EDN europe

NOVEMBER 2013
www.edn-europe.com
Issue 11

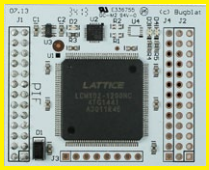

EDN europe

contents

NOVEMBER 2013

DEPARTMENTS & COLUMNS

DESIGNIDEAS

pulse

Cover image

This month's cover shows, not an abstract art installation but a new series of inductors from AVX; the unusual construction is part of what is required to maintain low losses up to the high-GHz region – **more on page 33.**

FTDI Chip

AVAILABLE NOW

DISPLAY MADE EASY

EVE EMBEDDED VIDEO ENGINE

A REVOLUTIONARY SOLUTION ENABLING HI-QUALITY, HUMAN MACHINE INTERFACES AT A LOWER COST

DEVELOPMENT SYSTEMS FROM

$35

DISPLAY AUDIO TOUCH

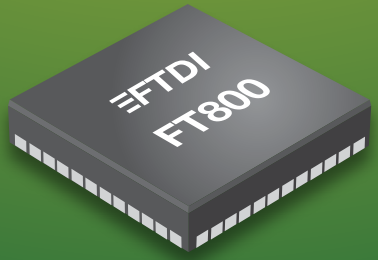

WATCH EVE DEMOS ONLINE

FTDICHIP.COM

EDN europe

NOVEMBER 2013
www.edn-europe.com
Issue 11

CONTACTS

PUBLISHER
André Rousselot
+32 27400053
andre.rousselot@eetimes.be

EDITOR-IN-CHIEF
Graham Prophet
+44 7733 457432
edn-editor@eetimes.be

Patrick Mannion
Brand Director EDN Worldwide

CIRCULATION & FINANCE
Luc Desimpel
luc.desimpel@eetimes.be

ADVERTISING PRODUCTION & REPRINTS
Lydia Gijsegom
lydia.gijsegom@eetimes.be

ART MANAGER
Jean-Paul Speliers

ACCOUNTING
Ricardo Pinto Ferreira

european business press

EUROPEAN BUSINESS PRESS SA
7 Avenue Reine Astrid
1310 La Hulpe
Tel: +32 (0)2 740 00 50
Fax: +32 (0)2 740 00 59
www.electronics-eetimes.com
VAT Registration: BE 461.357.437
RPM: Nivelles
Company Number: 0461357437

EDN-EUROPE is published 11 times in 2013 by European Business Press SA,
7 Avenue Reine Astrid, 1310 La Hulpe, Belgium
Tel: +32-2-740 00 50 Fax: +32-2-740 00 59
email: info@eetimes.be. VAT Registration: BE 461.357.437.
RPM: Nivelles.
It is is free to qualified engineers and managers involved in engineering decisions – see:
http://www.edn-europe.com/subscribe

SALES CONTACTS

Europe

Daniel Cardon
France, Spain, Portugal
+33 688 27 06 35
cardon.d@gmail.com

Nadia Liefsoens
Belgium
+32-11-224 397
n.liefsoens@fivemedia.be

Nick Walker
UK, Ireland, Israel, The Netherlands
+44 (0) 1442 864191
nickjwalker@btinternet.com

Victoria & Norbert Hufmann
Germany, Austria, Eastern Europe
+49 911 93 97 64 42
sales@hufmann.info

Monika Ailinger
Switzerland
+41-41-850 4424
m.ailinger@marcomedia.ch

Ferruccio Silvera
Italy
+39-02-284 6716
info@silvera.it

Colm Barry & Jeff Draycott
Scandinavia
+46-40-41 41 78
jeff.draycott@womp-int.com
colm.barry@telia.com

USA & Canada

Todd A. Bria
West
+1 831 477 2075
tbria@globalmediasales.com

Jim Lees
PA, NJ & NY
+1-610-626 0540
jim@leesmedia.com

Steve Priessman
East, Midwest, South Central & Canada
+1-630-420 8744
steve@stevenpriessman.com

Lesley Harmoning
East, Midwest, South Central & Canada
+1-218.686.6438
lesleyharmoning@gmail.com

Asia

Masaya Ishida
Japan
+81-3-6824-9386
MIshida@mx.itmedia.co.jp

Grace Wu
Asian Sources Publications Asia
(886-2) 2712-6877
wug@globalsources.com

John Ng
Asian Sources Publications Asia
(86-755) 8828 – 2656
jng@globalsources.com

THE IoT BANDWAGON GATHERS SPEED

Unless you have been on a very extended vacation or possibly, given the context, have spent some time with an index finger in each ear and intoning, “La la, can't hear you,” you cannot have failed to notice the growing momentum behind the expression “Internet of Things.” Buzz-words of-the-moment are hardly anything new, of course; if it's bad where you are, I invite you to contemplate what it's like in the world of PowerPoint presentations inhabited by technology editors. Part of the task has become separating today's reality from the rather more fanciful forecasting. Let me quote from a Texas Instruments white paper;

“The IoT of tomorrow: The hotel where I have a reservation knows I am coming and the approximate time of my arrival because I have allowed Apple and Google to track my location. It also knows that I am hot and sweaty from my trip because of the temperature and moisture sensors that are part of my smartwatch. The hotel room I will stay in is currently dormant (no lights, drapes closed, the temperature is at optimised dormant levels). Upon my arrival, the valet knows it is me. He opens my door and the car adjusts the seat because it detects the valet. My preference is to carry my own bag, so I am not accosted by the bell captain. Once in proximity of the hotel lobby, a secure key app is available on my smartphone. By the time I reach the elevator, the room temp has adjusted to coincide with my smartwatch sensors. The light level, music and privacy settings are to my requirements. Because I am hot and sweaty the room also prepares hot water for a shower I will probably take after entering into the room. As I approach, the secure key app unlocks the room door. Once settled for the night, the room detects the lights are turned out, it changes the temperature setting to my sleep preferences.

“In this scenario, every room in this particular hotel chain has multiple sensors and actuators. Every rental car has multiple sensors and actuators. I am wearing multiple sensors and actuators, like a watch vibration for alerts. I am not interacting with my smartphone touchscreen constantly to direct these connected things to take actions even though it is one gateway for my activity. There will be millions of people doing this every day. We will be living IN the data.”

This might, I suggest, sound to you eminently sensible and a joined-up use of technology: or, you may think it sounds like massive overkill, and what's the problem with flicking a few switches for yourself? I'm not going to try to call that argument; the interface between technology, mass-market expectations and fashion trends has become so fluid that it's now extremely hard to spot which items of futurology speculation become tomorrow's must-haves and which are destined to be still-born. (This may be particularly true if you're a European and not an American, and no power on Earth would persuade you to let a “valet” park your car. But I digress.)

There is a lot of this stuff about; and it's not in any way new. A little bit of Google and Wikipedia searching will reveal that the basic ideas of interconnected-everything go back a lot further than you might think; certainly, to well before landmark publications such as Neil Gershenfeld's 1999 book, “When Things Start to Think” which is often cited as a key moment in the growing awareness of where we might be headed. We now have the problem of stepping back from the futurology and asking what is useful, here and now, and what features we need to put in place to prepare for the era of greater connectivity that is to come. We also need to determine exactly what each commentator, and each potential supplier means by the IoT. For some, just having processing power plus connectivity is sufficient; while others only apply the term when there is some aspect of “big data” involved, where the network collates, analyses and then acts upon, information gathered by the nodes. But most seem agree that we'll start with locally-connected systems and evolve to opening those up to the wider IP world.

I don't, in any way, mean to suggest that TI (to take the example above) isn't firmly anchored in the reality of the here-and-now; as part of its programme to kick-start designers into working with connected systems, the company has recently introduced a hardware and software kit for lighting engineers, combining control with remote, or cloud connectivity. The same company has also added Cortex-M4 MCUs that have Ethernet PHY on-chip, to its range.

The same is true of many other suppliers in the field; to take just one other example, Freescale has approached the issue of Internet connectivity by forming partnership with Oracle (and ARM) to bring a standards-based and Java-based, approach to programming for connectivity. The intention is to bring in a platform that is already developed, uses widely-recognised standards, and is secure. Freescale recently held a day-long seminar at its application development centre in Scotland, for application developers and potential end-users. The company's Tim Summers told EDN Europe that delegates and presenters came from a wide range of backgrounds, from developers of connected medical devices, to municipal authorities with responsibilities for environmental monitoring, to machine/machine specialists and condition-monitoring engineers. Summers observes that there was a rapidly growing awareness among the attendees at the seminar awareness among the attendees at the seminar of the potential of connected devices, and of the possibilities of assembling larger data sets, correlating, analysing the data and going on the create real knowledge out of it – something that is beyond the capabilities of many of these organisations and engineers today.

pulse

XMOS' xCORE-XA blends multi-core MCU, ARM core and low-power operation

Fabless multi-core microcontroller maker XMOS is aiming to tap into the ARM world, with its resource of existing software and embedded applications code: and to take its technology into the low-power, battery-powered arena – in one move, with the introduction of its xCORE-XA (for "extended architecture") devices.

XMOS has concluded a deal with ARM to add a Cortex-M3 core into its multi-(virtual)-core architecture and with Silicon Labs/Energy Micro to bring in some of the "energy-friendly" techniques used in that company's Gecko chip series.

Initial offering will be a multi-chip package with two dice, a 7-core xCore device and a low-power Silicon Labs contribution that will add resources such as Silabs' low-power peripherals and autonomous peripheral-peripheral channel that allows many functions to complete without waking a central processor. The xCore part of the solution is a single-physical-core device that would normally run eight of XMOS' virtual cores, with one of the eight "slots" now devoted to the ARM core which will connect into the architecture via an AHB channel.

With this device, according to XMOS CEO Nigel Toon, you will be able to create a configurable system-on-chip for an embedded application, that you can program entirely in C, with multi-core capability and at low-power. XMOS presents its existing xCore devices as occupying a gap between conventional MCUs and FPGA-based designs; now it adds low-power to that equation. Toon asserts that FPGAs have developed to offer great flexibility, but with little focus on low-power. He says the XA will have five power modes, drawing on the multiple levels that Silabs/EM products offer, with the lowest offering as little as 100 nA in standby. It requires less than 1µA to run the integrated real-time clock and 32 kHz peripherals, for fast turn-on and time-polled operation. Differentiating factors will be XMOS' deterministic attributes; the types of peripherals on offer, which will include dual Ethernet ports and the option of USB, with the ability to run sophisticated functions such as motor controls, using separate cores.

ARM, in the person of Nandan Nayampally, vice president, Application Systems Marketing, endorses the concept; "ARM believes that xCORE-XA represents a significant step forward for embedded systems, allowing engineers to create an integrated SoC that is configured completely in high level software."

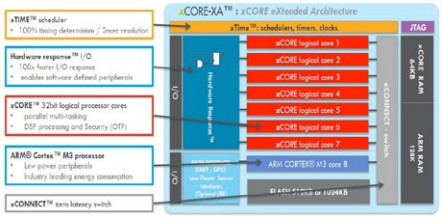

Multi-core MCU plus ARM core makes for "a completely new class of programmable IC", says XMOS.

The new architecture allows embedded system designers to use high-level software to configure a device with the exact set of interfaces and peripherals needed for their design, while at the same time re-using existing ARM binary code and taking advantage of ultra-low energy peripherals. Designers can also add real-time data-plane plus control processing and DSP blocks, using multiple xCORE processor cores, with the ARM available to run larger control plane processing software such as communication protocol stacks, standard graphics libraries, or complex monitoring systems.

Toon says that, for the first time, embedded system designers no longer have to choose between expensive, power hungry, programmable logic devices, inflexible fixed-function alternatives, or traditional microcontrollers that lack computing power and are constrained by a hardware-defined peripheral set.

The configurable xCORE multicore microcontroller technology features multiple deterministic 32-bit processor cores that execute high-level code concurrently. It allows customers to software-configure the exact combination of peripherals and interfaces required for their designs, and enables demanding hard real-time requirements with timing-precise execution. It also delivers advanced DSP and security processing. xCORE-XA extends these capabilities, offering designers access to the rich ARM ecosystem, including standard code libraries that can dramatically accelerate product design time.

The first device in this new family, the XA-U8-1024, features eight 32-bit processors (seven xCORE logical cores plus an ARM Cortex-M3 processor), 192 KB SRAM, and 1024 KB of Flash. The device includes a low energy USB interface, ultra low-energy peripherals and analogue functions including ADC, DAC, op-amps and capacitive sensing comparators. Future family members will include 6-core and 8-core products with Flash sizes ranging from 512 KB to 1024 KB, and device variants with or without the low-energy USB 2.0 interface.

XA-U8-1024 can make use of the range of xSOFTip software peripherals from XMOS's growing library and is supported with an integrated design-flow, by the xTIME composer Suite of tools that includes comprehensive design entry, compilation and debug support for both the ARM and multiple xCORE processor cores. It will sell for $16.39 in low volumes.

"The xCORE-XA extended architecture redefines what embedded developers can achieve using a programmable platform," says Toon, " We believe xCORE-XA represents a completely new class of programmable SoC. Now we can bring system level programmable configuration with hardware levels of real-time performance to low-energy battery-powered applications, configured and programmed completely in software."

XMOS
www.xmos.com

Toshiba steps up LED-lighting presence with 1W GaN-on-Si white LEDs

Toshiba says its second generation of 1W LEDs will drive down costs for general-purpose and industrial lighting.

Toshiba has announced the first devices in its second generation of LETERAS white light-emitting diodes fabricated using a gallium nitride-on-silicon (GaN-on-Si) process. The company is using its 200-mm wafer fabs in Japan and employing technology to grow the GaN on silicon wafers, that it developed along with Bridgelux. This, the company says, will give it a continuing cost advantage over competitors using GaN-on-sapphire substrates. Meanwhile, it promises further developments in luminous efficacy and colour performance. The 1W TL1F2 LEDs offer a cost-competitive alternative to current LED packages, allowing manufacturers of general purpose and industrial LED lighting to drive down costs.

Luminous efficacy of the TL1F2 white LEDs has been improved compared with the TL1F1 series by optimising the package and increasing the optical output power of the GaN-on-Si LED chips. The TL1F2 series offers a full correlated colour temperature (CCT) range from 2700K to 6500K, with minimum colour rendering index (CRI) values of 80 and 70 respectively. Typical luminous flux of the 1W LEDs ranges from 104 lumen to 135 lumen depending on colour temperature and CRI.

The new devices are supplied in a standard 6450 package measuring 6.4 by 5.0 by 1.35mm. Typical driving current (IF) is 350mA, with a typical forward voltage (VF) of 2.85V helping designers to reduce system power consumption. An operating temperature range of -40°C to 100°C makes the TL1F2 series suitable for both indoor and outdoor use in applications such as lamps, ceiling lighting, street lights and floodlights.

Toshiba Electronics Europe
www.toshiba-components.com

Single-chip software-defined radio handles multiple signal types

An RF Agile Transceiver provides up to three times the noise performance of alternative solutions and significantly reduces BOM cost; it comes with design kit, and FPGA rapid prototyping environment to reduce design times and risk.

Analog Devices' AD9361 is a monolithic software-defined radio chip that contains two complete transmit and receive channels, can operate anywhere in the spectrum from 70 MHz to 6 GHz, and with channel bandwidths from 200 kHz to 56 MHz. In designing it to enable programmable radio applications that operate over a wide range of modulation schemes and network specifications such as defence electronics, instrumentation equipment and communications infrastructure, ADI also describes it as an RF Agile Transceiver. An ADI spokesman disclosed that the device has in fact existed for some time, and is already the basis of products developed by a limited customer group, and that Analog has now chosen to make it available to a wider market. In part, the spokesman explains, this approach has been driven by the fact that applications for the SDR chip are so diverse that the company has had to grow an infrastructure and design-support package so as not to have its own design-support facilities overwhelmed by customer requests.

The AD9361 is supported by a wide range of design resources to expedite time to market including a software design kit and FPGA mezzanine card (FMC) to rapidly develop software defined radio solutions. One of the early customers quoted by ADI is Ettus Research, a National Instruments company; President Matt Ettus comments, "The AD9361 Agile Transceiver is a complete RF transceiver solution on a single chip – it's an RF architect's dream device...we combined it with a Xilinx Spartan-6 FPGA, USB 3.0 interface and comprehensive software support, to create easy-to-use and flexible software-defined radio solutions."

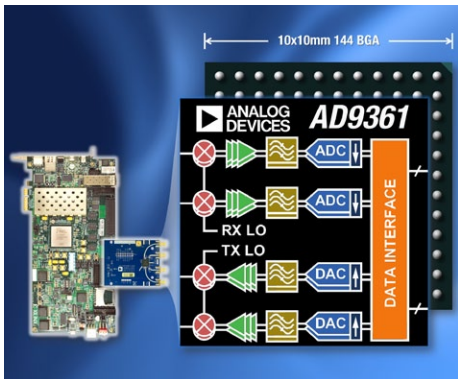

Selectable blocks such as front-end LNAs, combined with programmable analogue functions, comprise a complete SDR chip.

The transceivers integrate an RF front end with a selection of low-noise-amplifiers, flexible mixed-signal baseband section, frequency synthesisers, two analogue-to-digital converters and two direct conversion receivers. Functional blocks are combined, alternate blocks such as LNAs selected and parameters (such as local-oscillator frequencies – there are full fractional-N synthesis LO blocks on-chip) set according to the needs of any given signal, and everything can be changed and reconfigured on-the-fly. Among other capabilities, you can build a complete 2x2 MIMO configuration with the chip.

Two independent direct conversion receivers have a state-of-the-art noise figure and linearity. Each receive subsystem includes independent automatic gain control, DC offset correction, quadrature correction, and digital filtering, eliminating the need for these functions in the digital baseband. The AD9361 also has flexible manual gain modes that can be externally controlled. Two high-dynamic-range A/D converters per channel digitise the received I and Q signals and pass them through configurable decimation filters and 128-tap FIR filters to produce a 12-bit output signal at the appropriate sample rate. The transmitters use a direct conversion architecture that achieves high modulation accuracy with ultra-low noise.

Design resources include a Software Design Kit and FMC Board. Together with the FPGA mezzanine cards, ADI offers a range of AD9361 design resources including Gerber files, code references, Linux sample applications and drivers, and design support packages, which are available for download.

At 1000-quantities, the AD9361 costs $175 each; the AD-FMCOMMS2-EBZ FMC board costs $750.

Analog Devices
www.analog.com/ad9361rftransceiver

Altera/Enpirion DC-DCs aim to solve FPGA-power-supply challenges

Power converter configurations specifically targeted at providing the correct power feeds to large Altera FPGAs will improve system power efficiency by up to 35% while reducing board area by up to 50%.

Altera's acquisition of Enpirion, designers of small-outline high-efficiency DC/DC converter modules, was one of the more unexpected corporate merger-and-acquisition deals – now, the company has announced the first results from the joint operation that are specifically aimed at Altera FPGAs. According to Altera's Mark Davidson, marketing director, power business unit, the company's customers often face big challenges in providing correct power supplies to its FPGAs. There are multiple power rails, requiring precisely-controlled levels at low voltages and high currents: demanding transient response performance must be met; and in addition, correct sequencing at power-up and power-down is essential. Its customers, Davidson says, have had to apply, "engineering bandwidth to areas that are not key [to designing the actual functionalities of the FPGA]"

Accordingly, Altera has announced "Power-optimised FPGA reference designs" - both product and design process – to simplify the development of FPGA-based systems. Primarily for Cyclone V parts, there are four reference designs in this release. Davidson says Altera's parts built in 14-nm silicon technology will require ±3% voltage rail stability and as with all complex FPGA designs, actual currents will be dependent on the configuration and operating parameters of the FPGA so point-of-load regulation will need to handle large current variations. Up to five power rails are needed for core, I/O, PLLs, peripherals and SERDES functions, with appropriate sequencing. Davidson emphasises the need for voltage stability, "[when] supplying a SERDES [function], noise on the rail appears as jitter on the data stream."

"All FPGAs need significant bulk capacitance," he adds, before going on to detail how the solution reduces that need.

Enpirion brought expertise in three area, in its miniature-packaged-DC/DC designs; efficient high-frequency switching in CMOS; miniature magnetics; and the packaging itself. The four reference designs exploit those features, as "turnkey" power solutions that increase power efficiency by up to 35% (that is, reduce the losses by that percentage), reduce board area (used by the regulation) by up to 50%, and reduce overall bill-of-material (BOM) bulk capacitance costs by up to 50%. Multiple high-quality tantalum capacitors can amount to a significant cost in themselves. The designs are offered to customers as downloadable design packages and are demonstrated in hardware on development kits. A design package targeting Cyclone V SoCs is available for download, with additional design packages targeting 28-nm FPGAs available later in 2013. The designs include optimised board layouts and detailed schematics and have been verified by Altera to work correctly with its FPGAs and SoCs. Each design package contains the schematics for the FPGA's power circuits, application notes, bill-of-materials list, power tree layout guidelines and Gerber files for individual power components.

The first power reference design available for download targets the Cyclone V SoC. Power reference designs targeting the Stratix V, Arria V and Cyclone V FPGAs will be available later this quarter. There are four corresponding 28 nm development kits. Among the features of the designs is that regulation is good enough to enable provision to the SERDES function with a switching design; previously, usual practice has been to use a linear regulator in this function for noise reasons, with a corresponding efficiency penalty. "That design achieves the same, or lower, jitter, than the LDO version," Davidson says. You can use fewer tantalum capacitors, or an all-ceramic design may be feasible with attention to impedance issues; and the reference design will always work, irrespective of the FPGA configuration. Switching frequencies are in the range 2 to 9 MHz – the technology is capable of 18 MHz, but Davidson notes that this is not easy to use at present.

Going forward, Altera says that Enpirion's primary goal will be to develop power solutions for its FPGAs – but the unit will continue to provide product to the wider market and will support applications that have similar needs for precisely regulated power.

The current Enpirion/Altera designs do not have a bus connection, such as PMBus, to control/report to a system power monitor but Davidson says that, "the way the FPGA business is going, that is a high probability." In these designs, some parameters are set by resistor value selection, and some by digital register values. Davidson concludes, "Discrete power [solutions] are always cheaper, but here Altera/Enpirion has the opportunity to save customer hassle while selling a premium product."

Altera
www.altera.com/enpirion

Wireless charging receiver delivers 400mA across a 1.2cm air gap

Linear Technology's LTC4120 marks the company's entry to the wireless battery charging space, for which it has initiated a cooperation with new entrant PowerbyProxi. The LTC4120 combines a wireless power receiver with a constant-current/constant-voltage battery charger, functioning as the receive circuit component in a complete wireless power transfer system comprised of transmit circuitry, transmit coil, receive coil and receive circuitry.

The LTC4120 works with Linear Technology's simple discrete resonant transmitter reference design or with advanced off-the-shelf transmitters designed and manufactured by PowerbyProxi, a New Zealand-based wireless power solutions company. PowerbyProxi transmitters offer features including simultaneous charging of multiple receivers with a single transmitter and foreign object detection to prevent excessive heating during transmit faults.

The LTC4120 accepts a rectified 4.2V to 40V input from the receive coil to power a constant-current/constant-voltage 400mA battery charger that includes programmable charge current, a programmable 3.5V to 11V battery float voltage with ±1% accuracy, battery preconditioning with half-hour timeout, bad battery fault detection, NTC thermal protection, charge status and a 2-hour safety termination timer.

The LTC4120 features a technique that PowerbyProxi calls dynamic harmonisation control (DHC), that enables optimal wireless power transfer across a variety of conditions while providing thermal management and overvoltage protection. This technique modulates the resonant frequency of the receiver tank to provide lossless adjustment of the power received as well as the power transmitted to enable an efficient and robust solution for wirelessly charging battery-powered devices.

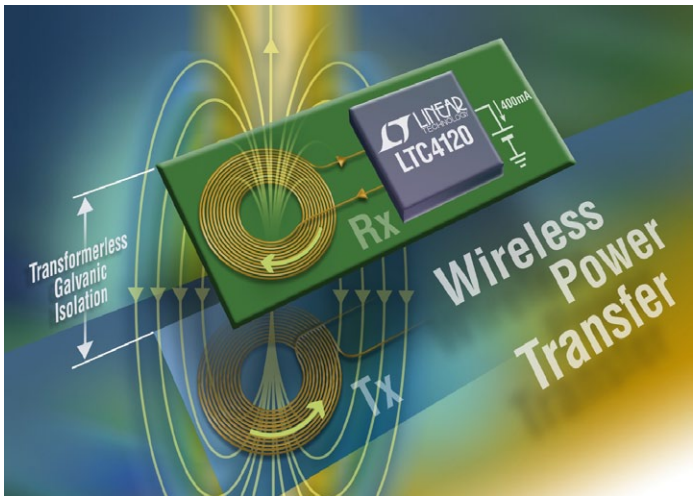

A partnership between Linear Technology and PowerbyProxi has produced a wireless charging system for which the companies claim ease of use as a key feature.

Wireless charging, Linear comments, enables battery powered devices to be recharged without expensive, failure-prone connectors. Products incorporating the LTC4120 may be contained within sealed enclosures, in moving or rotating equipment or used where cleanliness or sanitation is critical. Applications include handheld instruments, industrial/military sensors and similar devices in harsh environments, portable medical devices, physically small devices and electrically isolated devices. LTC4120-based systems offer solutions that, Linear claims, are much simpler than those implementing the Qi standard, with additional benefits, including greater transmission distance and no software required.
The LTC4120 is housed in a low profile (0.75 mm) 16-pin 3 x 3mm QFN package with backside metal pad for thermal performance. The device is guaranteed for operation from -40C to 125C, in both E and I grades. Pricing starts at $3.55 (1000).

Linear Technology,
www.linear.com/product/LTC4120
PowerbyProxi
http://powerbyproxi.com

Gecko range of "energy-friendly" MCUs gains ARM Cortex-M0+ Core-based chips

Following its acquisition of Energy Micro, Silicon Labs has introduced the EFM32 Zero Gecko MCUs that provide low-energy operation, with 32-bit processing, at prices down to $0.49 in high volume.
Calling them the industry's most energy-friendly 32-bit MCUs based on the ARM Cortex-M0+, Silabs says that EFM32 Zero Geckos are designed to achieve the lowest system energy consumption for battery-powered applications such as mobile health and fitness products, smart watches, activity trackers, smart meters, security systems and wireless sensor nodes, as well as battery-less systems powered by harvested energy. The family includes 16 variants; as with previous Gecko chips, its designers say that lowest overall power usage comes from attention to each of active power, standby power and the power needed to wake the device to an active from any of its four sleep modes. Support continues to include an energy-profiling tool that will accurately predict battery life for a given application.
In deep-sleep mode, Zero Gecko MCUs have 0.9 µA standby current consumption with a 32.768 kHz real-time-clock source, RAM/CPU state retention, brown-out detector and power-on-reset circuitry active. Active-mode power consumption scales down to 110 µA/MHz at 24 MHz using the example of real-world code (a prime number search algorithm) executed from flash. Current consumption is less than 20 nA in shut-off mode. The EFM32 MCUs further reduce power consumption with a 2-microsecond wakeup time from standby mode.
Carried over from previous parts is the Peripheral Reflex System (PRS) that helps system-level energy efficiency; it monitors complex system-level events and allows different MCU peripherals to communicate directly with each other and autonomously without involving (waking) the CPU. Also continued from Tiny Gecko, Giant Gecko and Wonder Gecko devices are low-energy analogue peripherals include an analogue comparator, a supply voltage comparator, an on-chip temperature sensor and a 12-bit ADC with 350 µA current consumption at a 1 MHz sample rate.

Positioned by Silicon Labs as the lowest-power "Gecko" range of MCUs to date, these devices host an ARM Cortex-M0 core.

New is a programmable current digital-to-analogue converter (IDAC), and Silabs says the EFM32 Zero Gecko devices are the only Cortex-M0+ MCUs with this facility on-chip. It generates a biasing current from 0.05-64 µA with only 10 nA overhead. The IDAC provides an accurate bias and/or control capability for companion ICs and other external circuits including amplifiers, sensors, Wheatstone bridges and resistor ladders, eliminating the need for external power amplifier components for many cost-sensitive applications.
The family is pin- and software-compatible with other EFM32 Gecko MCU products and provides a low-cost entry point to higher-performance Gecko devices including the Wonder Gecko based on the ARM Cortex-M4 core, with DSP instruction set and dedicated floating point unit (FPU). A full range of tools, including third-party ARM tool chains, supports the introduction. Silab's Simplicity Studio is a free download. In QFN and QFP packages with 4 to 32 kB flash sizes, high-volume prices (100,000) begin at $0.49.

Silicon Labs
www.silabs.com/zero-gecko

Drive appliance motors with 3-phase integrated PFC inverter

International Rectifier has added to its iMOTION design platform which integrates digital, analogue and power technologies; the IRAM630-1562F intelligent power module (IPM) has in-built Power Factor Correction (PFC) ahead of its inverter stage. Applications include air conditioners and washing machines.

Introducing its latest iMOTION module, IR's Marco Palma places it in a 10-year development story that is now changing at an exponential pace. Although different markets see different adoption rates, the application that was served by the induction motor now turns by default to the brushless motor. Algorithms that a short time ago were treated with caution are now accepted, he says, and the biggest growth sector of all for this technology is variable speed designs for (for example) washing machines and airconditioners. Sensorless rotor angle estimation derived from motor current is key to this, he adds.

Palma was launching the IRAM630-1562F that combines IR's low loss, trench IGBTs with a three-phase, high-voltage gate drive IC and PFC input stage, integrating more than 30 components into an isolated package. Built-in over-temperature/over-current protection, along with an integrated under-voltage lockout function, and built-in temperature monitor provide for a high level of protection and fail-safe operation. Other integrated features include bootstrap diodes for the high-side drive function and single polarity power supply.

It offers 15A current rating at a case temperature of 25C; delivers up to 2 kW to a motor; has a PFC peak current of 30A; and occupies an SIP2 package. The IPM's open emitter configuration enables multi-shunt current feedback for a sophisticated vector control loop, in a V/Hz control loop, with no circuit layout restrictions. Over-temperature is detected internally and triggers a fault condition. IR says that EMI emissions are also minimised due to shorter connection routing, component layout and internal shielding. The module costs $9.25 (10,000).

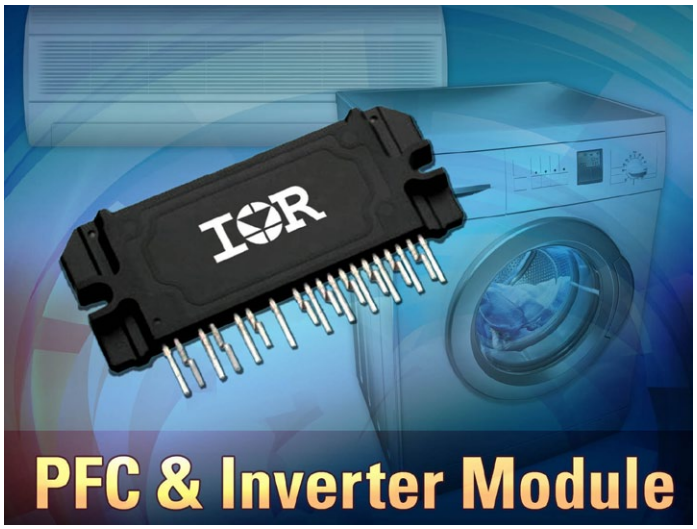

More compact and more efficient drives for appliance and light-industrial motors result from the integration in IR's latest iMOTION product.

International Rectifier
www.irf.com

Lowest-profile, highest density 180W PSUs have medical approvals

XP Power's ECP180 series of low-profile high efficiency "green power" open frame 180W AC-DC power supplies claims to deliver higher power outputs from a 2 x 4-in. standard outline. The power levels delivered by the units – 120W with convection-only cooling, and 180W with 10 cfm airflow – would more typically be expected of a 3 x 5-in. outline, relative to which the ECP180 takes a 46% smaller footprint, XP says. The units also have a low profile of 1 in. (25.4 mm.) XP believes there is, "nothing [in the market] that will give this much power from this volume."

The IT- and medically-approved units have an average active mode efficiency of 93% and a no load power consumption of less than 0.5 W. Their published efficiency curve is essentially flat from 30 – 100% of full load and declines only to around 84% at 10% of full load. With PFC and a universal input range of 85 to 264 VAC, the ECP180 is available in six output voltage models providing the popular nominal outputs of +12, +15, +24, +28, +36 or +48 VDC. A secondary 12 VDC / 0.5 A fan output is standard across the range. There is no derating of output until 50C ambient temperatures, or with lower output voltages; "many supplies in the market derate well below 50C and at voltages under 24V," according to the company. Extended operating temperature range from -20 to +70C is standard.

Meeting the ITE safety specifications of UL/IEC/EN 60950-1 and the latest 3rd edition medical safety standards ANSI/AAMI ES 60601-1 and IEC/EN60601-1, the ECP180 series is suitable for use in a wide range of information technology, industrial and medical equipment applications. They also have a 5000-m altitude rating. The units also comply with the internationally recognised EMC standards EN55011 and EN55022 Class B specifications for conducted noise emissions. The ECP180 has a 3 year warranty.

XP notes that in their portfolio, supplies for healthcare applications – in general, those requiring the medical approvals – is a steady growth sector, and is particularly strong in their North American sales figures; the company also notes that "green" factors are an increasing part of customer requirements in power supplies – not only the efficiency and standby performance of the units, but environmental considerations that extend to how the factories that make the products are operated.

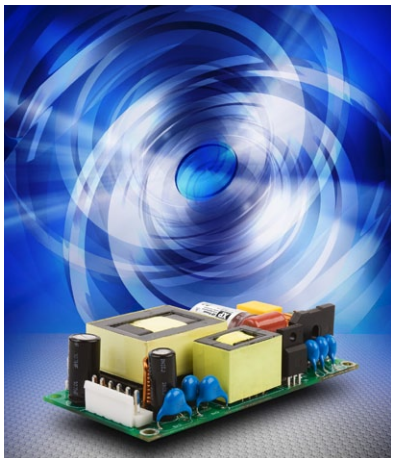

You cannot get more power out of any other PSU in the 2 x 4-in. standard outline, XP asserts.

XP Power
www.xppower.com

Cadence adds to IP offering for analogue in 28-nm, and in Tensilica audio blocks

Cadence Design Systems has recently been expanding its profile in IP (intellectual property) for SoC chip design and has added to its offering with three separate announcements. The company's Christian Malter, director of technology solutions and business development in the EMEA region, noted the growing importance of the IP offering and the company acquisitions that Cadence has made: he comments, "Ten years ago IP was not profitable – now it is." Malter was introducing new IP including high-performance data converters for 28-nm silicon processes; audio DSP for the Dolby Digital Plus audio stream; and, also in the audio space, an IP core for DTS Neural Surround.

The converter technology comprises two ADCs and two DACs that are ten-times faster than any other competing IP, Cadence claims. These analogue designs are intended for next-generation high-speed wired and wireless communications applications, such as WiGig (802.11ad), which runs on a 60 GHz spectrum with potential data throughput up to 7 Gbps, as well as LTE and LTE Advanced.

The data converter family includes: 7-bit 3 GSPS dual ADC and DAC; 11-bit 1.5 GSPS dual ADC; and a 12-bit 2 GSPS dual DAC. The data converter IP cores combine to form a complete analogue front end (AFE) IP solution for wired/wireless communications, infrastructure, imaging and software-defined radios.

The ADC IP cores are developed with a parallel Successive Approximation Array (SAR) architecture, producing extremely fast and scalable sample rates. High effective-number-of-bits (ENOB) values are achieved with a unique implementation and built-in background auto calibration, producing more accurate conversion and consistent performance, Cadence claims. The IP includes features such as differential data inputs, reference and timing generator, internal offset correction, and voltage regulators for improved supply noise immunity.

The DAC IP cores use a current switching architecture and include a digital multiplexer and FIFO for easy integration into an SoC. The DACs include digital gain control and all required reference circuitry.

All the IP includes multi-level power-down modes for additional power savings, a built-in analogue test bus for design testability, and single-ended CMOS or differential Current-Mode Logic (CML) clock inputs for a flexible clock interface. The IP provides matching dual channels for communication systems where these are desired, simplifying implementation and reducing risk, and a standard CMOS process target for easy manufacturing.

From the product line that was formerly Tensilica, Cadence has introduced audio DSP IP supporting Dolby Digital Plus with DS1, which will "bring clearer voice and richer sound to small speaker systems, from mobile devices to flat-panel TVs", enhancing audio quality in mobile devices and small speaker systems. The HiFi Audio/Voice DSP is the first IP core to offer a certified decoder for Dolby DS1 for Dolby Digital Plus, which Cadence notes as having been licensed by over 55 companies accounting for over 200 million HiFi DSP cores in smartphones, tablets, computers, digital televisions, home entertainment systems and other devices.

More; www.tensilica.com/products/audio.

Also from the Tensilica line, Cadence has an IP Core supporting DTS Neural Surround, which is part of the HD Radio specification and which aims to provide an immersive and natural surround sound for cars and A/V receivers, enhancing upmixing performance from MP3 and other compressed music formats.

DTS Neural Surround technology uses advanced techniques to encode up to 7.1 channels of audio to stereo while maintaining the surround queues from the discrete digital multi-channel track. The two-channel audio can then be broadcast in stereo or converted back up to 7.1 channels on the consumer end product enabled with Neural Surround, providing an experience claimed to be, "as close to the original discrete multi-channel audio as is available today."

Cadence
www.cadence.com

Qt package speeds embedded device creation with rich graphics

Digia's Qt Enterprise Embedded enables fast development and deployment with Android and Embedded Linux stacks; the fully-integrated solution enables users to get started quickly with software development on an embedded device with a tailored user experience.

Digia says that the package continues to support its approach of cross-platform development; write once, and deploy on multiple targets, with benefits for product quality, maintainability, and speed of development. It supports, the company adds, the ability to deliver the user experience that consumers have come to expect from products such as smartphones, to all classes of embedded device. Qt Enterprise Embedded comprises two main parts: a development environment and a software stack, called the Boot to Qt Software Stack. The self-contained development environment is installed and updated through a single online installer and features a full Qt Creator IDE. Development is done with Qt Enterprise libraries and device deployment can be achieved with one click directly from Qt Creator to a target connected via USB or network. Alternatively, the application can be deployed to an emulator on the host system that is running the same software stack as the target device. The Boot to Qt software stack is a lightweight, Qt-optimised, full software stack that is installed on the target hardware. The stack is supplied in embedded Android and embedded Linux formats.

Qt's uses OpenGL ES 2.0, the version of the graphics application programming interface designed specifically for embedded systems and mobile devices. This makes it easier to develop and deploy rich graphics with "velvet-like" animations and transitions as well as smoothly rendered 2D and 3D animations on devices with relatively limited performance. In addition to the versatile cross-platform Qt APIs, applications built with Qt can also use platform and device-specific APIs directly, giving the developer the freedom to integrate any functionality.

The Qt Enterprise Embedded development environment runs on Ubuntu Linux 64-bit 12.04 LTS or later. Target hardware supported as standard at time of launch includes: Google Nexus 7 tablet (Tegra 3, ARM Cortex-A9); Beagle Board xM (ARM Cortex-A8); Boundary Devices SabreLite (Freescale i.MX 6); Raspberry Pi Model B (ARM11); and BeagleBone Black (TI AM335x). Digia is also able to port the software stack to custom hardware.

A 30-day free trial offer from Digia is at;
http://qt.digia.com/qtenterpriseembedded /
http://qt.digia.com

BY YURIY KURTSEVOY AND STUART SMITH, MAXIM INTEGRATED

A HIGH-ACCURACY 4–20MA CURRENT-LOOP TRANSMITTER FOR TOUGH INDUSTRIAL REQUIREMENTS

The 4–20mA current loop has been widely used as an analogue communication interface in industrial applications. It facilitates transmission of data from remote sensors over a twisted-pair cable to a programmable logic controller (PLC) in a control centre. Simplicity, reliable data transfer over long distances, good noise immunity, and low implementation cost make this interface well suited for long-term industrial process control and automated monitoring of remote objects.

To no one's surprise, industry is evolving just like all electronic applications today. It has more stringent demands, new requirements for higher accuracy; lower power; reliable operation over an extended -40°C to +105°C industrial temperature range; added security and system protection; and implementation of the digital Highway Addressable Remote Transmitter (HART) protocol. Collectively, these requirements make the design of today's 4–20mA current loop quite challenging.

This article explains how to develop a 4–20mA current-loop transmitter, analyse its performance, and select the components that meet rigorous industrial requirements. Test data for error analysis, over-temperature characterisation data, schematics, and analysis software are provided.

KEY DESIGN PARAMETERS

We start by focusing on the new reference design. The block diagram in Figure 1 shows the high-performance, low-power, 4–20mA current-loop transmitter that reduces component count and yields the best results for price versus performance.

This reference design uses low-power, high-performance components that provide less than the 0.01% at 25°C and less than 0.05% over the temperature range for industry's most demanding 4–20mA current loop. The design features a low-power 16-bit DAC (U1), a zero-drift rail-to-rail input/output (RRIO), high-precision op amp (U2), a voltage reference (U3), and a 40V low-quiescent-current LDO (U4).

The U3 voltage reference provides low noise, low temperature drift of 5ppm/°C (max) and a precise 2.500V for U1. The smart sensor microcontroller commands U1 through a 3-wire SPI bus. The U1 output is divided and converted to the loop current by the Q1 power MOSFET, 10Ω (±0.1%) sense resistor (RSENSE), and U2. The U1, U2, and U3 devices are powered by U4, which is powered directly from the loop. There is a current-limiting circuit made with Q2, a BJT transistor, and sense resistor (R6). This circuitry limits the loop current to approximately 30mA, which prevents runaway conditions and any damage to an ADC on the PLC side. The Schottky diode (D1) protects a transmitter from reverse current flow.

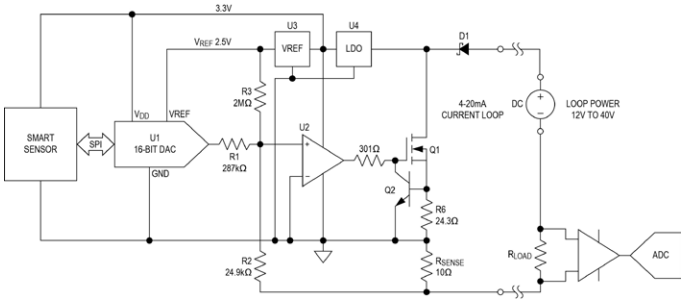

Figure 1. *Reference design for a 4–20mA loop-powered transmitter uses a MAX5216 16-bit DAC (U1), MAX9620 op amp (U2), MAX6133 voltage reference (U3), and MAX15007 LDO (U4).*

DESIGN PERFORMANCE

The reference design operates at low power. The maximum current consumption of the selected components is less than 200 µA at +25°C and less than 300 µA over the -40°C to +105°C temperature range. The U2 op amp has a 25 µV (max) zero-drift input offset voltage over time and temperature, so it is ideal for the accuracy and stability requirements of the application. The 10Ω current-sense resistor allows operation with a low loop-supply voltage; its smaller resistance dissipates less power and allows use of a smaller package, which further shrinks this transmitter. For example, if only a 10Ω RSENSE and 10Ω load are present, then the maximum voltage drop on them at 30 mA is 600 mV. The U4 LDO requires only 4V for proper operation with a 3.3V output, and the total minimum loop supply can be as low as 5V. However, if the PLC load is 250Ω, then the minimum loop supply must be 4V + 30 mA × (10 + 250)Ω = 11.8V.

Note that to determine a more accurate estimation of the minimum loop supply voltage, the loop cable resistance must also be considered.

During testing the output exhibited some noise at 10Ω. Increasing the value of the RSENSE resistor will increase power dissipation and a minimum loop supply voltage, but it will also

Table . 4–20mA Current-Loop Transmitter Error Analysis

	Tol (±%)	min	nom	max	
V_{REF}	0.04	2.4990	2.5000	2.5010	V
R1	0.1	286.71	287	287.29	kΩ
R2	0.1	24.88	24.9	24.92	kΩ
R3	1	1980.00	2000	2020.00	kΩ
R_{SENSE}	0.1	0.00999	0.0100	0.01001	kΩ
		min	nom	max	
Zero-scale DAC code		0	0	0	*dec*
Zero-scale I_{OUT}		3.07430	3.11250	3.15149	mA
4mA DAC code		2806	2682	2555	*dec*
4mA I_{OUT}		3.99984	4.00015	3.99999	mA
Full-scale DAC code		65535	65535	65535	*dec*
Full-scale I_{OUT}		24.69058	24.8024	24.91527	mA
20mA DAC code		51314	51025	50734	*dec*
20mA I_{OUT}		19.99988	20.00007	19.99995	mA
4mA error		-0.00410	0.00381	-0.00016	%FS
20mA error		-0.00061	0.00035	-0.00026	%FS
24mA I_{OUT} DAC code		63441	63111	62779	*dec*
24mA I_{OUT}		23.99989	24.00013	24.00002	mA

Table 1. *4–20mA Current-Loop Transmitter Error Analysis*

Table 2. Temperature Error Analysis of 4–20mA Current-Loop Transmitter

	TC (± ppm/°C)	min	nom	max	
VREF	5	2.4991	2.5000	2.5009	V
R1	10	286.7919	287	287.2081	kΩ
R2	25	24.8549	24.9	24.9451	kΩ
R3	100	1985.5000	2000	2014.5000	kΩ
RSENSE	10	0.00999	0.0100	0.01001	kΩ
		min	nom	max	
Zero-scale DAC code		0	0	0	*dec*
Zero-scale IOUT		3.08114	3.11250	3.14433	mA
4mA DAC code		2806	2682	2555	*dec*
4mA IOUT		4.00647	4.00015	3.99302	mA
Full-scale DAC code		65535	65535	65535	*dec*
Full-scale IOUT		24.69253	24.8024	24.91297	mA
20mA DAC code		51314	51025	50734	*dec*
20mA IOUT		20.00289	20.00007	19.99655	mA
4mA error		0.04047	0.00095	-0.04362	% FS
20mA error		0.01807	0.00044	-0.02157	% FS
24mA IOUT DAC code		63441	63111	62779	*dec*
24mA IOUT		24.00199	24.00013	23.99751	mA

Table 2. *Temperature Error Analysis of 4–20mA Current-Loop Transmitter*

reduce noise on the loop. This is a trade-off that the user can control.

The U2 op amp tracks the voltage drop across R2 and RSENSE, and maintains 0V at both of its input nodes. The following equations are used for this circuitry:

$$Iout = I(R2) * \frac{R2}{Rsense} \quad \text{(Eq. 1)}$$

$$I(R2) = I(R1) + I(R3) \quad \text{(Eq. 2)}$$

Where:

IOUT is the loop current

I(R2) is the current flowing through R2

I(R1) is the current flowing through R1

I(R3) is the current flowing through R3

In Equation 2 we assume that the input current to IN+ and IN- of U2 is 0. Following Equations 1 and 2, the initial loop current of 4 mA is set by the I(R3) current while I(R1) is 0. Therefore:

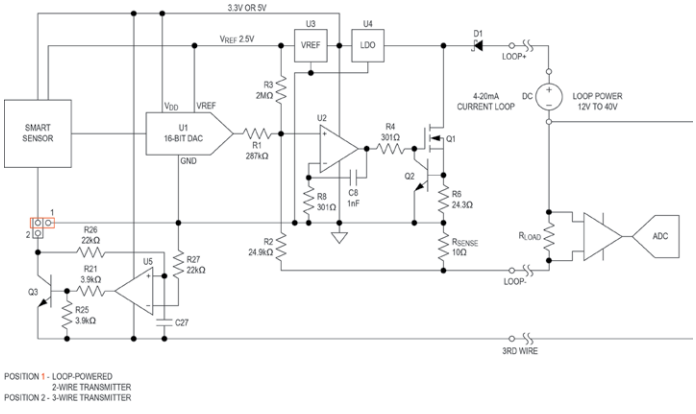

Figure 2. *Block diagram for a universal 2- or 3-wire smart sensor transmitter.*

$$Iout_{init} = I(R3) * \frac{R2}{Rsense} \quad \text{(Eq. 3)}$$

Current through R3 is equal to the U3 voltage reference output divided by R3. Equation 3 can be expanded as:

$$Iout\_init = \frac{VREF}{R3} * \frac{R2}{Rsense} \quad \text{(Eq. 4)}$$

According to the Namur NE43 recommendations for failure information transmitted over a 4–20mA current loop, the signal range for measurement information is from 3.8 mA to 20.5 mA, allowing for a small amount of linear overrange process readings. In some cases when additional failure conditions are defined, an even larger dynamic range is required for the loop current, for example, from 3.2 mA to 24 mA. Thus, selecting R2 = 24.9k, IOUT_INIT = 3.2mA, and solving Equation 4 for R3 yields:

$$R3 = VREF * \frac{R2}{Rsense * Iout\_init} = 2.5 * \frac{24.9 * 10\text{^}3}{10 * 3.2 * 10\text{^}(-3)} = 1.945 * 10\text{^}6(\Omega)$$

A 1.945 MΩ resistor is costly and, perhaps more important, not well suited either for automated production or for easy field calibration. Therefore, it is preferable to use a regular 1% tolerance resistor and regain accuracy by calibrating out the 4 mA offset current and the 20 mA full-scale current with the U1 DAC. In this case, some digital codes are needed for calibration to ensure the required accuracy. Thus, I(R1) = VDAC/R1, where VDAC is the U1 DAC output voltage. This can be rewritten as:

$$I(R1) = \frac{VREF * code}{65535 * R1} \quad \text{(Eq. 6)}$$

And:

$$I(R3) = \frac{VREF}{R3} \quad \text{(Eq. 7)}$$

Finally, the Equation 1 can be rewritten as:

$$Iout = VREF * \left[\frac{code}{65535 * R1} + \frac{1}{R3}\right] * \frac{R2}{Rsense} \quad \text{(Eq. 8)}$$

ERROR ANALYSIS AND PERFORMANCE OPTIMIZATION

Transmitter Error at +25°C

Table 1 presents the error analysis of the passive components and VREF in the 4–20mA current-loop transmitter at +25°C. Data are based on Equation 8. This table, as a spreadsheet calculator, is available for download and the designer is encouraged to use the What-If Analysis/Goal Seek… feature from the Data tab to find the appropriate codes for 4mA, 20mA, and 24mA IOUT.

Thus, having the standard 1% tolerance 2 MΩ R3 resistor and setting the U1 DAC to 2682 decimal code, the initial loop current of 4.00015 mA is maintained. Note that the total calculated error is much less than the tolerance of the individual components because their errors are calibrated out by the high-resolution U1 DAC.

The effective number of bits (ENOB) of a 4–20mA current-loop transmitter can be calculated as:

A HIGH-ACCURACY 4–20MA CURRENT-LOOP TRANSMITTER

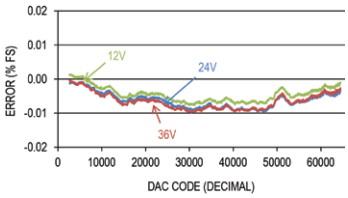

Figure 3. *Transmitter error at 25°C. Data for MAX5216 DAC.*

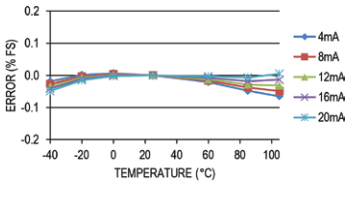

Figure 4. *Transmitter error change vs. temperature with a 12V loop supply.*

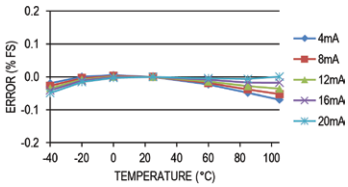

Figure 5. *Transmitter error change vs. temperature with a 24V loop supply.*

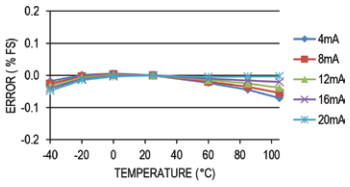

Figure 6. *Transmitter error change vs. temperature with a 36V loop supply.*

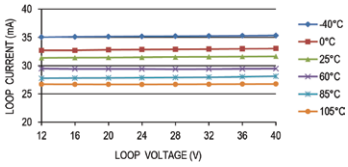

Figure 7. *Current limit vs. loop voltage with a 24.3Ω sense resistor.*

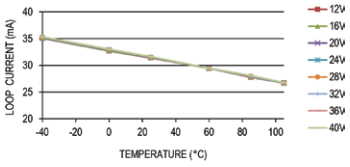

Figure 8. *Current limit vs. temperature with a 24.3Ω sense resistor.*

$$ENOB = \frac{LOG(20mA\ DAC\ code - 4mA\ DAC\ code)}{LOG(2)} \quad \text{(Eq. 9)}$$

Based on the data from Table 1, the ENOB is equal to 15.56 bits. So, dropping less than 0.5 bit of the total resolution allows the calibration process to be automated and lowers the number of expensive precision components.

The selected resistors in Table 1 cover the current loop's dynamic range from 3.2mA up to 24.6mA. Different combinations of R1, R2, R3, and RSENSE can shrink the dynamic range. Close attention should be paid to the temperature coefficients (TC) for each resistor.

TRANSMITTER ERROR OVERTEMPERATURE ANALYSIS

The overtemperature error analysis of the passive components and VREF is shown in Table 2.

The following formulae are used to calculate the minimum and maximum resistance drift:

$$R(T) = Rnom * (1 \pm \frac{TC * \Delta T}{2 * 10^6}) \quad \text{(Eq. 10)}$$

$$VREF(T) = VREFnom * (1 \pm \frac{TC * \Delta T}{2 * 10^6}) \quad \text{(Eq. 11)}$$

Where TC is the temperature coefficient in ppm/°C and ΔT is the total temperature range of 145°C.

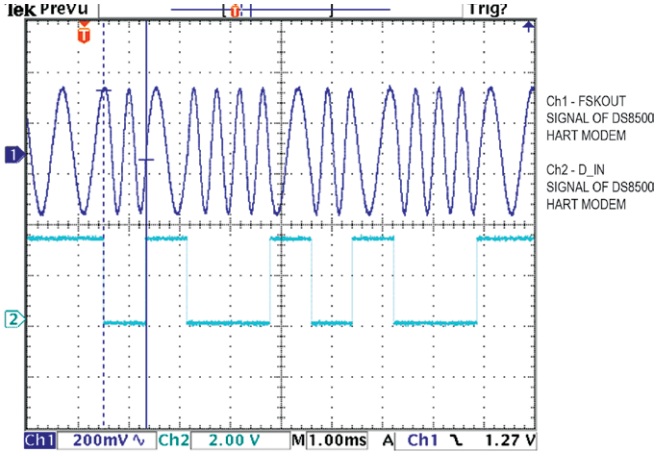

Figure 9. *HART communication over 4–20mA current loop.*

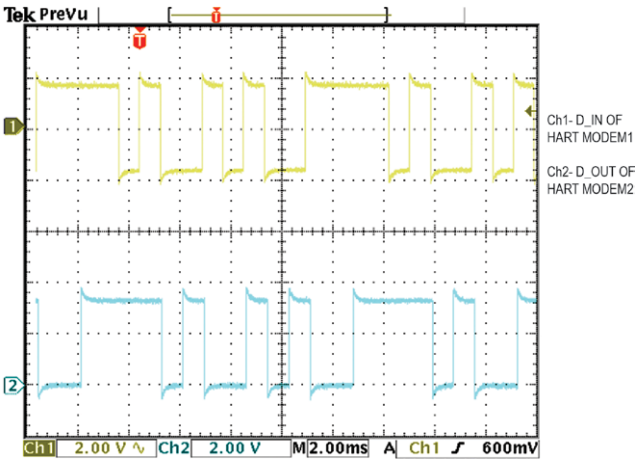

Figure 10. *HART communication between two modems.*

As can be seen from Table 2, the 0.05%FS error is achievable with the following TC for R1, R2, R3, and RSENSE:

R1 = 287 kΩ ±0.1%, 10ppm/°C
R2 = 24.9 kΩ ±0.1%, 25ppm/°C
R3 = 2 MΩ ±1%, 100ppm/°C
RSENSE = 10Ω ±0.1%, 10ppm/°C

Note that total error is the square root of the sum of the squares of each source of error: component's tolerance, component's tempco, measurements, etc.

If a smart sensor consumes more than 3.4 mA, it cannot be used as part of a loop-powered 2-wire transmitter. This happens, for example, when a microcontroller or ADC consumes more than 3 mA or when a sensing element requires a higher supply current to increase its dynamic range and/or resolution. In such cases, the extra current has to flow through an additional third wire. This configuration, called a 3-wire transmitter, can be modified as shown in Figure 2. This design makes it universal as a 2- or 3-wired smart sensor transmitter.

The U5 op amp and Q3 buffer in Figure 2 are sensing the virtual ground, continuously maintaining the common point for the smart sensor and keeping it at the constant voltage of the U4 output. The U5 op amp must be capable of accepting a maximum supply voltage of 12V with a PLC RLOAD/sense resistor value up to 250Ω. The C8 and R8 negative feedback network stabilises the loop current and assures stability for all normally expected loading conditions.

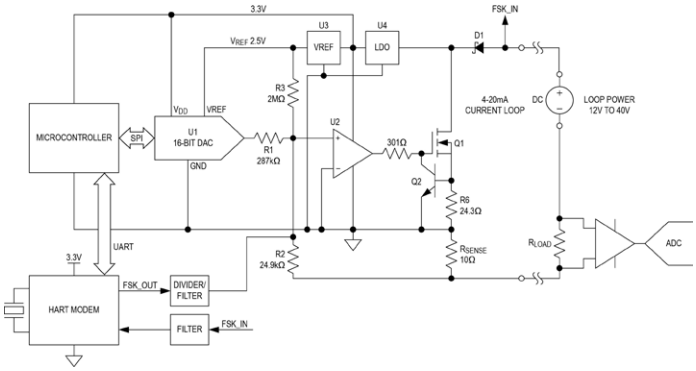

Figure 11. *Block diagram with HART modem*

POWER TRANSISTOR AND PROTECTION COMPONENTS

There are no special requirements for the Q1 power transistor. It could be either a MOSFET or bipolar power transistor which satisfies maximum safe operating area criteria. For example, if the loop power supply is 36V and the highest limiting current is 35 mA, then the maximum dissipation requirement is 1.26W. Close attention should be paid to proper layout, traces width, and the heatsink capabilities of the PCB.

The Schottky diode (D1) (see Figure 1) is a safety device to prevent any damage to the transmitter from reverse current flow. In addition, a transient voltage suppressor (D2, not shown in block diagram) can be added between the LOOP+ and LOOP- inputs to protect from overvoltage surge conditions. The requirements for D1 and D2 depend on the safety standards of the application.

TESTING THE DESIGN

A 4–20mA loop-powered transmitter evaluation (EV) kit, the MAX5216LPT, was built and characterised with a 1000ft 22-gauge shielded communication cable and load resistor of 249Ω ±0.1%. The loop current was measured with an Agilent HP3458A DVM as the voltage drop across that load resistor. The characterisation data from the MAX5216 DAC are presented in Figures 3 to 8.

This transmitter reference design also supports the HART protocol. It allows simple connection with a HART modem such as the DS8500 (see Figure 11). Figures 9 and 10 show HART signals over a 1000ft 4–20 mA current loop with a 249Ω load resistor.

About the Authors

Yuriy Kurtsevoy is a senior strategic applications engineer with the Industrial Measurement business unit in the Industrial and Medical Solution Group at Maxim Integrated. He joined Maxim in 2011. He holds an MS in honors electrical engineering from the East-Ukrainian National University, Luhansk, Ukraine.

Stuart Smith joined Maxim Integrated in 2011 as a Product Definer. He has worked for over 30 years as an analogue and mixed-signal IC design engineer and has received eight patents during that time. Mr. Smith has a BSC EE from Abertay University (Dundee, Scotland) and is a Chartered Engineer.

BY BONNIE BAKER

Simply an adjustable low-pass filter

A low-pass filter is the most common filter found in data acquisition systems. Typically this type of filter is used to reduce analogue-to-digital converter (ADC) aliasing errors and noise outside the signal bandwidth. A signal path requires this type of dedicated filter to match the signal's requirements. If the circuit has a front-end multiplexer, it is possible to have a variety of signals that reach the ADC where each signal source has its own set of filter requirements. Consequently, a variety of different filters and corner frequency requirements may be required in the circuit prior to the multiplexer. These filters use independent operational amplifiers (op amps) in combination with fixed resistors and capacitors.

An alternative filter design solution is to have one programmable filter after the multiplexer (Figure 1). The obvious advantage is reduction in chip count, from multiple op amps to one single amplifier. Subsequently, the cost is lower when a single filter serves many analogue inputs. You can use a dual digital potentiometer, two capacitors, and a single amplifier to configure a low-pass, second-order Butterworth response with a programmable corner frequency range of 1:100. Table 1 summarises the digital potentiometer programmed settings.

With this circuit, it is possible to program second-order Bessel or Chebyshev filters with a programmable corner frequency range of 1:100. Additionally, you can realise a combination of Butterworth, Bessel, and Chebyshev filters with the same circuit using a 1:10 corner frequency range.

Figure 1 shows the details of a single-supply, unity gain, second-order programmable low-pass Sallen-key filter. The OPA314 is a single-supply, rail-to-rail op amp. The filter implementation requires two resistors and two capacitors. The dual TPL0102-100, a 100 kΩ 8-bit digital potentiometer, replaces the two resistors in this circuit. The capacitors are hard-wired in. A reprogramming of the dual digital potentiometer in Figure 1 changes the filter's frequency cut-off and approximation (Butterworth, Bessel, vs. Chebyshev) of this second-order, low-pass filter.

You can calculate the appropriate resistance and capacitance with a little research, a sharp pencil, and a good eraser. An alternative to this tedious design exercise is to determine the capacitor and resistor values using TI's WEBENCH Filter Designer software.

In the input screen, enable the low-pass filter button and type in a filter bandwidth with Ao = 1 V/V, fc = 100, fs = 1000, and Asb = –35 dB. This will produce a second-order Butterworth filter. In this view, you can also select the supply requirements of Single Supply = +5 V.

Then press the green button, "Start Filter Design."

The next page displays a list of filter response options that meet your requirements. Select a second-order Butterworth from the list and click "Open Design."

The Filter Designer Design Summary view (page 3) allows the user to adjust the capacitor seed value to the desired value of C1 and C2 on the left side of the screen, under the "Filter Topology Specifications" section. Change to Capacitor seed value to 15 e-9 or 15 nF, then press the "update" button. When this capacitor seed is set, the software changes the resistors in the circuit to appropriate values.

To produce the remaining filters in Table 1, return to page one and set the conditions for the next filter. As you do this make sure that fs is ten times higher than fc.

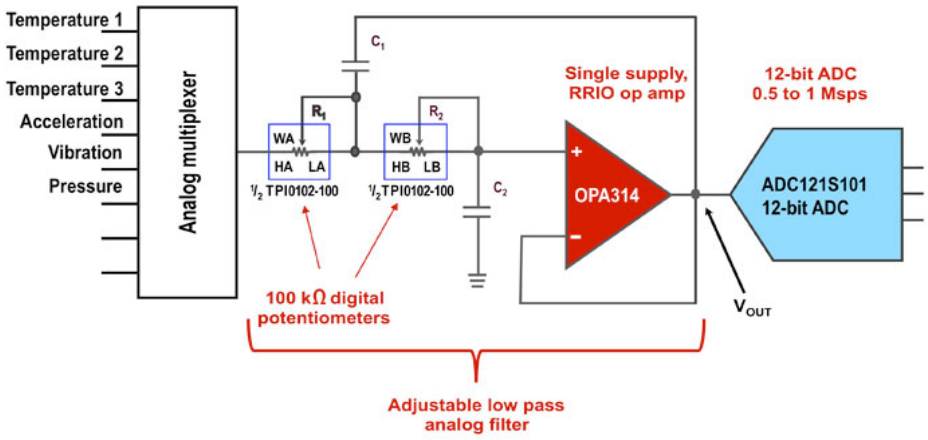

Figure 1. *Figure 1 A second-order, analogue filter using a dual digital potentiometer, two capacitors, and one operational amplifier reduces chip count.*

Cutoff Frequency, Hz	Calculated 1% R1 Value, Ohm	Closest Nominal Digital Pot, R1 Value, Ohm	Digital Pot. R1 Code, decimal	Calculated 1% R2 Value, Ohm	Closest Nominal Digital Pot, R2 Value, Ohm	Digital Pot. R2 Code
100	52.3K	52.3K	134	97.6K	97.7K	250
200	26.1K	26.2K	67	48.7K	48.8K	125
300	16.9K	16.8K	43	32.4K	32.4K	83
1k	5.23K	5.11K	13	9.76K	9.77K	25
2k	2.61K	2.7K	7	4.87K	4.69K	12
3k	1.69K	1.56K	4	3.24K	3.13K	8
10k	523	390	1	976	781	2

Table 1 *shows the digital potentiometer program setting for Butterworth filter with corner frequencies ranging from 100 to 10 kHz (C1 = 33 nF, C2 = 15 nF).*

BY WOLFGANG KATTERMAN, ALTERA

IMPLEMENTING HIGH AVAILABILITY NETWORK COMMUNICATIONS FOR THE SMART GRID AND INDUSTRIAL APPLICATIONS

Network communication for smart grid applications must be highly reliable, highly available and highly deterministic. The network must be resilient against node, link, and single points of failures. For this reason the IEC standardisation group has established the IEC 62439 standard as a way of validating communication network resilience and reliability.

Many new smart grid designs, especially automation equipment in transmission and distribution substations, offer end-to-end communications with IEC 61850-compatible intelligent electronic devices (IEDs), such as industrial switches, and protection relays. IEC 61850 over Ethernet is fast becoming the backbone of substation and utility automation designs.

However, the use of IEC61850 is not limited to Smart Grid applications but can be used anywhere where production interruptions or down times impact productivity. The considerations of high-availability and seamless redundant or parallel redundant topologies with PTP (precision time protocol) for clock synchronisation, apply anywhere there is a need for resilient communications. Examples are industrial automation and factory automation systems.

COMMUNICATIONS TOPOLOGIES BASED ON IEC 62439

Clause 4 of the standard describes the Parallel Redundancy Protocol (PRP), shown in Figure 1. The two networks are assumed to be fail-independent. The destination node will always receive at least one packet as long as either one of the two networks is operational. The destination node must detect the duplication of packets and discard the second packet upon its arrival to make sure only one packet with the same content will be successfully delivered. This provides zero-time recovery in case of a single failure, so no frames are lost. The downside of PRP is that the network cost is doubled when comparing with a single non-redundant network. This makes it more expensive to implement than the other redundancy protocols mentioned in the following paragraph.

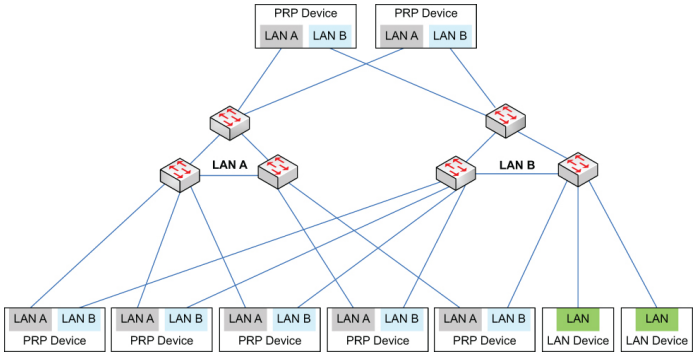

Figure 1. *PRP Network Topology with PRP acting destination nodes*

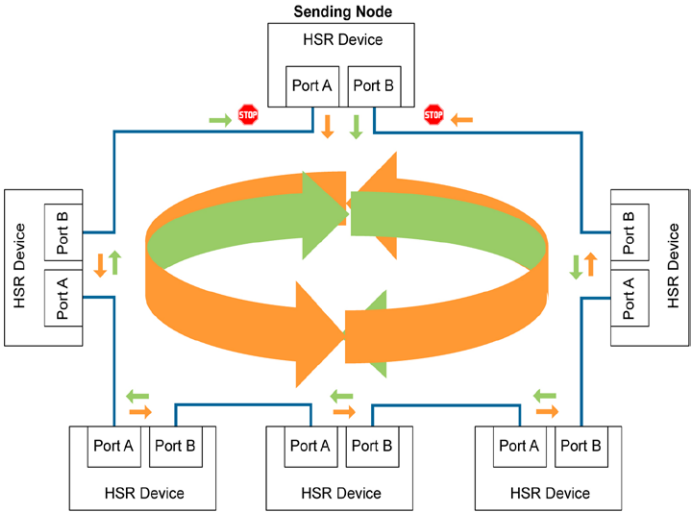

Figure 2. *HSR Network Topology with HSR acting as "Redboxes"*

Clause 5 of the IEC 62439 standard describes the High-Availability Seamless Redundancy (HSR) shown in Figure 2.

The typical HSR topology is a ring. The source node duplicates all the frames it has to send, and sends them using two different paths to their destination. If either one of the paths is broken, due to link or node failure, the frames are still able to reach their destination. End-nodes have two external Ethernet ports to connect to one HSR ring. They may also have one or more additional Ethernet ports, for example for maintenance access. When sending frames, an End-node sends two duplicates of each frame into the network, one to each of the directions in the ring. When receiving, the nodes accept the first copy and discard the second, thus eliminating the duplicate. Sensors such as IEDs (Intelligent Electronic Devices), cameras, and merging units in substation automation are examples of End-nodes.

"RedBox" (or Redundancy Box) is an entity that has three external Ethernet ports. Two of the ports are connected to an HSR ring and one port (interlink) is a traditional Ethernet port. "RedBoxes" are used to connect non-HSR nodes and non-HSR network segments to HSR rings. This means, for example, ordinary Ethernet switches and computers. When forwarding frames to the ring, a "RedBox" duplicates each frame and sends the two duplicates to the ring, one to each direction. When forwarding frames from the ring, RedBox forwards the first copy and removes the one that arrives later.

Currently the HSR ring is perhaps the most effective high availability Ethernet method. HSR is similar to Parallel Redundancy Protocol (PRP) in that it sends duplicate packets, however, unlike in PRP, there is no need for duplicate cables or switches in ring topology. This makes an HSR ring much more cost effective than the two local area networks of PRP. Compared with Rapid Spanning Tree Protocol (RSTP), the benefit of HSR is its zero-time recovery - a single network fault in the ring will not result in any frame loss; the recovery times of RSTP are simply not acceptable in all applications.

IMPLEMENTING HIGH AVAILABILITY COMMUNICATIONS

What means are available to accommodate all the above-mentioned standards whilst keeping flexibility for upcoming changes, as well as avoiding bottlenecks in performance? One simple approach is to use an FPGA combined with an ARM subsystem and with a dedicated IP block handling the hardware requirements of the communication standard. For example, the Cyclone V SoC is well suited for those applications. The device features one or two ARM Cortex-A9 processor cores capable of being clocked at 800 MHz, and other hardwired peripherals including embedded flash, RAM, caches, GPIO, and communications ports. The FPGA fabric affords opportunities for integration, performance acceleration, and upgradeability.

With these resources it becomes practical to evolve existing designs, or already installed systems, with changing PRP/HSR protocol standards. Communications with media redundancy in real time can be processed. In addition, increasing requirements to connect external I/O interfaces are easily maintainable and expandable. For example, Flexibilis, a company that has a focus on mission-critical communications, has developed an FPGA IP core that can be integrated into the end-product design without the need for an external switch.

The Flexibilis Redundant Switch (FRS) IP core is a triple-speed (10 Mbps/100 Mbps/1 Gbps) Ethernet Layer-2 switch that supports new protocols, providing seamless redundancy for Ethernet networking. FRS is compatible with IEC 62439-3 High-availability Seamless Redundancy (HSR) and Parallel Redundancy Protocol (PRP). The Redundant Switch also includes IEEE1588v2 Precision Time Protocol (PTP) transparent clock functionality. FRS has been explicitly to be implemented in FPGAs. The key feature of the FRS switch is that it is scalable from three to eight ports without the requirement for separate RedBoxes. It is full-duplex 1000 Mbps (GMII1) and 10/100 Mbps (MII2) on all ports. It supports wire-speed packet forwarding, non-blocking operations, a reliable store-and-forward operation with data integrity checking as well as compatibility with IEEE 1588 Precision Time Protocol (PTP) transparent clock.

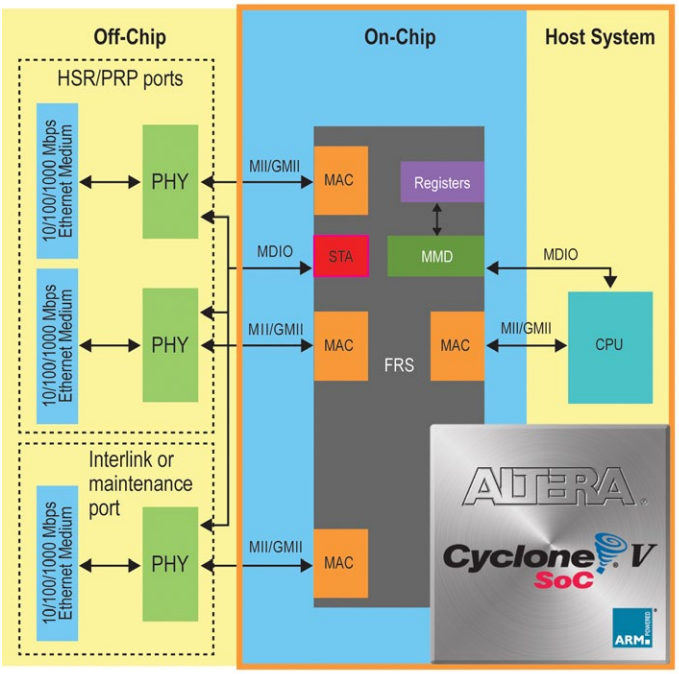

Figure 3 – *Implementing the Flexibilis Redundant Switch IP*

The IP block integrates a PRP/HSR switch implementation with other system functions on a Cyclone V FPGA or Cyclone V SoC. This saves development and integration time and mitigates risk. One other major advantage of the combined FPGA/Flexibilis solution is the avoidance of an external switch, since each HSR node supports multiple Ethernet connections (two for the ring and one or more for additional service ports). The FRS is a layer 2 switch that supports HSR/PRP which means that there are two redundant ports and up to six ordinary Ethernet switch ports. Therefore, the End Node/IED can be used for both HSR/PRP and standard switching requirements.

Another key element in industrial designs is maintainability and longevity of parts used in the system design. This goes beyond reliability and a commitment to provide the solution for the life of the product. The ability to reconfigure and upgrade products (in either manufacturing, production, or development) is critical, particularly when certain important standards evolve over time (for example, IEC 62351). FPGAs help mitigate this problem by providing scalability and re-configurability to implement product updates that go beyond a simple software change.

A large-scale commoditisation is under way in the industrial and smart energy market. Supporting a communication standard is no longer any sort of differentiating factor. In considering industrial Ethernet progression, slave devices and platforms have to be able to support all major industrial standards utilised in this market. Even the leading companies in the sector, who have developed their own industrial Ethernet implementations, are reducing their efforts in maintaining their proprietary solutions. Communications is a commodity product and does not add differentiation value.

As mentioned in the opening paragraphs above, HSR/PRP standards can be employed in areas other than in the Smart Grid. A migration or combination of real time industrial Ethernet standards using the HSR/PRP for availability and reliability reasons is a likely outcome.

Wolfgang Katterman is a market development engineer with Altera.

BY HENRY J. ZHANG, LINEAR TECHNOLOGY

BASIC CONCEPTS OF LINEAR REGULATOR AND SWITCHING MODE POWER SUPPLIES

This article explains the basic concepts of linear regulators and switching mode power supplies (SMPS). It is aimed at system engineers who may not be very familiar with power supply designs and selection. The basic operating principles of linear regulators and SMPS are explained and the advantages and disadvantages of each solution are discussed. The buck step-down converter is used as an example to further explain the design considerations of a switching regulator.

Today's designs require an increasing number of power rails and supply solutions in electronics systems, with loads ranging from a few mA for standby supplies to over 100A for ASIC voltage regulators. It is important to choose the appropriate solution for the targeted application and to meet specified performance requirements, such as high efficiency, tight printed circuit board (PCB) space, accurate output regulation, fast transient response, low solution cost, etc. Power management design is becoming a more frequent and challenging task for system designers, many of who may not have strong power backgrounds.

A power converter generates output voltage and current for the load from a given input power source. It needs to meet the load voltage or current regulation requirement during steady-state and transient conditions. It also must protect the load and system in case of a component failure. Depending on the specific application, a designer can choose either a linear regulator (LR) or a switching mode power supply (SMPS) solution. To make the best choice of a solution, it is essential for designers to be familiar with the merits, drawbacks and design concerns of each approach.

This article focuses on non-isolated power supply applications and provides an introduction to their operation and design basics.

LINEAR REGULATORS

How a linear regulator works; Let's start with a simple example. In an embedded system, a 12V bus rail is available from the front-end power supply. On the system board, a 3.3V voltage is needed to power an operational amplifier (op amp). The simplest approach to generate the 3.3V is to use a resistor divider from the 12V bus, as shown in Figure 1. Does it work well? The answer is usually no. The op amp's VCC pin current may vary under different operating conditions. If a fixed resistor divider is used, the IC VCC voltage varies with load. Besides, the 12V bus input may not be well regulated. There may be many other loads in the same system sharing the 12V rail. Because of the bus impedance, the 12V bus voltage varies with the bus loading conditions. As a result, a resistor divider cannot provide a regulated 3.3V to the op amp to ensure its proper operation. Therefore, a dedicated voltage regulation loop is needed. As shown in Figure 2, the feedback loop needs to adjust the top resistor R1 value to dynamically regulate the 3.3V on VCC.

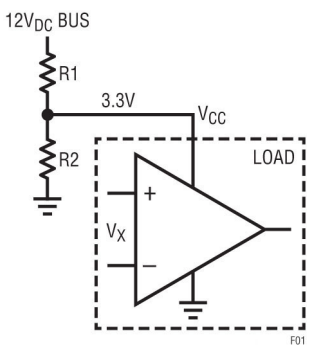

Figure 1. *Resistor divider generates 3.3 VDC from 12V bus input*

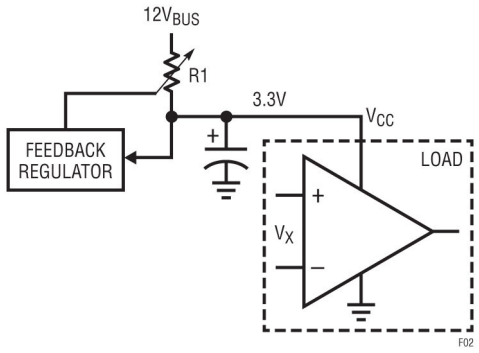

Figure 2. *Feedback loop adjusts series resistor R1 value to regulate 3.3V*

This kind of variable resistor can be implemented with a linear regulator, as shown in Figure 3. A linear regulator operates a bipolar or field effect power transistor (FET) in its linear mode. So the transistor works as a variable resistor in series with the output load. To establish the feedback loop, conceptually, an error amplifier senses the DC output voltage via a sampling resistor network RA and RB, and then compares the feedback voltage VFB with a reference voltage VREF. The error amplifier output voltage drives the base of the series power transistor via a current amplifier. When either the input VBUS voltage decreases or the load current increases, the VCC output voltage goes down. The feedback voltage VFB decreases as well. As a result, the feedback error amplifier and current amplifier generate more current into the base of the transistor Q1. This reduces the voltage drop VCE and hence brings back the VCC output voltage, so that VFB equals VREF. On the other hand, if the VCC output voltage goes up, in a similar way, the negative feedback circuit increases VCE to ensure the accurate regulation of the

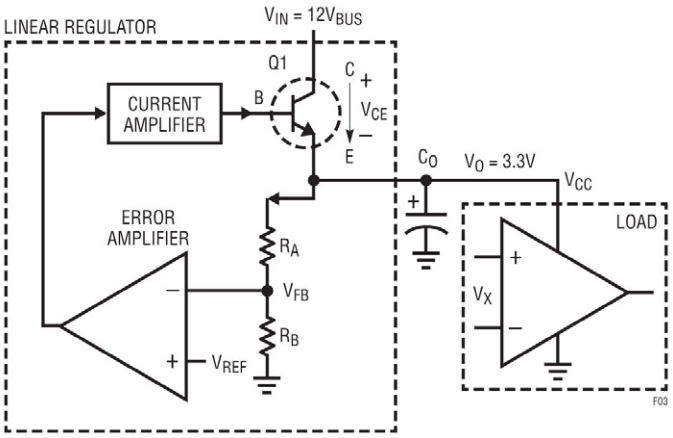

Figure 3. *A linear regulator implements a variable resistor to regulate output voltage*

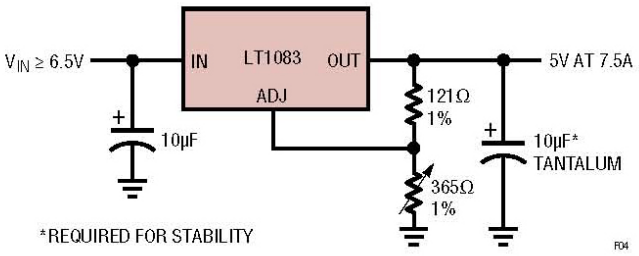

Figure 4. *Integrated linear regulator example: 7.5A Linear regulator with only three pins*

3.3V output. In summary, any variation of VO is absorbed by the linear regulator transistor's VCE voltage. So the output voltage VCC is always constant and well regulated.

WHY USE LINEAR REGULATORS?

The linear regulator has been widely used by industry for a very long time. It was the basis for the power supply industry until switching mode power supplies became prevalent after the 1960s. Even today, linear regulators are still widely used in a wide range of applications. In addition to their simplicity of use, linear regulators have other performance advantages. Power management suppliers have developed many integrated linear regulators. A typical integrated linear regulator needs only VIN, VOUT, FB and optional GND pins. Figure 4 shows a typical 3-pin linear regulator, the LT1083, which was developed more than 20 years ago. It only needs an input capacitor, output capacitor and two feedback resistors to set the output voltage. Almost any electrical engineer can design a supply with these simple linear regulators.

One drawback – a linear regulator can burn a lot of power. A major drawback of using linear regulators can be the excessive power dissipation of its series transistor Q1 operating in a linear mode. As explained previously, a linear regulator transistor is conceptually a variable resistor. Since all the load current must pass through the series transistor, its power dissipation is PLOSS = (VIN – VO) • IO. In this case, the efficiency of a linear regulator can be quickly estimated by:

$$\eta_{LR} = \frac{P_{OUTPUT}}{P_{OUTPUT}+P_{LOSS}} = \frac{V_O \bullet I_O}{V_O \bullet I_O + (V_{IN}-V_O)\bullet I_O} = \frac{V_O}{V_{IN}} \quad (1)$$

So in the Figure 1 example, when the input is 12V and output is 3.3V, the linear regulator efficiency is just 27.5%. In this case, 72.5% of the input power is just wasted and generates heat in the regulator. This means that the transistor must have the thermal capability to handle its power/heat dissipation at worst case at maximum VIN and full load. So the size of the linear regulator and its heat sink may be large, especially when VO is much less than VIN. Figure 5 shows that the maximum efficiency of the linear regulator is proportional to the VO/VIN ratio.

On the other hand, the linear regulator can be very efficient if VO is close to VIN. However, the linear regulator (LR) has another limitation, which is the minimum voltage difference between VIN and VO. The transistor in the LR must be operated in its linear mode. So it requires a certain minimum voltage drop across the collector to emitter of a bipolar transistor or drain to source of a FET. When VO is too close to VIN, the LR may be unable to regulate output voltage any more. The linear regulators that can work with low headroom (VIN – VO) are called low dropout regulators (LDOs).

It is also clear that a linear regulator or an LDO can only provide step-down DC/DC conversion. In applications that require VO voltage to be higher than VIN voltage, or need negative VO voltage from a positive VIN voltage, linear regulators obviously do not work.

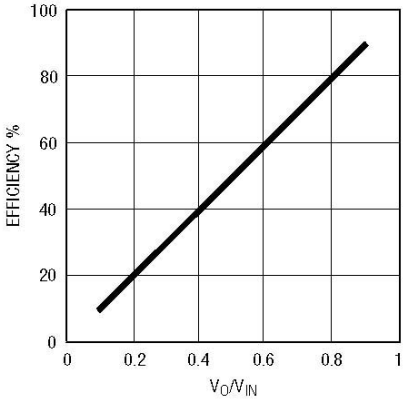

Figure 5. *Maximum linear regulator efficiency vs. VO/VIN ratio*

LINEAR REGULATOR WITH CURRENT SHARING

For applications that require more power, the regulator must be mounted separately on a heat sink to dissipate the heat. In all-surface-mount systems, this is not an option, so the limitation of power dissipation (1W for example) limits the output current. Unfortunately, it is not easy to directly parallel linear regulators to spread the generated heat.

Replacing the voltage reference shown in Figure 3 with a precision current source, allows the linear regulator to be directly paralleled to spread the current load and thus spread dissipated heat among the ICs. This makes it possible to use linear regulators in high output current, all-surface-mount applications, where only a limited amount of heat can be dissipated in any single spot on a board.

The LT3080 is the first adjustable linear regulator that can be used in parallel for higher current. As shown in Figure 6, it has a precision zero TC 10 µA internal current source connected to the non-inverting input of the operational amplifier. With an external single voltage-setting resistor RSET, the linear regulator output voltage can be adjusted from 0V to (VIN – VDROPOUT).

Figure 7 shows how easy it is to parallel LT3080s for current sharing. Simply tie the SET pins of the LT3080s together; the two regulators share the same reference voltage. Because the operational amplifiers are precisely trimmed, the offset voltage between the adjustment pin and the output is less than 2 mV. In this case, only 10 mΩ ballast resistance, which can be the sum of a small external resistor and PCB trace resistance, is needed

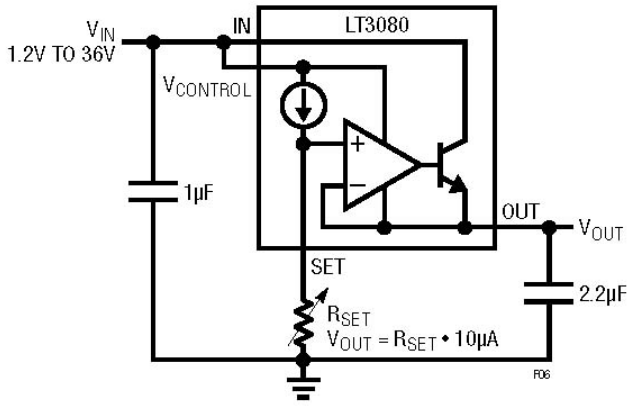

Figure 6. *Single resistor setting LDO LT3080 with a precision current source reference*

to balance the load current with better than 80% equalised sharing. Need even more power? Even paralleling 5 to 10 devices is reasonable.

WHERE LINEAR REGULATORS ARE PREFERABLE

There are many applications in which linear regulators or LDOs provide superior solutions to switching supplies, including:

1. Simple/low cost solutions. Linear regulator or LDO solutions are simple and easy to use, especially for low power applications with low output current where thermal stress is not critical. No external power inductor is required.
2. Low noise/low ripple applications. For noise-sensitive applications, such as communication and radio devices, minimising the supply noise is very critical. Linear regulators have very low output voltage ripple because there are no elements switching on and off frequently and linear regulators can have very high bandwidth. So there is little EMI problem. Some special LDOs, such as Linear Technology's LT1761 LDO family, have as low as 20 µVRMS noise voltage on the output. It is almost impossible for an SMPS to achieve this low noise level. An SMPS usually has mV of output ripple even with very low ESR capacitors.
3. Fast transient applications. The linear regulator feedback loop is usually internal, so no external compensation is required. Typically, linear regulators have wider control loop bandwidth and faster transient response than that of SMPS.
4. Low dropout applications. For applications where output voltage is close to the input voltage, LDOs may be more efficient than an SMPS. There are very low dropout LDOs (VLDO) such as Linear's LTC1844, LT3020 and LTC3025 with from 20 mV to 90 mV dropout voltage and up to 150 mA current. The minimum input voltage can be as low as 0.9V. Because there is no AC switching loss in an LR, the light load efficiency of an LR or an LDO is similar to its full load efficiency. An SMPS usually has lower light load efficiency because of its AC switching losses. In battery powered applications in which light load efficiency is also critical, an LDO can provide a better solution than an SMPS.

In summary, designers use linear regulators or LDOs because they are simple, low noise, low cost, easy to use and provide fast transient response. If VO is close to VIN, an LDO may be more efficient than an SMPS.

SWITCHING MODE POWER SUPPLY BASICS

Why Use a Switching Mode Supply? A quick answer is high efficiency. In an SMPS, the transistors are operated in switching mode instead of linear mode. This means that when the transistor is on and conducting current, the voltage drop across its power path is minimal. When the transistor is off and blocking high voltage, there is almost no current through its power path. So the semiconductor transistor is like

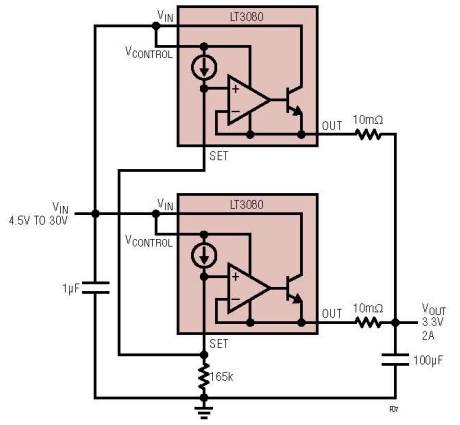

Figure 7. *Paralleling of two LT3080 linear regulators for higher output current*

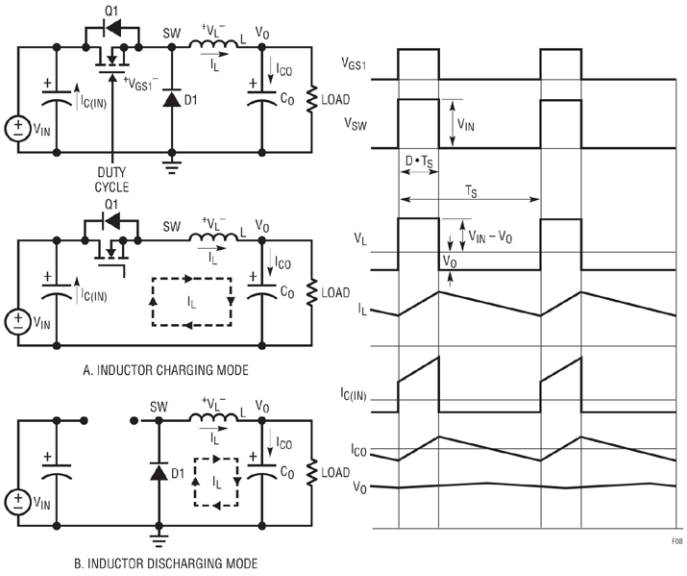

Figure 8. *Buck converter operating modes and typical waveforms*

an ideal switch. The power loss in the transistor is therefore minimised. High efficiency, low power dissipation and high power density (small size) are the main reasons for designers to use SMPS instead of linear regulators or LDOs, especially in high current applications. For example, nowadays a 12 VIN, 3.3 VOUT switching mode synchronous buck step-down supply can usually achieve over 90% efficiency vs. less than 27.5% from a linear regulator. This means a power loss or size reduction of at least eight times.

THE BUCK CONVERTER

Figure 8 shows the simplest and most popular switching regulator, the buck DC/DC converter. It has two operating modes, depending on if the transistor Q1 is turned on or off. To simplify the discussion, all the power devices are assumed to be ideal. When switch (transistor) Q1 is turned on, the switching node voltage VSW = VIN and inductor L current is being charged up by (VIN – VO). Figure 8(a) shows the equivalent circuit in this inductor-charging mode. When switch Q1 is turned off, inductor current goes through the freewheeling diode D1, as shown in Figure 8(b). The switching node voltage VSW = 0V and inductor L current is discharged by the VO load. Since the ideal inductor cannot have DC voltage in the steady state, the average output voltage VO can be given as:

$$V_{O(DC)} = AVG[V_{SW}] = \frac{T_{ON}}{T_S} \bullet V_{IN} \qquad (2)$$

Where TON is the on-time interval within the switching period TS. If the ratio of TON/TS is defined as duty cycle D, the output voltage VO is:

$$V_{O(DC)} = \frac{T_{ON}}{T_S} \bullet V_{IN} = D \bullet V_{IN} \qquad (3)$$

When the filter inductor L and output capacitor CO values are sufficiently high, the output voltage VO is a DC voltage with only mV ripple. In this case, for a 12V input buck supply, conceptually, a 27.5% duty cycle provides a 3.3V output voltage.

Other than the above averaging approach, there is another way to derive the duty cycle equation. The ideal inductor cannot have

DC voltage in steady state. So it must maintain inductor volt-second balance within a switching period. According to the inductor voltage waveform in Figure 8, volt-second balance requires:

$$(V_{IN} - V_O) \bullet D \bullet T_S = V_O \bullet (1 - D) \bullet T_S \quad (4)$$

$$\text{Hence, } V_O = V_{IN} \bullet D \quad (5)$$

Equation (5) is the same as equation (3). The same volt-second balance approach can be used for other DC/DC topologies to derive the duty cycle vs. VIN and VO equations.

POWER LOSSES IN A BUCK CONVERTER

DC conduction losses; With ideal components (zero voltage drop in the ON state and zero switching loss), an ideal buck converter is 100% efficient. In reality, power dissipation is always associated with every power component. There are two types of losses in an SMPS: DC conduction losses and AC switching losses.

The conduction losses of a buck converter primarily result from voltage drops across transistor Q1, diode D1 and inductor L when they conduct current. To simplify the discussion, the AC ripple of inductor current is neglected in the following conduction loss calculation. If a MOSFET is used as the power transistor, the conduction loss of the MOSFET equals IO2 • RDS(ON) • D, where RDS(ON) is the on-resistance of MOSFET Q1. The conduction power loss of the diode equals IO • VD • (1 – D), where VD is the forward voltage drop of the diode D1. The conduction loss of the inductor equals IO2 • RDCR, where RDCR is the copper resistance of the inductor winding. Therefore, the conduction loss of the buck converter is approximately:

$$P_{CON\_LOSS} = I_O^2 \bullet R_{DS(ON)} \bullet D + I_O \bullet V_D \bullet (1 - D) + I_O^2 \bullet R_{DCR} \quad (6)$$

For example, a 12V input, 3.3V/10AMAX output buck supply can use following components: MOSFET RDS(ON) = 10 mΩ, inductor RDCR = 2 mΩ, diode forward voltage VD = 0.5V. Therefore, the conduction loss at full load is:

$$P_{CON\_LOS} = 10^2 \bullet 10 \bullet 10^{-3} \bullet 0.275 + 10 \bullet 0.5 \bullet (1 - 0.275) + 10^2 \bullet 2 \bullet 10^{-3}(W) = 0.275W + 3.62W + 0.2W = 4.095W \quad (7)$$

Considering only conduction loss, the converter efficiency is:

$$\eta_{BUCK\_CON} = \frac{P_{OUTPUT}}{P_{OUTPUT} + P_{CON\_LOSS}} = \frac{3.3V \bullet 10A}{33W + 4.095W} = 88.96\% \quad (8)$$

The above analysis shows that the freewheeling diode consumes 3.62W power loss, which is much higher than the conduction losses of the MOSFET Q1 and the inductor L. To further improve efficiency, diode D1 can be replaced with a MOSFET Q2, as shown in Figure 9. This converter is referred to as a synchronous buck converter. Q2's gate requires signals complementary to the Q1 gate, i.e., Q2 is only on when Q1 is off. The conduction loss of the synchronous buck converter is:

$$P_{CON\_LOSS} = I_O^2 \bullet R_{DS1(ON)} \bullet D + I_O^2 \bullet R_{DS2(ON)} \bullet (1 - D) + I_O^2 R_{DCR} \quad (9)$$

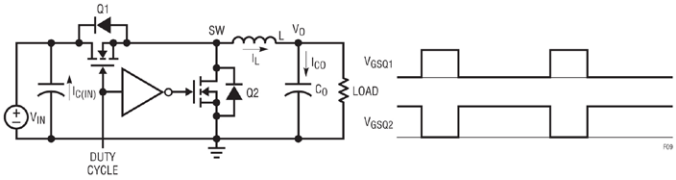

Figure 9. *Synchronous buck converter and its transistor gate signals*

If a 10 mΩ RDS(ON) MOSFET is used for Q2 as well, the conduction loss and efficiency of the synchronous buck converter are:

$$P_{CON\_LOS} = 10^2 \bullet 0.015 \bullet 0.275 + 10^2 \bullet 0.015 \bullet (1 - 0.275) + 10^2 \bullet 2 \bullet 10^{-3}(W) = 0.275W + 0.725W + 0.2W = 1.2W \quad (10)$$

$$\eta_{BUCK\_CON} = \frac{P_{OUTPUT}}{P_{OUTPUT} + P_{CON\_LOSS}} = \frac{3.3V \bullet 10A}{33W + 1.2W} = 96.45\% \quad (11)$$

The above example shows that the synchronous buck is more efficient than a conventional buck converter, especially for low output voltage applications where the duty cycle is small and the conduction time of the diode D1 is long.

AC SWITCHING LOSSES

In addition to the DC conduction losses, there are other AC/switching related power losses due to the non-ideal power components:

1. MOSFET switching losses. A real transistor requires time to be turned on or off. So there are voltage and current overlaps during the turn-on and turn-off transients, which generate AC switching losses. Figure 10 shows the typical switching waveforms of the MOSFET Q1 in the synchronous buck converter. The charging and discharging of the top FET Q1's parasitic capacitor CGD with charge QGD determine most of the Q1 switching time and related losses. In the synchronous buck, the bottom FET Q2 switching loss is small, because Q2 is always turned on after its body diode conducts and is turned off before its body diode conducts, while the voltage drop across the body diode is low. However, the body diode reverse recovery charge of Q2 can also increase the switching loss of the top FET Q1 and can generate switching voltage ringing and EMI noise. Equation (12) shows that the control FET Q1 switching loss is proportional to the converter switching frequency fS. The accurate calculation of the energy losses EON and EOFF for Q1 is not simple but can be found from MOSFET vendors' application notes.

$$PSW\_Q1 = (EON + EOFF) \bullet fS \quad (12)$$

2. Inductor core loss PSW_CORE. A real inductor also has AC loss that is a function of switching frequency. Inductor AC loss is primarily from the magnetic core loss. In a high frequency SMPS, the core material may be powdered iron or ferrite. In general, powdered iron cores saturate softly but have high core loss, while ferrite material saturates more sharply but has less core loss. Ferrites are ceramic ferromagnetic materials that have a crystalline structure consisting of mixtures of iron oxide with

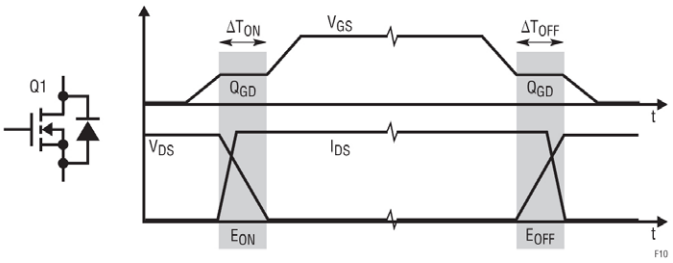

Figure 10. *Typical switching waveform and losses in the top FET Q1 in the buck converter*

either manganese or zinc oxide. Core losses are due mainly to magnetic hysteresis loss. The core or inductor manufacturer usually provides the core loss data for power supply designers to estimate the AC inductor loss.

3. Other AC related losses. Other AC related losses include the gate driver loss PSW_GATE, which equals VDRV • QG • fS, and the dead time (when both top FET Q1 and bottom FET Q2 are off) body diode conduction loss, which is equal to (ΔTON + ΔTOFF) • VD(Q2) • fS.

In summary, the switching-related loss includes:

$$P_{SW\_LOSS} = P_{Q1\_SW} + P_{CORE\_SW} + P_{DRV} + P_{DEADTIME} \quad (13)$$

The calculation of switching related losses is usually not easy. The switching related losses are proportional to switching frequency fS. In the 12 VIN, 3.3 VO/10 AMAX synchronous buck converter, the AC loss causes about 2% to 5% efficiency loss with 200 kHz – 500 kHz switching frequency. So the overall efficiency is about 93% at full load, much better than that of an LR or LDO supply. The heat or size reduction can be close to 10x.

PART TWO of this article will cover design considerations of the switching power components, discuss the feedback loop, PCB layout and other important aspects and topologies of switching supply design.

Henry Zhang is Applications Engineering Manager, Power Products, at Linear Technology Corp. Milpitas, CA, USA.

Linear offers free power supply design & simulation tools at: http://www.linear.com/designtools/software/

VICOR

Solutions from AC to Point of Load

Our Latest Products

Picor Cool-Power Isolated ZVS DC-DC Converters

Simple to Use

- Complete Isolated DC-DC converter with zero-voltage switching module
- Three input voltage ranges for communication, industrial, rugged/M-Grade applications
- Communication converters have max peak input voltage of 100 V / 100 ms (non-operating)
- 2,250 V input to output isolation

High Density

- 50 W / 60 w output power (dependent upon converter model)
- Surface mount 22mm x 16.5mm x 6.7mm high-density package
- 900 kHz switching frequency, miminizes input filtering and reduces output capacitance

Rich Feature Set

- On/Off Control, positive logic ENABLE
- Wide trim range +10/-20% Trim
- Temperature monitor™ & Over-temperature Protection (OTP)
- Input UVLO & OVLO and output OVP
- Over current protection with auto restart
- Adjustable soft-start
- Output voltage sensing without opto coupler use for higher reliability

Resources

Blog Post: The technology behind the Cool-Power ZVS DC-DC Converters

Video: An Introduction to Vicor's Cool-Power ZVS DC-DC Converters

Cool-Power	Input	Output Set	Output Range	I_{OUT}Max
Communications (-40°C to 125°C)				
PI3101-00-HVIZ	36 – 75 Vin	3.3 V	3.0 to 3.6 V	18 A
PI3105-00-HVIZ		12 V	9.6 to 13.2 V	5 A
PI3110-01-HVIZ	41 – 57 Vin	18 V	16.2 to 19.8 V	3.3 A
Industrial (-40°C to 125°C)				
PI3109-01-HVIZ	18 – 36 Vin	5 V	4.0 to 5.5 V	10 A
PI3106-01-HVIZ		12 V	9.6 to 13.2 V	4.2 A
M-Grade (-55°C to 125°C)				
PI3109-00-HVMZ	16 – 50 Vin	5 V	4.0 to 5.5 V	10 A
PI3106-00-HVMZ		12 V	9.6 to 13.2 V	4.2 A
PI3111-00-HVMZ		15 V	12 to 16.5 V	3.3 A

Two-Part Web Seminar Series About EMI

Watch Part 1

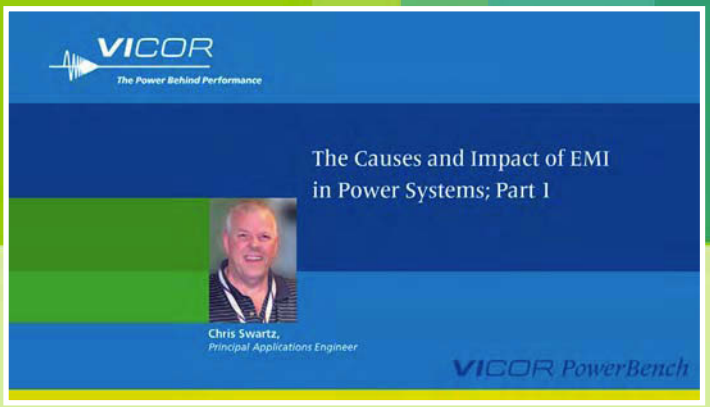

The Causes and Impact of EMI in Power Systems

Watch Part 2

Practical Design Considerations and Solutions for the Reduction of Conducted EMI and Filter Size in Power Systems

VI Brick AC Front End

Overview

- Universal Input: 85 – 264 Vac
- Output: 48 Vdc - isolated, regulated (SELV)
- Power: 330 W - over entire input voltage range
- Isolated AC-DC converter with active Power Factor Correction (PFC)
- Integrated rectification, filtering and transient protection
- Peak efficiency: >92%
- EN55022, Class B EMI conducted emissions with a few components
- EN61000-3-2 harmonic limits
- -55 to 100°C baseplate operation

Features

- Low profile, 9.55 mm height above board
- Small footprint, size of a business card
- Flanged aluminum package for secure mounting and thermal management
- Consistent high efficiency across the worldwide mains (flat efficiency curve)
- Reduced power loss and cooling requirements
- Module includes PFC, regulation, isolated 48 V output (SELV), filtering, rectification, transient protection, agency approvals, simplified thermal management
- Simple design, requires few external components
- Module power density, 121 W/in³
- Complete solution including hold-up capacitors, 54 W/in³

Basic Application

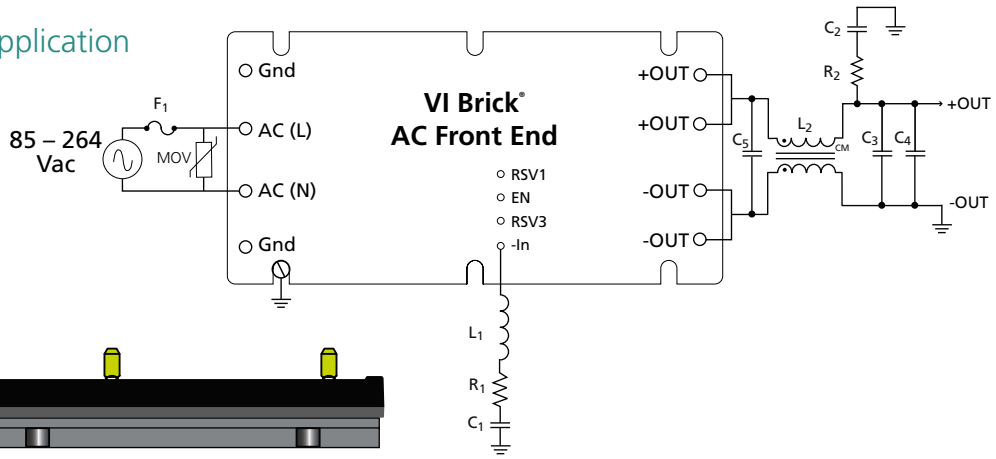

Resources

- An Introduction to the Vicor AC Front End Module
- Webinar: Designing High Performance AC-DC Power Systems Using a Power Component Approach
- AC Front End Product Information

Part Number	Input Voltage	Output Voltage	Output Power	Operating Temperature
FE175D480C033FP-00	85 – 264 Vac	48 Vdc	330 W	-20 to 100°C
FE175D480T033FP-00	85 – 264 Vac	48 Vdc	330 W	-40 to 100°C
FE175D480M033FP-00	85 – 264 Vac	48 Vdc	330 W	-55 to 100°C

Replace the "–00" suffix in the part number with "–CB" to order an evaluation board.

Consistent High Efficiency

Over Line, Load, Temperature

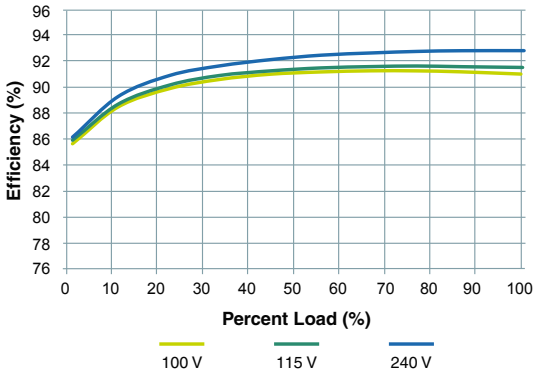

Picor Cool-Power ZVS Buck Regulators

Wide Operating Range

- Wide V_{IN} (8 – 36 V) and wide V_{OUT} (1 – 16 V)
- 12 V-optimized performance with PI34xx Series
- -40°C to 125°C operating range

Simple to Use; Fast Development Time

- Internal compensation - few external components
- No additional design or additional settings required

High Efficiency

- Up to 98% peak efficiency (19 V_{IN} to 15 V_{OUT})
- PI34xx Series optimized for 12 V_{IN} with even higher efficiency
- Light and full load high-efficiency performance

Flexible and Rich Feature Set

- Paralleling and single wire current sharing
- Frequency synchronization
- User adjustable soft-start & tracking
- Power-up into pre-biased load
- Optional I²C functionality & programmability

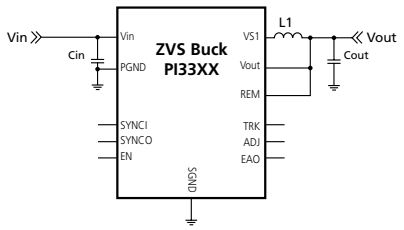

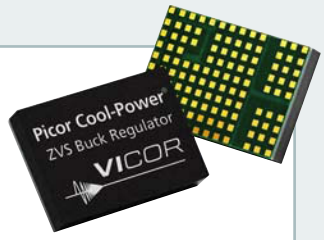

	Cool-Power Model Number	Output Range Set	Output Range Trim Range	I_{OUT} Max
8 – 36 Vin	PI3311-00-LGIZ	1.0 V	1.0 V to 1.4 V	10 A
	PI3318-00-LGIZ	1.8 V	1.4 V to 2.0 V	10 A
	PI3312-00-LGIZ	2.5 V	2.0 V to 3.1 V	10 A
	PI3301-00-LGIZ	3.3 V	2.3 V to 4.1 V	10 A
	PI3302-00-LGIZ	5.0 V	3.3 V to 6.5 V	10 A
	PI3303-00-LGIZ	12 V	6.5 V to 13.0 V	8 A
	PI3305-00-LGIZ	15 V	10.0 V to 16.0 V	8 A
	Higher Current Versions			
	PI3311-01-LGIZ	1.0 V	1.0 V to 1.4 V	15 A
	PI3318-01-LGIZ	1.8 V	1.4 V to 2.0 V	15 A
	PI3312-01-LGIZ	2.5 V	2.0 V to 3.1 V	15 A
	PI3301-01-LGIZ	3.3 V	2.3 V to 4.1 V	15 A
	I²C Functionality and Programmability			
	PI3311-20-LGIZ	1.0 V	1.0 V to 1.4 V	10 A
	PI3318-20-LGIZ	1.8 V	1.4 V to 2.0 V	10 A
	PI3312-20-LGIZ	2.5 V	2.0 V to 3.1 V	10 A
	PI3301-20-LGIZ	3.3 V	2.3 V to 4.1 V	10 A
	PI3302-20-LGIZ	5.0 V	3.3 V to 6.5 V	10 A
	PI3303-20-LGIZ	12 V	6.5 V to 13.0 V	8 A
	PI3305-20-LGIZ	15 V	10.0 V to 16.0 V	8 A
	PI3311-21-LGIZ	1.0 V	1.0 V to 1.4 V	15 A
	PI3318-21-LGIZ	1.8 V	1.4 V to 2.0 V	15 A
	PI3312-21-LGIZ	2.5 V	2.0 V to 3.1 V	15 A
	PI3301-21-LGIZ	3.3 V	2.3 V to 4.1 V	15 A
	12 V Optimized Option			
8 – 18 Vin	PI3420-00-LGIZ	1.0 V	1.0 V to 1.4 V	15 A
	PI3421-00-LGIZ	1.8 V	1.4 V to 2.0 V	15 A
	PI3422-00-LGIZ	2.5 V	2.0 V to 3.1 V	15 A
	PI3423-00-LGIZ	3.3 V	2.3 V to 4.1 V	15 A
	PI3424-00-LGIZ	5.0 V	3.3 V to 6.5 V	15 A

I²C is a trademark of NXP Semiconductors

Resources

- Video: Interview with ECE Europe about ZVS Regulators
- Webinar: ZVS Point-of-Load Regulation – Enabling High Performance On-Board Power Solutions
- Webinar: Design Considerations For High Performance On-Board Power Design
- Cool-Power ZVS Buck Regulators Product Information

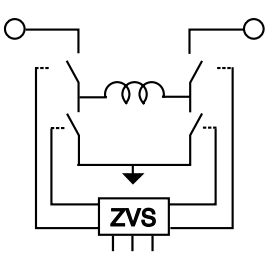

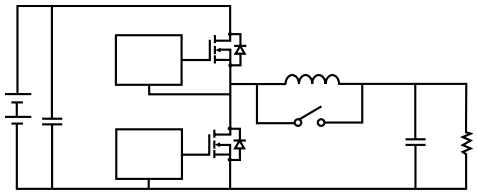

- Reduces Q1 turn-on losses
- Reduces gate drive losses
- Reduces body diode conduction

Benefits of **Zero-Voltage-Switching Topology**

VI Chip PRM Module

Simple to Use
- Point-of-load, Buck-Boost regulation
- Factorized Power Architecture
- Minimal external components

High Density
- Up to 1,700 W/in³, with 500 W in 1.1in² package

Wide Vin Optimized for 48 Vout
- 24 Vin, 18 – 36 Vin range
- 36 Vin, 18 – 60 Vin range
- 45 Vin, 38 – 55 Vin range
- 48 Vin, 36 – 75 Vin range

High Efficiency
- Full chip 500 W: 97.8%
- Half chip 250 W: 96.7%

Flexible
- Regulation: Remote sense, local loop, adaptive loop
- Parallel capabilities

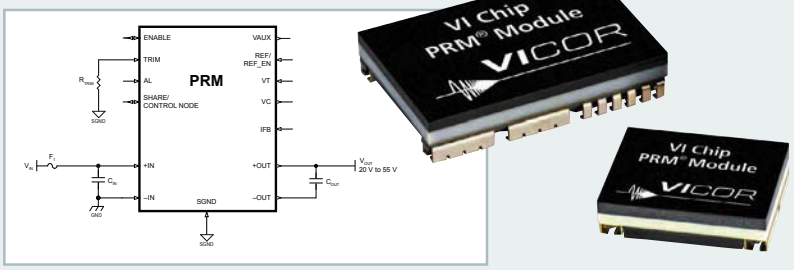

PRM Modules Model Number	Input Voltage Nom. (V)	Input Voltage Range (V)	Output Voltage Voltage Range (V)	Output Power Max.	Output Current Max.	Package Size
P024F048T12AL	24 V	18 – 36 V	26 – 55 V	120 W	2.5 A	Full
P036F048T12AL	36 V	18 – 60 V	26 – 55 V	120 W	2.5 A	Full
P045F048T17AL	45 V	38 – 55 V	26 – 55 V	170 W	3.5 A	Full
P045F048T32AL	45 V	38 – 55 V	26 – 55 V	320 W	6.67 A	Full
P048F048T12AL	48 V	36 – 75 V	26 – 55 V	120 W	2.5 A	Full
P048F048T24AL	48 V	36 – 75 V	26 – 55 V	240 W	5.0 A	Full
PRM48BH480T200A00	48 V	38 – 55 V	5 – 55 V	200 W	4.17 A	Half
PRM48BF480T400A00	48 V	38 – 55 V	5 – 55 V	400 W	8.33 A	Full
PRM48AH480T200A00	**48 V**	**36 – 75 V**	**20 – 55 V**	**200 W**	**4.17 A**	**Half**
PRM48AF480T400A00	**48 V**	**36 – 75 V**	**20 – 55 V**	**400 W**	**8.33 A**	**Full**
PRM48BH480T250A00	**48 V**	**38 – 55 V**	**20 – 55 V**	**250 W**	**5.21 A**	**Half**
PRM48BF480T500A00	**48 V**	**38 – 55 V**	**20 – 55 V**	**500 W**	**10.42 A**	**Full**

These PRM modules can be further configured to meet your exact needs.

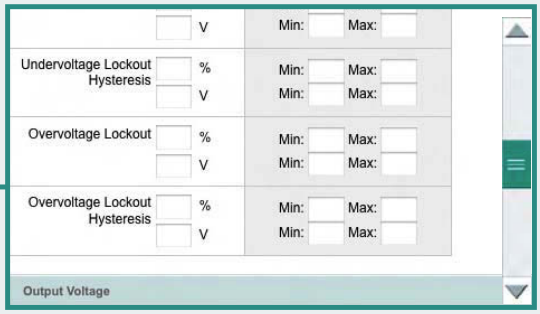

Resources

- Video: Overview of Vicor's VI Chip PRM Module
- PRM Product Information
- Configure a PRM for your application's requirements

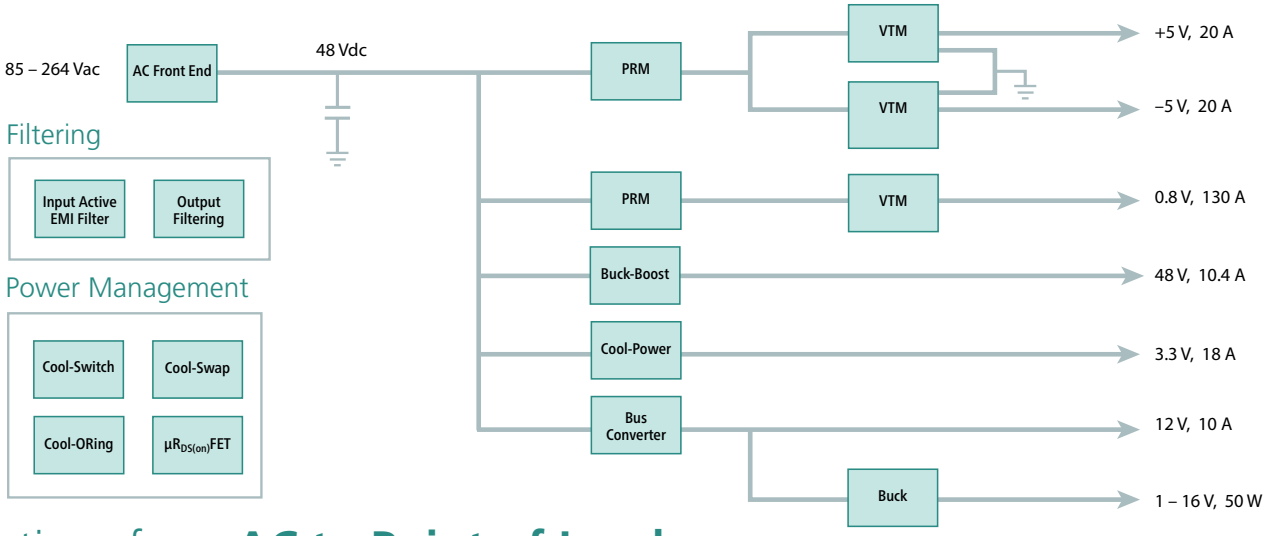

Solutions from **AC to Point of Load**

Introducing... The Growing ChiP Lineup

"Converters housed in Package" Technology

Recent Press Articles

- EDN: Vicor reveals new ChiP power package at APEC 2013
- Power Electronics: New Thermal-Design Options Drive Power Density

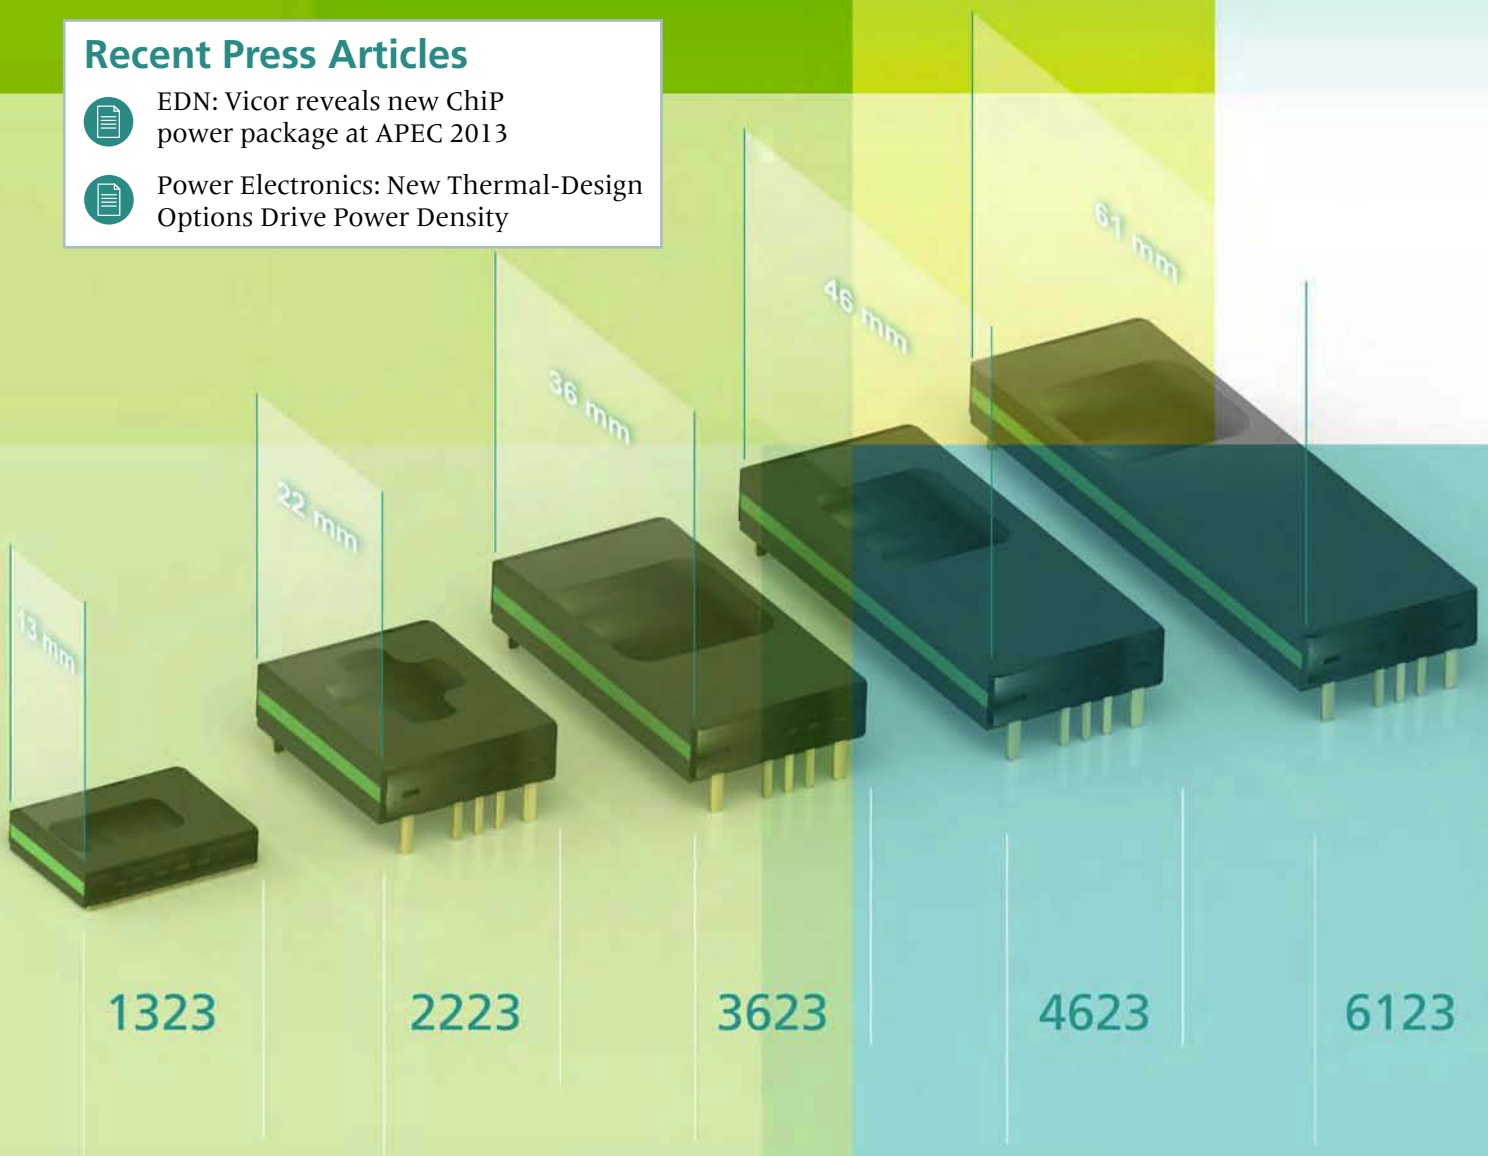

Resources

Vicor's CEO discusses ChiP technology at APEC

An introduction to ChiP technology

Online Design Tools

Online Simulator

- Simulate electrical and thermal behavior
- User defines line and load conditions, input and output impedance and filters
- Simulations include start-up, steady state, shutdown, Vin step and load step, as well as thermal.
- Electrical and thermal performance showed in charts and tables

- Determine trim resistors for fixed and variable output voltage trimming
- Calculate required bus capacitance for VI-ARM, FARM, and ENMod modules
- Thermal calculator for heat sink selection

Filter Design

- Select attenuation and frequency
- Choose from five different topologies
- Supports regulated and unregulated converters

Design Calculators

Resources

- Video: Using Vicor's online simulator
- Online simulator: VI Brick IBCs
- Online simulator: BCMs
- Online simulator: PRMs
- Filter design tool
- Design calculators

Calculation of Components for Parallel Damped Filter

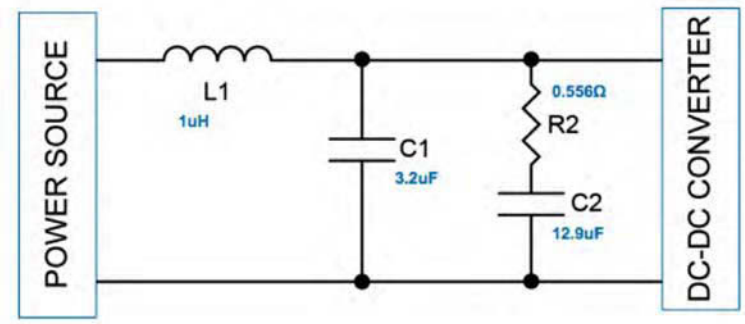

Configure Your Product

PowerBench™ online design center

- Design your own DC-DC converters to meet your application's requirements
- Or use hundreds of predefined designs
- Online registration allows designs to be saved

VI Chip® PRM® Module

- Point-of-Load Buck-Boost regulation with remote sense
- Full Chip (up to 500 W in 1.1 in²)
- Half Chip (up to 250 W in 0.57 in²)

Other DC-DC Converters

- Maxi, Mini, Micro Series: Full (160–600 W), Half (100–300 W), Quarter (50–150 W)
- VI-200 / VI-J00 Series: Full brick (50–200 W), Half brick (25–100 W)
- ComPAC, VIPAC Arrays and chassis-mount MegaMods

AC-DC Converters

- VIPAC - Autoranging input with filtering, multiple output, cold plate chassis,
- FlatPAC - Multiple output and autorange input with heat sink or conduction-cooled models

Complete power systems

- Westcor custom AC-DC
- High power density, small size, high efficiency
- Fan-cooled, slide-in assemblies

Resources

- PowerBench online design center
- Design calculators
- Webinar: Modeling, Simulation, and Selection Techniques in Power Design

VI Chip **PRM Module Configurator**

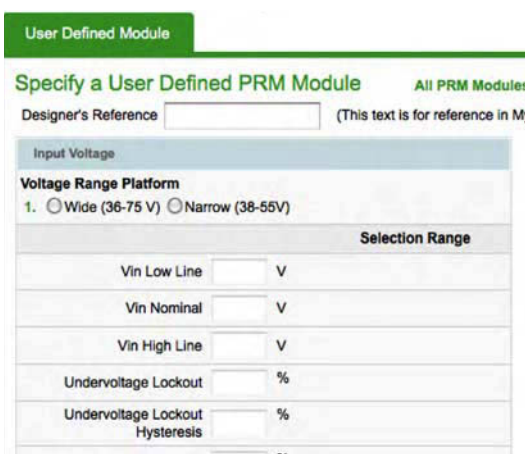

Enabling Our Customers' Competitive Advantage

At Vicor, we enable customers to efficiently convert and manage power from the wall plug to point of load.

We master the entire power chain with the most comprehensive portfolio of high efficiency, high-density, power distribution architectures addressing a broad range performance-critical applications.

Vicor's holistic approach gives power system architects the flexibility to choose from modular, plug-and-play components ranging from bricks to semiconductor-centric solutions.

By integrating our world-class manufacturing and applications development, we can quickly customize our power components to meet your unique power system needs.

Focus Performance-Centric Markets /Applications

Communications

> 400 VDC Power Distribution
> Datacom
> Netcom
> Telecom Infrastructure

Computing

> Data Centers
> High Performance Computing
> Network Servers

Industrial

> ATE
> Lighting
> Process Control
> Transportation

Automotive

> Electric Vehicles
> Hybrid Vehicles

Defense/Aerospace

> Aircraft Test Equipment
> Ground Vehicles
> Radar
> Telemetry
> Unmanned Vehicles

Phone:
Europe: 00 800 8426 7000
US: 1 800 735 6200

VICOR

BY JOHN BLEVINS, DIRECTOR OF PRODUCT MARKETING, LYNUXWORKS

SECURING THE INTERNET OF THINGS

Traditional network and endpoint protection mechanisms are struggling today to protect our corporate IT infrastructure and computers. Adding the projected billions of nodes that will constitute the Internet of Things will exacerbate the problem dramatically.

The advantage that the embedded industry has over the traditional computer world, is that the devices within it are generally being built for a specific purpose, and as part of the design of these systems, security and protection can be a primary consideration. However, adding security of any form to a resource-constrained device that isn't secure by nature, is almost impossible. Providers of the traditional embedded software platform, the real-time operating system, are in a prime position to address this need by providing the security functionality, as an integral part of their offering, that developers need to protect connected embedded devices and allow them to be designed to be secure.

NSA's "Protection Profile for Single-level Operating Systems in Environments Requiring Medium Robustness" (SLOSPP), specifies the security requirements for commercial-off-the-shelf (COTS) general-purpose operating systems in networked environments that contain sensitive information. LynuxWorks' LynxOS 7.0 is amongst the first to implement these standards, allowing developers to embed military-grade security directly into their devices. Security features include discretionary access control, audit, roles and capabilities, identification and authentication, cryptography, quotas, self-test, residual information protection, and local trusted path. The following paragraphs take a deeper look at how these features can help protect networked embedded devices from malicious attacks.

Discretionary Access Controls (DAC) are a means of restricting access to objects (such as files, applications, directories, and devices), based on the identity of the user or group to which they belong. LynxOS implements DAC using Access Control Lists (ACL) as defined by the Posix.1e standard. ACL's provide for much finer grain control over who can access an object as compared to the traditional UNIX "user/group/others" security implementations. ACLs allow authorised users to specify which resources may be accessed by specific users and groups of users. Essentially, DAC gives the networked device the ability to carefully control who can access and execute files and data on the device.

Audit facilities allow an embedded device to capture significant system events and perform security monitoring of these events. Fine-grained event auditing tracks events such as login, logout, object accesses, and administrative tasks, which can all be logged to an audit trail. The audit trail contains invaluable information that can be used to

- Review security-critical events
- Discover attempts to bypass security mechanisms
- Track usage of privileges by users
- Provide a deterrent against attempted attacks since audit logging captures intrusion
- Perform forensic analysis

Comprehensive audit records contain information such as the audit event, the user that caused the event, and success or failure of the event. The OS can provide utilities to monitor the live audit events as they are occurring and to analyse the audit records written to the audit trail helping to make connected devices much more secure.

Roles and Capabilities provide finer grain privilege levels for users over the traditional UNIX model. Historically, in UNIX and similar systems, privilege to perform administrative tasks is reserved to the root user, who has complete privilege to see or modify the entire system, whether or not those privileges were required by the original task at hand. Normal users, on the other hand, had no administrative privileges whatsoever. A normal user could do only what was allowed based on his or her user ID and group ID.

In general, the all-encompassing privileges of root makes it susceptible to abuse, whether unintended or malevolent. This problem is resolved in POSIX.1e by decomposing the root privileges into a set of fine-grained privileges or capabilities. The idea is that instead of having an all-powerful root user, one can create multiple administrative "roles" each with a subset of privileges (called "capabilities") specific to the task at hand. For example, one might create a "netadm" role and only give it the capabilities related to manipulating the network (CAP_NET_ADMIN). One might create a second role called "auditadm" and assign to it the capabilities of controlling the Audit logs (CAP_AUDIT_CONTROL) and writing to those logs (CAP_AUDIT_WRITE).

Embedded devices that implement distinct roles, each with limited capabilities, will be more secure against attacks which target and compromise a specific "user" account on the system.

Identification and Authentication refers to the process whereby a device on the network recognizes a valid user's identity and can verify (authenticate) the claimed identity of the user. LynxOS 7.0 identifies users with the traditional POSIX.1 user id / group id model, where a user has a unique user id and is a member of one or more groups. This strong identification and authorisation mechanism significantly reduces the risk that unauthorised users will gain access to embedded devices.

Cryptography provides techniques for secure communication in the presence of potential adversaries. Cyrptographic software uses sophisticated algorithms to turn plaintext data into nonsensical strings and vice versa based on encryption keys (public, private) and digital signatures. The FIPS 140-2 certified crypto algorithms and random number generators are based on the certified OpenSSL and OpenSSH libraries. These libraries provide powerful open standard security protocols for the Secure Sockets Layer (SSL), Transport Layer Security (TLS), and a full strength cryptography package.

Embedded devices on the network can take advantage of the FIPS 140-2 certified crypto algorithms to protect sensitive data and passwords.

Quotas are implemented to protect the device from running out of disc, memory, or CPU. As an example, if a denial of service attack is made against the system with the intent of over-utilising the CPU, an unprotected device could essentially appear to freeze up. By being able to set limits on the amount of critical resources, such as CPU, that a task is able to consume, the overall system is protected from such an attack.

Likewise, if a process tries to allocate memory or disk resources and the usage of the user or group exceeds the quota then an error will be returned and the allocation will fail.

Self-Test provides the ability to run a set of tests during the device's initial start-up, periodically during normal operation, or at the request of an authorised administrator, to demonstrate the correct operation of the device. Self-tests are designed to provide confidence that the device's memory, and file systems, are in the expected state(s).

Residual Information Protection is a feature that ensures that whenever a resource is allocated or freed, the content of the resource can be made unavailable to others' processes. LynxOS 7.0 makes user memory and file system data unavailable to others when it is freed resulting in a more secure system. Whenever a file system object is removed, all the blocks allocated to that file system object will be filled with zeros. The OS can be configured to fill memory allocated or freed by a process with zeros.

Trusted Path is a mechanism that provides confidence that the user is communicating with the entity the user intended to communicate with, ensuring that attackers can't intercept or modify whatever information is being communicated.

A simple and elegant mechanism called Secure Attention Key Sequence provides a trusted path because normal programs can't intercept this key pattern: this approach creates a trusted path between the user and the OS computing base. The OS also provides a modular framework called Trusted Menu Manager (TMM) for all trusted operations. Trusted operations such as user login, role login, and password change operations can be performed from this menu; screen locking/unlocking operation can also be performed.

The trusted path mechanisms combine to ensure that a user who supplying sensitive data, such as login information will not have to worry about their login information being stolen by spyware programs.

With these powerful RTOS security features available to embedded developers they can design their systems to be more secure before they are connected to the internet. Once the applications are secured, developers can take advantage of the OS's networking support for both long haul networks using TCP/IPV4, IPV6, 2G / 3G / 4G cellular, and WiMax communication stacks as well as short-haul networks common with M2M applications such as 802.11 Wi-Fi, ZigBee wireless mesh and Bluetooth.

For systems that require the use of a more general purpose OS, such as Windows or Android to provide end-user familiarity, security becomes more difficult to design into the operating system itself, and other protection or isolation mechanisms need to be used to help protect against cyber attacks.

A further technology allows embedded developers to build security into their systems, regardless of which operating system is used. Originally designed and developed to meet the exacting security needs of the DoD, a **separation kernel** is a technology that provides isolation of devices and memory on a microprocessor based system. It is particularly well suited to embedded devices as it is typically a very small and efficient kernel implementation offering real time properties as well as isolation. In DoD systems the separation kernel is used to help provide domain isolation for applications running at different security levels on a single hardware system, and protected domains from seeing each other's data, network and applications.

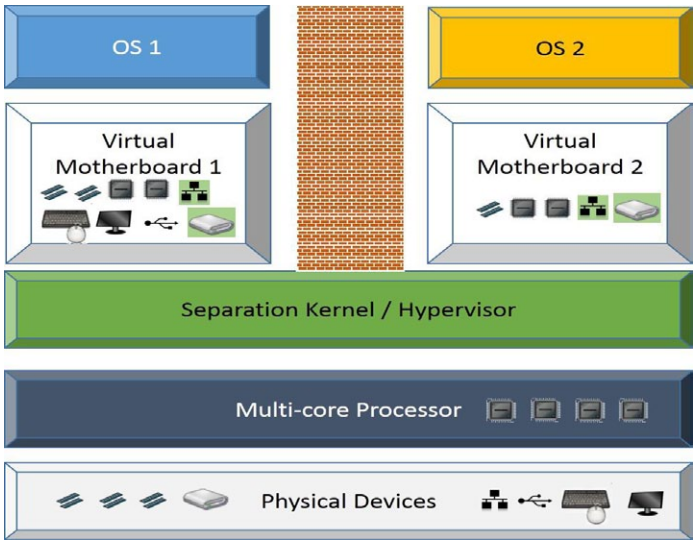

Figure 1 – Type-0 hypervisor offering two secure virtual motherboards to guest operating systems

Adding virtualisation to an embedded system is now particularly interesting as more embedded processors contain two or more cores, and using virtualisation is a good way of segmenting a multi-core system in an efficient way, and offer consolidation of multiple physical systems onto a single multi-core system. A flexible separation kernel/hypervisor solution will allow the embedded developer to look at the requirements of each guest OS, and allocate the appropriate number of processors. This could involve sharing processors across multiple OSs, dedicating a single core to a single OS, and allowing OSs that support symmetric multi-processing to access multiple cores.

As embedded systems use more traditional computer OSs and then link to the open internet, they become are every bit as vulnerable to stealthy cyber threats as our normal computers. However, unlike our normal computers, these embedded devices could be controlling our critical infrastructure and when infected, could pose a dramatic security breach, even affecting nation states. The latest super stealth cyber weapons include rootkits and bootkits that infect the system below the OS, residing on hard discs, in devices or in memory and wait for instructions from a command and control center communicating over the internet before actually starting their attack. Traditional network and endpoint protections are quite poor at detecting them, and general purpose operating systems are very susceptible to them, and hence our new connected embedded systems are vulnerable to them.

The LynxSecure separation kernel/hypervisor has accordingly been equipped with a rootkit detection feature that has been designed to detect and alert when the rootkit has entered the system and is looking for a place to hide itself.

In conclusion, the propagation of connected embedded systems, as predicted as part of the "internet of things" is opening up a potential security hole for cyber criminals and cyber terrorists to take advantage of. Embedded developers now have new military-grade tools available to design security into their devices at the RTOS level as well as by using a secure separation kernel/hypervisor. For smart devices these threats can be mitigated and detected long before the attack has really started.

John Blevins is the Director of Product Marketing at LynuxWorks. He has over 25 years of software development experience in the embedded industry working primarily on software compilers, debuggers, IDEs and related tools. He can be reached at jb@lnxw.com

Solid-state circuit breaker for microcontrollers

By Didier Juges

Many MCU designs are intended to control power to devices such as motors, actuators, or other electronic subsystems. This power switching usually involves using a power device such as a bipolar transistor or MOSFET that the MCU controls, through a level shifter if necessary.

This MCU-based power control has a potential weakness. If the device to be controlled is defective or fails in a way that overloads the system power supply, the microcontroller's own input power may drop, causing the MCU to reboot. When the processor reboots, the default condition is typically one of no power to the controlled device, and so everything seems to work again – that is, until the device is powered again.

This kind of self-correcting failure mechanism is usually a recipe for difficult troubleshooting. To avoid the situation, it is desirable to maintain power to the processor at all times to help with system troubleshooting, while protecting the power-controlled circuit against excessive currents.

The typical solution to this problem is to use one of those chips designed for hot-plugging (hot swap controller). One of my favourites is the LT1910 fault-protected high-side driver from Linear Technology. I like it because it works over a wide voltage range (not just 3V or 5V) and can restart autonomously or be controlled with one pin. However, in my opinion, there are several issues with using such chips: They are single-sourced and not inexpensive, they still require a number of external components, and if you want to monitor the chip (i.e., detect when faults occur) as well as control it, you will need two lines to the MCU (although the LT1910 can be configured to use just one).

Before such chips were available, I designed my own solid state circuit breaker for just such applications, and it still offers many advantages. The design uses two bipolar transistors, a P-channel MOSFET, and a handful of passive components. The LT

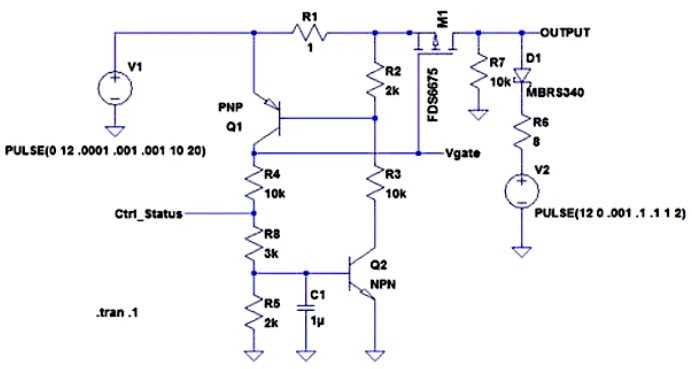

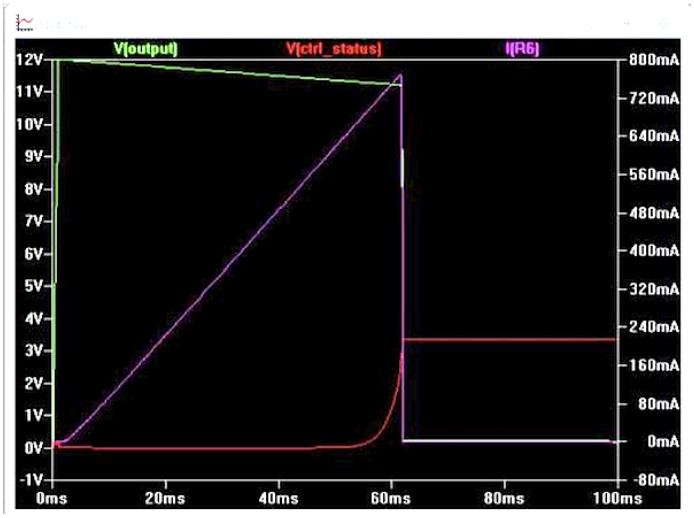

Spice model schematic below shows values suitable for a 12V supply voltage, approximately 700 mA maximum load current, and a 5V logic interface to the MCU. You can adjust R4 and R8 for different supply or logic voltages, and change R1 to alter the load current. (Note that V2, D1, R6, and R7 are parts of the load simulation and not part of the circuit breaker itself.)

The main disadvantage of this circuit compared to integrated hot-swap controllers is the relatively large voltage drop through the sense resistor R1, which is sized to drop about 0.7V at the current value where the circuit is supposed to trip. Another disadvantage is the trip point's variability. Because the trip point depends on the Vbe of Q1 it is temperature sensitive, changing at a rate of about 1.8mV/K (0.25%/K). If your application can tolerate these modest shortcomings, however, the circuit is inexpensive, uses widely available components, and is quite flexible.

An advantage over most hot-swap controllers is that the circuit provides status and control through a single connection: the Ctrl_Status pin. After turning the circuit breaker ON, the processor changes the Ctrl_Status pin to an input and then needs only to periodically check the pin's status to ensure the breaker has not tripped. Here is how it works:

The Ctrl_Status point ties to an I/O pin of a conventional microcontroller. Unless you want the power-controlled device to come up automatically at power-up, you should use an I/O that can be configured to provide a high logic level at power-up. The high level ensures that Q1 and Q2 will be ON and the gate of the P-channel MOSFET M1 will be near the source potential,

keeping the MOSFET in the OFF state at power-up.

To turn the breaker on and apply power to the load, the processor momentarily sets the Ctrl_Status pin to a low state (long enough to discharge C1), then sets the pin as an input (high impedance). The low pulse turns off Q2, which then turns off Q1 and lets M1's gate go to ground, turning the MOSFET ON. As long as the Ctrl_Status pin remains in high impedance, Q1/Q2 will both remain OFF and the MOSFET will remain ON.

If an overload occurs, Q1 will turn ON, which will turn Q2 ON and both Q1/Q2 will keep each other ON (and the MOSFET OFF) until the Ctrl_Status pin is brought low again.

If the load includes filtering capacitors, the output current may momentarily exceed the trip threshold during a normal turn on. To prevent such parasitic triggering at power-on with capacitive loads, you only need to lengthen the low pulse on the Ctrl_Status pin. As long as the Ctrl_Status pin is maintained low, Q1 will keep the MOSFET in current limit, but Q2 won't be able to close the latch. The current limit action will avoid a large current spike on the power supply.

Capacitor C1 provides an additional filtering in cases where the load may draw short current spikes during normal operation. During those spikes, Q1 will provide current limit, but the circuit will not latch unless those spikes exceed a duration that the value of C1 determines.

The LT Spice simulation output below shows the circuit breaker tripping. You can see the current ramping up and the output voltage dropping to zero when the current reaches about 750mA.

I have used this circuit in many designs and have been quite happy with it.

Voltage-to-period converter improves speed, cost, and linearity of A-D conversion

By Jordan Dimitrov

Designers often use VFCs (voltage-to-frequency converters) to perform A-to-D conversion in data acquisition systems that require strict monotonic response, high resolution, reduced noise and moderate speed. The VFC produces a pulse train with frequency proportional to the input voltage. Then a microcontroller or logic converts frequency into a number by opening a gate for a fixed amount of time and counting how many pulses go through the gate for that time.

The main drawback of this approach is that to increase speed, designers have to run the VFC at high frequency, which degrades linearity.

This Design Idea reverses things. A circuit converts input voltage into a proportional time interval; then, the micro uses that interval to count pulses coming from its internal clock. The results are impressive:

- Good linearity as the voltage-to-period converter runs at low frequency
- Faster A-to-D conversion due to the high value of the clock frequency
- Potentially simpler program or logic, as it only has to count clock pulses, gated by the circuit
- Low price

The key is that increasing the count frequency does not affect linearity of the A-to-D conversion, while increasing the frequency of the VFC always means worse linearity.

Figure 1 shows the circuit. It is a modified VFC (Ref 1), where the input voltage VIN and the reference voltage VREF swap their roles. The R1-R2 network shifts the input voltage so it is always more positive than the reference voltage and maintain proper operating conditions. The circuit uses all switches of the 4066 part: two in parallel build S1 to reduce the effect of imperfect switch flatness on linearity, one switch goes for S2, and the last switch is part of the start-up circuit, paralleling CINT, and controlled by the logic during initialisation, or as shown in Ref 1.

When the circuit gets power, the start-up circuit shorts CINT for a while and the one-shot is reset by its internal circuitry. Switch S1 opens and S2 closes. The left side of RINT connects to the ground and the right side gets the potential of 0.2V. After the start-up time is over, the switch in parallel to CINT opens and the capacitor start charging. The integrator makes a rising ramp. When the ramp reaches the ~2.5V threshold, the one-shot triggers. Switch S1 closes and switch S2 opens. The right side of RINT gets a positive potential that depends on the input voltage, but is always greater than 0.2V. The current through RINT reverses direction, and CINT starts discharging. When the one-shot interval is over, the cycle repeats.

As the input voltage changes from 0 to 5V, the output period changes from 78 to 578 µsec. Integration capacitor CINT and the threshold level of the one-shot's Schmitt input do not participate in the period-vs.-voltage relation. Filling the period with 10 MHz clock pulses generates numbers from 780 to 5780 – one count per millivolt. Linearity is one count or ±0.02%, which is not a surprise when the maximum frequency is only 12.8 kHz. The maximum time of the A-to-D conversion is 578 µsec. This is 8.65 times faster compared to the case of a 1 MHz VFC, where it would take 5,000 µsec to count 5,000 pulses of 1µsec. The interface program is short and simple; very similar to what is

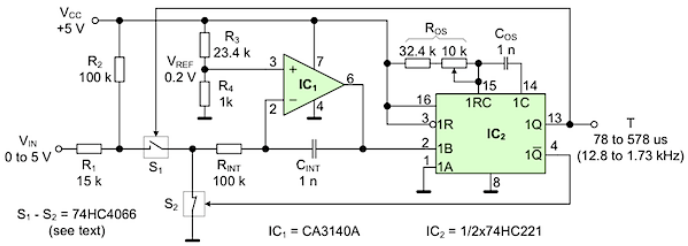

Figure 1. *The circuit uses a modified VFC architecture where the input voltage applies to the reference voltage terminal and vice versa.*

Parameter	V-to-F converter	V-to-P converter	Unit
Max frequency	1000	12.8	kHz
Linearity	0.1	±0.02	%
Conversion time	5000	78 to 578	µs
Program interface	Define counting interval and count pulses	Count pulses	
Power supply	Dual	Single	
Price	18 to 24	3-4	USD

Table 1. *Performance of a 5,000 count A-to-D conversion with 1 MHz VFCs and the proposed circuit.*

presented in Ref 2.

Calibration involves some back and forth due to the shift of the input voltage: adjust sensitivity to 100 µsec/V using the trim-pot of the one-shot. The nominal duration of the one-shot pulse is 26 µsec. Cancel the 780 count offset in the controller.

Table 1 shows that the V-to-P approach is significantly better than the V-to-F one (Refs 3, 4). Surprisingly, no chip-maker offers this type of converter.

REFERENCES:

1. Dimitrov J., Inexpensive VFC features good linearity and dynamic range. EDN, Design Ideas, Dec 1, 2011, pp.47-48.
2. Dimitrov J., Linearize optical distance sensors with a voltage-to-frequency converter. EDN, Design Ideas, Apr 19, 2012, pp.47-48.
3. AD650 voltage-to-frequency and frequency-to-voltage converter.
4. VFC320 voltage-to-frequency and frequency-to-voltage converter.

Jordan Dimitrov is an Electrical engineer in Measurement and Instrumentation. After 30 years of experience he features 2 patents, 70+ papers and a PhD degree. Currently he teaches electrical and electronics courses at a large community college in Toronto.

White noise source flat from 1 Hz to 100 kHz

By Steve Hageman

White noise is very useful in testing many types of circuits. When combined with an FFT analyser, a flat noise source can make for quick and easy gain plots of circuits. If the noise into a circuit is flat and of known quantity, then the gain of the output circuit is easy to determine, even visually. This method has been used at least as far back as 1978 on the HP3582A Low Frequency Spectrum Analyzer [Ref. 1].

The "modern" way to generate white noise is to use a digital shift/feedback register arrangement in a CPLD or FPGA. Some authors have even created a parallel arrangement of microcontrollers to generate Gaussian white noise.

The implementation below is all analogue, and the parts are even available in through-hole varieties if desired, making prototype construction easy.

It is well known that Zener diodes are a good source of wideband noise. The trick is, as always, to find a diode that is flat over the frequency range you need. Jim Williams used a conventional 6.8V Zener in his 5MHz wideband noise generator. Also common is to reverse bias the base to emitter junction of a NPN transistor and use this as a noise diode.

The design goal here was to generate a large amount of noise that is flat over the frequency range of 1 Hz to 100 kHz for FFT testing purposes. The 6.8V Zener diode designs do in fact have wideband noise, but I found them to have large amounts of 1/f noise at low frequencies, and are not flat down to DC. In this design, I used instead the trusty 12V Zener as a noise source. I have found that these diodes are flat over the frequency range, have a large amount of intrinsic noise, and operate very well over the discharge life of the 9V batteries used [Ref. 2].

The intrinsic noise of the diode chosen when it is biased to 18V through 1 MΩ is approximately 20 mV RMS. This was fortuitous scaling, as the P-P value is approximately five times

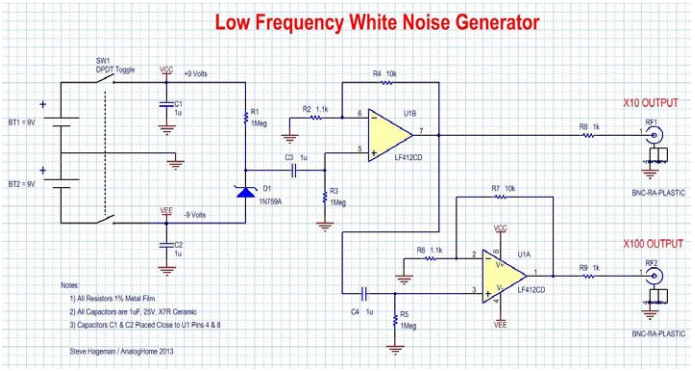

Figure 1 *This low frequency noise generator produces 1 or 10 VP-P over a frequency range of 1 Hz to 100 kHz.*

this, or 100 mV P-P.

To keep the DC offset errors in check I used the venerable LF412 dual JFET-input op-amp to amplify the diode noise in two x10 steps.

Figure 1 shows the resulting circuit. The Zener diode is biased from the two 9V batteries in series (18V) through 1 MΩ. The LF412 is operated split-supply between the 9V batteries. The low input current and low offset voltage allow the elimination of any bulky output coupling capacitors, as the output's DC value is within millivolts of ground. While the LF412 is not a low noise op-amp, its noise level is still well below the intrinsic noise level of the diode and is not a concern here.

The Zener Noise is amplified by U1A and U1B in two x10 steps, to approximately 1V P-P and 10V P-P. If this is too much noise for the circuit under test, then the 1 kΩ R8 & R9 can be used to advantage in making a voltage divider to lower the noise level to any desired value.

Figure 2 shows the resulting noise of the circuit, which is flat from 1 Hz to 100 kHz. The slight high frequency roll-off at 100 kHz in the x100 output is less than 0.5 dB and could be ac-

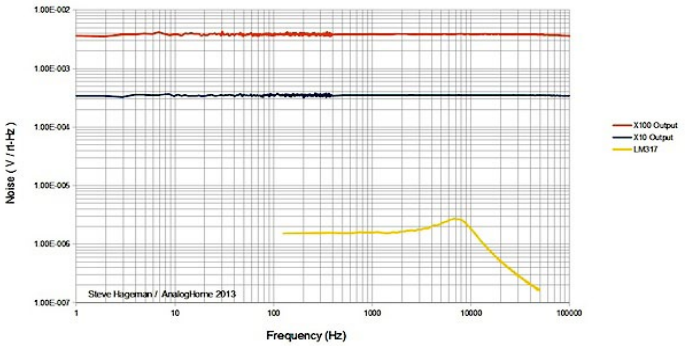

Figure 2 *The power spectral density of the noise output is very flat from 1 Hz to 100 kHz. As a comparison, the noise of an LM317 regulator is also plotted, as this regulator is normally thought of as being very noisy.*

counted for by adding some frequency dependent gain to U1B if desired, but for my purposes this was unneeded. For comparison, Figure 2 also shows the noise plot of an LM317 regulator operating with minimum capacitors, as this configuration is normally regarded as a "Very Noisy Regulator", but it is nothing compared to a 12V Zener diode and some gain.

Due to the use of stable, low power FET amplifiers and ceramic coupling capacitors, the 1/f noise induced by temperature gradients from stray air currents is kept to a minimum. Still, the design should be operated in some sort of enclosure, and kept away from circulating air currents, for maximum stability.

The circuit draws just 4 mA, and if the batteries are used down to 7V, the expected battery life from regular alkaline 9V batteries is greater than 100 hours. The noise of the circuit as-designed varies by approximately 15% over the life of the batteries; if desired, a more complex and stable biasing arrangement for the Zener could be employed to improve this.

REFERENCES:

1. Pendergrass, N., Farnbach, J., "A High Resolution, Low Frequency Spectrum Analyzer", Hewlett-Packard Journal, September 1978.
2. I tested 5 each 1N759A and 1N4742A diodes with essentially the same results, based on the samples I had either would work in this application.

Steve is the owner of AnalogHome - A Custom Electronics Design Resource. Steve has designed and provided high quality electronic hardware and software solutions from DC to 20 GHz for a diverse set of markets; www.AnalogHome.com

Double-speed interpolation

By Michael Mahon

It is sometimes necessary to perform linear interpolation or "blending" of two vectors. For example, this is done to "mix" two audio waveforms or blend two images. The equation is:
out = a•in1 + (1-a)•in2
where out is the output vector, **in1** and **in2** are the two input vectors, and **a** is the fraction of **in1** to be blended with the complementary fraction of **in2.**
For each element of the vectors, this computation usually requires two multiplications and an addition.
While this is not particularly onerous for a processor with a fast multiply operation, it can be a significant performance issue for processors that must synthesise multiplication by repeatedly shifting and adding. This Design Idea simplifies the computation on processors that do not have a fast multiply by evaluating the complete equation at the cost of just one shift-and-add multiply loop, as shown below in 6502 code for 8-bit operands:

```
*****************************************
*              BLEND
*
*  Called with blend fraction (0..1-) in Accumulator
*         X register = vector subscript i
*     (IN1 and IN2 are input vector base addresses)
*
*  Upon exit, Accumulator = OUT(i)
*         factor = low 8 bits of OUT(i) (unused)
*
*****************************************

blend    lsr           ; Low bit of factor to CARRY
         sta  factor   ;  and save shifted result.
         lda  IN2,x    ; Init sum to IN2(i)
         ldy  #8       ; 8 iterations
blender  bcc  addIN2 ; bit off: add IN2(i)
         clc           ; bit on:  add IN1(i)
         adc  IN1,x    ;  (after clearing CARRY)
         jmp  shift

addIN2   adc  IN2,x    ; add IN2(i) (CARRY already clear)
shift    ror           ; (shifts in any carry out)
         ror  factor   ; Next bit of factor to CARRY
         dey
         bne  blender; 8 iterations
         rts           ; Return with OUT(i) in Accumulator.
```

The following is a brief summary of the 6502 operations used in this algorithm (with the exception of the Store, Load, and Jump instructions, which do just what you'd expect):

lsr; Logical Shift Right; Shift the accumulator right one bit, shifting a zero into the MSB and shifting the LSB out into the CARRY toggle.

bcc; Branch if Carry Clear; Branch if CARRY toggle is 0

clc; CLear Carry; CARRY = 0.

adc; ADd with Carry; Accumulator = Accumulator + operand + CARRY. Therefore, CARRY is typically cleared prior to doing 8-bit adds.

Ror; ROtate Right; Shift the operand (either a memory byte or the accumulator if no operand is specified) one bit to the right, shifting the CARRY toggle into the MSB and the LSB out into the CARRY toggle.

dey; DEcrement Y; Decrement the Y register by 1, setting the Equal toggle if result is 0.

bne; Branch Not Equal; Branch if the Equal toggle is set.

rts; Return from SubRoutine; Returns from a subroutine that has been entered by a Jump to SubRoutine instruction.

The hash prefix "#" is used to indicate a literal value.

The operation of the algorithm is straightforward. For every bit of the blend factor that is one, the first input is added to the sum, and for every bit that is zero, the second input is added. Because **a** and **(1-a)** are complements, the loop directly computes the sum of the two products in the time normally required by a single multiply loop.
The blend fraction, **a**, is an 8-bit unsigned fraction. **IN1**, **IN2**, and **OUT** are all 8-bit unsigned integers.

Note that the Accumulator is initialized to **IN2(i)** so that if the blend factor is zero, the value returned is equal to **IN2(i)**.

Michael Mahon is a retired Hewlett-Packard computer system architect. During his career, he developed and later managed labs developing compilers, operating systems, and processor architectures. His hobbies include retrocomputing (Apple II), digital photography, astronomy, flying, and sailing. Apple II projects are documented on his website: http://home.comcast.net/~mjmahon/

productroundup

Offline LED driver eliminates inductors

A FLOATING SWITCH, AC switched matrix technique will replace flyback, buck and boost converters, according to Texas Instruments. The floating switch architecture will simplify the offline linear drive of LEDs in lamps, downlights and fixtures. The AC switched matrix technique features TI's TPS92411 floating MOSFET switch and is a new approach to producing low-ripple LED drive current without magnetic components. The architecture provides excellent compatibility with legacy wall dimmers and delivers high power factor and low total harmonic distortion (THD). Its features include; 100-V, 2-Ω floating MOSFET switch with 350 mA current capability; LED drive performance comparable to traditional flyback, buck and boost converters; greater than 0.95 power factor; less than 15% total harmonic distortion; low LED current ripple resulting in efficient LED utilisation as compared with "tapped linear" drivers. Its low-frequency, slew-rate controlled switch action produces very little EMI noise, and it supports lighting designs up to 70 W.

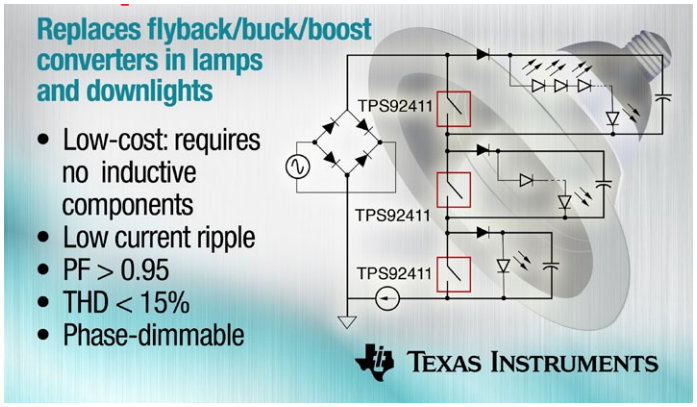

TPS92411 comes in a 5-pin SOT-23 package at $0.23 (1,000) An evaluation module costs $75.00.

Texas Instruments; www.ti.com/tps92411-pr

Haptics platform for small touchscreens

FROM REDUX LABS, the customisable Bullfinch platform mimics the sensation of real-world keys on a flat screen, imitating keyboard, double-press camera button, home key, and volume sliders. With Bullfinch, said to be the first 4G-haptics platform optimised for a smartphone's small form factor, 4G haptic responses can be customised according to where a user touches the screen, what application they're in and how hard they press. Bullfinch is able to accurately reproduce, on a flat

screen, a wide range of customisable haptic responses that mimic real world key press sensations, be it a home key, a volume slider, a 2-level (focus/shoot) camera button or even a phone's qwerty keyboard. The inclusion of pressure sensing technology also enables the user to feel the key before they press down. Bullfinch-D uses two Redux small-form-factor transducers, positioned at either end of the touchscreen, to deliver this feedback.

More here, or; Redux Labs; www.reduxst.com

20V battery capacity monitor with I²C

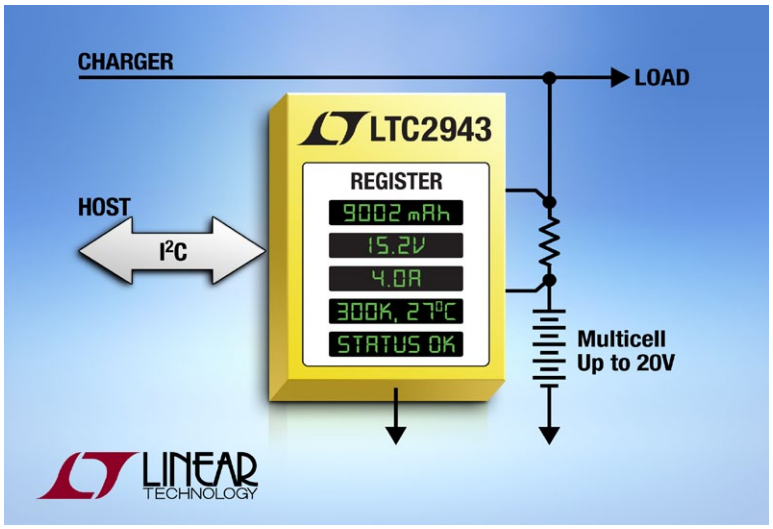

Linear Technology's LTC2943 is a multicell battery 'gas gauge' which makes direct measurements of 3.6V to 20V battery stacks. No level shifting circuitry on the supply and measurement pins is required to interface with multicell voltages, so total current consumption is minimised and measurement accuracy is preserved. The LTC2943 is a true high voltage gas gauge that measures charge, voltage, current and temperature to within 1% accuracy — all of the essential parameters required to accurately assess battery state of charge (SoC). Battery current is measured by monitoring the voltage across an external, high side sense resistor and integrating this information to infer charge. A bidirectional analogue integrator accommodates either current polarity (battery charge or discharge), and a programmable prescaler allows for a wide range of battery capacities. Charge, voltage, current and temperature information are communicated to the host system over an I²C/SMBus-compatible 2-wire interface that is also used to configure the gas gauge.

More here, or from; Linear Technology; www.linear.com/product/LTC2943

3A transistors in a 1.1-mm² package

NXP HAS INCREASED current/board area values with high Ptot MOSFET and bipolar transistors with benchmark values for low RDSon and VCEsat for power management and load switches in space-critical applications. Claimed as the first transistors in a 1.1 x 1 x 0.37 mm low-profile DFN (discrete flat no-leads) package, the new portfolio consists of 25 types including low RDSon MOSFETs, as well as low saturation and general purpose transistors that boost current capabilities up to 3.2 A.

More here or;
NXP; www.nxp.com

Cortex-M4 MCU with Ethernet MAC+PHY

TEXAS INSTRUMENTS has added to its Tiva C Series microcontrollers with the Tiva TM4C129x MCUs which are the first ARM Cortex-M4-based MCUs with Ethernet MAC+PHY, enabling customers to create a new class of sophisticated, highly connected products that connect to the cloud for Internet of Things (IoT) applications.

More here, or;
TI; www.ti.com

Precision linear current-sensor IC

ACS770 FROM ALLEGRO MICROSYSTEMS is a fully integrated Hall-effect based linear current sensor IC incorporating thermal enhancement for high precision. The device consists of a precision, low-offset linear Hall circuit with a copper conduction path located near the die. The total output accuracy of the ACS770 is achieved by using a new piecewise linear digital

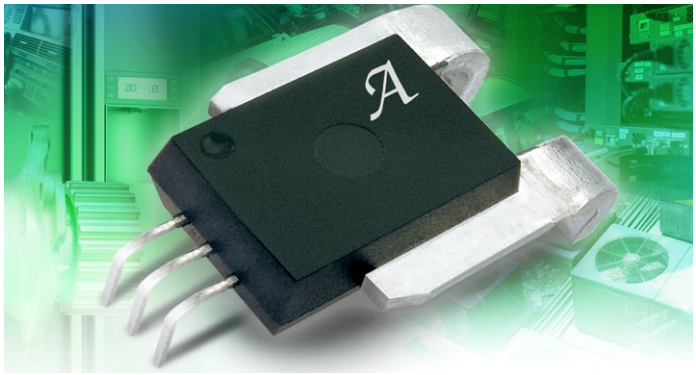

temperature compensation technique for offset and sensitivity. This greatly improves the IC's accuracy and temperature stability without influencing the high-bandwidth (125 kHz) operation of the analogue output.

More here, or;
Allegro Microsystems;
www.allegromicro.com

SMT inductors operate to 40 GHz

AVX' GLM AND GLN series passives provide low insertion loss and excellent matching for ultra-broadband DC decoupling networks and bias-tee applications. The GLM and GLN Series provide extremely low insertion loss and excellent matching over multiple octaves of the frequency spectrum. Developed to deliver repeatable and reliable ultra-broadband performance up through 40 GHz, these surface mount broadband inductors provide maximum inductance in a compact physical space, suiting them for applications including: semiconductor data communications, receiver optical sub-assemblies, transimpedance amplifiers, and test equipment. In addition to ultra-low insertion loss they offer excellent return loss, and extreme part-to-part performance and repeatability. Available in two sizes, M and N, the GLM and GLN Series are rated for -55°C to +125°C and offer flat frequency response and a rugged powdered iron core

AVX; www.avx.com

Software platform builder from Altium/Tasking for ARM Cortex-M

A SOFTWARE TOOLSET aims to provide a time-efficient way to implement middleware; Altium's Software Platform delivers rapid prototyping and development for ARM Cortex-M based microcontrollers, at a much-reduced cost level. The Software Platform includes a range of middleware functionalities, such as an RTOS, CAN, USB, TCP/IP, I²C, HTTP(S), file systems, graphical user interface, and touch panel control. The Software Builder is integrated into the TASKING VX-toolset, which includes a C/C++ compiler, debugger and an Eclipse based IDE. All Software Platform functionalities, including the software development tools, are made available at the cost of a traditional development toolset. The first release of the Software Builder provides support for the STMicroelectronics STM32 series. Pricing starts at $1,795 for the TASKING VX-toolset Standard Edition and $2,995 for the Premium Edition. Existing customers with a maintenance contract will get the update for free.

More here, or;
Altium; www.altium.com

Software-radio interface for FPGAs

LIME, AZIO AND E-ELEMENTS are co-operating to launch the Myriad RF connector board for multiple FPGA platforms: a low cost mezzanine board connect field programmable RF boards

with both Altera and Xilinx FPGA boards. The twin-standard mezzanine board supports both DE0-Nano and FMC FPGA interfaces, enabling it to link the majority of field programmable RF boards – including Lime's development platform, Myriad RF – with both Xilinx and Altera FPGA developer kits. The launch enables wireless networks of almost any size and complexity to be created on any mobile standard or frequency. The board also integrates USB functionality enabling you to connect it straight to a PC or Raspberry Pi, enabling the PC itself to act as the baseband. The ESO-001 is priced at US$199 (£123/€145), said to be a fraction of competing technologies. When combined with a low cost platform, such as Myriad RF ($299/£184/€219) it creates a low cost, professional grade RF development platform capable of developing networks of almost any complexity, which can be configured to work with or around the available whitespace in any given geographical market.

More here, or;
Lime Micro; www.limemicro.com

Development kit for 3D gesture systems

THE MGC3130 HILLSTAR DEVELOPMENT KIT for 3D gesturing systems provides designers with a step-by-step approach to develop 3D gesturing systems with Microchip's MGC3130 and electrodes that meet their specific space requirements in different form factors. Microchip's MGC3130 is an E-field-based 3D gesture controller; this reference system includes an MGC3130 Module, I²C-to-USB bridge module and reference electrode; other reference electrodes and the Aurea Graphical User Interface Software Development kit are also available. The MGC3130 Hillstar Development Kit enables you to integrate an advanced 3D hand position and gesturing user interface into virtually any electronic product using a modular approach with complete documentation, reference layouts, hardware and a software GUI. The Hillstar Development Kit (DM160218) is priced at $179.00.

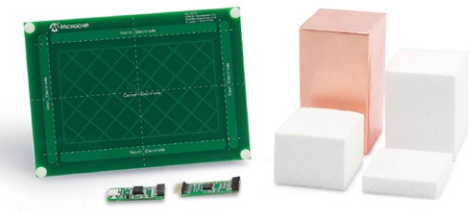

Microchip;
www.microchip.com/get/E0AA

IP core for mixed FPGA designs

REFLEX CES (PARIS), provider of Modified-Off-The-Shelf (MOTS) solutions for embedded and complex systems, has announced its FPGA Aurora-like 64B/66B IP core, a turnkey solution enabling designers of complex systems to interconnect Xilinx and Altera high speed transceiver FPGAs, giving

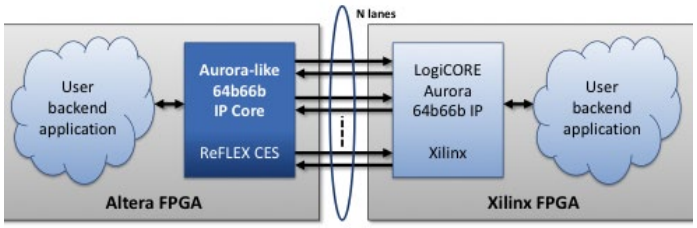

users freedom to choose the best FPGA solution. Based on Altera FPGAs, the new IP core supports 64B/66B encoding and high-speed interfaces up to 14.4 Gbps, enabling interoperability between Xilinx and Altera FGPAs, with an effective bandwidth of up to 97%. Targeted specifically at complex, embedded military and telecommunications and networking applications, the 64B/66B FPGA IP core allows designers to choose the most appropriate FPGA solutions to meet their needs, from single or multiple vendors. The IP core offers a fully compliant implementation of the Xilinx Aurora 64B/66bB scalable, link-layer protocol for high-speed serial communication, and allows for communication between FPGAs through a backplane. Based on this open standard protocol used to transport data with higher connectivity performance for chip-to-chip and board-to-board architecture, the Reflex CES Aurora-like 64B/66B IP core allows designers to move data from point-to-point across 1 to 16 serial lanes at up to 14.1 Gbps.

More here, or;
Reflex CES (Custom Embedded Systems);
www.reflexces.com/en/

Power modules drive up to 300W

INTERNATIONAL RECTIFIER'S highly integrated, ultra-compact, µIPM (micro-integrated-power-module) series now adds the IRSM808-105MH and IRSM807-105MH, for high efficiency appliance and light industrial applications with motor power up to 300W. In an compact 8 x 9 x 0.9mm PQFN package, the half-bridge power modules offer up to a 60% smaller footprint than existing 3-phase motor control power ICs, IR says, to provide a compact heatsink-free solution. Both the new modules are rated at 500V and 10A DC output (at 25C) (max), delivering 1.1A to a motor if not heatsinked, or 1.3A with a heatsink, for motor power of 285/390W. (Motor currents are RMS, with Fc=16 kHz, 2-phase PWM operation, ΔTCA=70°C, TA ≈ 25°C). They are priced at $1.44 (10,000).

More here, or;
International Rectifier; www.irf.com

Automotive boost/step-down DC-DC pairing

A NON-SYNCHRONOUS BOOST CONTROLLER with automatic wake up and shutdown functions supplies a minimum output voltage during start-stop conditions in order to counteract any sag in the vehicle's battery voltage. Running off a

2V to 44V input voltage, the NCV8876 is enabled when the supply voltage drops below 7.2V, then boost operation is initiated once this voltage goes under 6.8V, with the IC driving an external N-channel MOSFET. A quiescent current of 11 µA is drawn when the device is in sleep mode so that power consumption is minimised. Paired with this device is the NCV896530, a 2.1 MHz switching frequency, dual channel step-down DC-DC converter targeted at driver assistance systems. Both of its channels are externally adjustable (covering 0.9 to 3.3V) and can source currents up to 1600 mA. Synchronous rectification enables the device to offer enhanced system efficiency. Inclusion of features such as integrated soft-start, cycle-by-cycle current limit and thermal shutdown protection mean that this device is highly optimised for deployment in harsh automotive environments. The NCV8876 and NCV896530 both support a junction temperature range of–40 °C to 150 °C. The NCV8876comes in an SOIC-8 package and the NCV896530 in a DFN-10 package, at $0.90 and $1.20 (3,000)

More here, or;
ON Semiconductor; www.onsemi.com

Miniature low noise LDO

WITH A HIGH PSRR of 75 dB at 1 kHz and a low-noise performance of 60 µVrms over the 10 Hz to 100 kHz range, the AP7340 low-dropout regulator from Diodes Incorporated en-

sures a high-quality supply for analogue circuits in noise-sensitive audio, video and RF applications. With its 1.0 x 1.0 mm DFN1010-4 package, the AP7340 suits high-density portable product designs. This miniature, low-noise LDO produces a 150 mA output current for a wide input voltage range, starting from as low as 1.7V and extending up to a maximum of 5.25V. It comprises a voltage reference, error amplifier, current limit and an enable input for output on/off control. Its integrated resistor network produces fixed output voltages in 0.1V steps between 1.2V and 3.3V. The AP7340's output voltage accuracy is higher than other solutions, with an accuracy of ±1.5% over the entire operating temperature range of -40ºC to +85ºC. This low-power LDO offers a typical quiescent current of 35 µA. It costs $0.10 (5000).

Diodes; www.diodes.com

In-circuit authentication chip

ORIGINAL OR COUNTERFEIT? OPTIGA Trust Security Solutions from Infineon protects IP in small electronic devices and accessories. Infineon's OPTIGA Trust authentication chip helps manufacturers of electronic accessories and replacement parts

protect their businesses against damages caused by counterfeiting. OPTIGA Trust is a complete solution consisting of a chip and software that can be integrated into headphones, cartridges and electronic replacement parts to authenticate them when connected to the end-device. The end-device can recognise authorised products and reject counterfeits. Features include; strong cryptographic security due to asymmetric encryption (ECC Elliptic Curve Cryptography); space and cost savings because a single chip replaces a typical two-chip solution; minimal development and integration effort; a complete solution consisting of a chip and software; easy-to-integrate single-wire host interface; and a small chip size makes it suitable even for small devices such as in-ear headphones. The OPTIGA Trust SLS 10ERE authentication chips come in the USON-3-Package (2 x 3mm).

Infineon; www.infineon.com/optiga-trust

Medically-approved 200W PSUs

XP POWER'S CCB200 series of 200W single output AC-DC power supplies is aimed at a broad range of medical and industrial applications, that require convection cooling. Using an open-frame format, the series has a typical efficiency of 94% and a maximum of 95%, with a flat efficiency curve across the entire operating load range above 20% of full load. This can amount, XP says, to reduction of losses (waste heat) of up to 50% compared to typical 200W supplies. The high efficiency has enabled the convection-cooled units to be packaged in an industry standard 76.2 x 127.0 x 36.32 mm (3.0 x 5.0 x 1.43 in.) outline, usually used for forced-cooled power supplies of this rating. No load power consumption is less than 0.5W. The series conforms to the ANSI/AAMI ES60601-1 medical 3rd edition safety standard. The full operating range of the CCB200 series is 80 – 300 VAC, with full power available from 90 – 264 VAC. A total of five models is available providing the popular nominal outputs from + 12 to +48 VDC.

More here, or;
XP Power; www.xppower.com

Programmable Hall sensor measures currents to 200 kHz

A PROGRAMMABLE and high-speed linear Hall IC, with a sensitivity range from 5 to 150mV/mT, is suitable for measuring DC and AC currents up to 200kHz. The uprated MLX91209CA can be applied in all current sensing situations, Melexis says, as the fastest Hall current sensor available; its underlying response time is improved to under 2 µsec from the 6 µsec of its predecessor. The sensor has internal, switchable filtering but even with this applied, it is as fast as previous parts without filtering. It features thermal and lifetime stability as a foundation for fully programmable critical characteristics - allowing one single part to meet a wide variety of current sensing applications. It achieves very small offset drift of under 10 mV (-40-125°C), and very small sensitivity drift of typically under 100 ppm/°C (-40-125°C). Calibration is done using Melexis' PTC (Programming Through Connector) protocol. This modulates the supply voltage and does not require any additional pin for programming.

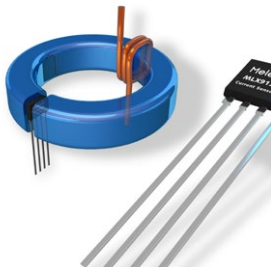

More here, or;
Melexis; www.melexis.com/91209

Flex-PCB with Altium Designer release 14

ALTIUM HAS RELEASED a new version of Altium Designer, focused on core PCB design technology, and extending features for 3D PCB design. Altium Designer 14 concentrates on core PCB design technologies, and now has support for flex and rigid-flex design, including schematic capture, 3D PCB layout, analysis and programmable design – all in a single, unified solution. Flex and Rigid Flex PCB design support in this release includes the ability to handle advanced layer stack management. Support for embedded PCB components allows greater design miniaturisation opportunities by placing standard components on an inner layer of the circuit board during fabrication.

More here, or;
Altium; www.altium.com/en/products/altium-designer/features

JESD204B FPGA software debug for designs with fast converters

FREE, ON-CHIP, 2-D statistical eyescan software enables fast in-system verification of high-speed data converter-to-FPGA signal integrity, in systems with ADI's data converters and Xilinx programmables; Analog Devices has issued an FPGA-based reference design with software and HDL code that helps verify high-speed systems incorporating JESD204B-compatible data converters. Called the JESD204B Xilinx Transceiver Debug Tool, it supports the 312.5-Mbps to 12.5-Gbps JESD204B data converter-to-FPGA serial data interface, Xilinx 7 series FPGAs and Zynq-7000 All Programmable SoCs. It is available at no cost with ADI converters and provides an on-chip, 2-D statistical eyescan that helps designers of radar arrays, software-defined radio and other high-speed systems more quickly verify the signal integrity of JESD204B data converter-to-FPGA designs using gigabit transceivers. The reference design is a download here and there is also a video outlining the tool's usage. ADI's reference design gathers data directly from the on-chip Rx margin analysis feature in the 7 series IBERT core and manages the data locally inside the FPGA or one of the ARM dual-core Cortex-A9 MPCore processors, displaying the data on an HDMI monitor or over Ethernet to a remote monitoring station. Typically, says ADI, other scanning tools measure signals off-chip and require costly test and measurement equipment or transfer the data back over JTAG to be viewed on a host/development PC in the lab.

"The Analog Devices JESD204B Xilinx Transceiver Debug Tool provides on-chip eye-scanning that augments the test and measurement process by statistically determining signal integrity inside the FPGA," says Revathi Narayanan, High Speed I/O product manager, Xilinx. "Where other techniques probe the outside of the FPGA package and acquire the signal before it's been processed by Xilinx's automatic gain control and equaliser blocks, ADI's approach yields a more accurate result by utilising the Xilinx transceiver on-chip eyescan feature to allow developers to monitor the signal integrity and design margin on their JESD204B links inside the FPGA."

More from Xilinx; www.xilinx.com/jesd204
Analog Devices; www.analog.com

Ultra-low power FPGAs for always-on sensors

WITH ITS ULTRA-LOW-DENSITY iCE40 FPGAs, Lattice Semiconductor claims to be delivering the world's most flexible, single-chip sensor solutions for making a new generation of context-aware, ultra-low power mobile devices possible. Delivered in a 16-ball 1.48 x 1.40mm wafer level chip scale package (WLCSP), the iCE40LP integrates an RGB LED driver and is flexibly programmable to implement advanced functions such as IrDA, barcode emulation (requiring a dedicated LED driver), service LED, and more in a single chip with available logic for additional customer defined functions. With hard IP for strobe generators, I^2C and SPI interfaces, the iCE40LM FPGAs deliver near-zero latency to the mobile market, enabling context-aware systems with the real-time capturing of user and environmental inputs with minimal delay or error, all from a WLCSP measuring 1.71 x 1.71 x 0.45mm (0.35 mm pitch). Because the iCE40LM draws just under 1 mW in active mode, it can be used to constantly monitor sensors, buffering data for up to a dozen of seconds before it wakes up the application processor for context-specific responses.

More here, or; Lattice Semiconductor; www.latticesemi.com/iCE40

650V silicon carbide Schottkys

MICROSEMI HAS ADDED to its silicon carbide (SiC) Schottky product family with 650V diodes targeted at high-power industrial applications including solar inverters, offering zero reverse recovery, temperature independent behaviour, higher voltage capability and higher temperature operation. This 650V SiC Schottky diode product portfolio includes:

- APT10SCD65K (650V, 10A, TO-220 package)
- APT10SCD65KCT (650V, 10A, common cathode TO-220 package)
- APT20SCD65K (650V, 20A, TO-220 package)
- APT30SCD65B (650V, 30A, TO-247 package)

Microsemi; www.microsemi.com

PIF – an FPGA daughter card for the Raspberry Pi

THE PIF FPGA BOARD for the Raspberry Pi (RPi) adds an instant-on, non-volatile, high performance FPGA from Lattice Semiconductor's MachXO2 range to the Pi; use it as a target for experimenting with FPGA design, for developing design options, and for implementing custom hardware functionality. It can support peripheral interfacing, adding I²C and SPI ports, controlling specialised ICs, buffering and pre-processing the data from sensors, and precisely controlling output signals. A modern FPGA with wide internal memories and lots of logic capability is just what you need for experimenting with cryptographic routines, developing your own CPU architecture, or designing novel signal processing algorithms, according to its designers.

More here, or;
Bugblat; www.bugblat.com/products/pif

Gesture recognition on ARM-based devices

ISRAELI COMPANY EYESIGHT, producer of intuitive gesture control technology, has completed extensive work in partnership with ARM to optimise its gesture recognition solution for use on ARM's Mali T600 Graphics Processing Units. Manufacturers using ARM Mali GPUs can use eyeSight's natural-feeling, advanced gesture control capabilities, using GPU Compute for improved robustness, accuracy and energy efficiency. The improved efficiency of gesture computation through use of the GPU will also improve a variety of new use cases, such as face and emotion detection, long distance finger tracking, and even 3D motion recognition (such as finger pointing for selection).

More here, or;
ARM; www.arm.com
eyeSight; www.eyeSight-tech.com

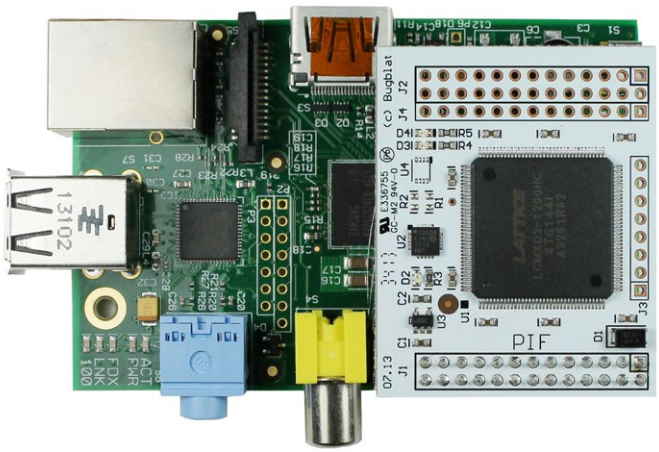

Point taken

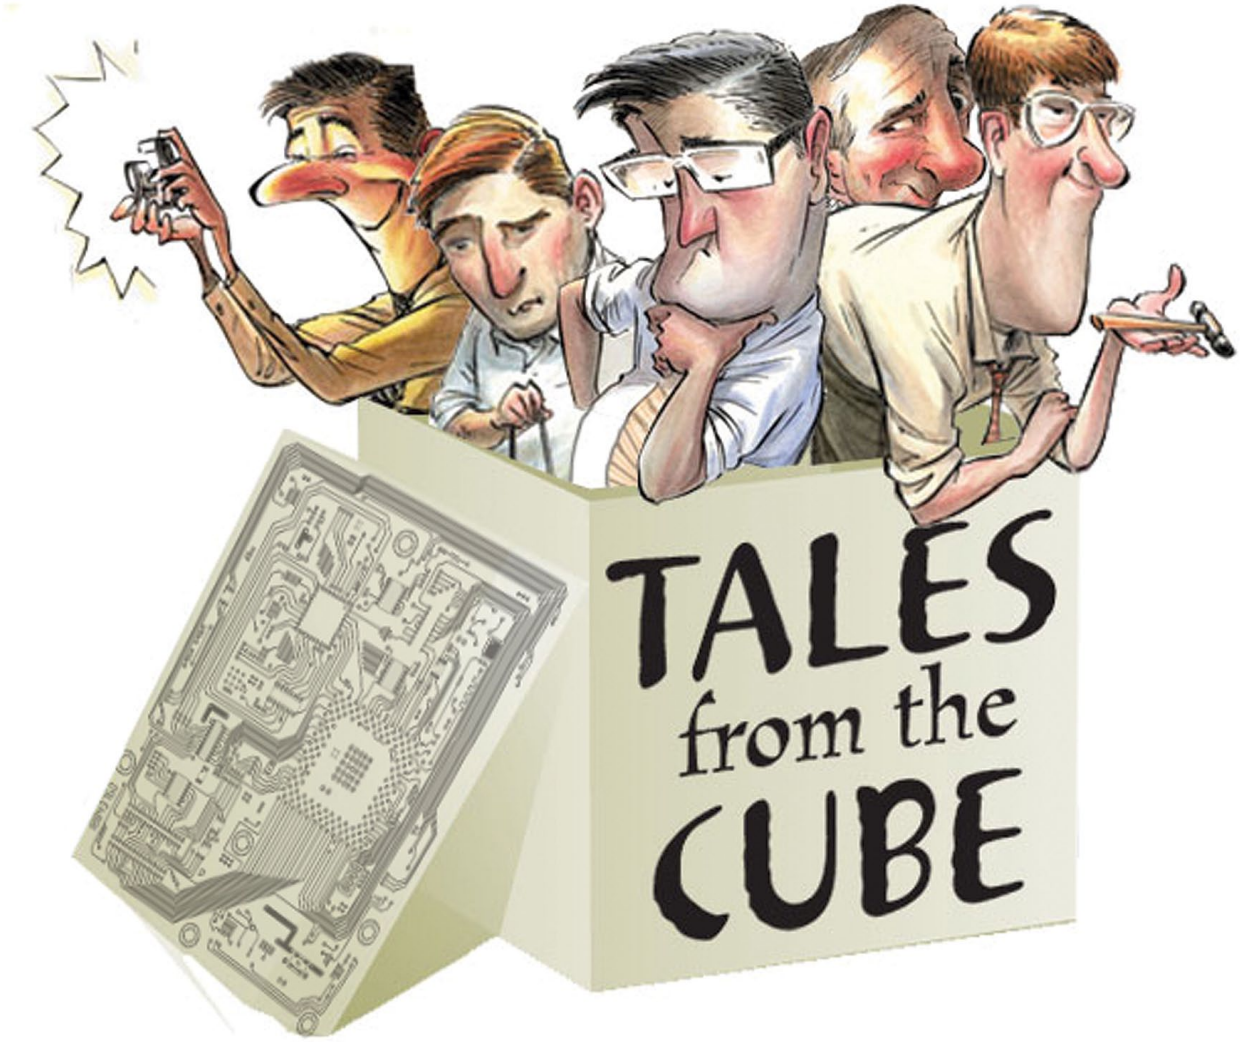

We designed an embedded processor board that ran under Windows 7. When the hardware was deployed to the software developers, all was running fine until they had the platform software "sit" in a bootloader menu, which consisted of a list of options on the screen and a command prompt – sort of like a DOS shell but not running Windows.

After a while, say 20 seconds to a few minutes, it was reported that the unit would shut down and the power manager chip (PMIC) would enable a red error light that indicated a power rail fault. The strange thing was the unit did everything it was supposed to do in Windows, which was the actual customer application. At first, we thought that in the menu mode, the watchdog timer was timing out or there was some other software bug, but we couldn't figure out why the time to fail was so variable. At one point we thought it would happen when the mouse was installed in the USB port.

The software engineer discovered that when in the bootloader menu, all the processor clocks were enabled, causing more power consumption than when in Windows mode. Thinking now that the problem may be related to heat in some way, I took a working board and put it into a temperature chamber while monitoring 5V main power. As the board temperature increased, the voltage started decreasing. Hmmm, this is not supposed to happen.

Now I went back to the schematics and reviewed the power supply, which was a very basic 24 VDC input, 5V output switching supply. I looked at the main inductor and it was 4.7 µH. Going back to the power supply controller data sheet, the inductor should have been 47 µH, so the board was operating with an inductor one order of magnitude lower than it should have been. As the temperature and the load increased, the inductor began to saturate, causing the 'B' field to drop off and the voltage to drop, which, in turn, caused the PMIC to indicate a voltage rail problem.

I think the moral to the story is to watch for inductor saturation and always environmentally test your designs. Also, I have found that most problems have a simple solution rather than an exotic one like sunspots, gamma particles, or idle computer mice.

Michael Gambuzza is a senior design engineer for GE Sensing in Billerica, MA. He has over 33 years experience as a mixed-signal designer and has made his share of technical mistakes.

Keep abreast of the latest industry news with our newsletters

european business press

EDN europe

MICROWAVE engineering europe

EE Times europe LEDLighting

EE Times europe AUTOMOTIVE

EE Times europe ANALOG

EE Times europe TEST & MEASUREMENT

EE Times europe POWER MANAGEMENT

EE Times europe

EE Times europe Technical papers

www.electronics-eetimes.com/newsletters

EET Search

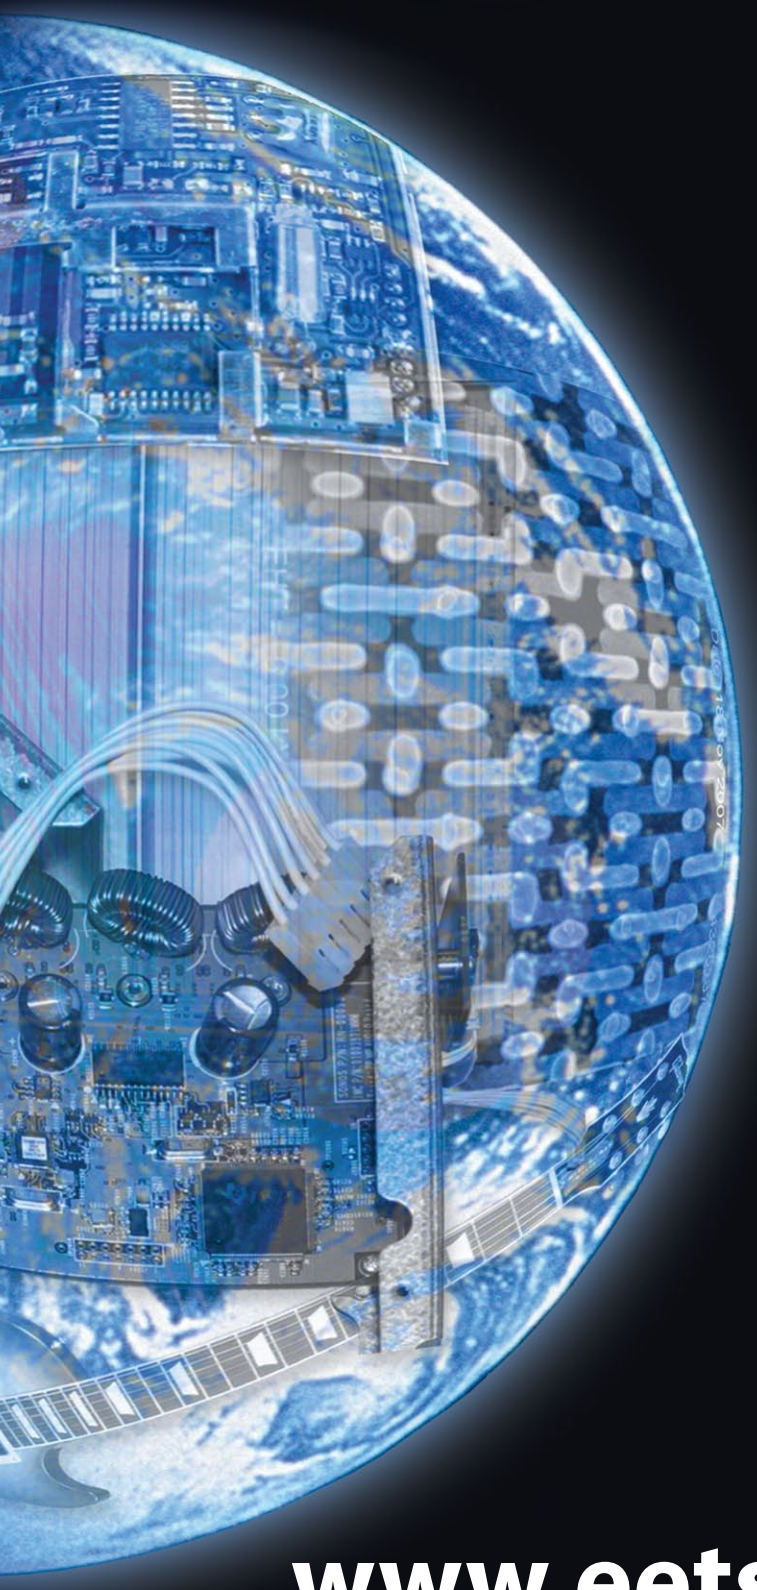

- searches all electronics sites

- displays only electronics results

- is available on your mobile

www.eetsearch.com