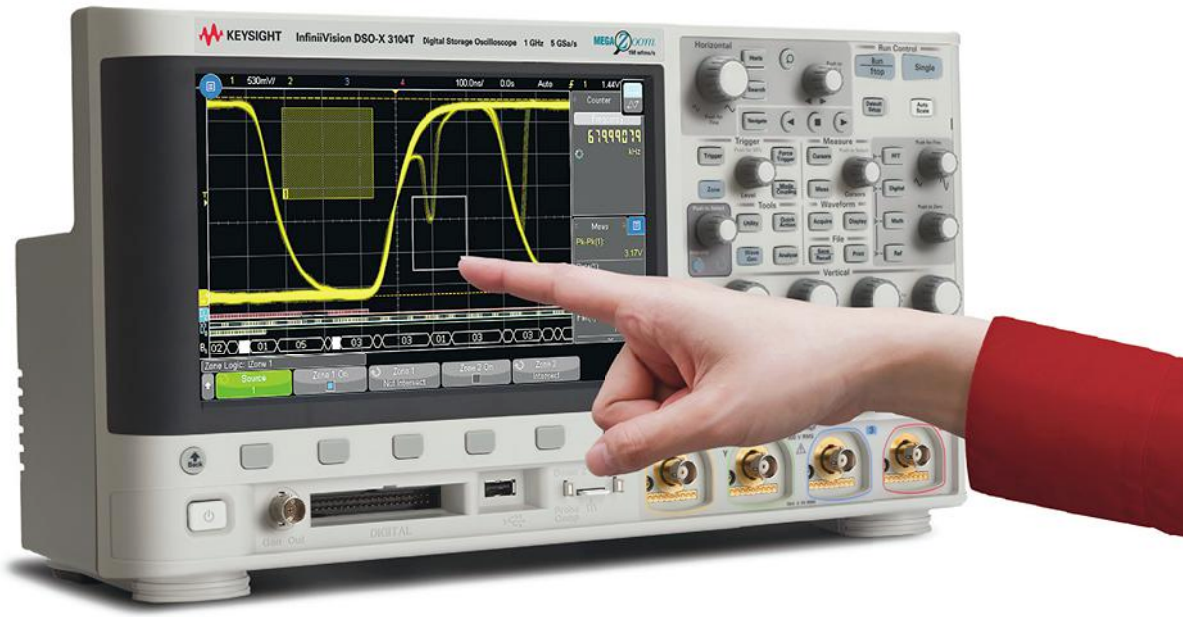

Oscilloscopes redefined.

We've improved on the proven. The new Keysight 3000T oscilloscope is the next-generation of the InfiniiVision X-Series. With its zone touch triggering, you can trigger on any signal in just two steps. So you can isolate a signal in seconds—much faster than with any competing scope. The 3000T is also a 6-in-1 instrument. Along with your oscilloscope, you can get an MSO, WaveGen function generator, protocol analyzer, DVM and counter. Get in touch with the future of scopes. Take the Trigger Challenge today.

Models	Keysight 3000T X-Series*	Tektronix MDO3000 Series**
Bandwidth	100 MHz-1 GHz	100 MHz-1 GHz
Update Rate	Uncompromised 1M wfms/s	Up to 280K wfms/s
Touch Screen	8.5-inch Capacitive	No
Zone Touch Triggering	Yes	No
Sample Rate	5 GSa/s	2.5 GSa/s (>= 500 MHz) 5 GSa/s (1 GHz)

Take the Trigger Challenge at:
www.keysight.com/find/triggerchallenge

USA: 800 829 4444 CAN: 877 894 4414

KEYSIGHT TECHNOLOGIES

Unlocking Measurement Insights

* Refer to Keysight document 5992-0140EN for product specs, and 5989-7885EN for update rate measurements.
** Competitive oscilloscopes are from Tektronix publication 48W-30020-3

NON-VOLATILE NRAM Nanotubes Deliver Endurance and Performance

NANTERO'S nanotube-based NRAM technology has an impressive feature list. It has the read/write speed of DRAM with flash memory's non-volatility, but with significantly higher endurance. It can retain data for more than 1,000 years at normal temperatures. It is also amenable to the latest CMOS technology scaling below 5 nm as well as supporting MLC (multi-level cell) and 3D structures. Best of all, it works with current CMOS fab technology.

NRAM is based on film of crossed nanotube (CNT) deposited on a silicon substrate. Each cell can be switched between a low and high resistive state (*Fig. 1*) that is stable in excess of 300° C.

The CNT layer is etched using standard lithography techniques. A single cell has hundreds of nanotubes (*Fig. 2*). They are scattered, but proper alignment is required for it to act as a programmable storage device. Changing the cell's state requires application of a programming voltage for a few picoseconds. Experimental chips have reached 1012 write cycles and 1015 read cycles.

NRAM addresses almost all aspects of storage, given its potential scalability and performance. It is especially well suited for embedded applications because it is compatible with conventional CMOS technology, allowing it to be employed in microcontrollers. Its wide operational temperature range will allow its use

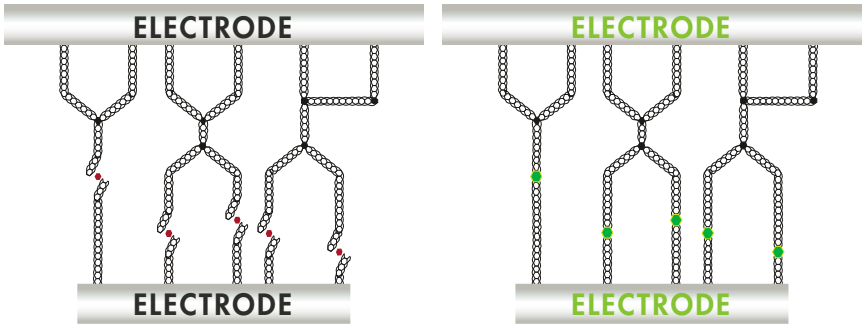

1. The crossed nanotube (CNT) fabric can be changed between high and low resistive states.

in rugged applications like automotive and aviation.

NRAM—like previous technologies such as FRAM, MRAM, and phase-change memory—has the potential to unseat the existing DRAM/flash combination that currently dominates computer technology. It remains to be seen whether it will be a winner or a niche player, but it is looking good so far. Nantero is currently licensing its technology. Initial chips are expected to be compatible with DDR4 interfaces. ■

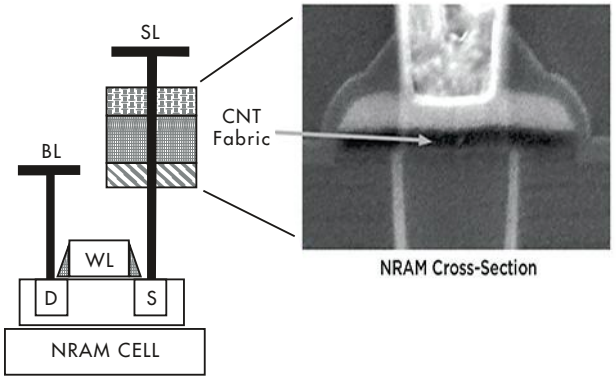

2. A typical NRAM cell has hundreds of nanotubes.

edX, QUALCOMM Launch Mobile Education Initiative

edX, A NON-PROFIT ONLINE EDUCATION ENTERPRISE, has partnered with QualComm Education Inc., a subsidiary of QualComm Inc., to expand its open-source courses onto mobile platforms.

Founded by Harvard and the Massachusetts Institute of Technology, edX provides online courses from accredited universities on an open-source platform, which allows the distribution of educational software to students. These programs, called Massively Open Online Courses, also feature a social component so that students are able to connect with each other for educational purposes.

edX is attempting to give students instant mobile access to both educational materials and social learning tools. In addition to supporting traditional methods of instruction, edX will pursue a more interactive, game-like approach to education.

According to Anant Agarwal, CEO of edX and professor at MIT, "The ability for edX students to take their coursework with them, access it as they go about their day, use collaborative tools to enhance student engagement, and mobile devices to create new learning experiences, is the next frontier for learning worldwide." ■

HYPER-REGISTERS Boost Throughput in Monolithic FPGA Chip

ALTERA'S STRATIX 10 IS its top-of-the-line FPGA, built using Intel's 14 nm Tri-Gate 3D transistor technology. It incorporates a host of new technologies designed to deliver twice the performance in a more secure fashion while using less power per function.

The Intel connection is not surprising given Altera's and Intel's previous work such as the Intel E600C system-on-chip. It packed a 40-nm Altera Arria II FPGA with an Intel Atom core into a 37.5 mm × 37.5 mm BGA package. And now, of course, Altera is going to be inside Intel (*see "Intel Really Set to Buy Altera FPGAs" on electronicdesign.com*) if everything pans out.

Stratix 10 can be delivered with more than 5.5 million logic elements (LEs) or adaptive logic modules (ALMs) in Altera-speak. This monolithic chip can deliver 10 TFLOPS of computation using hard DSP blocks and its interfaces can deliver 1.2 Tbits/s.

One way for Altera to hike performance was to spin the high-speed interfaces onto separate silicon tiles. The 3D System-in-Package (SiP) uses an embedded multi-die interconnect bridge (EMIB), also known as a silicon interposer, that has through-silicon vias (TSVs). TSV technology has been used in other FPGAs and ASICs. It tends to be used in larger chips and ones where different technologies are needed. Multichip packages are another alternative, but the interconnects tend to be less numerous and less efficient. The approach allows Altera to utilize 144 SERDES operating at speeds up to 30 Gbits/s. The approach is not limited to a particular SERDES. Look for future tiles with Ethernet, PCIe Gen 4, 56 Gbit/s, PAM-4, and optical transceiver technologies.

Altera's new HyperFlex technology (*see figure*) highlights another change in design to support high throughput. The change is critical because chips are larger and a signal can rarely get from one side of the chip to another as speeds increase. It is just a physical limitation. HyperFlex adds a host of Hyper-registers to facilitate pipelining of data through the FPGA connection fabric. Stratix 10 has 10 times the number of Hyper-Registers in the interconnect fabric than the previous FPGA designs.

Hyper-Registers are simply buffers that can be placed along a signal path to save data as it flows from a source to a sink. Additional buffers increase the number of steps, but allow the overall clock rate to increase, thereby increasing overall system throughput.

Software optimizations using this technique include Hyper-Retiming, Hyper-Pipelining, and Hyper-Optimization, which are part of the new SpectraQ engine in Altera's Quartus II FPGA design tool. The addition of more Hyper-Registers increases system complexity, but the tools make their use transparent to designers so design time is actually reduced.

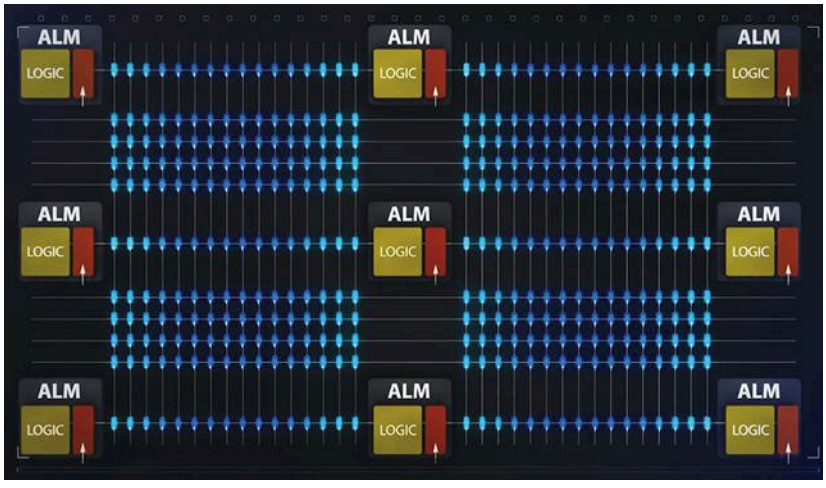

Altera's Stratix 10 has more than 10 times the number of Hyper-Registers in the interconnect fabric than the previous FPGA designs.

The approach can also help reduce power requirements. Some applications see as much as a 70% power reduction due to this and other enhancements in the Stratix 10 family. For example, sometimes the bus width can be cut in half while doubling the clock frequency, resulting in a 50% space reduction.

The Secure Device Manager (SDM) is part of Altera's security support. This includes sector-based authentication and encryption allowing sections of the FPGA to be programmed using different security methods and keys. SDM supports multi-factor authentication and incorporates Physically Unclonable Function (PUF) technology. ARM Trust Zone technology is included with hard-core processor support.

Stratix 10 targets high-performance applications so the 1.5 GHz, quad-core Cortex-A53 is a good match for programmatic application aspects. The 64-bit cores include NEON SIMD and floating point support. Four cores share a 1 Mbyte L2 cache with ECC support.

The FPGAs are also designed to work with Altera's Enpirion power solutions. The EM1130 PowerSoC supports the Stratix 10 SmartVoltage ID power management. ■

The Original Electronics Parts Search & Procurement Tool

Search Parts

Part # | Manufacturer

SEARCH PARTS FAST

Your Trusted Source

www.SourceESB.com

WILLIAM WONG | Technology Editor

New technologies address the universal challenge of having enough storage space.

High-Density Storage

One can never have enough storage, and chip designers have been working hard to deliver as much storage as possible. Almost every avenue is being pursued, from Nantero's nanotube-based NRAM to AMD's 3D high-bandwidth memory (HBM) chip architecture (*Fig. 1*). New technologies like NRAM, MRAM, and FRAM deliver non-volatility to challenge not only flash memory, but DRAM.

Conventional single-die chip storage continues to be replaced by 3D multi-die (*Fig. 2*). It pays to limit interconnect runs as speeds and capacity increase rapidly. Moving storage closer to the processor is important to performance, and it is occurring at all levels. For example, Diablo Technologies' Memory Channel Storage moves flash memory onto the DRAM memory channel.

3D ARCHITECTURES

Using multiple dies in the same package allows a designer to easily mix different chip technologies. It is also a way to increase the packing density to improve capacity. Furthermore, it can boost performance by reducing interconnect length and allowing higher-speed transfers, since designers have better control over the connections.

Multiple dies on a chip have been commonly available for many years, but they were often used in areas where their higher cost could be justified by the added performance and reliability. A chip carrier would typically be wired internally in a fashion similar to connecting a die to external pins.

Die stacking has also become common because of the tight spaces required by smartphones and other mobile devices. Microcontrollers incorporate flash memory and RAM on-chip, but these are often separate chips for more powerful microprocessors. Stacking a flash memory and DRAM die on top of a microprocessor die cuts down on space while providing a more reliable platform. Again, the wiring between chips is similar or identical to the chip-to-carrier connection.

The use of silicon interposer layers with through-silicon vias (TSVs) lets designers drop down multiple chips, but the approach is more akin to a silicon circuit board than the multichip carriers. The latter requires wiring to connect the chips together and to the outside world.

The use of interposers has been pushed by FPGA vendors such as Xilinx and Altera. Xilinx originally called the technology 2.5D (*see "10,000 Connections Between FPGA Slices" on electronicdesign.com*),; however, these days everyone refers to it as 3D. Xilinx initially used this approach to provide more logic elements using multiple FPGA die. The company also paired it with high-speed SERDES die that employed different silicon technology.

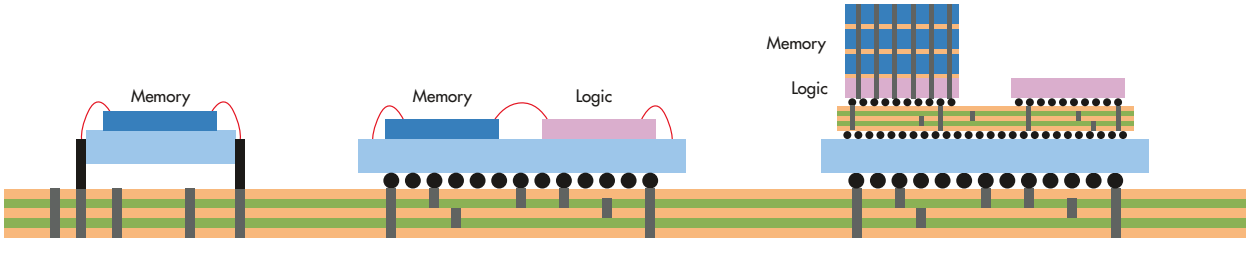

2. Solid-state storage has moved from memory chips (left) to multicarrier packages (middle) and now to 3D chips (right). The 3D chips essentially incorporate a multilayer circuit-board approach to provide an on-chip interconnect.

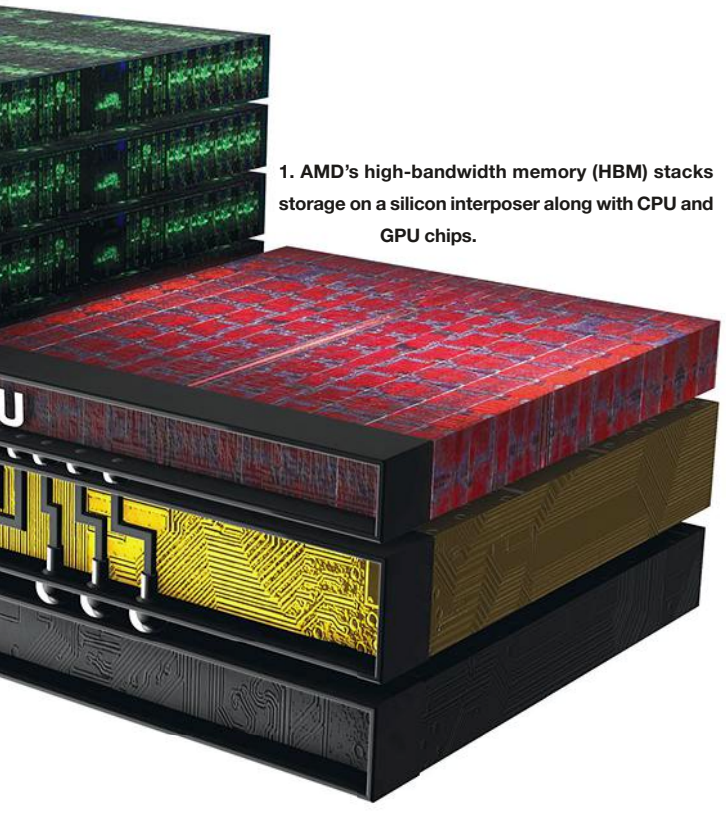

1. AMD's high-bandwidth memory (HBM) stacks storage on a silicon interposer along with CPU and GPU chips.

Altera's Stratix 10 uses a similar approach, but delivers 5.5 million logic elements and a quad-core Cortex-A53 complex on a single die. Altera's SERDES die are connected using the interposer layer that the company calls the Embedded Multi-die Interconnect Bridge (EMIB). The main difference is that the EMIB is a small set of interposers that link the FPGA die to adjacent dies referred to as Tiles (*Fig. 3*). Altera plans on having different I/O Tiles for application-specific interfaces.

A number of vendors are doing 3D on the vertical side. The Hybrid memory cube (HMC) addresses the RAM side of storage (*see "Hybrid Memory Cube Shows New Direction for High-Performance Storage" on electronicdesign.com*). It stacks die using TSV connections. The base layer is the controller that has multiple, high-speed SERDES linking it to one or more hosts. The HMC specification addresses a range of configurations from a simple single-port connection to one host to multiport connections (*Fig. 4*). It targets enterprise servers, where the demand for RAM is insatiable.

There's an equally unquenchable appetite for flash storage, too. Single-level cell (SLC) has given way to 2-bit multi-level cell (MLC) and 3-bit triple-level cell (TLC). The tradeoffs between SLC, MLC, and TLC tend to be space and capacity versus write speed and write lifetime.

Samsung V-NAND, also known as vertical NAND, is a 32-layer TLC device (*Fig. 5*). It provides the highest storage density at this point (*see "Mass Production Underway for 32-Layer 3D V-NAND Flash" on electronicdesign.com*). Identical dies are stacked on top of each other. Samsung's SSD 850 EVO line is the initial home for V-NAND. This includes an M.2 and mSATA versions. The latter has a capacity up to 1 Tbyte using a 6-Gb/s SATA interface.

AMD's HBM combines the interposer and vertical stacking. CPU and GPU dies are single-layer types, but the memory is a stack of multiple RAM die. HBM's advantage over HMC is that it does not have to transfer data using high-speed SERDES. Instead, data can be moved in parallel.

HBM will first find a home in AMD's Radeon R9 Fury X GPU card. This will feature 4 GB of HBM plus the latest Fiji GPU from AMD with 4,096 stream processors, 64 GCN (graphics core next), 128 render output units, and 256 texture mapping units.

The memory interface is 4,096 bits wide. The GPU delivers 8.6 FLOPS while only having a TDP of 300 W. This is only 10 W more than the earlier Radeon R9 290X, yet performance improves by more than 50%.

NONVOLATILE MEMORY TECHNOLOGIES

Packaging technologies are one way to improve storage density and performance. Changing the storage technology is another that is complementary to the packaging. Quite a few technologies have been around a while and found a niche, as have emerging technologies. All look to take away all or part of the market from SRAM/DRAM and flash memory. These technologies include various resistive RAM approaches, FRAM, and MRAM.

Nantero's NRAM is a new resistive memory technology that uses nanotubes (*Fig. 6*). It meshes well with current CMOS fab technology and offers impressive retention and performance specifications. It should retain data at room temperature for over 1,000 years. It can handle extended temperature ranges with lower retention capability. It should scale to 5 nm and can handle 3D structures and MLC support. It is based on a crossed nanotube (CNT) film deposited on a silicon substrate.

Crossbar's resistive RAM is similar (*see "RRAM Challenges Flash Storage" on electronicdesign.com*). It uses a switching medium that includes metal filaments with metal ions migrated from one end to another to change state. It has a 10-year retention and support for MLC and TLC techniques as well as 3D stacking. It is 20 times faster than flash, and features endurance support of over 10,000 cycles.

FRAM has been showing up in quite a few of Texas Instruments' (TI) MSP430 microcontrollers. It has replaced SRAM and flash storage on some models. TI also has a set of routines

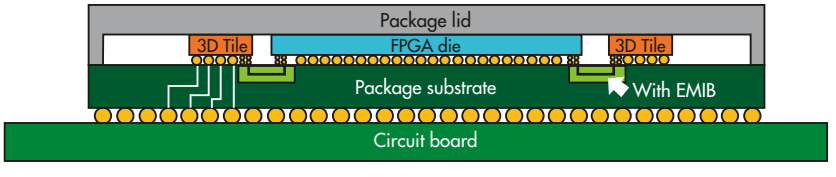

3. High-end FPGAs, such as Altera's latest Stratix 10, are putting the interface logic on separate TILEs. These die are connected using an interposer layer akin to a silicon circuit board.

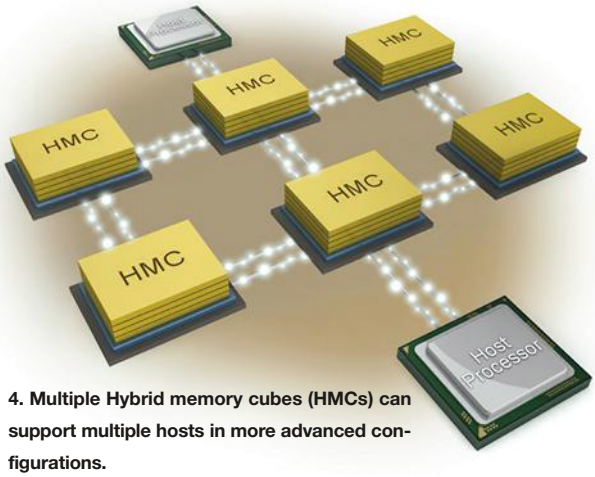

4. Multiple Hybrid memory cubes (HMCs) can support multiple hosts in more advanced configurations.

that can save and restore system settings to FRAM that are not stored in the memory. This can include system settings like status and control registers for peripherals.

The Compute Through Power Loss (CTPL) utility API allows an almost instant-on capability, since the application stack and status are kept in the nonvolatile FRAM. A sleeping system automatically restores these settings upon power-up.

Everspin MRAM is similar to FRAM functionality, but uses a much different storage mechanism for its magnetic tunnel junction (*see "Q&A: Everspin Takes MRAM to the Mainstream" on electronicdesign.com*). It has a data retention of more than 20 years with a 35-ns read/write time. It offers unlimited endurance, making it a good choice for a number of applications like RAID controller cache. The latest products employ Spin-Torque MRAM (ST-MRAM), which allow it to be used in DDR3 systems. Everspin also has parallel and SPI MRAM products. Chip capacity is 64 Mbits.

Phase change memory (PCM) originally utilized behavior of chalcogenide glass. Micron has phased out its commercial PCM line based on technology acquired from Numonyx. Research is ongoing, so the technology may resurface in the commercial space in the future.

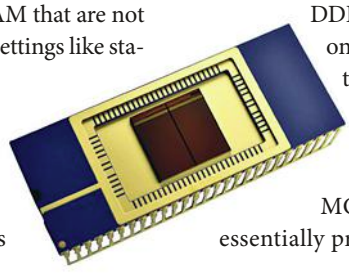

5. Samsung's V-NAND chip uses 32 layers of flash-memory die stacked on top of each other.

NEW FLASH CONFIGURATIONS

DRAM and flash memory in general remain the dominant force in storage. Flash memory has pushed rotating magnetic media into more niche markets as end-user devices have moved to mobile platforms requiring more compact and power-sipping flash storage. Large hard disks are still part of enterprise solutions, set-top box recorders, and PC backup, but flash memory has moved from its niche into the primary nonvolatile storage mechanism in most areas. Flash memory's ability to fit into almost any space makes it a much better choice than a hard disk for a growing number of applications, including wearable tech (*see "Wearing Your Technology" on electronicdesign.com*).

Diablo Technologies' Memory Channel Storage (MCS) moves flash memory to the memory channel of high-performance microprocessors. It licenses its technology to companies ranging from SanDisk to IBM that deliver DDR3 and DDR4 modules competing with DRAM for slots on the motherboard. The DIMM modules use the same interface, but the transactions are more like flash-memory storage devices rather than DRAM since the write speeds for flash are significantly slower.

MCS storage can be utilized as a block device, essentially providing a faster transfer rate to the system since the memory channel has a higher bandwidth than other I/O interfaces and there's less system overhead. MCS storage has a smaller transaction size than block devices that are designed to accommodate blocks of data on rotating magnetic media. The C2 version of MCS has a feature called NanoCommit that moves the minimum transfer size to a CPU cache line (*see "Memory Channel Flash Storage Provides Fast RAM Mirroring" on electronicdesign.com*). This allows mirroring of key data structures, opening the door to significant performance improvements for applications like database servers.

DRAM still tends to be faster and larger than other technologies to date. It is one reason that nonvolatile DRAM has emerged in enterprise environments where reliability is critical. Viking Technology's ArxCis-NV is a hybrid DIMM with DRAM and flash memory on-board (*see "Nonvolatile DIMMs and NVMe Spice Up the Flash-Memory Summit" on electronicdesign.com*). The flash memory is used to retain the DRAM contents when the system is powered down. A supercap provides power in the event of system power loss, enabling the on-board controller to copy DRAM to flash memory. DDR3 and DDR4 versions are available.

There is never a dull moment when it comes to storage technology. The mix of new storage technologies and packaging techniques provides designers with a wide range of options.

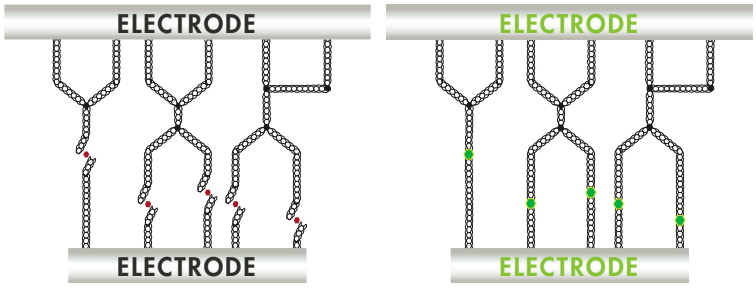

6. Nantero's NRAM uses nanotubes that can be modified to different resistances.

10 MHz Rubidium Frequency Standard

- **5 MHz and 10 MHz outputs**
- **Ultra-low phase noise (<–130 dBc/Hz at 10 Hz)**
- **0.005 ppm aging over 20 years**
- **Built-in distribution amplifier (up to 22 outputs)**
- **1 pps input and output**

The FS725 Benchtop Rubidium Frequency Standard is ideal for metrology labs, R&D facilities, or anywhere a precision frequency standard is required.

It has excellent aging characteristics, extremely low phase noise, and outstanding reliability. A 1 pps input is provided for phase-locking to GPS, providing Stratum 1 performance.

With a built-in 5 MHz and 10 MHz distribution amplifier, the FS725 is the ultimate laboratory frequency standard.

Allan Variance vs. Time

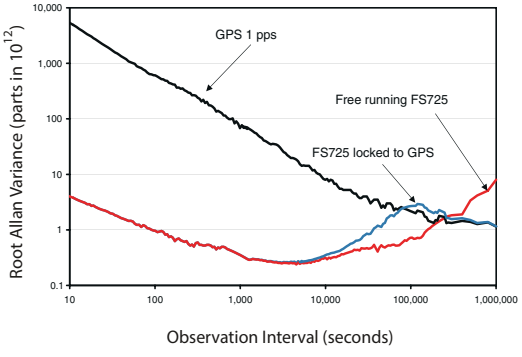

FS725 Single Sideband Phase Noise

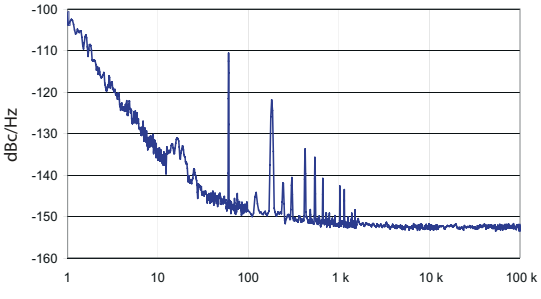

FS725 rear panel

SRS Stanford Research Systems

Phone: (408)744-9040
www.thinkSRS.com

RON LOWMAN | Synopsys

Optimized IP Fosters Energy-Efficient IoT Chip Design

By adopting IP tailored specifically for IoT applications, design teams can reduce power, add connectivity, improve sensory and communications interfaces, and ultimately achieve design goals faster, with less risk.

The growing Internet of Things (IoT) market continues to drive significant investment in new products, including smarter versions of existing products like the connected light bulb, or completely new products such as connected drones for videography, surveillance, and much more. IT research firm Gartner expects 10 billion shipments of these types of products will flood the market by 2020; however, the market remains fragmented.

One common trend in this arena involves adoption of new techniques to lower power consumption of the principle ICs providing the intelligence for these systems. These techniques go beyond those used in today's mobile, auto, and PC peripheral systems-on-a-chip (SoCs). The new levels of reduced power make it possible to fit into smaller spaces, shrink overall system cost, and increase overall battery life.

IoT POWER SOURCES

The trend toward larger mobile-phone form factors enables the use of bigger, higher-capacity batteries, relieving some of the pressure on design teams to optimize the SoC for the lowest possible power consumption. The automotive market has the benefit of using a 12-V lead-acid battery as a power source, and USB-tethered devices are just that, tethered. Unfortunately, many IoT applications don't have this luxury.

IoT applications are looking at either 25 years of battery life or very size-constrained spaces with limited room for batteries, while still meeting the expectation of adding elements such as wireless connectivity. So, while battery technology is slowly improving, minimization of the SoC's power consumption as well as the system architecture will be the keys to making IoT designs "special." For instance, a 570-mAh battery is estimated to keep Google Glass running for about a day, a 300-mAh source powers a Samsung smart watch for about two days, and a similar battery is good for up to 22 days when paired with a hearing aid.

SAVING ENERGY IN IoT DEVICES

Figure 1 below compares a typical energy use case for a generic IoT application, shown in green, with the desired profile, shown in purple. The goal is to reduce the entire area under the rectangles. This requires both reducing the peak usage of each function (amount of real work) and the time required to perform the task.

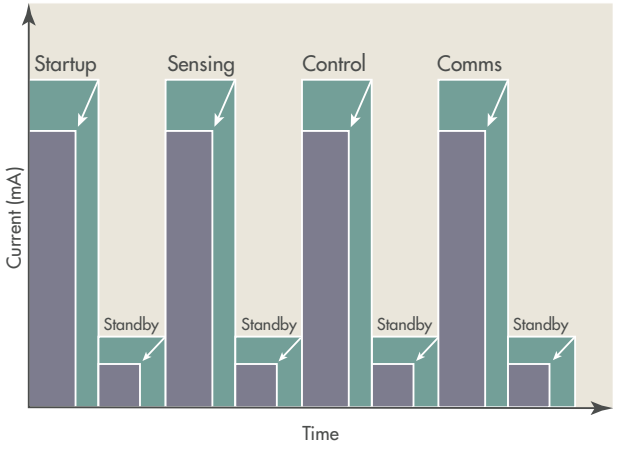

1. Function usage and the time required to perform a task are key factors in this common IoT energy use-case profile.

It's clear that IoT designs must go far beyond the advances made thus far by the mobile market to meet the power requirements of next-generation products. However, different applications require very different power use cases. Some use cases are dominated by "always on" functions, such as voice activation and power supplies. Others are off 99% of the time, such as motion, fire, and other environmental detection, which means a preponderance of standby power usage, whereas others fall somewhere in between these two extremes. Each use case requires different considerations to minimize power.

If the device operates in such a way that dynamic power is the biggest contributor, an effective strategy is to reduce V_{DD} as much as possible. If static power dominates, controlling the leakage per operation becomes more important.

New investments are already benefiting SoCs for IoT applications. For instance, IoT-specific process technologies developed by foundries support lower leakage and lower voltages, yet still meet processing performance demands.

On the design side, the mobile community and others have taken advantage of techniques such as clock and power gating, multi-voltage domains, dynamic voltage and frequency scaling (DVFS), and back-biasing. IoT teams are adding smart biasing, sub-threshold, and near-threshold design, and making more extensive use of multi-voltage, shutdown, and power domains, to accomplish very aggressive goals. Effective tools are required to ensure consistent interpretation of power intent, power network synthesis, and in-design rails, as well as deliver optimal low-V_T usage and comprehensive power-aware formal and static verification.

ENERGY-EFFICIENT IP FOR THE IoT

Advances, innovations, and new investments beyond the mobile market in both process technologies and design techniques will absolutely drive new IoT products closer to reaching power-consumption goals. However, it will be differentiated IP and integrated IP subsystems developed specifically for IoT applications that will truly create viable, competitive SoCs. Among these:

Logic libraries, embedded memories, and non-volatile memories. In the mobile domain, low-power logic libraries and memory compilers are essential ingredients in producing efficient SoC implementations. In the IoT world, design teams supplement this with advanced low-power options, including always-on libraries, lower-voltage memories and libraries, multichannel lengths, multi-bit flip-flops, and power-optimization kits.

Ultra-low-power non-volatile memory (NVM) IP is being used to supplant more power-hungry options. ROMs now store mature code, beyond just boot functions such as USB and Bluetooth stacks, thus minimizing code storage costs and lowering power versus alternative options. A number of integrated test-and-repair solutions now support embedded flash to speed development, which increases test coverage, provides in-field diagnostics, and decreases costs

Efficient processor cores and IP subsystems. One of the most effective ways to reduce power usage is to shorten the time it takes to perform tasks. Using co-processors and dedicated hardware to perform specific functions will shrink the overall time.

Frequency and the amount of memory required to support functions also contribute heavily to power usage. Higher-performance processors and tightly integrated subsystems, in terms of the amount of work completed per cycle, can improve overall system and design costs. Co-processors, more efficient processors, and customized hardware will provide huge benefits by reducing frequency and memory requirements of functions such as voice, vision, motor control, power conversion, filtering, and audio.

Today, designers tackling IoT applications make more use of integrated

IP subsystems that significantly reduce wait states within the architecture and optimize the overall die area needed to perform the tasks. For applications that require advanced math and very tight timing control, or must support "always-on" functions, improved processing capabilities (beyond just increasing frequency and pipelines) can dramatically enhance power consumption. Application-specific analog interfaces also are able to be used to manage functions for metering, medical, and power-management applications. They help better control the system and increase efficiency via higher resolution and higher conversion rates.

Interface IP. Standard interfaces designed for mobile and networking markets are being effectively used by IoT design teams. These include low-power RF solutions, DDR, MIPI for camera and display, Ethernet, and USB supporting battery charging and power-down features.

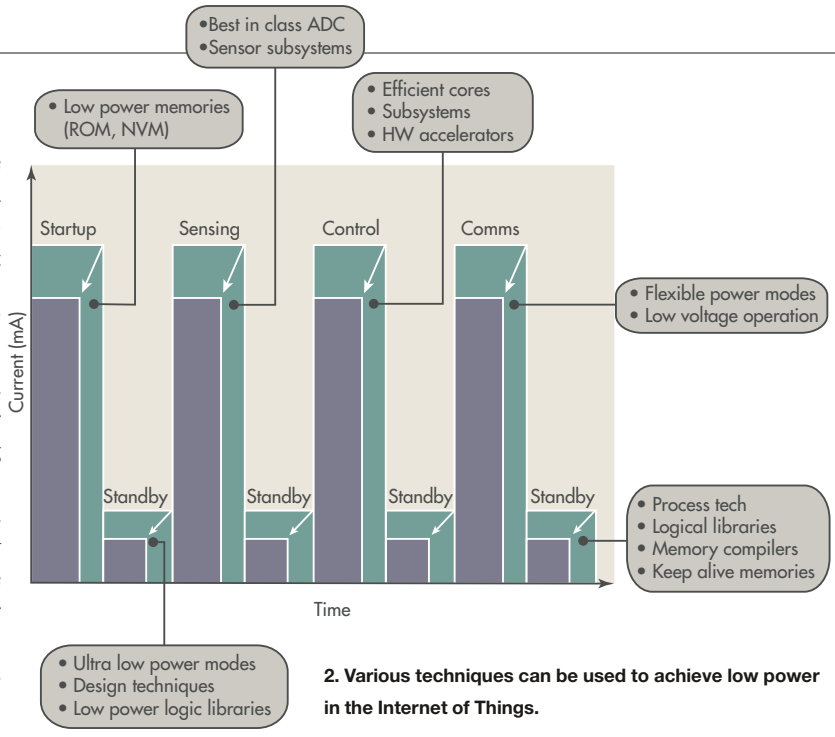

2. Various techniques can be used to achieve low power in the Internet of Things.

To further reduce peak current levels and task time, the industry is implementing more new wake-up features, flexible sleep modes, and interoperable low-bandwidth wireless radios that minimize the overall payload required to transmit.

Figure 2 summarizes some key techniques for achieving lower power in specific IoT functions:

- Support boot loaders via ROMs, rather than fetching from traditional methods.
- Support options during startup to lengthen the time to load memory to reduce startup currents.
- Support of more flexible low-power modes ensures that only the necessary SoC modules are operating, since this can change during the use case. For instance, during communication functions, sensing and measurement may not be feasible or desired and vice versa.
- During sensing and measurement stages, reduce the number of cycles and wait states. This is accomplished via improved analog measurement capabilities (higher resolution and/or higher conversion rates) and a reduced number of cycles to process the data (tightly coupling these interfaces with efficient and flexible processing cores in integrated subsystems). Processing can always be done by using a processor with a longer pipeline or increasing the frequency, but both of these options increase power consumption. Therefore, it's best to use processors that complete the most work per cycle at the lowest frequency.
- Process complex math at minimum frequencies, in the least amount of time. When combined with enhanced sensing capabilities, this efficient processing reduces system mechanical costs on top of the PCB-related system savings. A good example is the implementation of vector-oriented control for motors, which has improved control and reduced the size of

the motors over the past few decades. It should benefit new applications such as quadcopter drones, as well as traditional pumps found in appliances, utilities, and factories.

- Communicating at the proper time and with the appropriate amount of data payloads enables the application to right-size the power usage. This is being done with newly introduced low-bandwidth wireless standards.
- Employ the proper sleep states, with only the necessary "always-on" circuits, and minimize leakage with a low-power IP methodology.

Solving power issues isn't the only driver for IoT applications. Design teams also face challenges such as increasing integration to lower system costs, adding connectivity and security, and simplifying supply costs, while improving ease of use, meeting tighter project schedules, and getting to market first.

Though IP requirements can vary greatly depending on the specific application, it's clear that IoT design teams need low-power, robust, and proven IP solutions that help them achieve silicon success and quickly get their products to market.

SUMMARY

The Internet of Things is spawning hundreds of new products and applications for wearable and machine-to-machine markets. These applications demand innovative design techniques, the use of established and advanced process technologies, and most importantly, innovative low-power IP to reduce power consumption and increase battery life.

In turn, many chip companies are exploring and/or investing in new business opportunities offered by the IoT market. However, these investments need to include the adoption of IP designed specifically to meet IoT application requirements. This will accelerate design development and get products to market faster, with significantly less risk. Risk management of SoC power usage is critical for IoT devices.

By adopting IP tailored specifically for IoT applications, design teams are able to reduce power, add connectivity, improve sensory and communications interfaces. In the end, they will be able to achieve design goals at a faster clip, and with less risk. ed

RON LOWMAN is the strategic marketing manager for IoT at Synopsys. Prior to joining Synopsys, he spent 16 years at Freescale, where he was most recently the industrial business development manager for the company's Microcontroller Division as well as the product marketing manager for industrial MCUs and digital signal controllers. Ron holds a Bachelor's of Science in electrical engineering from The Colorado School of Mines and a Master's in business administration from the University of Texas in Austin.

PARKER DORRIS | Senior Applications Engineer
Microcontroller Products, Silicon Labs, www.silabs.com

Smart Capacitive-Sensing Design with 8-bit MCUs

While drop-in, fixed-function, capacitive-sensing solutions are readily available, general-purpose 8-bit MCUs with cap-sense capabilities can provide a more versatile approach that helps reduce power consumption, board size, and BOM cost.

Adding capacitive sensing to a consumer or industrial product can be a daunting challenge that requires developers to maintain sensor robustness and responsiveness while minimizing current consumption and addressing other system-level design priorities. Navigating this embedded development process without the help of tools and support can lead to lengthy trial-and-error design iterations, both in software and hardware, resulting in slipped release schedules and suboptimal solutions.

While fixed-function capacitive-sensing solutions can ease some of these design burdens, those devices are rarely pure drop-in solutions, as many still rely on developer-side configuration and calibration. In addition, fixed-function solutions prevent developers from giving the sensor additional responsibilities in a system, which can help reduce system current consumption, board size, and bill-of-materials (BOM) cost.

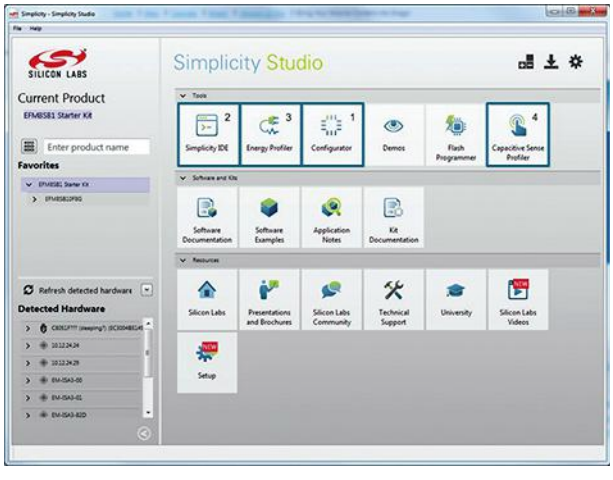

1. Annotations on the launcher window show the basic development path followed by many users.

Capacitive-sensing solutions based on general-purpose 8-bit microcontrollers (MCUs) can help designers through each step of the development process, enabling them to generate projects using capacitive-sensing firmware libraries, add and debug features with an intuitive and powerful integrated development environment (IDE), and view real-time capacitive-sensing performance in-system using capacitive-sensing visualization tools.

CAPACITIVE-SENSING DESIGN BASICS

Capacitive-sensing technologies measure the capacitance of an electrode connected to the sensor's input. The electrode being measured is often a printed-circuit-board (PCB)-designed round pattern of solid copper about 10 mm in diameter that's isolated from the board's ground. A thin overlay, usually made of a type of plastic or glass, is sometimes adhered to the PCB, and the resulting system can function as a touch interface that doesn't require mechanical buttons.

When a user touches the area of the overlay covering the electrode, the capacitance of the human body affects the capacitive coupling of the electrode, resulting in a change in the on-chip capacitive sensor's measured capacitance. Post-sample processing on the stream of samples will detect this change in capacitance and qualify it as a "touch."

A common goal among all capacitive-sensing technologies is the ability to sense a relatively small change in an analog signal while operating in a mixed-signal system that has the potential to be electrically noisy. As a result, designing a robust and reliable sensing interface can be challenging.

In addition to hardware considerations regarding board layout and electrode design, the sensor's post-sample processing operations must update state variables while responding dynamically to system-level events. Touch-sensing strategies usually involve maintaining a baseline, which is the expected output of the capacitive sensor when the electrode is in an untouched state, as well as setting one or more touch thresholds relative to that baseline.

Processing compares these thresholds against sensor output to determine the occurrence of touch and release events. In addition to these operations, the firmware must be designed to operate as efficiently as possible to minimize both code size and average current draw.

To help ease the burden of firmware development, many capacitive-sensor vendors offer fixed-function devices, which provide an integrated circuit that samples capacitance and outputs touch qualification through either a serial interface or port-pin logic states. While these products may appear attractive to developers because of their ease of use, the lack of programmable flash and limited customizable feature sets can force system developers to place more responsibilities for touch qualification onto more power-hungry host processors.

For example, a fixed-function device (FFD) that becomes susceptible to false-positive touch events in high-interference environments could force a host processor to ignore touch-qualification information from the FFD and instead read the FFD's raw data and develop its own touch-qualification algorithms. Similarly, an FFD erroneously sending frequent touch-event signals to a host processor will cause that processor to wake from its lower power state to examine touch events, driving up average current draw for a system.

A capacitive-sense-enabled microcontroller, such as the EFM8SB1 8-bit MCU from Silicon Labs, when used with a capacitive-sensing firmware library, combines a fixed-function device's ease of use with the flexibility of a general-purpose mixed-signal MCU. The precompiled library provides touch-qualification algorithms, state variable maintenance, and a simple application programming interface (API) for retrieving touch-event information. The remaining flash memory and numerous on-chip peripherals give developers the opportunity to add more features and responsibilities to the MCU.

A cap-sense firmware library for an 8-bit MCU leverages the device's power management unit (PMU) and on-chip real-time clock. This allows the developer to create a sophisticated, mode-switching MCU that can achieve less than 1-μA average current draw in most use cases.

SIMPLIFY CAP-SENSE DESIGN WITH AN IDE

An IDE, such as Silicon Labs' Simplicity Studio platform, provides end-to-end development support for capacitive-sensing-enabled embedded-system designs. This ranges from configuring and importing the capacitive-sensing firmware library, to in-system debugging and real-time visualization of capacitive-sensing output. Many IDEs provide a launcher window that dynamically populates tiles depending on the MCU being used in development. It will show the tools offered to support every step of development.

2. A configurator tool provides users with a graphical interface and properties windows. Customers can designate port pins as capacitive-sensing inputs and configure library performance settings.

The annotations on the launcher (*Fig. 1*) show the basic development path followed by most users:

- Configurator generates the capacitive-sensing library and allows developers to configure hardware peripherals.
- The IDE provides a code development and debug environment where additional features can be added to a firmware project.
- An energy profiler tool built into the IDE provides real-time current-draw measurements of a project that's been downloaded to a starter-kit evaluation board.
- A capacitive-sense profiler tool displays runtime capacitive-sensing data raw values and algorithm-derived values such as touch state, baselines, and thresholds.

CONFIGURING CAPACITIVE-SENSING PROJECTS

A project's development starts with the configurator tool (*Fig. 2*), which provides users with a graphical interface and properties windows. In this tool, customers can designate port pins as capacitive-sensing inputs and configure library performance settings (e.g., scan period, low-power functionality). If a customer's configuration choices require implementation of other configuration settings, a "problems perspective" view lists the relevant warnings and errors. These items can be clicked on to jump to specific settings that must be changed.

As the user configures the project, appropriate files are generated and organized in a project explorer window. Clicking on one of these files will cause the configurator perspective to switch to the IDE perspective. Users can go back to the configurator and make changes or additions to a configuration anytime by clicking on the configurator file, which is also listed in the project explorer window.

CODE DEVELOPMENT AND DOWNLOAD IN THE IDE

A comprehensive Eclipse-based IDE offers all of the features that developers expect in a development environment, including code completion and formatting tools. The IDE can provide deep integration with the debug circuit on evaluation boards as well as USB debug adapters, enabling full read/write access to MCU registers and variables.

When developers include the capacitive-sensing firmware library in their projects, the configurator provides a build-ready project that will scan enabled sensors, qualify touches, and even output information through the serial interface. This feature lets customers generate a project without any additional development in the IDE. In other words, if a user wants to test out capacitive sensing on a board, the IDE can simply be used to build and download the image, without any coding required.

CHARACTERIZING PERFORMANCE WITH PROFILERS

Once code has been downloaded to the 8-bit MCU, the developer can use the IDE's energy profiler and capacitive-sense profiler to see how configuration choices made in the configurator and any additional features created in the IDE affect current consumption and capacitive-sensing performance. For capacitive-sensing applications, the capacitive-sense profiler *(Fig. 3)* is particularly useful, because users are able to examine performance in-system, complete with product overlays and other late-stage-development system components in place.

The profiler supports advanced features, such as standard deviation and signal-to-noise-ratio calculations, to characterize board performance. The tool also allows the captured data to be exported into a text file for further examination in other programs like Microsoft Excel.

DEVELOPMENT WORKFLOW FLEXIBILITY

One of the challenges in designing capacitive-sensing systems concerns managing a project through an iterative development process—build a project, test performance, and return to development to further optimize performance. The tools included in Silicon Labs' Simplicity Studio platform are designed to support these kinds of workflow cycles.

For example, a developer might create a capacitive-sensing project designed for thinner 1/16-in. overlays, build the project, and test performance

in the IDE's capacitive-sense profiler. If the product's requirements change, modifying the overlay's thickness to 1/8 in., the developer need only click on the configurator file for that project and adjust the touch-threshold settings. After making that change, the developer just has to rebuild and download the project, and performance with the thicker overlay can be seen in real-time through Profiler.

Adding other firmware features to a project as it develops provides another example of iterative development. Let's say that a developer has created a project with capacitive sensing, and checked that all performance and current-consumption requirements for the project are met using the IDE's diagnostic tools. The developer then starts work on a serial interface that controls other system components based on capacitive-sensing input. This new component could impose new requirements on current draw and capacitive-sensing responsiveness.

With Simplicity Studio, for example, the new component can be added to the project's code base without breaking the link between the source code and the configurator tool. Once the project is built, the developer might find that the functionality of the new component requires a modification to the capacitive-sensing component.

For example, the system might need to stay in an active mode and scan for button touches for a longer period of time than

3. When using the capacitive-sense profiler, designers can examine performance in-system with product overlays and other late-stage-development system components in place.

was first configured, because the serial interface to the rest of the system must stay active and optimally responsive to input. The developer can go to the configurator, adjust the amount of time the system stays awake before entering a low-power sleep state between touch sessions, and then rebuild code.

KEY-SEQUENCE EXAMPLE

The ability to generate code and minimize the amount of low-level capacitive sensing code that needs to be written leads to many use cases that may add application-specific sens-

ing features to the sensing MCU. These use cases distinguish general-purpose MCU solutions that support capacitive-touch capabilities from fixed-function solutions in which all application-specific functionality must be added to more power-hungry processors in the system. A key-sequence detector feature clearly illustrates the benefits of using a general-purpose MCU in capacitive sensing.

In many products with control panels, the end user must enter a key sequence or password to unlock product functionality. In a system that uses a fixed-function device, the main processor would have to wake up upon the first keystroke, wait in an active state while the user enters the key sequence, and then process the sequence to determine if it matches the correct password. This might not seem like a significant amount of time for a high-performance MCU to be active, but let's examine the total power budget.

Say, for instance, that this battery-powered control panel sees an average of 100 15-second user interaction sessions per day. The sessions have two phases: a validation phase where the key sequence is entered and checked against a password; and a response phase in which the control panel issues commands elsewhere in the system. In a design with a fixed-function device, the main processor must be awake for both the validation and response phase; here, it's 15 seconds per session.

If an 8-bit microcontroller, such as the EFM8SB1 MCU, is used for the validation phase, the main processor only needs to be awake for the response phase, or about 7.5 seconds. Let's also say that the main processor uses 10 mA when active, a common trait of many 32-bit processors, and the 8-bit MCU would use 20 µA when active, which is achievable thanks to the sophisticated power-management state machine in the capacitive-sensing firmware library. In this example, we're simplifying the current draw a bit and not considering the current used between user interaction sessions.

In situations where an FFD requires that the main processor be active during validation, the system's average current draw would be about 17 µA. If the lower-power 8-bit MCU is used for part of the interaction session, the average current draw would be about 9 µA. If this battery-powered system were running on a coin-cell battery, such as a CR2032 with a typical capacity of around 225 mAh, an 8-bit design based on an EFM8SB1 MCU would run without a battery change for almost three years. The system with an FFD would require a battery change in about 1.5 years—a significant reduction in run time. ed

PARKER DORRIS, a senior applications engineer supporting Silicon Labs' microcontroller product line, joined the company in 2003 He holds a BSEE from the University of Texas at Austin.

LINEAR TECHNOLOGY

DESIGN NOTES

36V Input, 8.5A Buck-Boost µModule Regulator Easily Parallels for Increased Power

Design Note 540

Andy Radosevich

Introduction

The LTM®8055 is a buck-boost µModule® regulator with a 5V to 36V input range that can be easily paralleled to extend load current capability. Its 4-switch buck-boost topology features high efficiency while allowing the input voltage to be below, at or above the output—with smooth transitions between regulation modes. The LTM8055 can regulate a constant-voltage (CV) and a constant-current (CC) output, with analog programming of output current. Its out-of-phase clock output, combined with switching frequency synchronization, enables easy phase interleaving in parallel operation. A complete solution fits a space not much larger than the 15mm × 15mm footprint of the LTM8055, including input capacitors, output capacitors and current sense resistors as needed.

High efficiency is an inherent benefit of a 4-switch buck-boost converter design, due mainly to low power losses in NMOS switches. 2-switch converters are

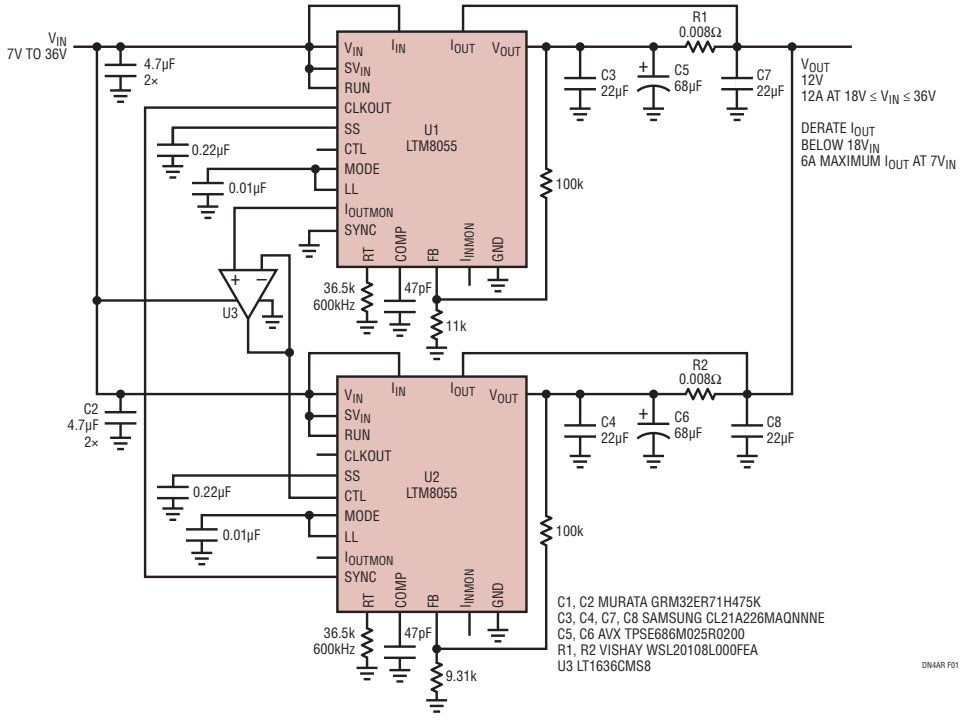

Figure 1. 144W Buck-Boost Regulator Circuit: Two LTM8055s Can Be Easily Paralleled to Increase Output Power. This Parallel Circuit Is Capable of 12A of Output Current When the Input Is Between 18V and 36V

similarly efficient but limited to buck or boost topologies. Furthermore, DC/DC supplies that use a single switch and a rectifier diode, instead of two switches, are less efficient because of power losses in the diode.

Parallel Operation

It is easy to parallel multiple LTM8055s to increase output power capability. Figure 1 shows a 12V output regulator composed of two parallel LTM8055s, one CV master regulator U1 and one CC slave regulator U2, capable of 12A of output current when the input is from 18V to 36V. Figure 2 shows that the efficiency of the Figure 1 circuit is as high as 93% at 24V input. The output current is derated at input voltages lower than 18V.

CV operation of the master regulator and CC operation of the slave regulator assures balanced current sharing. I_{OUTMON} monitors the output current, while the CTL pin programs the output current of the current source regulator. The master is programmed to the desired output voltage while the slave is programmed to a higher output voltage, since it normally regulates current and not voltage. The master regulates the output voltage and its I_{OUTMON} pin connects to the CTL pin of the slave which forces the LTM8055s to share the load current equally. A unity gain buffer allows the master I_{OUTMON} pin to drive the CTL pin of the slave. Figure 3 shows load sharing of the two parallel regulators in Figure 1, where the I_{OUTMON} outputs remain equal during a 6A to 10A load transient test.

The two switching cycles are synchronized 180° apart by simply connecting CLKOUT of U1 to the SYNC pin of U2. Out-of-phase synchronization reduces the required capacitance on the input and output. Synchronization prevents any beat frequency noise that could be caused by a random switching cycle relationship between the LTM8055s.

More than two LTM8055s can be paralleled by connecting the I_{OUTMON} pin of the master to the CTL pins of the slaves using a single unity gain buffer. An external clock generator with two or more phases may be necessary for proper channel interleaving.

Current Regulation and Monitoring

Parallel operation is not the only use for the LTM8055's current regulation feature. It can be used for battery and supercapacitor charging or to protect the output from overcurrent and short faults. The LTM8055 also provides input current regulation and monitoring. Input current limit can prevent the LTM8055 from overloading its input supply.

Conclusion

Minimal design effort is required to produce a compact and high performance converter with the LTM8055. This 36V buck-boost µModule regulator can be easily paralleled for increased output power, with 2-phase designs requiring no external clock.

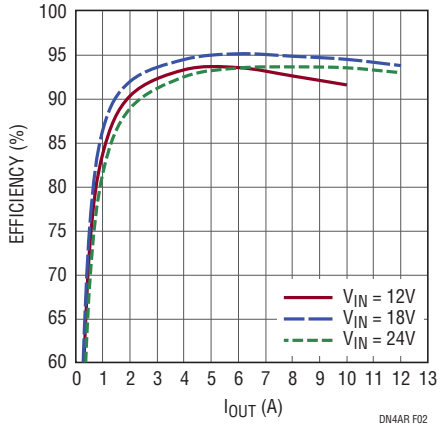

Figure 2. The Efficiency for the Figure 1 Circuit Is as High as 93% at 24V Input. The Output Current Is Derated at Input Voltages Lower Than 18V

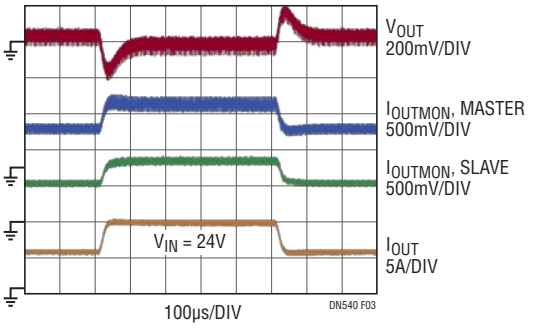

Figure 3. The I_{OUTMON} Outputs of the Master and Slave Regulators of Figure 1 Show Balanced Current Sharing During a 6A to 10A Load Transient

Data Sheet Download

www.linear.com/LTM8055

For applications help,
call (408) 432-1900, Ext. 3568

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

dn540f LT/AP 0715 111K • PRINTED IN THE USA

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

Ethernet Continues to Push Faster and Add More Flexibility

It was hard to imagine Ethernet's staying power when it first appeared on the market almost 35 years ago. It used thick coax. It pumped data at 10 Mb/s. An entire coax-based Ethernet network could easily be brought down by one bad connection.

These days, an RJ45 with a twisted-pair cable connected to an Ethernet switch or hub is pretty much the norm. Only two pairs of the four are used for copper Ethernet connections up to 10 Gb/s. Cat 5e is needed for 1 Gb/s Ethernet. At 10 Gb/s, Ethernet standards eliminate hubs and half duplex operation. Switches are necessary to maintain throughput. Cat 6a or Cat 7 cables are needed for copper connections. These higher-speed connections tend to be found in top-of-rack (ToR) switches.

PAM-4 has four distinct levels to keep cables at four 50 Gbaud connections to deliver 400 Gbit/s Ethernet (as shown on Keysight's real-time oscilloscope running N8827B PAM-4 analysis software).

Higher-speed Ethernet is also available, including 25-, 40-, and 100-Gb/s connections. The 25-Gb/s speed matches the addition of 2.5-Gb/s Ethernet that increases the signaling speed. The 25 Gb/s is a 10-Gb/s system increased in the same fashion. By 2017, 400-Gb/s Ethernet should be arriving, and it will require even more ambitious changes.

Ethernet started with an 8b/10b encoding scheme that has a 25% overhead; 10-Gb/s Ethernet switched to 64B/66B with a 3.125% overhead. Pulse amplitude modulation (PAM) has also been used to increase throughput. The 1000Base-T specification uses PAM-5, a five-level scheme. The potential 400-Gb/s Ethernet standard is looking to use PAM-4 running at 50 Gbaud with four channels to reach the necessary speed (*see the figure*).

Fiber use has become more important and more common to accommodate higher-speed Ethernet. It supports much longer runs than copper and it is inherently less susceptible to noise. It tends to be more economical for 10-Gb/s Ethernet and above compared to copper.

Higher speeds are critical in the enterprise, but slower versions are very useful in embedded applications. A large number of microprocessors and microcontrollers come with 10/100-Mb/s Ethernet interfaces. Industrial Ethernet spans a number of standards to provide real-time Ethernet communications running at 1 Gb/s or slower (*see "Industrial Automation Relies on Ethernet" on electronicdesign.com*). There is also the IEEE 1588 Precision Time Protocol (PTP) standard.

The Audio Video Bridging (AVB) standard allows replacement of point-to-point multimedia connections with an Ethernet network. The AVB standard addresses precision synchronization, traffic shaping, and admission controls.

Another standard worth mentioning is Power over Ethernet (PoE). This targets copper Ethernet cabling using 1-Gb/s Ethernet or lower, providing dc power on the unused RJ45 connections. Most Ethernet cables contain all eight wires, although only four are used for communication at these speeds.

The original IEEE 802.3af-2003 standard promised up to 15.4 W of power. The newer IEEE 802.3at-2009 increased this to 25.5 W. This amount of power is sufficient for a wide range of applications from Voice-over-IP (VoIP) telephones to wireless access points and network video cameras. It reduces device cost, since an additional power connection is not required nor is additional cabling needed. Many PoE switches can be configured to limit the amount of power as well as provide it on a scheduled basis that may be necessary for some applications.

The trend for wired communication is clearly Ethernet, even if the plethora of standards can be overwhelming.

BRIAN ROBERTS | Senior Designer
Dawn VME Products, 510-657-4444, www.dawnvme.com

Solve VPX Backplane Problems Before You Deploy

Connector-less micro-overlays simplify and automate topology customization to quickly optimize off-the-shelf VPX backplanes.

VPX backplanes commonly implement high-speed signal standards such as PCI Express, Serial Rapid I/O, SATA, SAS, and 10-Gbit (XAUI) Ethernet. It's important to know that when VPX backplanes use these types of signal paths, they require point-to-point connectivity from slot to slot to maintain signal integrity and speed of communication.

Connecting multiple plug-in cards, such as CPU processor boards, graphics cards, GPU math processors, and the like via a VPX backplane can be problematic. That's because the nature of the extremely high frequency signals used means simple "busing" between multiple card slots no longer works effectively.

High-performance, mission-critical backplanes need more flexibility to meet the wide variations in point-to-point signal-connection standards. Fabric mapping modules—ball-grid-array (BGA), signal mapping overlays—present an effective solution with the necessary signal-integrity improvements to meet these challenges, and their use makes it possible to solve many application problems in the design phase.

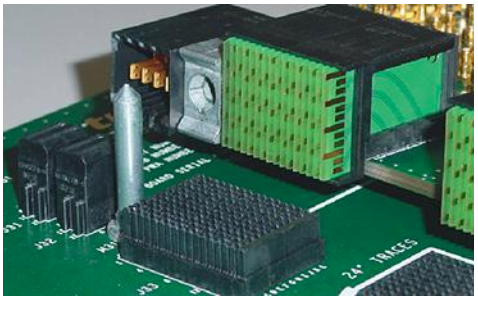

1. The MultiGig RT2 connector, shown here with an alignment pin, can be used in 3U and 6U VPX backplanes.

FROM VME TO VPX

Before we discuss the fabric-mapping-module (FMM) overlay techniques, let's get more in-depth about VPX backplanes themselves. In the older VME-bus systems (the predecessor of VPX), connectors were able to accommodate parallel data buses. The VME International Trade Association (or VITA) eventually developed new standards to accommodate switched fabrics that use differential signaling at multi-gigabit speeds; this necessitated a new connector technology. The differential pairs in switched fabrics use pairs of pins that are close to one another and shielded from other signals by ground pins.

The MultiGig RT2 from Tyco was developed to meet the design requirements of VITA 46, now known as VPX. Among its most important features are quad-redundancy in pin connectivity, and the potential for electrical customization: Wafers can be manufactured for differential or single-ended signal paths; and impedance, propagation delay, and crosstalk specifications can be altered as per customer requirements.

Consider some examples of backplanes that are built to accommodate the MultiGig RT. There's a 3U VPX board with two 16-column, 7-row RT2 connectors and one 8-column, 7-row RT2 connector. And a 6U VPX board incorporates six 16-column, 7-row RT2 connectors and one 8-column, 7-row RT2 connector. VPX boards have alignment keys that also supply a safety grounding contact. The 6U board has three such keys, while the 3U has two (*Fig. 1*).

The density of the RT2 connector is such that when applied to a 6U board, there are 464 signaling contacts that can be allocated as follows:

- 384 differential pins that can be implemented as 192 high-speed differential pairs for core fabric
- 40 single-ended pins for customer I/O
- 28 pins for system utilities

Even with the impressive performance and capabilities of the MultiGig connector, its impedance variations can create challenges when using the standard overlay techniques of circuit-board design. Printed-circuit boards (PCBs) that support high-speed signal standards, such as PCI Express, Serial Rapid I/O, SATA, SAS, and 10-Gbit Ethernet, do so by providing point-to-point direct signaling paths to maintain signal integrity. Clearly, then, multiple plug-in cards on a backplane cannot

"share" a signal path, because communication no longer occurs via busing between card slots—the differential-pair nature of the connectors keeps that from being an option.

As an example, consider PCI Express generation 1. In differential-pair transmission technology, the signal from one chip's output pairs links directly to the input pairs of the receiving chip by way of two traces that form a 100-Ω transmission line. The MultiGig VPX connector consists of small "blades" that are actually PCBs themselves, and they continue the 100-Ω transmission line path. The connector's "leaf-spring" contacts are designed so that any "stubs," or short trace lengths, are as short as possible to minimize any signal reflection noise.

Any stub present on a signal path will, at some frequency, cause an additional signal wave front that's out of phase. This causes destructive interference with the original signal transmission, diminishing the signal seen by the receiving chip. The design goal is to optimize the path between the transmitting chip and receiving chip by eliminating stubs and impedance discontinuities. These factors are further enhanced by the PCI Express chip standard implementation of pre-emphasis and equalization to make the signal transmission as optimal as possible.

Another key aspect of VPX backplane design to consider is the adjacencies of the plug-in cards that go into the backplane itself. The closer the source and destination VPX cards are to each other, the faster the communication links between them. That design principle, along with the difficulties associated with balanced-pair transmission-line implementation, explains why VPX backplanes need more flexibility to meet the wide variations in point-to-point signal-connection standards. Micro-overlays that are external to the backplane itself can solve the challenge of linking these two requirements by delivering the necessary signal-integrity improvements.

USING SIGNAL-MAPPING MICRO-OVERLAYS

Micro-overlays use BGA solder-connection technology to interface a PCB-based differential pair matrix with compatible backplanes. They reduce transmission-line impedance variations and stubs associated with connector-based interfaces by employing two techniques:

- Stacked dielectric with alternating etched copper layers implementing signal shield plane and controlled-impedance copper-etched signal lines.
- Sub-picosecond matched signal path length. The thickness of the dielectric stack-up is tuned for optimal impedance matching to the signal standard requirements. Off-signal-path stub lengths are minimized by this overlay connection technique.

These techniques improve the signal integrity between system cards beyond the requirements of PCI Express Gen 3, Serial Rapid I/O, and the 10-Gbit Ethernet standards. In addition, micro-overlays can facilitate rear transition modules and low-profile connector interface systems when normal rear transition modules don't mechanically fit the system application.

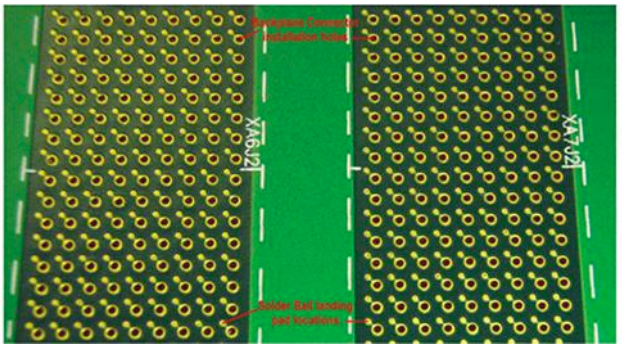

2. Ball-grid-array connection pads are shown adjacent to connector holes.

FABRIC-MAPPING MODULES FOR VPX BACKPLANES

FMMs are micro-overlays that provide benefits to designers of systems incorporating VPX backplanes. Use of FMMs allows the backplanes to be semi-customized and quickly reconfigured according to user needs. Consider a backplane with power and other "utility" planes that serve multiple plug-in boards. Those planes need not change or be redesigned when a VPX-based backplane is used for a design iteration, since FMMs can be utilized to make iteration-specific changes. In this way, FMMs add flexibility to the system design (*Fig. 2*).

One example is a primary defense contractor that needed a four-slot backplane to fit inside a conduction-cooled cube-shaped form factor meant for use in multiple vehicle types. In each case, the payload plug-in cards are different, but the power-supply plane remained the same. Processor cards and graphics cards from different manufacturers were used. Here, the changes in plug-in cards did not require a complete redesign of the backplane. Instead, FMMs were employed, which accommodated the necessary changes on the backplane that correspond to the differences in plug-in cards and front I/O pin out.

Another example of FMM usage involves the major supplier of VPX-backplane-based deployable radio systems. The company uses processor cards, storage cards, switch cards, and radio cards that go in to different plug-in slots in the backplane, depending on the specific configuration. They rely on different FMMs to make the configuration-specific connections.

On top of that, the supplier implements connectors on the FMMs themselves to pick off signals such as Ethernet and Serial ATA, and then routes those signals directly to the front panel of the system. None of these changes require backplane redesigns because of the inherent flexibility derived from using FMMs.

In general, fabric mapping modules give system designers the flexibility needed to meet the wide variations in point-to-point signal connection requirements. ED

BRIAN ROBERTS, senior designer at Dawn VME Products, has 24 years' experience developing electronic products for commercial and defense applications.

VICTORIA FRAZA KICKHAM | DISTRIBUTION EDITOR victoria.kickham@penton.com

When Easily Sourced Items Aren't So Easy to Find

Pricing and availability remain steady for many electronic components, but buyers are keeping an eye on proposed interest-rate increases and other factors that could affect both issues in 2015.

BRIDGET McCREA | CONTRIBUTING EDITOR

KAREN SAVAGE IS no stranger to long lead times on electronic components, but when she recently learned that a specific switch was going to take 10-12 weeks from order to delivery, she couldn't believe her ears.

"It was just a simple switch that we hadn't had any issues with in the past," says Savage, senior buyer with A&D Technology Inc., Ann Arbor, Mich.

As a manufacturer and designer of engine testing solutions, the company relies on a steady stream of circuit boards, switches, and computer equipment. When the lead times on such products are long—in this case, by nearly three months—Savage either has to source the items elsewhere or risk project delays.

While Savage was grappling with the switch issue, other component makers were dealing with growing pains of their own.

"It was a compound issue," says Savage. And, as it happens, this "perfect storm" occurred right around the time that some of A&D's customers were scrambling to spend the remaining pieces of their fiscal-year budgets.

"Many of our customers have a 'use it or lose it' mentality when it comes to budgeting," says Savage. "We can't just tell them that we're going to deliver in February on an order that they placed with us for December; it doesn't work that way."

Savage then started looking to other sources for the switch in question. "The item didn't seem that special to me," she says, "so I looked to brokers all over the world and no one had any stock, consignment stock, or even resale stock at all on this switch."

In the end, Savage says A&D was able to get the switches within an eight-week time frame: "I was ecstatic when we received them," noting that in some cases the specialized nature of the manufacturer's business comes into play when sourcing components. "We don't do a huge volume of any specific item because we're so specialized, but in this case our supplier was able to push the order through and make the switches for us within a shorter time frame."

Looking at the broader global purchasing picture, Savage is keeping an eye on proposed interest-rate increases and other issues that could affect product availability and pricing in 2015.

"Even the Federal Reserve can't figure out whether the economy is growing fast enough and/or if the growth needs to be slowed down," says Savage. "If you're a chip builder right now—and if you want to run 50,000 units of a specific chip—you're probably not going to make the investment if there aren't any customers waiting for those items. Without a good forecast, you may wind up building 10,000 instead of the intended 50,000 units."

That scenario could result in less quantity and higher prices, particularly if system builders don't "stock up" on components.

"With technology, you don't want pieces sitting on the shelf in [inventory]," says Savage. "When you don't know that there's a sufficient demand, no one wants to go out on a limb and say 'I'm ready to take an order and I'm going to build these up right now,' because it just doesn't make sense."

In terms of current component pricing, Savage says that costs are mostly either remaining steady or rising slowly: "I don't see anything dropping unless it was for a specific spot buy. For the most part, I have my ordering set up for continuous orders in lot size, so I'm keeping my buys at routinely the same price."

Finally, she's started reviewing invoices and agreements to ensure good future supply or find alternative sources of supply. "As buyers, we have to watch and assess bill health to make sure that when the market does pick up, we're talking to our manufacturers about issues like component availability," says Savage. ■

The Power of Distribution

Let America II Power Your Supply Chain

America II is a leader in component distribution and serves as a primary supplier for OEM and EMS companies around the world. We offer the most complete range of products, from active, passive and electromechanical components to memory and specialty products. In addition, America II supplies components for nearly every application including automotive, industrial, military/aerospace, telecommunications, medical, consumer electronics, the Internet of Things (IoT), and so much more.

800.955.5302
www.americaii.com

AMERICA II ELECTRONICS®
The power of distribution℠

Universal Footprint RF Duplexers

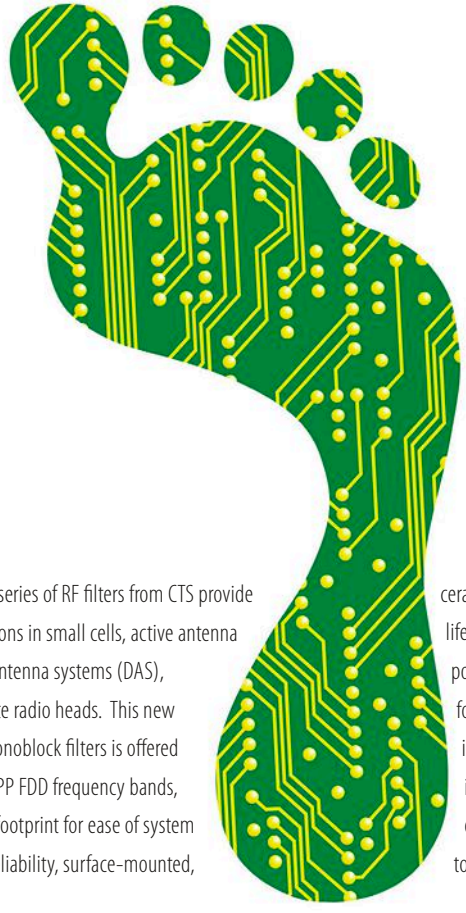

Universal-footprint series of RF filters from CTS provide better communications in small cells, active antenna arrays, distributed antenna systems (DAS), repeaters and remote radio heads. This new breed of ceramic monoblock filters is offered for all the major 3GPP FDD frequency bands, sharing a common footprint for ease of system design. Our high-reliability, surface-mounted, ceramic filters support 20-year operating lifetimes and as much as 200 W peak input power handling capability. Universal footprint filters provide industry-leading insertion loss and rejection. For more information on our ceramic universal duplexers and bandpass filters go to www.ctscorp.com.

Filters with Industry-leading insertion loss and rejection.

	UMB/UMD Pair of Bandpass Filters For Metro Cell	USD DUPLEXER For Small Cell	UPD DUPLEXER For Pico Cell
Input Power Rating	20W Avg 200W Pk	6W Avg 60W Pk	1.5W Avg 15W Pk
Insertion Loss (5MHz AVG)	2.2dB	2.6dB	3.0dB
Rx Band Isolation*	80dB	72dB	63dB
Tx Band Isolation	74dB	66dB	57dB
Universal Footprint Size (mm)	62 x 44	63 x 18	44 x 18
Operating Temp Range	-40 to +85°C	-40 to +85°C	-40 to +85°C

* Note: "Difficult" bands may have 2dB lower worst case Rx band isolation.

CTS®
ELECTRONIC COMPONENTS

www.ctscorp.com

Design Balanced Op-Amp Circuits for Performance and Simplicity

DIETER KNOLLMAN | dknollman@gmail.com

IN A BALANCED OP-AMP circuit, the op-amp inputs see equal impedances looking into the circuit. Balanced circuits provide better performance by cancelling some input errors and have a simple R_F/R_I gain formula that works for all cases, where R_F is the feedback resistor that connects from the output to the inverting input and R_I is the input resistor that connects the input voltage source to an op-amp input.

It's simple enough. The gain is positive if R_I connects to the non-inverting input and negative if R_I connects to the inverting input.

To design a balanced circuit, select R_F (such as 100 kΩ), and calculate R_I for each input by dividing R_F by the magnitude of the gain. If the gain is positive, connect R_I to the non-inverting input. If it's negative, connect it to the inverting input. Finally, add a balance resistor to create equal impedances for the op-amp inputs. The balance resistor connects from the op-amp input with the higher impedance to ground.

Alternatively, you can use Daisy's Theorem to get the balance resistor, which states that the sum of the gains in a linear circuit must be equal to +1. The balance resistor adds a ground input to the circuit. Select the ground input gain to satisfy Daisy's Theorem and you have a balanced circuit.

Consider a differential op-amp circuit that has a positive input with gain of +5 and a negative input with gain of −5 *(Fig. 1)*. R_F was selected to be 100 kΩ. The input resistors are 100 kΩ/5 = 20 kΩ. You can get the balance resistor by noting that the inverting input sees R_F (100 kΩ) and R_{In} (20 kΩ) while the non-inverting input has R_{IP} (20 kΩ). It needs a 100-kΩ balance resistor. Alternatively, the signal gains add to zero and need a +1 ground gain to satisfy Daisy's Theorem.

The procedure works for all combinations of positive and negative inputs. If you have been using the $1 + R_F/R_I$ formula, you may have some misconceptions that should be cleared up.

For negative gains, the $-R_F/R_I$ formula is the same as the classical op-amp formula. For a positive gain, the R_F/R_I gain formula differs from the $1+ R_F/R_I$ formula, which only applies to the non-inverting amplifier circuit and only for an input connected to the non-inverting op-amp input. The R_F/R_I formula applies for all positive inputs in a balanced circuit.

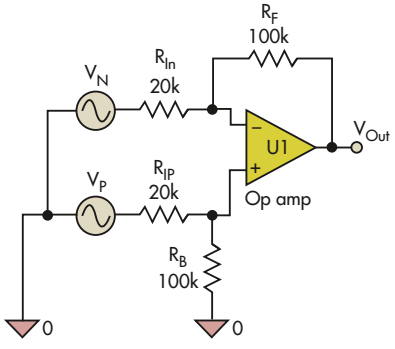

1. Analysis begins with a basic differential op-amp circuit with two inputs: a positive one with gain of +5 and a negative one with gain of –5. The signal gains sum to zero and need a ground gain of +1 to satisfy Daisy's Theorem.

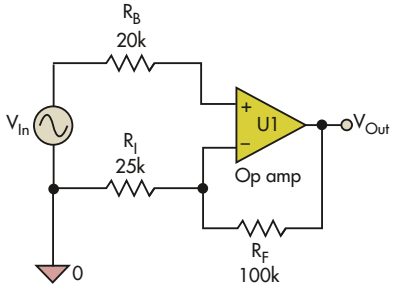

2. In the non-inverting amplifier circuit, balance resistor R_B makes the analyses from both perspectives agree, and protects the op amp.

A look at a non-inverting amplifier circuit clarifies the source of the confusion *(Fig. 2)*. To show that both formulas create the same result, we need to add a balance resistor. The standard circuit has the input connected directly to the non-inverting op-amp input.

To protect the op amp and provide a way to balance the circuit, a balance resistor R_B has been added. If you select $R_B = R_F$ in parallel with R_I, you will have a balanced circuit, and R_B does not affect the circuit gain. Since the non-inverting input impedance is infinite, no current flows through R_B, and the non-inverting op-amp input voltage is V_{In}.

The R_F and R_I labels come from the standard circuit. Note that R_I is not a signal input resistor. It is the ground input resistor. The V_{In} input resistor is missing in the standard circuit. It has been added above and labeled R_B.

You can analyze the circuit in many ways. With the $1 + R_F/R_I$ formula, the gain is 1 + 100 kΩ/25 kΩ = 5. With the balanced circuit formula, the gain is R_F/R_B = 100 kΩ/20 kΩ = 5. With Daisy's

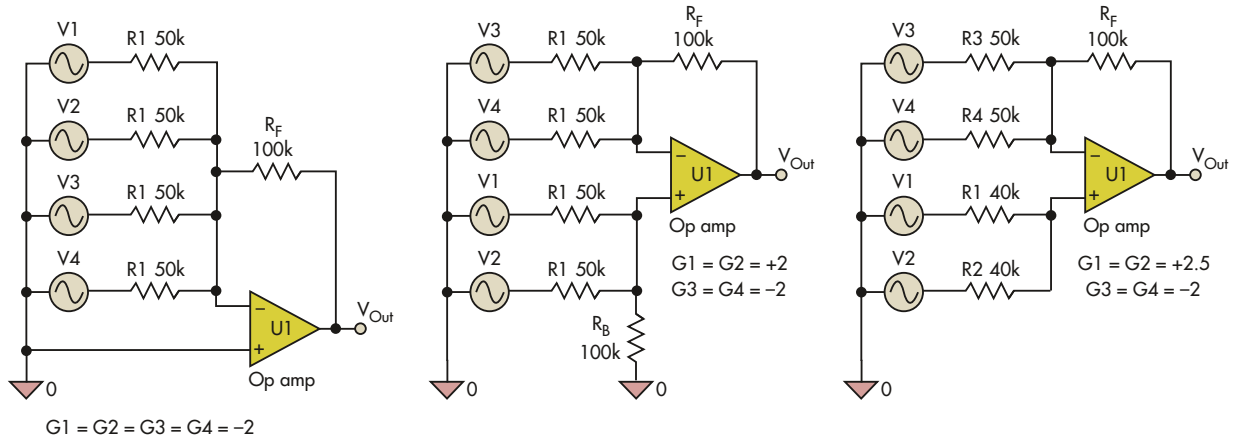

3. A four-input audio mixer can be configured in three ways: as an inverting-amplifier mixer circuit (left); as a balanced mixer with two positive and two negative inputs (middle); and as a mixer with different positive and negative gains (right).

Theorem, the ground gain is $-R_F/R_I = -4$ and the input gain needs to be +5 to have the gains add to 1.

The results are the same. The confusion is created by the labeling. It comes from how the equations are derived. Classical analysis derives the non-inverting formula using circuit assumptions. If the op-amp gain is large, both op-amp inputs must possess the same voltage in order to create a finite output voltage. Such an assumption makes the analysis of the circuit rather simple.

In summary, the $1 + R_F/R_I$ formula works only for the non-inverting circuit. The R_F/R_I formula works for all positive inputs in a balanced circuit. There is no conflict. The balanced circuit formula is derived from Plato's Gain Formula:

$$V_{Out}/V_{In} = p \times Z_F/Z_I$$

K9 Analysis represents a circuit component as an impedance Z, where Z can be an ideal component, a two-terminal network for a non-ideal component, or a resistor. Plato's Gain Formula is the unified gain formula for a single op-amp circuit. The gain is a constant p times the feedback impedance (Z_F) divided by the input impedance Z_I.

For a circuit with equal impedance at the positive (+) and negative (–) op-amp inputs, the gain magnitude is simply Z_F/Z_I, or R_F/R_I for resistors. The gain is positive for inputs connected to the positive op-amp input and negative for inputs connected to negative input.

Constant p is equal to –1 for negative gains and Z_P+/Z_P- for positive gains. Z_Pn is K9 notation for the parallel combination of the impedances that connect to node n. Z_P+ is the parallel combination of the positive input impedances that connect to the non-inverting op-amp input.

Since the inputs are ideal sources, this is the impedance seen by the non-inverting op-amp input. Z_P- is the parallel impedance of the negative input resistors and the feedback impedance Z_F. The Z_P terms will vary with the circuit and depend on the number of circuit inputs. For a balanced circuit, the Z_P terms cancel and the magnitude of the gain is Z_F/Z_I.

Plato's Gain Formula is derived via K9 Analysis using

nodal analysis. If you're still skeptical, compare Plato's Gain Formula to the gain formula in a textbook solution. They should match. If the positive input is connected to ground, add an input resistor to avoid singularities. (This may require a bit of algebra.)

The intent of K9 Analysis is to simplify analog circuit design. With the balanced-design technique, op-amp circuit design becomes simple, and you are no longer constrained to existing op-amp circuits.

If you are using multiple op amps to implement a linear circuit with a single output, you may be able to reduce the circuit to a circuit with a single op amp. Just get the circuit gains and design a balanced, single op-amp circuit. Keep in mind that one op amp can implement any linear circuit equation.

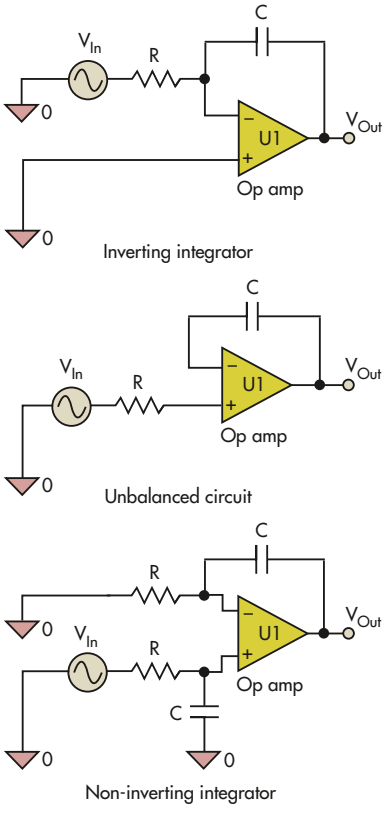

4. The balanced principles are not restricted only to simple amplification. They can be adapted for use in various integrator circuits, the second most commonly used function (after amplification, of course).

No longer is a single op-amp design with multiple positive and negative gains difficult. You do have to create a balanced design by adding a balance resistor, which determines the ground gain. Since ground is zero voltage, you may think that the ground input does not contribute to the op-amp output. That's wrong! Ground is never zero voltage. There is always a little ground noise. Reducing the ground gain will reduce the output noise.

Let's look at a four-input audio mixer, where we want each input to have a gain of 2 *(Fig. 3)*. Since this is audio, the sign of the gain may not matter, and a gain of +2 or –2 is acceptable. Set R_F = 100 kΩ, and each input resistor will have a value of R_F divided by two, or 50 kΩ. The circuit on the left is an inverting-amplifier mixer circuit. The middle one is a balanced mixer with two positive and two negative inputs. The circuit on the right has the V1 and V2 input gains changed to +2.5.

What's the ground gain in the inverting mixer circuit? The inputs add to a gain of –8. Daisy's Theorem has the ground gain at +9. Since all of the gains are negative, the circuit does not need a balance resistor. The non-inverting op amp input can be connected to ground.

If you can afford an extra resistor, you can create the balanced design. Two of the inputs are moved from the inverting input to the non-inverting input. What's the ground gain now? The input gains add to zero and the ground gain is +1. This is a big noise reduction.

You can even do better by making the sum of the signal gains equal to 1, in the third variation, where two of the input gains are +2.5 and the other two are –2, so the gains add to +1. The ground gain is 0 and no balance resistor is needed. This circuit adds no ground noise to the output, which is nice.

Up to this point, we have discussed how to reduce your circuit and eliminate ground noise by employing a balanced design. K9 Analysis utilizes impedances. If we use a capacitor for Z_F, we can build an

integrator circuit. Figure 4 shows the textbook integrator circuit on the left. The integrator has a negative gain, but are we able to design an integrator with positive gain?

In the circuit in the middle, the input has been moved to the non-inverting input. This circuit is not balanced, as the inverting input sees the feedback capacitor, while the non-inverting input sees the input resistor.

To balance the circuit, we need to add a resistor equal to the input resistor to the inverting op-amp input and a capacitor equal to the feedback capacitor to the non-inverting input. Thus, both inputs see equal impedances and are balanced.

The balance components were connected to ground. However, you can connect them to a voltage-source input to create a different circuit, such as a differential integrator, by returning the inverting input-balance resistor to a second input, or an ac amplifier by returning the balance capacitor to an input. These circuit techniques allow you to convert inverting textbook circuits to non-inverting circuits, as long as you maintain balance.

But there are caveats. The simplified positive-gain formula hides the fact that positive gains are very interactive. Changing any input impedance can change all of the positive gains. You need to have a balanced circuit for the simple positive formula to apply. Most changes require a change in the balance impedance. In contrast, negative gains are independent and don't interact. If all signal gains are negative, you do not need a balanced circuit.

Also, op-amp circuit voltages are limited. You need to verify that inputs do not cause the output to saturate. For positive inputs, you also must verify that the inputs do not violate the common-mode input range. (If all signal gains are negative, the op-amp inputs are at ground potential and the common-mode input range is not an issue). You also need to verify that the circuit is stable. The design equations are solutions of circuit equations. Although the solutions are correct, there's no guarantee that the circuit will remain in the solution state.

SUMMARY

Balanced op-amp designs are relatively easy. Keep in mind that op amps prefer balanced design, so consider a balanced op-amp design for your next linear circuit.

Plato's Gain Formula is Z_F/Z_I for balanced circuits. Just pick a feedback impedance and select the input impedances for the desired gains. Connect the input impedance to the non-inverting op-amp input for a positive gain or the negative op-amp input for a negative gain. Create a balanced cir-

cuit by adding balance impedances.

For amplifier designs, use Daisy's Theorem to find the ground gain. If the circuit's requirements allow, explore changing the gain signs to reduce the ground gain and

DIETER KNOLLMAN is a retired member of the technical staff at Bell Labs. He has a PhD from New York University and about a dozen patents. He created K9 Analysis to simplify analog analysis (Lucent has given permission to share the K9 analysis) and taught circuit design at a state university in the 1980s.

Well-Controlled Audio-Band Noise Source Uses Basic Microcontroller Filtering

DENNIS SEGUINE | CYPRESS SEMICONDUCTOR CORP. seg@cypress.com

A WHITE-NOISE SOURCE can be constructed based on the random noise of a Zener diode or reverse-biased base-emitter junction. While these circuits work, they are temperature-sensitive and not predictably calibrated.

Another approach is to use a pseudorandom sequence (PRS) generator that, when filtered, yields a noise with characteristics such as flat in frequency, Gaussian in amplitude distribution, and as stable in amplitude as the reference that

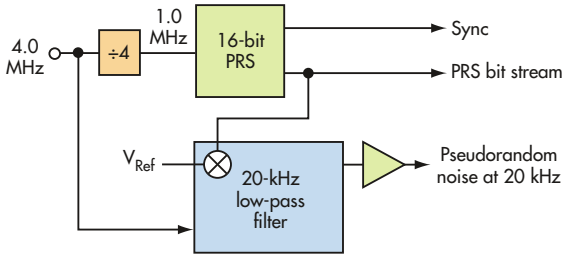

1. The basic pseudorandom noise generator based on a linear feedback shift register is simple, but an effective circuit requires attention to filtering and parameter settings.

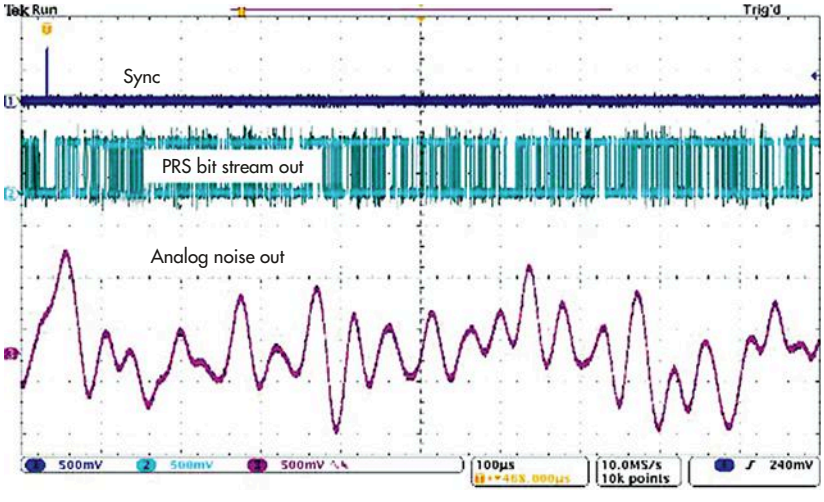

2. With proper filtering, the pseudorandom-noise-generator digital-output waveform is converted into a close representation of an analog noise pattern.

is used. The noise source, including PRS, filter, and driver, can be implemented entirely within a programmable system-on-chip (PSoC) device *(Fig. 1)*.

The PRS is derived from a modular linear feedback shift register (LFSR) with the output fed back to specific taps that are XORed with the data stream. The tap values are available from a variety of sources, including the PSoC device's own design software. The PRS generator in the PSoC is a modular design that's available in bit lengths from 2 to 32.

The bit-stream output is a digital signal that's converted to analog noise with a filter. A sync pulse is available, triggering once per sequence *(Fig. 2)*. The bit-stream output of the PRS repeats every 2^{n-1} clock pulses. For a 1.0-MHz clock, a 16-bit PRS repeats its pattern every 65 ms. If the system is slow and repeating patterns might interfere, just use a longer sequence. A 32-bit sequence yields a repeat pattern at 1.19 hours.

The bit-stream digital output can connect directly to the input of an on-chip, switched-capacitor, low-pass filter. This yields a signal that's proportional to the supply voltage. A supply-independent source can be generated by setting the filter input to a reference, then connecting the PRS bit stream to the filter's modulator input. This multiplies the reference fed to the filter by +1 or –1 and provides an amplitude-stable source.

The analog noise output and the bit stream are synchronized as shown in the waveforms of Figure 2. The bit-stream amplitude is constant and doesn't look very Gaussian because it's always driven to the supply rails. Making the signal Gaussian requires limiting the bandwidth and keeping the output from hitting the rails. This can be accomplished by setting the filter corner frequency at less than 5% of the PRS clock rate. This implementation is a white-noise generator with a bandwidth of 20 kHz, sampled at 1.0 MHz.

Figure 3 shows the noise spectra. *(See Listing 1 in the spreadsheet available with the online version of this article at electronicdesign.com for the raw file of the spectra in Fig. 3.)* The PRS noise (red trace) has a sin(x)/x shape with the first null at the sample rate. Filtered output noise (blue trace) is as flat in the band as allowed by the 20-kHz filter. The finished noise source has an observed noise output of about 1.5 V p-p and measured amplitude of 275 mV rms.

Small clock spurs occur at some submultiples of the clock rate, but they are easily filtered out with an external

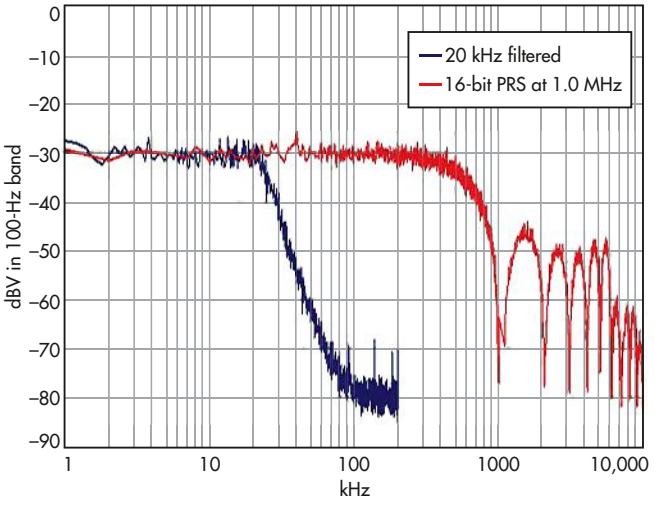

3. The frequency spectrum of the filtered noise shows how closely it mimics a true random signal, within defined limits.

single-pole low-pass RC filter. Higher corner frequencies can be achieved by reducing the oversample ratio and increasing the clock rates. The practical corner-frequency limits for a PSoC low-pass filter are 300 Hz to 120 kHz.

IDEAS FOR DESIGN WANTED

Send us your Ideas For Design. We'll pay you $150 for every Idea For Design that we publish. In addition, this year's top design as selected by our readers will earn an additional $500, with two runners-up each receiving $250. You can submit your Ideas For Design via:

- E-mail: roger.engelke@penton.com

OR BY

- Postal mail to:
  Ideas For Design
  Electronic Design
  1166 Avenue of the Americas, 10th Floor
  New York, NY 10036

Go to *www.electronicdesign.com* for our submission guidelines.

DENNIS SEGUINE is a member of the technical staff at Cypress Semiconductor Corp. He has been an applications engineer for Cypress Semiconductor since 2000, following many years of analog, embedded system, and software design for the underwater, instrumentation, and medical industries.

Intelligent I/O Module with IMU Simplifies Robot Design

THE ROBOTEQ IO eXtender Module (RIOX) is an intelligent I/O expansion card, with optional Inertial Measurement Unit (also known as AHRS) that seamlessly integrates with PC's via its USB interface, PLCs via its CANbus or RS485 interface, and most SBCs, such as Aduino, Raspberry PI, or Beagle Boards. The module includes a 32-bit ARM MCU for processing and buffering the I/O, and managing communication with the host PC/PLC/SBC, motor controllers, and other components of the system. The processor can read the 3-axis accelerometer, gyroscope, magnetometer, and perform the fusion algorithm in the AHRS version. A Basic-like programming language is built into the card and allows users to write programs that process the I/O in real time. Each of the 12 inputs can be individually configured as digital, 0-5 V analog or pulse inputs, and can also be configured as a driver for RC servos or as interface for ultrasound distance sensors. The 16 digital outputs are capable of driving resistive loads such as lights, or inductive loads such as relay, solenoids, or motors, up to 1 A each at 40 V.

The RIOX module is available now at $175 and $295 in the Basic and AHRS version, respectively.

ROBOTEQ

www.roboteq.com

Smart-Card Keyboard Designed for Access Control Security

ZF ELECTRONIC Systems' CHERRY JK-A0100EU smart card keyboard is designed to control logical access to networks in order to maintain a secure user base. The secure keyboard is ideally suited for regulated industries such as government, finance, and healthcare. However, it can also be used by any organization that wants reliable access control to enhance the security of their network.

The wired keyboard requires the secondary authentication of a smart card before allowing access to the connected system. In addition to the smart-chipset card reader, the keyboard also supports secure PIN entry. The convenient placement of the smart card reader allows for one-handed operation.

The CHERRY JK-A0100EU smart-card keyboard also features a thin profile and weighs 1.85 lbs. FIPS 201 certification is in process.

ZF ELECTRONICS

cherrycorp.com

Automotive-Grade High-Current Power Inductors Enhance Reliability

EATON'S COILTRONICS HCMA Series of high-current power inductors are designed to withstand temperature and shock stresses found in automotive applications, supporting greater reliability in electronics for advanced drive-assistance systems, vehicle control, body control, lighting, infotainment, cluster controls, heating and air conditioning, chassis and safety systems. The power inductors meet and exceed the AEC-Q200 Grade 3 standard with a maximum total temperature operation range from minus 55°C to 125°C, while operating reliably in temperatures ranging from minus 40°C to 85°C. The inductor series comes in SMT packaging with a range of compact heights from 3 up to 7 mm, and covers inductance values from 0.10 to 68.0 µH with current ranges from 1.8 A to 118 A. Coiltronics HCMA Series inductors are supplied in tape and 13-in. reel packaging at prices starting at $1.31 each/1,000.

EATON

www.cooperindustries.com

For more information on products or services visit our website www.electronicdesign.com, menu item Reader Service. The advertisers index is prepared as an extra service. Electronic Design does not assume any liability for omissions or errors.

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

Java Turns 20—
So What Else is New?

Oracle's Java turns 20 this year. C++ hit the big 3-5 last year, while C is over 40. Microsoft's C#, the baby of the group, is 15. All of these continue to command a significant following even in the embedded space. All have significant corporate support that helps with their staying power, and they have been continually updated with new features. For example, C++14 allows units *(see "Bjarne Stroustrup Talks about C++14" on electronicdesign.com)*, and lambda expressions are now part of Java, C#, and C++.

Of course, new programming languages keep cropping up, including ones designed for embedded applications. Many look to displace existing programming languages, with others targeting new or emerging platforms like IoT or mobile devices. In theory, they try to simplify the programmer's job and help make applications safer and more secure.

For example, Google Go is a new language from Ken Thompson, Rob Pike, and Robert Griesemer. Ken, as you may know, helped create C *(see "Thompson, Ritchie, And Kernighan: The Fathers Of C" on electronicdesign.com)*.

Apple's Swift looks to replace Objective-C, the company's current programming language. Objective-C has primarily been used for Apple platforms, and Swift is following suit—so if you aren't interested in those platforms, Swift won't be, either. This is very similar to Microsoft's approach for C#, although C# has found some open-source, third-party support.

Swift does not expose pointers, and its syntax has moved toward languages like C++ and Java. Objective-C syntax was closer to Smalltalk. It supports named parameters for methods, while espousing features like protocols (think Java interfaces), categories, and closures.

Mozilla's Rust is another general-purpose programming language that targets highly concurrent and safety-critical applications. It has a syntax similar to C/C++, but it guarantees memory and thread safety. While Rust uses manual memory management, its control of pointers is worth looking at—especially for embedded applications. Its "boxed pointers" have a unique reference to a heap object that will be deleted when the pointer goes out of scope. In addition , there are "borrowed pointers" designed to prevent dangling references. Data referenced using these pointers cannot be deleted, while a borrowed pointer references it. A "borrow checker" provides compile-time checking.

Of course, like almost any practical programming language these days, it has C bindings. What I find interesting are Rust's trait-based generics and pattern-matching support. The latter is not unique, as Scala *(see "If Your Programming Language Doesn't Work, Give Scala A Try" on electronicdesign.com)* and other programming languages have this feature.

Go, Swift, and Rust still have to prove themselves. These and other languages will not have a lot of impact on C, C++, C# and Java in the near term.

C, C++, and Java are still the primary languages for embedded applications, with C leading the charge. The reason is simple: Every processor coming out of the foundry has a set of development tools with a C compiler. Assemblers have essentially disappeared; they still exist, but board support packages (BSPs) and runtime support are almost always in C these days. Even if a developer uses another programming language, it will need a C interface to gain access to base services.

Other tried-and-true languages are available without having to try out a new language. Dynamic scripting languages like Python and Lua are already used in embedded applications, ranging from robot control to wireless Internet-of-Things (IoT) sensor devices. Then, of course, we have JavaScript. While primarily used in web browsers, it also works as a standalone platform.

Duktape is just one open-source C implementation of JavaScript that can be embedded. JavaScript, like Lua and Python, will run on many Cortex-M microcontrollers. The Google Chrome V8 JavaScript engine is written in C++.

Sometimes I wonder if we should not give Ada or Lisp a chance. Ada is 35 years old and has built-in multitasking, plus contracts that make most languages look lame when it comes to safety *(see "Ada 2012: The Joy of Contracts" on electronicdesign.com)*. Its object-oriented support matches that of the newer languages. Lisp has been around for 60 years and is where a lot of the "borrowed" ideas came from for these other languages. Think lambda.

Follow Bill Wong on Twitter @AltEmbedded

Become a fan at facebook.com/ElectronicDesign

DUAL 42V, 2MHz Sync Buck

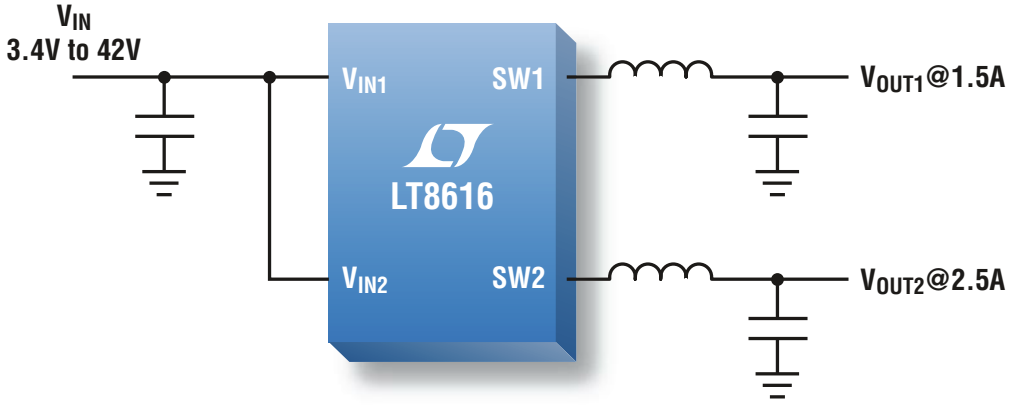

96% Efficient, 2.5A and 1.5A Outputs, 6.5µA I_Q

The LT®8616 brings a new level of performance for 42V input monolithic synchronous buck converters by incorporating dual outputs in a FMEA compliant 28-pin package. Channel 1 is capable of delivering 1.5A of continuous output current, while channel 2 can deliver 2.5A, both with outputs as low as 0.8V. Each channel operates 180° out of phase to further reduce input ripple. Burst Mode® operation requires only 6.5µA of quiescent current (both channels on), thereby extending battery life in battery-powered systems. For high step-down ratio applications, the LT8616 bucks from 16V down to 1.5V at a switching frequency of 2MHz, enabling a compact solution footprint.

Features

- Input Voltage Range: 3.4V to 42V
- Dual Outputs: 2.5A and 1.5A
- Ultralow Quiescent Current: 6.5µA (Both Channels On)
- Low Output Ripple: < 10mV_{P-P}
- Adjustable/Synchronizable Frequency: 200kHz to 3MHz
- High Step-Down Ratios: 16V to 1.5V@2MHz
- 28-Lead 3mm x 6mm QFN and FMEA Compliant TSSOP-28E Packages

Efficiency with $12V_{IN}$ to $5V_{OUT}$ and $3.3V_{OUT}$

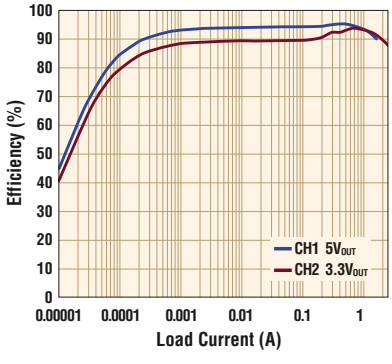

Info & Free Samples

www.linear.com/product/LT8616
www.linear.com/LT86XX
1-800-4-LINEAR

LINEAR TECHNOLOGY