EDN europe

MAY 2016

www.edn-europe.com

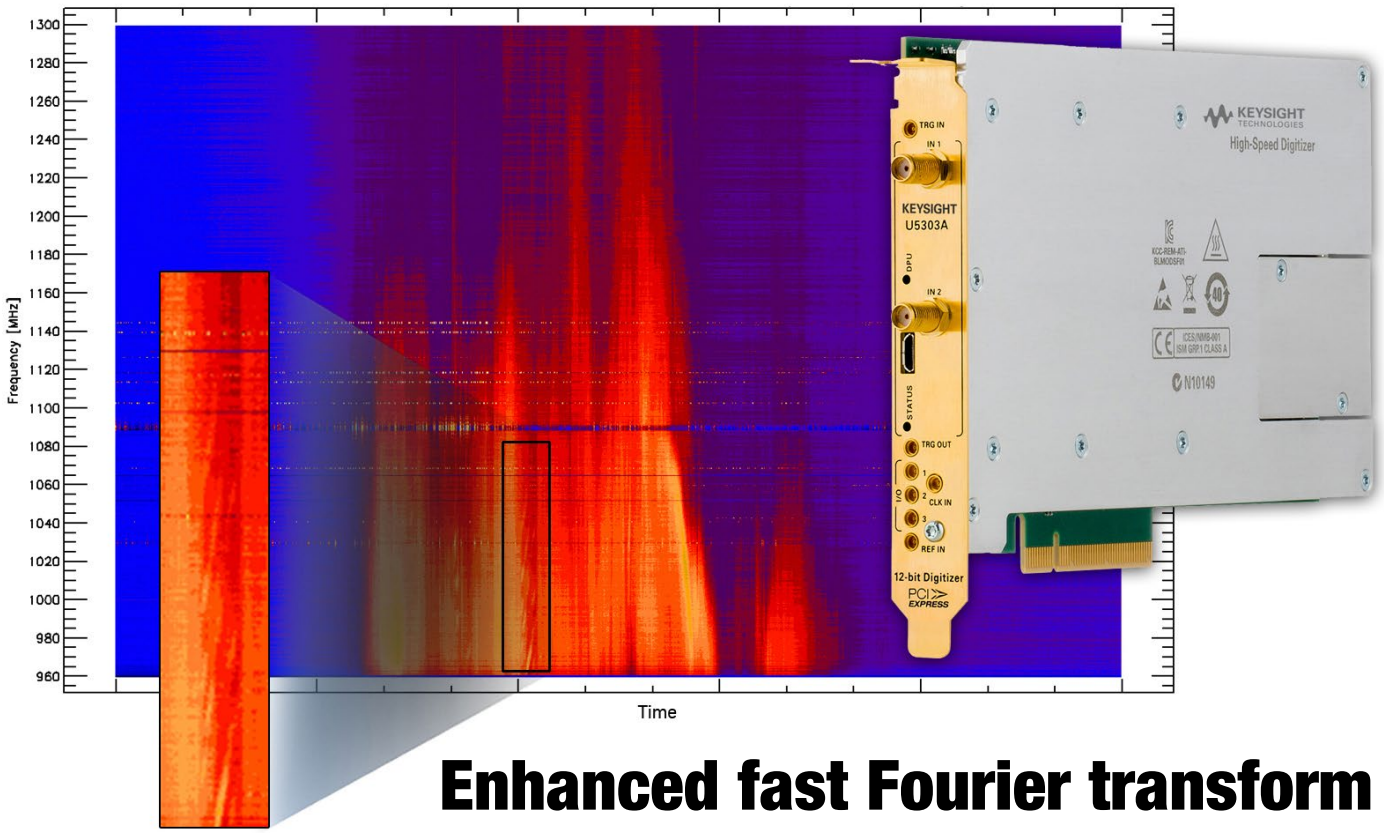

Enhanced fast Fourier transform application aids radio astronomy

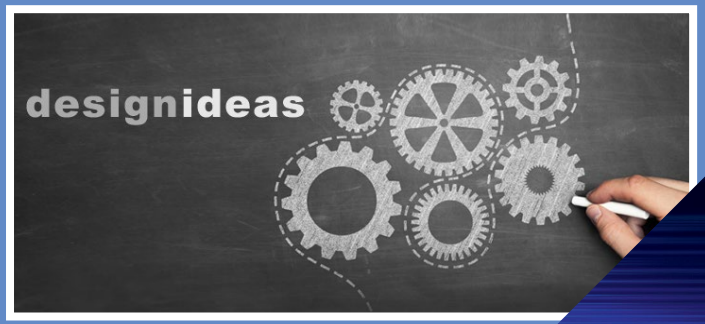

Enabling Your Innovation™

From Concept to Production™

FREE SHIPPING
ON ORDERS OVER €65!*

4.8 MILLION PARTS ONLINE

OVER 650 INDUSTRY-LEADING SUPPLIERS

100% AUTHORIZED DISTRIBUTOR

Digi-Key® ELECTRONICS

FIND CONTACT AND ORDERING INFORMATION FOR YOUR REGION AT

DIGIKEY.COM/EUROPE

*A shipping charge of €18.00 (£12.00) will be billed on all orders of less than €65.00 (£50.00). All orders are shipped via UPS for delivery within 1-3 days (dependent on final destination). No handling fees. All prices are in euro and British pound sterling. If excessive weight or unique circumstances require deviation from this charge, customers will be contacted prior to shipping order. Digi-Key is an authorized distributor for all supplier partners. New product added daily.

ecia MEMBER ecsn member CEDA

COVER

Enhanced fast Fourier transform application aids radio astronomy

The image on the cover of this issue shows a radio astronomy signal capture; specifically, a “Solar burst with ‘tilted fibres’” observed with a 12-bit 12-bit PCIe digitizer plus FFT option, by Keysight Technologies, taken at the Radio Astronomy Observatory of ETH Zürich, Switzerland.

Keysight recently introduced an improved version of its FFT option for the U5303A PCIe 12-bit high-speed data acquisition card. This card features two channels and a sampling rate from 1 Gsample/sec to 3.2 Gsample/sec, yielding bandwidth of DC up to 1.4 GHz, and resolution of 97 kHz per spectrum line. It features, the company adds, greatly improved fixed point arithmetic and is aimed at applications in astronomy, physics and environmental measurements. While providing continuous real-time FFT at full sampling rate, this option also enables spectra accumulation capability, a key requirement for radio-astronomy.

With DC up to 2 GHz bandwidth nominal, the card converts even low frequencies of around 0 Hz not observable with an AC front-end; moreover, AC coupling would introduce BF noise, whereas DC coupling provides measurement fidelity. By increasing the resolution from 8- to 12-bit, this option provides computation of a 32k points FFT for single or dual channels. Low-noise capabilities allow users to detect phenomena difficult to observe, such as tilted fibres (a phenomenon associated with flux lines in the sun’s corona) in a solar burst. Longer item/links.

FEATUREARTICLES

ONLINE THIS MONTH

EDN's columns

OUTSIDE-THE-BOX THINKING

This column is being written as the 2016 PCIM power electronics conference approaches. This year's event is the first since the conclusion, earlier in the year, of the Google/IEEE "Little Box Challenge", that was conceived to stimulate innovation in inverter design. With a $1million prize, the Challenge was to build a significantly more compact inverter than current standard-practice: Google's contention was that proliferation of distributed renewable energy installations needs a step-change in size and, to a lesser extent, efficiency. From, "the size of a cool-box to the size of a small laptop" as the originators put it. The Challenge website is still up, if you didn't see it at the time; among other resources it archives the technical submissions of the leading contenders.

The Challenge – and the $1million – was won by power specialist CE+T Power (Wandre, Belgium) with the award being made at the ARPA-E Energy Innovation Summit in Washington D.C. on 29th February 2016, after its winning design was chosen over 18 other finalists and some 2000 initial submissions from teams around the world.

I spoke at the time to CE+T's Olivier Bomboir: part of my interest was to gain some insight into whether the design exercise needed some exotic departure into fundamentally new architectures of power conversion. Recall that the competition's original specification was pitched at 50W/in³: CE+T Power's design went far beyond that, achieving 145 W/in³. The 2 kVA design is only 13.77 in³ (225 cm³) – smaller than a block of post-it notes, which is far smaller than the brief requested – and it uses technology already commercially available.

Therefore; no, there is no exotic, "skunk-works" technology here – but every aspect of the design, electrically and thermally, was refined to get to the final result.

CE+T has experience in inverter design, and in particular, modular inverters. From their initial study it was clear, Bomboir says, that the Challenge could be done, but [even] at 50 W/in³, not with silicon IGBTs or MOSFETs; it would need high-frequency switching and 95% efficiency or better. Simple arithmetic tells us that 95% efficiency at 2 kW still leaves 100W to dissipate, so given the target size, advanced thermal design was clearly going to be needed.

CET+T used switches from GaN Systems, exploiting that company's planar packaging style in its thermal design and to achieve very low-inductance paths. The topology used is what the team terms a five-leg inverter; five half-bridge switch pairs, two of which generate the outgoing neutral, and two the live, lines. The fifth half-bridge acts as an active filter; the control and generation of the (soft) switching waveforms is fully digital, and the final half-bridge is used to feed compensating waveforms on to the outputs to reduce the need for passive filtering. The software/firmware uses an adaptive (learning) algorithm for filtering.

The thermal design brings together a remarkable array of materials and techniques; honeycomb pattern copper, thin "sandwich" connections to power devices, silicone foam, ultra-thin PCBs, and more. Once again, however, nothing "sci-fi" - but existing technology, developed and refined for the application. All electrical and thermal aspects were modelled and simulated, with the result that the design was essentially first-time-right.

Olivier Bomboir identifies a number key design aspects, and lessons learned; the use of GaN switches is essential to this exercise but – as we have heard from other sources – they must be applied with care. They switch so fast, Bomboir confirms, that transient potentials can easily exceed the ratings of most available driver devices; CE+T developed its own driver configurations for the design, which it intends to patent. Establishing that drive capability, at the frequencies required, was a major focus. The other key aspect was the software; the fully-variable frequency drive algorithms involve a high level of DSP computation, and a change to an uprated processor was needed along the way. You can read a more detailed description of the architecture of the winning inverter here, and our initial report of CE+T's success is here.

Which leaves me with this thought; what other functions or designs could – with similarly rigorous and meticulous use of today's design techniques, and already-available technology – be driven forward by the same step-change that CE+T achieved in this case? And, do we necessarily need a body like Google to provide the spur to make the attempt?

pulse

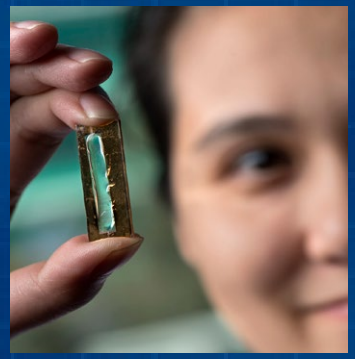


embOS
OS

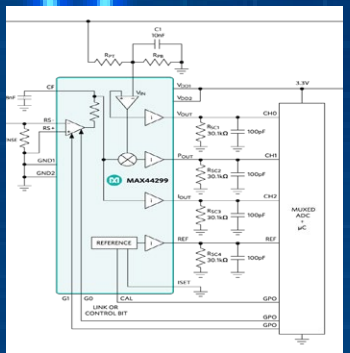
MAX44299
REFERENCE
MUXED ADC I²C

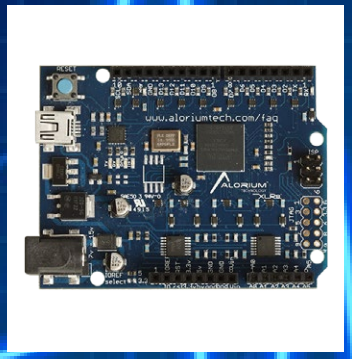
www.aloriumtech.com/faq
ALORIUM

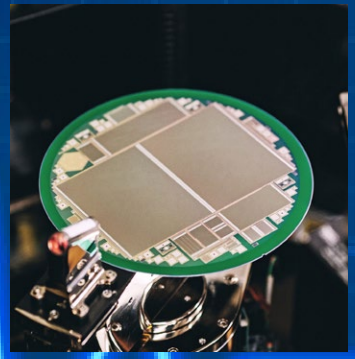


Integrated 600-V GaN FET power stage
• Easy-to-design
• Twice the power density, half the power loss
• Enables new topologies
GaN FET Power Stage
LMG3410
Texas Instruments

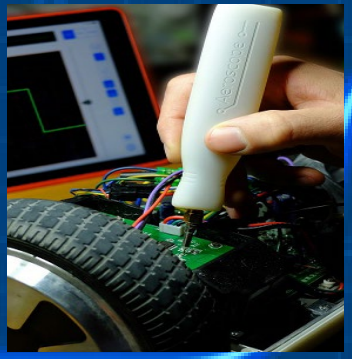

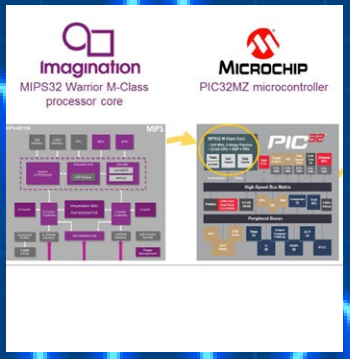
Imagination
MIPS32 Warrior M-Class processor core
MICROCHIP
PIC32MZ microcontroller
MIPS
PIC32

Pico expands USB scope approach to multifunction T&M, in handheld format

Pico technology has become well-known for its display-less, USB-connected and -powered, compact modular oscilloscopes. The PicoScope 2000 Series oscilloscopes – intended for general purpose fault-finding use up to 100 MHz bandwidth – develop that concept by packaging the functions of six instruments in a portable package.

PicoScope two-channel, four-channel and mixed-signal models have the functionality of an oscilloscope plus a logic analyzer (on MSO models), spectrum analyzer, function generator, arbitrary waveform generator, and serial bus analyzer with support for 15 protocols included as standard. They are USB-powered and come in an portable package that can be easily transported in a laptop bag.

PicoScope 2000A 2- and 4-channel models are aimed at technicians, trainers, students, and hobbyists doing fault-finding on signals up to 25 MHz. The 2205A MSO (mixed-signal oscilloscope) has two analogue plus 16 digital channels for viewing and making measurements on digital or mixed technology designs.

PicoScope 2000B Series models are equipped with deep buffer memory from 32 to 128 M samples, and bandwidths of 50, 70 or 100 MHz supported with sampling speeds to 1 Gsample/sec, and hardware acceleration, to deliver over 80,000 waveforms per second update rates. Like the “A” models they are available in 2- and 4-channel models, plus 2+16 channel MSO versions. The “B” series models are packaged in the same size portable enclosure as the “A” Series, suiting them for use in the lab or on the move, or for fast low-cost shipment to wherever they are needed.

Complete article, here

MISRA clarifies safe and secure uses of the C language

MISRA is releasing new documents to clarify use of the MISRA C Guidelines in developing any application with high integrity or high reliability requirements – both safety related and security-related. The release took place at a MISRA workshop at the Device Developer Conference in Cambridge UK on April 27th, 2016.

The MISRA C Guidelines (The latest version is MISRA C:2012 Guidelines for the Use of the C Language in Critical Systems), the organisation notes, are internationally accepted as setting out a subset of C for use in critical systems. Generally this is understood to mean for use in safe systems. However the guidelines are equally appropriate for secure systems, a topic of increasing concern with the growth of the Internet of Things. ISO/IEC JTC1/SC22/WG14 (the committee responsible for maintaining the C Standard) has published its C Language Security Guidelines (ISO/IEC 17961:2013). MISRA has carried out a coverage comparison between this and MISRA C:2012 and is publishing the resulting coverage matrix as MISRA C:2012 Addendum 2. Alongside MISRA C:2012 Amendment 1 “Additional security guidelines for MISRA C:2012”, which includes a small number of additional guidelines, to improve the coverage of the security concerns highlighted by the ISO C Secure Guidelines, particu-

Simple, Easy Solutions

Design 101: Use MPS

MPQ4423A

Excellent Efficiency, 3A, 36V Synchronous Step-Down Converter

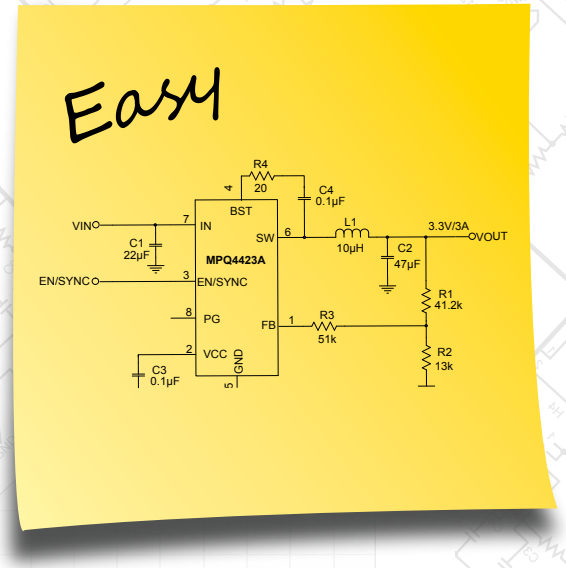

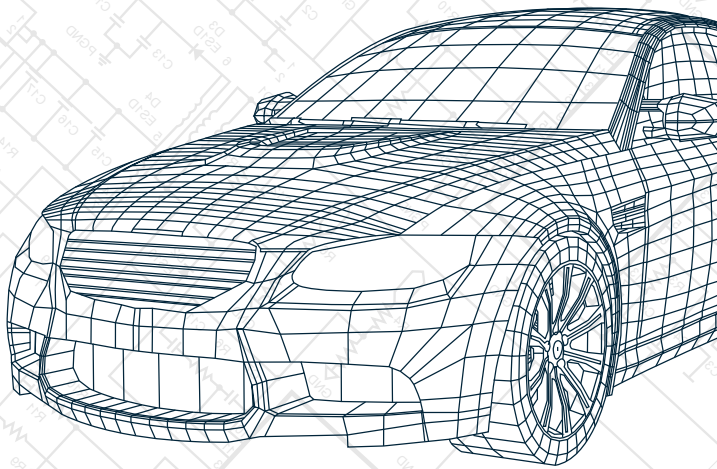

Key Highlights

- AEC-Q100 Qualified
- Wide 4V to 36V Continuous Operating Input Range
- 95% Efficiency Synchronous Mode Operation
- Good EMI, Compliant with CISPR25 Class 5
- Low Dropout Voltage for Automotive Cold Crank

www.monolithicpower.com

Follow us on:

larly in the use of "untrustworthy data", MISRA C is demonstrably suitable for both safe and secure applications.

"Anyone using the C language for system development, particularly for systems that have to safe and/or secure should be using the MISRA C Guidelines," said Andrew Banks Chairman of the MISRA C Committee. "The coverage matrix, plus the new rules, provides reassurance that code will be of high quality."

Complete article, here

MIT researchers refine single WiFi access point to deliver tens-of-cm location

A system designed by MIT's Computer Science and Artificial Intelligence Lab enables a single WiFi access point that can locate users within tens of centimetres; the team offers the prospect of safer drones, smarter homes and password-free WiFi. In a new paper, a research team led by Professor Dina Katabi present a system called Chronos that enables a single WiFi access point to locate users to within tens of centimetres, without any external sensors.

Chronos locates users by calculating the "time-of-flight" that it takes for data to travel from the user to an access point. The system is 20 times more accurate than existing systems, computing time-of-flight with an average error of 0.47 nsec. MIT's Deepak Vasisht presented the paper at a recent USENIX Symposium on Networked Systems Design and Implementation (NSDI '16).

Existing localisation methods have required four or five WiFi access points. This is because today's WiFi devices don't have wide enough bandwidth to measure time-of-flight, and so researchers have only been able to determine someone's position by triangulating multiple angles relative to the person. What Chronos adds is the ability to calculate not just the angle, but the actual distance from a user to an access point, as determined from time-of-flight.

The team demonstrated the system in an apartment and a cafe, while also showing off a drone that maintains a safe distance from its user with a margin of error of about 4 centimetres. "From developing drones that are safer for people to be around, to tracking where family members are in your house, Chronos could open up new avenues for using WiFi in robotics, home automation and more," says PhD student Deepak Vasisht, who is first author on the paper alongside Katabi and former PhD student Swarun Kumar, who is now an assistant professor at Carnegie Mellon University. "Designing a system that enables one WiFi node to locate another is an important step for wireless technology."

A video demonstration can be viewed here.

Complete article, here

OS extension brings Java platform into ARM Cortex-M design sphere

Segger's real time operating system embOS is now supporting MicroEJ's platform, thus opening the world of ARM Cortex-M based embedded applications to Java developers.
The package provides a complete Java platform, including a Virtual Machine and other components. The virtual machine is a 32-bit processor that manages the Java threads. It is executed as a task controlled by Segger's embOS kernel, thereby combining all advantages of both ANSI-C and Java languages on a single embedded target. Developers can focus on their Java applications and do not need to have any deeper knowledge of ANSI-C.
"We are very happy to be able to leverage the usability of embOS to Java users, and thereby giving them the opportunity to produce embedded applications in an economically viable way," says Dirk Akemann, Partnership Marketing Manager of Segger. "A further group of developers can now profit from our high-performance RTOS with zero interrupt latency which has been optimised for the limited resources of microcontrollers."
"MicroEJ aims at offering developers a unified multi-language programming platform for microcontroller-based systems on a rich variety of hardware and software foundations," said Fred Rivard, CEO of MicroEJ. "We're happy to offer our MicroEJ platform to embOS developers, combining the predictive behaviour and precision of Segger's RTOS with the ease of programming and portability of our platform."

Complete article, here

ST's free software packages detect activity from sensors in wearables

STMicroelectronics has introduced three additions to its Open.MEMS portfolio of free software libraries for the development of motion-sensing applications. The software allows the detection of human activities from data acquired by inertial sensors embedded in the end-user equipment. Optimised to minimise power consumption, they are particularly suited for fitness and healthcare applications in portable or wearable platforms that monitor human physical activities in real time over long periods.
The three new software packages are:
- The osxMotionAR Activity Recognition package is a high-performance algorithm that identifies the user activity from a wide range of movements and transportation scenarios such as stationary, walking, fast walking, jogging, cycling, and driving. Exploiting the high precision of ST's LSM6DS3, LSM6DS3H, and LSM6DSL inertial modules, the Activity Recognition algorithm manages the data acquired from the sensors at a low sampling frequency and returns the identified activity in real time with a very low power consumption.

- The osxMotionCP Carry Position package detects how the device containing the motion sensors is being carried. For example, the algorithm can detect whether a portable device such as a mobile phone is placed on a desk, held in hand to view the display or in

Looking to simplify design and reduce microcontroller overhead?

DACs with Integrated EEPROM Save Settings at Power-down

With integrated EEPROM, Microchip's MCP48FEBXX Digital-To-Analog Converters (DACs) retain settings at power-down to eliminate the need for external memory and reduce microcontroller overhead.

Available in single and dual channel versions with Serial Peripheral Interface (SPI), the family offers various shut-down modes to significantly reduce power consumption and is ideal for battery powered applications. In addition, these devices also feature low Differential Nonlinearity (DNL) error to sustain monotonic output and low Integral Nonlinearity (INL) error for better linearity.

The same feature set, without the integrated EEPROM, is offered by the MCP48FVBXX family as lower-cost alternatives for applications that don't require memory.

microchip DIRECT
www.microchipdirect.com

MICROCHIP

www.microchip.com/eumcp48fxbxx

a swinging arm, near the user's head, or put in a shirt or trouser pocket. To minimise power consumption, sensor data is acquired at a low sampling frequency (50 Hz).
- The osxMotionGR Gesture Recognition package recognises the actions carried out on a mobile or handheld device, including pick-up, glance, or wake-up, which allows designers to develop controls for different functions on the device. This algorithm acquires data from inertial modules with a sampling frequency of 100 Hz and recognises the gestures carried out by the user platform in real time.
Based on the STM32Cube software development tool, Open. MEMS libraries are part of the X-CUBE-MEMS1 expansion software package designed to run on the X-NUCLEO-IKS01A1 motion MEMS and environmental sensor expansion board.

Complete article, here

Low-side I/V sensing IC, with internal power calculation

Maxim Integrated has posted details of the MAX44299, a low-side current, voltage, and power monitoring circuit that provides an analogue output current proportional to the measured current, voltage, and the internally calculated instantaneous power. The power monitor offers high precision and integration in tiny size. Instantaneous power is calculated internally by multiplying the load current and a fraction of the load voltage set by an external resistive divider. All three outputs are scaled to a full-scale current of 100 µA. An additional output current of 100 µA is available at the reference (REF) output; this current can be used to create a reference voltage for the ADC that is being used to measure the power, voltage, and current signals. By providing the ADC with both the measured signals and the input reference voltage, the ADC can make a ratiometric measurement, allowing improved accuracy. The use of currents, rather than voltage, to convey the measured signals to the ADC eliminates any errors caused by voltage drops across the parasitic resistance of PCB, which can be significant for high-current systems. To allow full-system calibration, the CAL bump provides a way to calibrate gain and offset for the ADC.

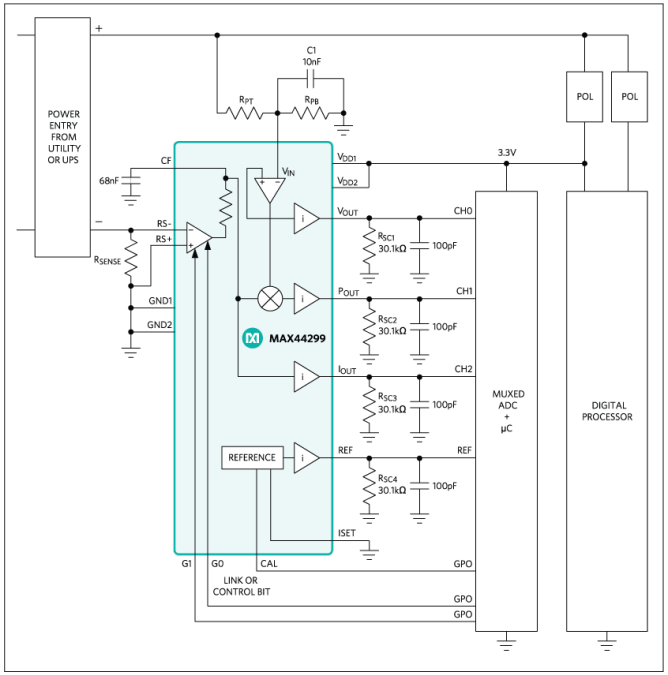

The device measures load current by using a precision, auto-zeroed current-sense amplifier (CSA), which due to its ultra-low offset voltage allows precise measurement of full-scale voltages of 5 mV, 10 mV, and 20 mV. The load voltage is measured via a user-selectable resistive network, dividing the input voltage down to a full scale of 1.00V.
The wide supply voltage range of 3V to 5.5V allows the simple sharing of supplies with either the ADC or a microcontroller. The device can be powered down and the outputs will then go high impedance. The device is available in a 2.4 x 2.4 mm, 16-bump wafer-level package (WLP) and is specified for the 0°C to +85°C temperature range.

Complete article, here

60 GHz chipset for WiGig outdoor wireless infrastructure

Canadian fabless chip maker Peraso recently announced production of its X610 Wireless Infrastructure Chipset; now the company has added that it is able to support higher production volumes due to the introduction of a fast and comprehensive production testing regime at 60 GHz.
Peraso Technologies' (Toronto) 60 GHz WiGig ICs are, the company believes, the first WiGig solution in production for the infrastructure market, and is suited for end-use applications such as wireless broadband for internet service providers (WISP), small cell backhaul, video surveillance, private networks and new hybrid fibre-wireless.
The X610 chipset includes the PRS1126 direct conversion 60 GHz RFIC and the PRS4601 baseband IC to implement complete RF to baseband functionality compliant with the IEEE 802.11ad (WiGig) standard. The chipset provides multi-gigabit per second throughput and operates across the industrial -40°C to 85°C temperature range. It incorporates a high-speed USB 3.0 system interface to integrate with common low-cost network processor ICs.
The X610 chipset uses 60 GHz V-band spectrum to enable rapid, licence-free deployment in most major geographic markets. The 60 GHz band offers a large contiguous operating band to deliver multi gigabit per second throughput, and is particularly well suited to avoid the increasing interference problems suffered by traditional 2.4 GHz and 5 GHz WiFi bands.
Now, Presto Engineering and Peraso Technologies have announced a collaboration in developing a test solution for Peraso's 60 GHz semiconductor devices., that enables testing at production volumes for infrastructure applications and provides foundation for high-volume testing for consumer applications.
See also; IDT bets on Peraso's 60GHz chips

Complete article, here

Californian researchers create battery mechanism with 'indefinite' cycle life

A team at the University of California, Irvine report – almost inadvertently - inventing nanowire-based battery material that can be recharged hundreds of thousands of times, moving us closer to a battery that would never require replacement.
The basic structures being investigated are charge-storage mechanisms based around nanowires; the 'wires' in question are deposited strutures (wires only in the sense they are long, thin conductors). In the work reported, the wires are gold, and they are coated or sheathed in manganese dioxide (MnO_2). The device is essentially a capacitor, with the charge-storing capacitance being formed between interleaved arrays of the wires, all suspended in an electrolyte medium.

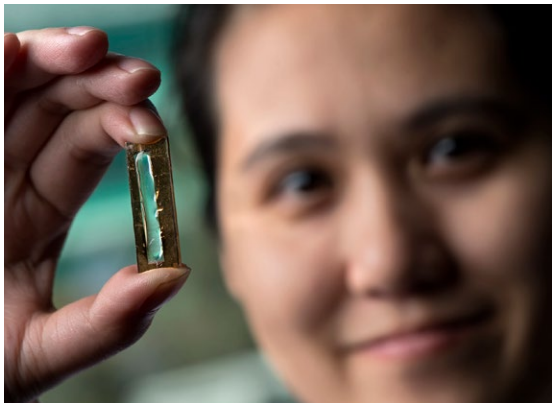

The use of nanowires is known to offer higher charge storage than films of identical materials. Thousands of times thinner than a human hair, they're highly conductive and feature a large surface area for the storage and transfer of electrons. The gold wires (laid down on glass) are 240 nm across and 35 nm thick; various shell thicknesses between 143 and 300 nm were evaluated. Prior work has found that these

filaments are extremely fragile and don't hold up well to repeated discharging and recharging, or cycling. In a typical lithium-ion battery, they expand and grow brittle, which leads to cracking.
UCI researchers have solved this problem by coating a gold nanowire in a manganese dioxide shell and encasing the assembly in an electrolyte made of a Plexiglas-like gel: a gell of poly-methyl-methacrylate and LiClO4 (lithium perchlorate, source of the lithium ions to form a lithium ion battery). The combination is reliable and resistant to failure.
The study leader, UCI doctoral candidate Mya Le Thai, cycled the testing electrode up to 200,000 times over three months without detecting any loss of capacity or power and without fracturing any nanowires.

Complete article, here

Alorium Technology's FPGA dev board, Arduino compatible, in distribution

Mouser Electronics has announced a global distribution agreement with Alorium Technology (Eau Claire, Wisconsin, USA), producers of the XLR8 development board which integrates an ATmega328-compatible microcontroller core and custom accelerator blocks on a field-programmable gate array (FPGA).
The board, Mouser says, allows developers to capitalize on the ease-of-use of the Arduino ecosystem with the hardware acceleration of an FPGA. The XLR8 is an Arduino-compatible board that hosts an Altera MAX 10 FPGA as its main processing chip. The FPGA implements an Atmega328-compatible MCU core and leaves room for developers to implement custom logic as Xcelerator Blocks (XBs) that can be integrated into the reconfigurable FPGA and interfaced directly to the processor. The XLR8's FPGA-based hardware acceleration and offload results in significantly improved performance in the same physical footprint, using the same tool chain as standard Arduino Uno boards. Sketches designed for Arduino-compatible boards will also run on XLR8, so developers can load code directly to XLR8 via the Arduino integrated development environment (IDE) and use Arduino-compatible shields.

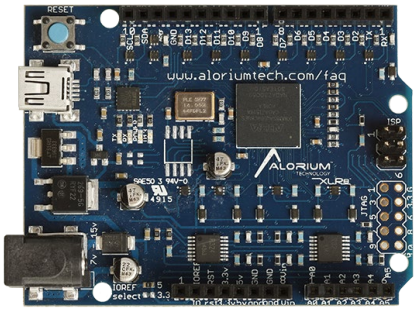

Complete article, here

Infineon builds 8-in. particle detectors, one per wafer, for CERN

At CERN, a unique sensor chip may contribute to proving the existence (or otherwise) of dark matter: the detector is eight inches or 15 cm x 10 cm and was developed jointly by Infineon Technologies Austria and the Austrian Academy of Sciences' Institute of High Energy Physics (HEPHY).
Tens of thousands of these silicon components could be used at CERN in the near future. The detectors improve on prior devices in several respects. They are not only more economical to produce than previous sensors, which measured up to six inches. The components also stand up better to constant radiation and thus age slower than the previous generation. Planned experiments, such as those probing the highest energies for evi-

dence of dark matter, will scarcely be possible without resistant sensors.
Two of the detectors for which the use of the Infineon sensors is currently being tested are named ATLAS (A Toroidal LHC Apparatus) and CMS (Compact Muon Solenoid). Particle physics experiments are huge cameras: When particles penetrate the silicon detectors, it registers them. With twenty meters (ATLAS) and, respectively fifteen metres (CMS) height both experiments are located 100m below ground. They have been in almost round-the-clock operation for years, carrying out 40 million individual experiments each second. The participants are currently discussing possible production of chips with a total area of up to 1,000m².

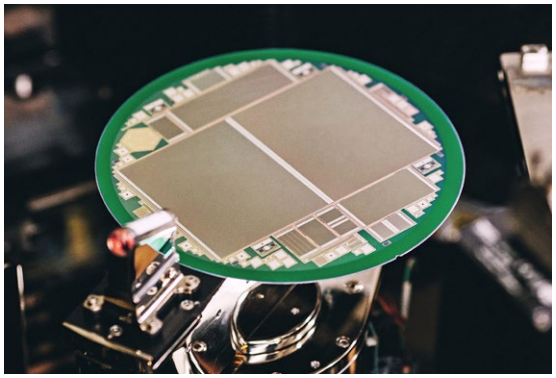

Infineon speculates that the technology developed for CERN could help cancer patients in less than ten years: several groups of researchers are currently testing proton computed tomography. The medical imaging procedure [that would be required for that technology] is based on the same fundamentals as the chip technology for CERN. Large silicon detectors like the ones Infineon and HEPHY are developing could supply tomographic images during therapeutic radiation. This would better determine the position of the tumour, allowing less damage to be done to healthy tissue than is possible with conventional X-rays. It would reduce the radiation load by a factor of 40.

Complete article, here

600-V GaN FET, integrated power conversion stage, from TI

Texas Instruments is sampling a 600-V gallium nitride (GaN) 70-mΩ field-effect transistor (FET) power-stage; TI claims to be the only semiconductor manufacturer to publicly offer a high-voltage driver-integrated GaN solution. The 12-A LMG3410 power stage is intended for use with TI's analogue and digital power-conversion controllers to create smaller, more efficient and higher-performing designs compared to silicon FET-based solutions, especially in isolated high-voltage industrial, telecom, enterprise computing and renewable energy applications.

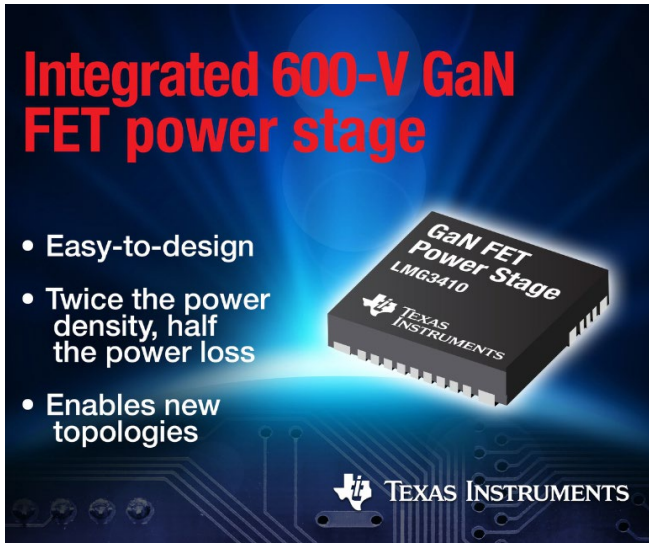

With its integrated driver and features such as zero reverse-recovery current, the LMG3410 provides reliable performance, especially in hard-switching applications where it can dramatically reduce switching losses by as much as 80%. TI claims a higher level of ease-of-use compared to stand-alone GaN FETs; the part integrates built-in intelligence for temperature, current and under-voltage lockout (UVLO) fault protection. It is the first IC to include GaN FETs manufactured by TI; the company creates its GaN devices in a silicon-compatible factory and qualifies them with practices that are beyond the typical Joint Electron Device Engineering Council (JEDEC) standards to ensure the reliability and robustness of GaN for demanding use cases. Easy-to-use packaging will help in-

crease the adoption of GaN power designs in applications such as power factor controller (PFC) AC/DC converters, high-voltage DC bus converters and photovoltaic (PV) inverters. The 600-V power stage delivers 50% lower power losses in a totem-pole PFC compared with state-of-the-art silicon-based boost power-factor converters. The reduced bill of materials (BOM) count and higher efficiency enable a reduction in power-supply size of as much as 50%.

Complete article, here

Former Agilent engineers developing a wireless oscilloscope

By Martin Rowe

A pair of engineers formerly with Agilent (now Keysight) Technologies are working on an oscilloscope that will connect to an iOS device for its user interface. The Aeroscope, conceived by Alexander Lee and Jonathan Ward, looks like the kind of device that you'd use on home projects or in field locations. I recently spoke with Lee about the design. At this time, Lee expects to ship in the Fall of 2016. The Aeroscope will be a one-channel device where the electronics are in a "pen" form factor. An SMA connector will let you attach mini-grabbers or other probe tips. Analogue bandwidth is 100 MHz with 500 Msamples/sec sample rate. It should store 10 ksamples of waveform data. "We designed the hardware from scratch," said Lee, focusing on power and the analogue front end." According to Lee the Aeroscope will have a ±40 V input range. It's based around an 8-bit ADC and a programmable gain amplifier. An FPGA handles triggering and memory management. An ARM-based processor handles communications to the BLTE (Bluetooth Low Energy) transceiver.

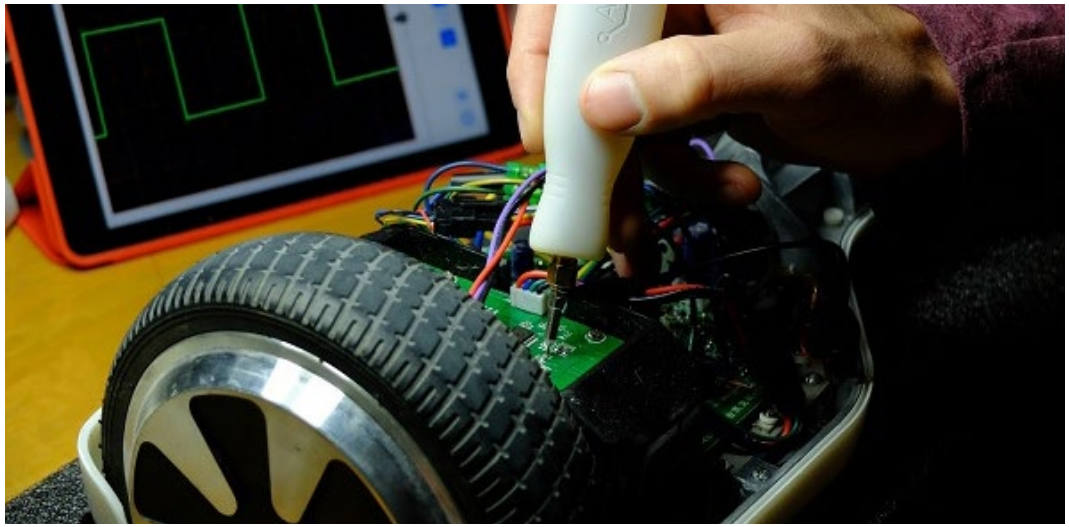

Lee and Ward are developing the Aeroscope software for iOS devices, at least for now. Lee claims that unlike the LabNation SmartScope , you won't have to jailbreak your device to use the software. That's critical, for would you jailbreak your iPhone or iPad to use a piece of test equipment? Lee didn't rule out the possibility of developing Android or Windows versions, but for now it's all iOS. A necessity for iOS-based oscilloscope is pinch gestures for setting vertical and horizontal sensitivity. Lee confirmed that the Aeroscope will have that ability. In addition, the software will let you change offset settings by dragging your finger. You can also set up triggers by sliding a slide control with your finger. "If you don't want to pinch to zoom," said Lee, "you can still use traditional menus."

Although Lee and Ward have designed the Aeroscope with one channel, you can have a single iOS device communicate with more than one unit. For the initial release, each unit's data will display in the same screen but the waveforms will not appear in the same grid.

Complete article, here

Academics offered intro-to-IoT curriculum by Imagination, Microchip & Digilent

Imagination Technologies and Microchip Technology, together with Digilent Inc. have created the “Connected MCU Lab”, a course developed through the companies' respective university programs.

The single-semester (half-academic-year) curriculum, available to universities worldwide, is designed to be an introductory and first microcontroller (MCU) class taken by undergraduate electronic engineering and computer science students. It delivers an interactive and compelling start to connected embedded systems education - covering MCUs and input/output (I/O), real-time operating system concepts, advanced MIPS processor architecture and cloud connectivity - all presented in a fresh and jargon-free style.

The Connected MCU Lab takes a hands-on approach, making use of a Wi-Fi enabled development board, tools, software, and cloud services - everything needed to design innovative Internet of Things (IoT) solutions. Lessons are based around the chipKIT Wi-FIRE board from Digilent Inc., which uses Microchip's PIC32MZ MCU incorporating a 32-bit MIPS M-Class CPU from Imagination. A chipKIT Basic I/O Shield is used for expansion along with a PICkit 3 In-Circuit Debugger from Microchip. Teachers and students have free access to professional software tools including MPLAB X Integrated Development Environment, MPLAB XC32 C compiler, and MPLAB Harmony Software Development Framework from Microchip, as well as Imagination's cloud technologies.

Imagination
MIPS32 Warrior M-Class processor core

MICROCHIP
PIC32MZ microcontroller

DIGILENT
chipKIT Wi-FIRE development board

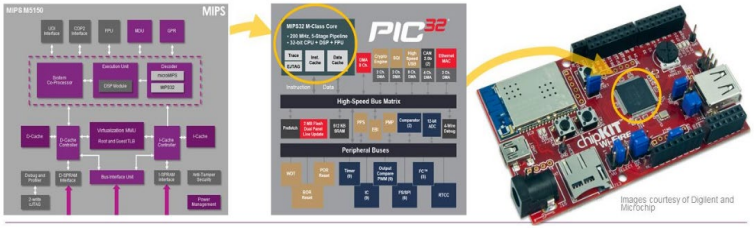

The Connected MCU Lab curriculum, written by Dr. Alexander Dean of North Carolina State University, includes presentation slides for each module, student guide, exercises, tests, solutions, and an Instructor's Guide.

Complete article, here

PREDICTING POWER SUPPLY RELIABILITY: AN ART OR A SCIENCE?

By Paul Baker, Future Power Solutions

The datasheet is a wonderful repository of sound, tested and verified information about the performance of a component, module or system. In the case of a Power Supply Unit (PSU), the datasheet tells engineers about a huge variety of performance parameters, including ripple and noise, efficiency, accuracy of regulation, isolation voltage and electro-magnetic emissions. The range and detail of the information on offer enable the user to characterise with great confidence the expected behaviour of the unit in any given application.

But what about one other important performance parameter: the reliability of the power supply? In truth, today's PSUs from reputable manufacturers offer extremely long lifetimes. The lifetime is precisely predictable when operated in the test conditions specified by reliability standards such as MIL-HDBK-217 or Telcordia. What is more, experience shows that high-quality PSUs also offer long lifetimes outside these strictly defined parameters.

There is a problem for system designers, however: how confidently can they predict the average lifetime when operating the PSU outside these test conditions? A wide variety of common factors can exceed these conditions: heat, shock and vibration, transient fluctuations in the supply voltage, and the ageing of electrolytic capacitors can all give rise to premature failure. The datasheet's standard lifetime rating, then, is rarely applicable to a real-world product.

At the same time, failure to manage the end product's reliability is hardly acceptable. The brand's reputation is a valuable asset. The environmental and financial cost of disposal and repair are also damaging.

So how can a system design engineer confidently estimate the reliability of a Commercial Off-The-Shelf (COTS) PSU? And which are the most effective ways to maximise this level of confidence?

The limitations of manufacturers' reliability data

The most commonly provided value expressing the lifetime of a new COTS PSU is the Mean Time To Failure (MTTF) or Mean Time Between Failures (MTBF) value. MTTF is normally specified in thousands of hours at a constant operating (ambient) temperature.

Of course, MTTF gives no indication about the time at which any single unit, chosen at random from a large population of units, will fail: MTTF is an average value. Some units will last longer than the specified MTTF value, and some will fail prematurely. In fact, assuming a constant failure rate (which might be an unrealistic assumption in the context of the operation of electronic equipment) the probability that an individual unit will last as long as the MTTF value is just 37%. Put another way, half of the units will have failed after 0.69 of the MTTF has elapsed (see Figure 1).

This is because the failures for a constant failure rate are characterised by an exponential factor, as expressed by the equation for calculating the probability of a component not failing after a given time:

$$R_{(t)} = e^{-\lambda t}$$

where λ is the average failure rate of the component.

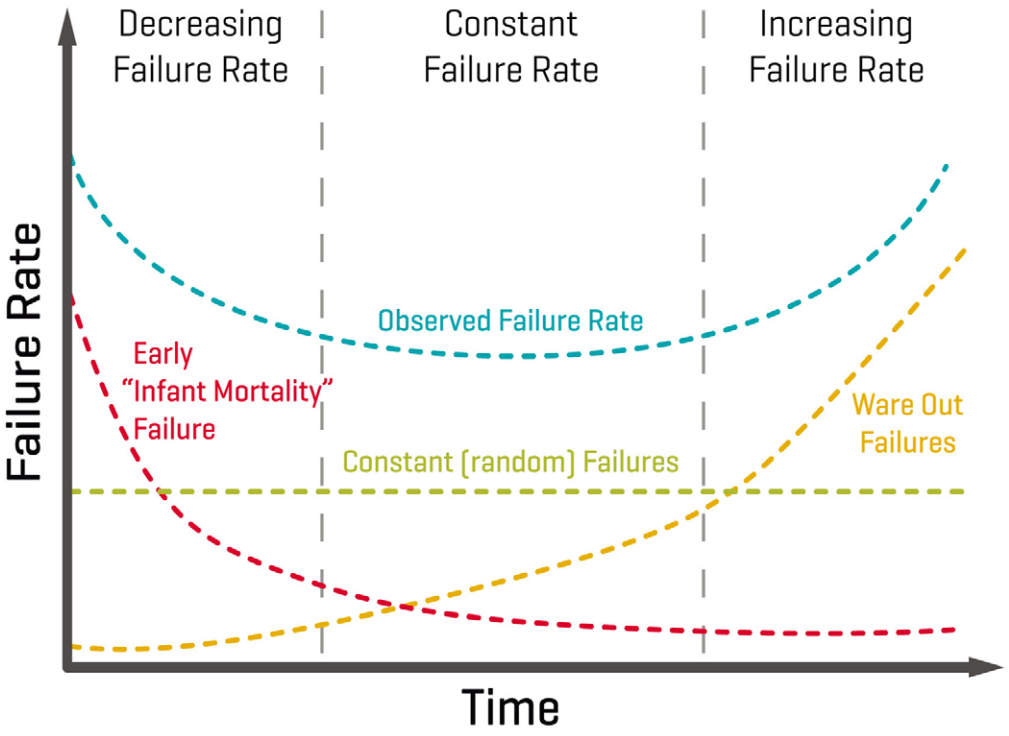

Figure 1. *Curve showing the probability that a unit with a given MTTF is still operational after a given multiple of the MTTF. (Source: CUI, 'Reliability Considerations in Power Supplies')*

PSU manufacturers employ models based on highly accelerated tests in order to predict the failure rate of their products. They cannot run a test population of PSUs under normal operating conditions and wait to observe failures, because it would take many years to gather statistically significant data. So they subject their products to excessive temperature, vibration, current and voltage stresses in order to rapidly impair them.

Clearly, a sound methodology should underlie the models that convert the results of accelerated tests into a datasheet's MTTF value; reputable PSU manufacturers carefully verify and refine their methodology to ensure it reflects users' experience in the real world.

Article continues, click below for pdf

DESIGNING FOR LOW-NOISE FEEDBACK CONTROL WITH MEMS GYROSCOPES

By Mark Looney, Analog Devices

MEMS gyroscopes offer a simple way to measure angular rate of rotation, in packages that easily attach to printed circuit boards, so they are a popular choice to serve as the feedback sensing element in many different types of motion control systems.

In this type of function, noise in the angular rate signals (MEMS gyroscope output) can have a direct influence over critical system behaviours – such as platform stability – and is often the defining factor in the level of precision that a MEMS gyroscope can support.

Therefore, "low-noise" is a natural, guiding value for system architects and developers as they define and develop new motion control systems. Taking that value (low-noise) a step further, translating critical system-level criteria, such as pointing accuracy, into noise metrics that are commonly available in MEMS gyroscope datasheets, is a very important part of early conceptual and architectural work. Understanding the system's dependence on gyroscope noise behaviours has a number of rewards, such as being able to establish relevant requirements for the feedback sensing element or, conversely, analyzing the system-level response to noise in a particular gyroscope.

Once system designers have a good understanding of this relationship, they can focus on mastering the two key areas of influence that they have over the noise behaviours in their angular rate feedback loops: (1) developing the most appropriate criteria for MEMS gyroscope selection and (2) preserving the available noise performance throughout the sensor's integration process.

Motion control basics

Developing a useful relationship between the noise behaviours in a MEMS gyroscope and how it impacts key system behaviours often starts with a basic understanding of how the system works. Figure 1 offers an example architecture for a motion control system, which breaks the key system elements down into functional blocks. The functional objective for this type of system is to create a stable platform for personnel or equipment that can be sensitive to inertial motion. One example application is for a microwave antenna on an autonomous vehicle platform, which is manoeuvring through rough conditions at a speed that causes abrupt changes in vehicle orientation. Without some real-time control of the pointing angle, these highly-directional antennas may not be able to support continuous communication, while experiencing this type of inertial motion.

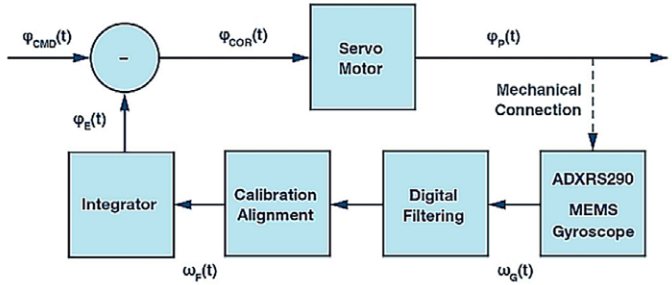

Figure 1. *Example motion control system architecture*

The system in Figure 1 uses a servo motor, which will rotate in a manner that is equal and opposite of the rotation that the rest of the system will experience. The feedback loop starts with a MEMS gyroscope, which observes the rate of rotation (ω_G) on the "stabilized platform." The gyroscope's angular rate signals then feed into application-specific digital signal processing that includes filtering, calibration, alignment and integration to produce real-time, orientation feedback, (ϕ_E). The servo motor's control

signal (Φ_{COR}) comes from a comparison of this feedback signal, with the "commanded" orientation (Φ_{CMD}), which may come from a central mission control system or simply represent the orientation that supports ideal operation of the equipment on the platform.

Example application

Moving from the architectural view of a motion control system in Figure 1, valuable definitions and insights also come from analyzing application-specific, physical attributes. Consider the system in Figure 2, which offers a conceptual view of an automated inspection system for a production line. This camera system inspects items that move in and out of its field of view on a conveyor belt. In this arrangement, the camera attaches to the ceiling through a long bracket, which establishes its height (See "D" in Figure 2), in order to optimize its field of view for the size of the objects it will inspect. Since factories are full of machinery and other activity, the camera can experience swinging motion (see "$\omega_{SW}(t)$" in Figure 2) at times, which can cause distortion in the inspection images.

The red dotted lines in this diagram provide an exaggerated view of total angular error ($\pm\Phi_{SW}$) that comes from this swinging motion and the green dotted lines represent the level of angular error that will support the system's image quality objectives ($\pm\Phi_{RE}$). The view in Figure 2 defines the key system-level metric (image distortion) in terms of linear displacement error (d_{SW}, d_{RE}) on the inspection surface. These attributes relates to the camera height (D) and the angular error terms (Φ_{SW}, Φ_{RE}) through a simple trigonometric relationship in equation 1.

$$d_{SW} = D \times \sin(\varphi_{SW})$$

$$\varphi_{SW} = a\sin\left(d_{SW}/D\right)$$

$$d_{RE} = D \times \sin(\varphi_{RE})$$

$$\varphi_{RE} = a\sin\left(d_{RE}/D\right)$$

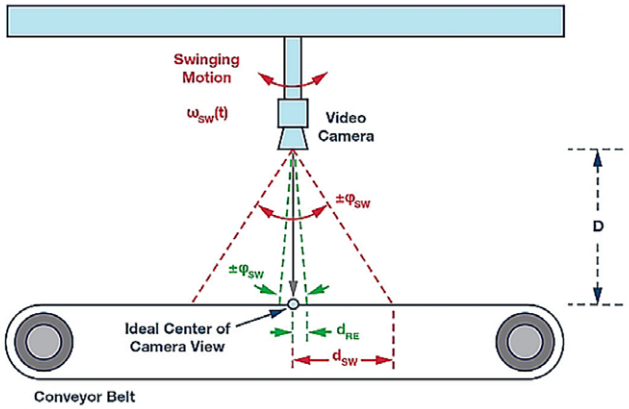

Figure 2. *Industrial camera inspection system*

The most applicable motion control technique for this type of system is known as Image Stabilization. Early image stabilization systems used gyroscope-based feedback systems to drive servo motors, which adjust the orientation of the image sensor during the time that the shutter is open. The emergence of MEMS technology helped reduce the size, cost and power of these functions in a revolutionary manner, which lead to wider-spread use of this technique in modern-day digital cameras. Advances in digital image processing techniques, which still use MEMS-based angular rate measurements in their algorithms, have led to elimination of the servo motor in many applications.

Whether the image stabilization comes from a servo motor or through digital post processing of image files, the fundamental function (feedback sensing) of the gyroscope remains the same, as does the consequence of its noise. For simplicity, this discussion focuses on the classic approach (servo motor on the image sensor) to explore the most relevant noise fundamentals, and how they relate to the most important physical attributes of this type of application. *Click for continuation in pdf*

THREE PREDICTIONS ON THE STATE OF SIGNAL INTEGRITY

BY ERIC BOGATIN

Based on discussions at DesignCon 2016, and since, I have three predictions about major changes ahead for high speed serial link systems.

1. Roll out of 28 Gbps systems will be slower than expected

I hear that the semiconductor companies producing the CMOS devices – ASIC, FPGA or custom – are doing fine producing the silicon with acceptable performance at 28 Gbps. Figure 1 is an example of a very clean eye from a 28 Gbps TX (transmitter).

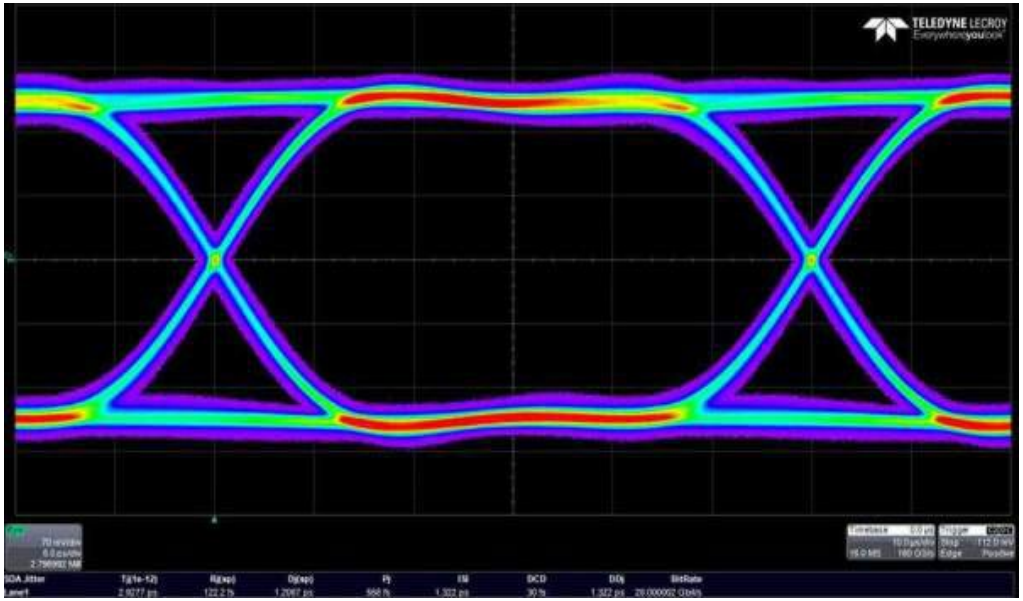

Figure 1. *Today's silicon can produce clean signals at 28 Gbps, at least at the transmit end.*

Semiconductor manufacturers' ability to sell to end users designing and manufacturing systems operating with 28 Gbps links is, however, limited by the their ability to support these customers.

A link operating at 28 Gbps, NRZ (non-return-to-zero), has to be designed with everything working almost perfectly. This data rate pushes the limits such as: low Df materials, smoother copper, wide enough lines, equalisation tuned to the limit of recovering -25 dB of insertion loss, minimal reflections, via stubs shorter than 15 mils [0.015 in. or about 0.4 mm], channel-to-channel cross talk less than -50 dB, and line-to-line skew less than 6 psec over as long as 20 in (50 cm).

By themselves, each item is possible to engineer, but all of them at the same time in the same channel requires solid engineering and analysis. Not every design team is capable of this task. When the channel does not work, who do they call? The silicon provider.

I hear that with a limited number of experienced support application engineers, the silicon providers are focusing on their large, high-end OEM customers and are limiting their sales based on which customers they have the resources to support. This may be a business opportunity for consulting engineering teams to work with silicon providers to support their customers and increase the design wins and sales of 28 Gbps-capable silicon.

2. There is a potential roadblock ahead for 56 Gbps PAM4 systems

A number of channels have been demonstrated operating at 56 Gbps with PAM4. The basic tools needed for PAM4 systems are in place. Most of the high-end software vendors have shown design tools for simulating PAM4. All the high-end oscilloscope and BERT (bit-error-rate tester) manufacturers have shown instruments able to measure and characterize PAM4 systems. Figure 2 shows the measured eye for a 56 Gbps PAM4 link.

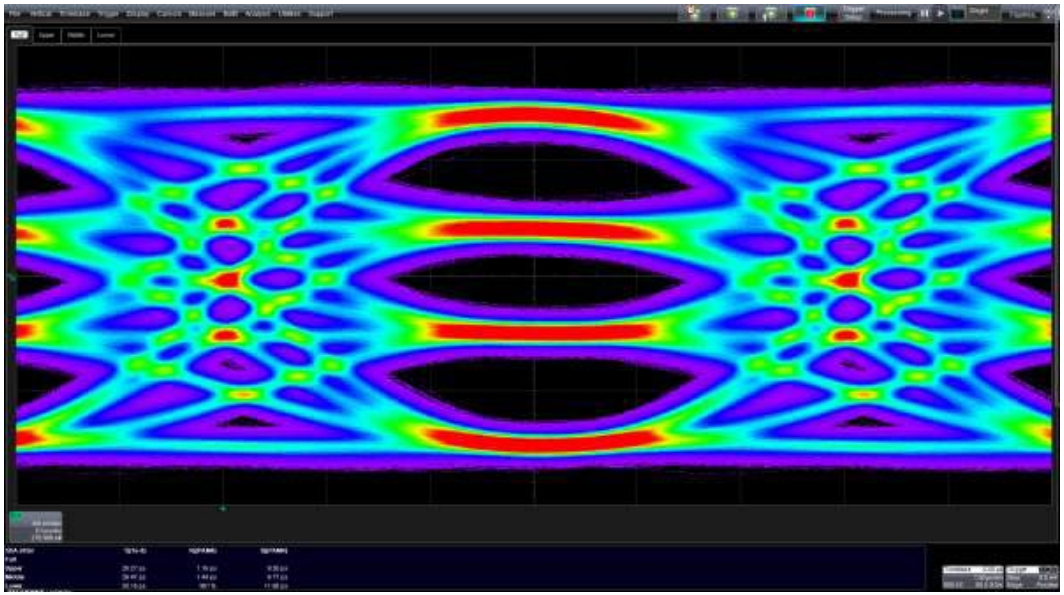

Figure 2. *At the transmitter, a PAM4 signal is clean enough for all three eyes to be visible.*

It's widely believed that the advantage of going to PAM4 for 56 Gbps is so that we are only dealing with signals with an equivalent bandwidth of 28 Gbps signals. If we can design a channel for 28 Gbps at PAM2, we should be able to design one for 56 Gbps at PAM4.

Not so fast: for there is one significant difference with PAM4. By dividing up the signal into three levels plus zero, we dropped the signal level for one bit by 1/3. The signal voltage level we have to measure is smaller. If we need a particular SNR (signal-to-noise ratio) at the receiver for an NRZ-PAM2 signal, and the signal level dropped by 10 dB, the acceptable noise level has to drop by 10 dB in PAM4. But wait: we're not done.

In NRZ-PAM2, we need about -50 dB isolation between a channel and all other aggressors for a SNR of 20 dB. With a lower noise floor required in PAM4, this means an isolation of –60 dB. When it comes to crosstalk, we still have high level signals corresponding to the fourth bit level sometimes coming out of the TX. This means the signal on the aggressor can be 3x higher than the signal of the second bit. To keep the same noise on the victim line when the aggressor has 10 dB higher strength, we need another 10 dB more isolation. This means an isolation of as low as -70 dB between the victim channel and all other aggressor channels.

I hear that the weak link in achieving this low level of isolation is in the via field under the BGA. At this low level of crosstalk required, issues such as differential to differential coupling in the via field under the BGA and common noise to differential noise conversion in all via fields, in connectors and in channel to channel cross talk, can be showstoppers. While it may be possible, with good engineering practices and optimized pad stack design to reduce cross talk to the -50 dB level, getting to -70 dB is a major engineering effort.

At this level, as well designed as a via area is [may be], manufacturing variations in the fabricated board can push a system into too much cross talk. There are some fundamental limitations to what can be done at the board level if the package footprint is poorly designed. This puts a larger burden on the silicon providers to design the package footprint with channel-to-channel crosstalk at the board level via field in mind. This does not play to their strengths.

While getting one channel operating at 56 Gbps PAM4 is possible, getting hundreds of channels operating, in close proximity, at acceptable bit error ratio, may require heroic efforts.

3. All is not doom and gloom

I did hear of one innovation that may be the saviour for high-speed serial links in copper-based interconnects. Given the increasing challenges to get a long channel operating at 28 Gbps in PAM2-NRZ or a 56 Gbps channel operating at PAM4, there may be an intermedi-

Figure 3. *Samtec's Firefly interconnect system merges optical and electrical connections to improve signal integrity.*

ate fix available. Every large connector company I spoke with has a practical plan to implement cabled interconnects integrated with the board to supplement laminated backplane and motherboard routing.

The advantage of a cabled system is lower loss and less channel-to-channel cross talk. The larger circumference in the round conductors means lower conductor loss per length in a cable than on a board. While there may be lower crosstalk in the cable interconnects, the crosstalk in the connector and its board footprint still needs to be considered, but many of the connector companies seem very good at this.

These solutions involve a connector system to mate between the board and an array of cables and back to the board. The idea is to route long distance, high bandwidth signals off the board, through cables and then back to the board. Figure 3 shows an example: the Firefly product from Samtec. A nice feature of the Samtec system is the integration of optical cables as well as copper cables to ease the transition to board-level optical interconnects. [Editor's note; Samtec describes it thus; "The FireFly Micro Flyover System is the first inside-the-box interconnect system that gives designers a choice of using either micro footprint high performance active optical engines or low-cost copper interconnects.]

This sort of approach, with a much lower loss at 14 GHz and 28 GHz, maybe the short-term fix to enable both a robust 56 Gbps (28 Gbaud) PAM4 or an PAM2-NRZ 56 Gbps system without the headaches of extremely high isolation requirements of a PAM4 system.

This sort of backplane architecture moves the interconnect roadmap onto a different trajectory and may give additional headroom to copper interconnects into the next generation of data rates. With the option of also including fibre optics, it may be the route into the long touted optical backplane architecture of the future.

Eric Bogatin is the Dean of the Signal Integrity Academy, www.beTheSignal.com, and an adjunct professor in the ECEE department at the University of Colorado, Boulder. He received his BS in Physics from MIT and PhD in Physics from the University of Arizona. He can be reached at eric@beTheSignal.com

HIGH VOLTAGE CHARGE PUMPS DELIVER LOW EMI

By Tony Armstrong, Linear Technology

Switching regulators are a popular choice for many power conversion supplies due to their size, output flexibility and efficiency advantages. Conversion efficacy of these supplies can now attain 98% levels, depending on operating conditions. Nevertheless, despite these advantages, they necessitate compromising on other parameters, one of the most difficult of which is noise.

[In the full version of this article, the author interposes a discussion of noise sources in switching regulators, and their mitigation – Ed.]

In some noise-sensitive applications, power supply designers simply do not like to use inductor-based regulators due to their associated EMI emissions. At the same time, the use of a linear regulator (aka LDO) may be precluded due to its relatively low conversion efficiency and need for heat sinking. As a result, they turn to a common alternative known as a charge pump.

Charge pumps have been around for decades, and they provide DC/DC voltage conversion, using a switch network to charge and discharge two or more capacitors. The basic charge pump switch network toggles between charge and discharge states of the capacitors. As shown in Figure 1, C1 the "flying capacitor" shuttles charge, and C2 the "reservoir capacitor" holds charge and filters the output voltage. Additional "flying capacitors" and switch arrays enable multiple gains.

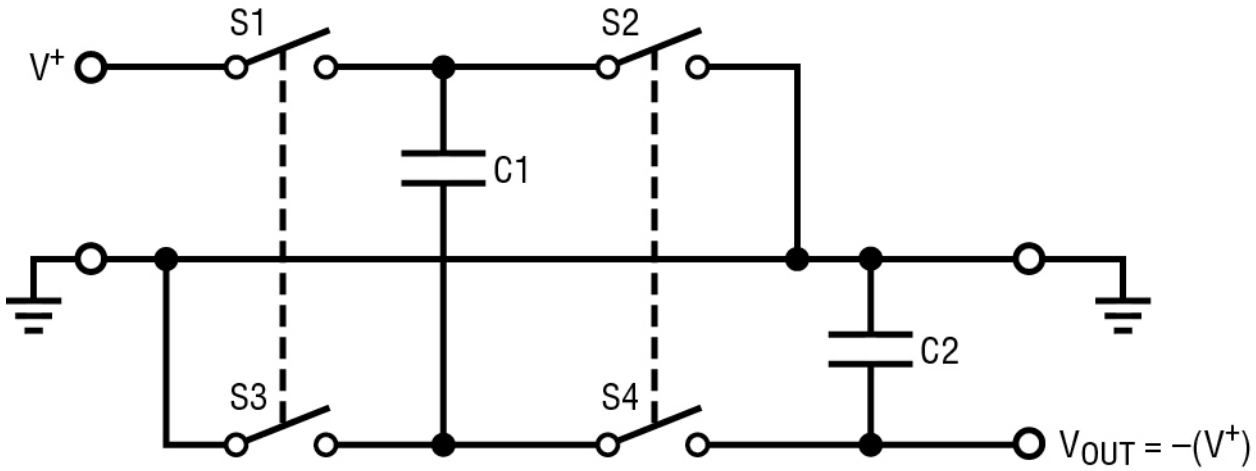

Figure 1. *Simplified charge pump block diagram of a voltage inverter*

When switches S1 and S3 are on, or closed, and switches S2 and S4 are off, or open, the input power supply charges C1. During the next cycle, S1 and S3 are off, S2 and S4 are on, and charge transfers to C2, generating V_{OUT} = - (V+).

However, until recently, charge pumps have had limited input and output voltage ranges, which has limited their use in industrial and automotive applications where inputs up to 40V or 60V are commonplace. However, this is now changed with the introduction of high voltage charge pumps.

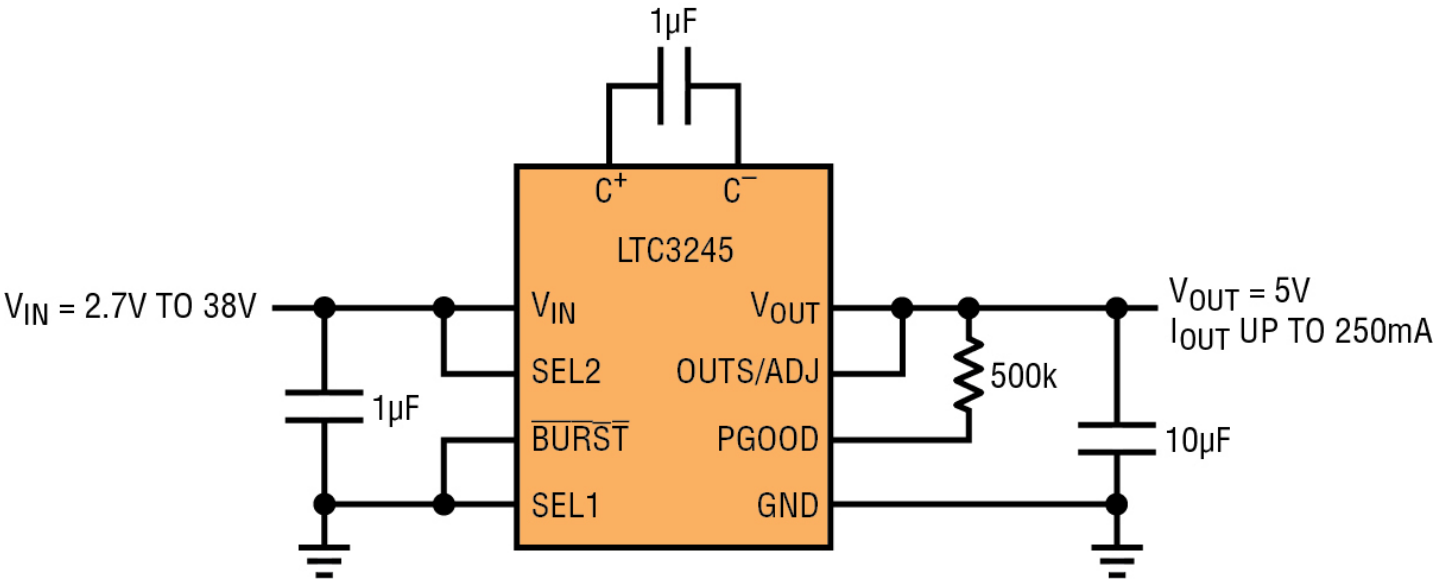

Figure 2. *LTC3245 schematic delivering a fixed 5V output from a 2.7V to 38V input*

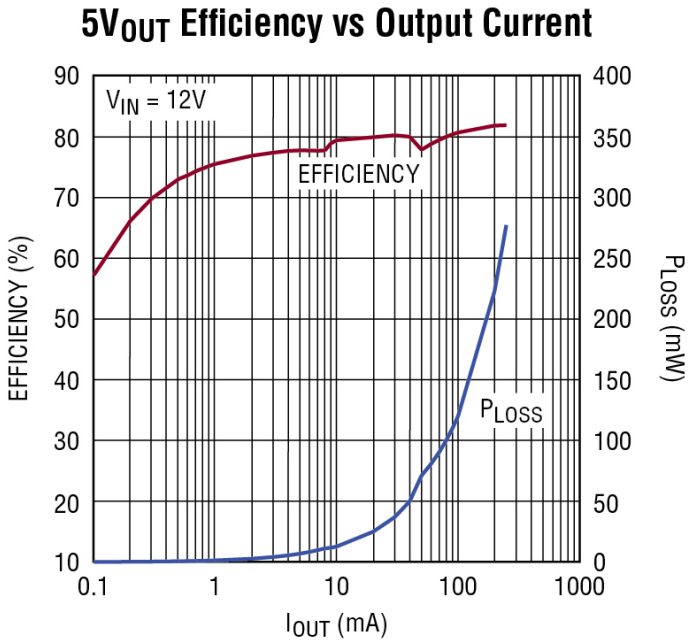

Figure 3. *LTC3245 efficiency/power loss curve for 12V input to 5V out*

High-voltage charge pumps

The part shown in Figure 2 is a buck-boost regulator that dispenses with the traditional inductor and uses a switched-capacitor charge pump instead. Its input voltage range is 2.7V to 38V, and it can be used without a feedback divider to produce one of two fixed output voltages, 3.3V or 5V, or programmed via a feedback divider to any output voltage from 2.5V to 5.5V. Maximum output current is 250 mA. The device is capable of regulating a voltage above or below the input voltage, allowing it to satisfy automotive cold crank requirements, for example. Figure 2 shows its complete schematic.

This charge pump is capable of achieving efficiencies of 80% when delivering 5V at 100 mA from a 12V source, almost twice than that of a linear regulator, making it possible to avoid the space and cost requirements of an LDO with a heat sink. It generates almost three times lower power loss at full load versus an LDO. See Figure 3 for its efficiency and power loss curves.

The article continues with results for efficiency/power loss, and for radiated and conducted EMI from this circuit configuration; it concludes with a circuit arrangement to provide symmetrical +/- output rails, using the same approach – click for pdf

MEMS GYROSCOPE BIAS STABILITY VS. SENSITIVITY

BY BY IAN BEAVERS, ANALOG DEVICES

Two key performance selection metrics of a high performance gyroscope are its sensitivity and bias stability. Distinguishing between the two will help engineers choose the best solution for their systems.

The sensitivity of a gyroscope describes its dynamic response to angular rotation. For a gyro with an analogue output, example specification units may be mV/degree/sec which describes how much the output voltage changes for the angular velocity of the component over time. When a precision analogue to digital converter (ADC) is internally used to sample the sensor's analogue output, the common sensitivity units are LSB/degree/sec. This nomenclature describes the number of ADC digital output codes within a degree of rotation over time. Higher sensitivity, a larger voltage output or larger number of LSBs per unit of rotation, is preferred in order to resolve a more accurate angular velocity measurement.

Although a multi-axis sensor may have similar sensitivity per axis, each can be quantified separately. This allows the design engineer to consider each axis of the gyroscope's performance against its most advantageous position within the system.

Bias stability, or rather bias instability, is often considered a key parameter in sensors. It defines the average output from the gyroscope when it is not subjected to angular movement. Unwanted bias error is any non-zero sensor output when the input movement is held at zero. Ideally, the output for this case should also be zero. The stability of the bias indicates how it varies in units of degrees over a given time interval, usually noted in seconds or hours.
For a gyroscope, the angular deviation per hour vs. the averaging time or desired sampling interval (Tau) is typically shown within the datasheet in units of degrees per hour. The Allan variance is often used as a method to determine the intrinsic noise as a function of Tau. All variables such as angular rotation, temperature, and the voltage at the supply domains are held static during the observation time. A smaller degree of angular deviation per hour is desired as it specifies the drift of the gyro. The in-run bias stability is often the headline performance metric of the gyro.

Many gyroscope datasheets specify the impact of temperature on sensitivity and bias. A system designer can then calibrate their application of interest to account for any temperature induced error. Industrial grade Isensor inertial measurement units (IMU) from Analog Devices are calibrated across their operating temperature range at the factory to alleviate this burden on the user. The compact six degrees of freedom inertial sensor ADIS16460 provides a calibrated triaxial digital gyroscope and triaxial digital accelerometer.

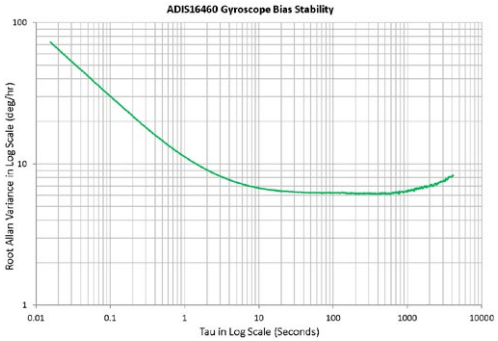

Figure 1. *The ADIS16460 gyroscope bias stability is plotted against time constant Tau. At a sample interval of one hour (3,600 sec.), the angular deviation is 8 deg/hour.*

About the author
Ian Beavers [ian.beavers@analog.com], a staff engineer for the Digital Video Processing Group at Analog Devices (Greensboro, NC), is a team leader for HDMI and other video interface products. With over 15 years' experience in the semiconductor industry, he has worked for ADI since 1999. He holds a bachelor's degree in electrical engineering from North Carolina State University and an MBA from the University of North Carolina at Greensboro.

MEASURING PHASE NOISE WITH A REAL TIME SAMPLING OSCILLOSCOPE – PART 2

By Mark Johnson, Application Physicist, Keysight Technologies

In the first part of this article we described how phase noise information can be extracted from real-time scope waveform acquisitions using two different techniques to demodulate the phase. In this article we'll take a look at the potential accuracy of the serial data clock recovery technique, what kinds of signals can reasonably be analyzed and some ways to improve such measurements.

Accuracy

In order to check that the scope phase noise measurement is accurate we can use a clean signal source with a broadband random phase modulation source built-in. By injecting a relatively large PM amplitude over broad frequency we can verify the noise level by comparison with a measurement made on a Signal Source Analyzer such as the Keysight E5052B.

In the measurement below (Figure 1), the SSA result is in blue and the scope measurement (Keysight MSOS804A) result is in green. There is excellent agreement over the range of injected PM. Above 2 MHz the SSA's lower noise floor is the reason for the separation in the curves.

Measurement noise floor

The measurement floor of a jitter measurement on a real-time sampling oscilloscope is affected by both vertical (voltage) accuracy and timing accuracy. Vertical noise in the sampling system, stability of the timebase, the phase noise of the scope's own oscillator and imperfections in the interleaving architecture of the scope will all contribute to errors in the jitter measurements and thus the measured phase noise.

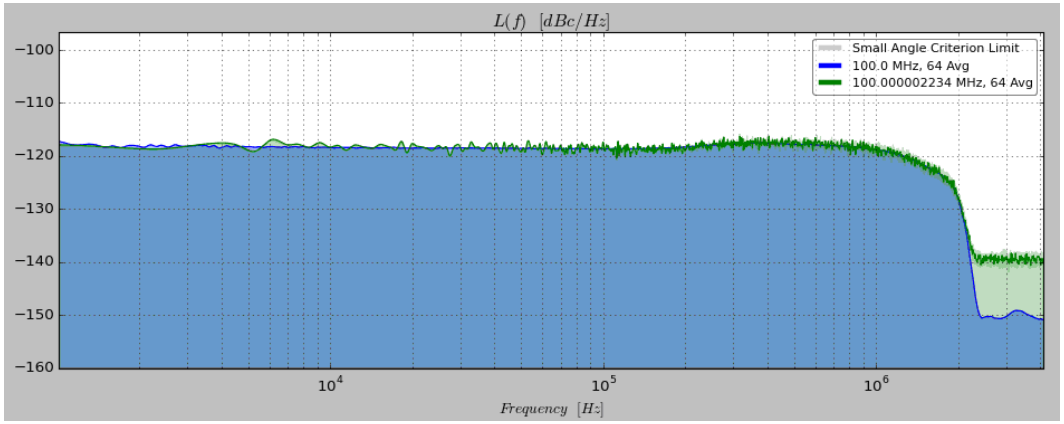

Figure 1.

An example of an oscilloscope jitter measurement floor specification is:

$$TIE = \sqrt{\left(\frac{Noise}{Slew\ Rate}\right)^2 + Intrinsic\ Jitter^2}\ [seconds\ peak]$$

The Intrinsic Jitter portion is dependent on the stability of the internal timebase reference. For the highest performance scopes such as the 63 GHz Keysight Z-series this can be as low as 50 fsec but it must be noted that this value is often only valid for fairly short acquisition times. To measure close-in phase noise we need to capture long acquisition times and the Intrinsic Jitter of the scope will increase due to its own phase noise.

Noise floor & signal slew rate

In most cases the first term in the equation dominates the jitter measurement floor. Both signal and oscilloscope vertical noise combine with the finite slew rate of the signal to create apparent horizontal displacement of edges,

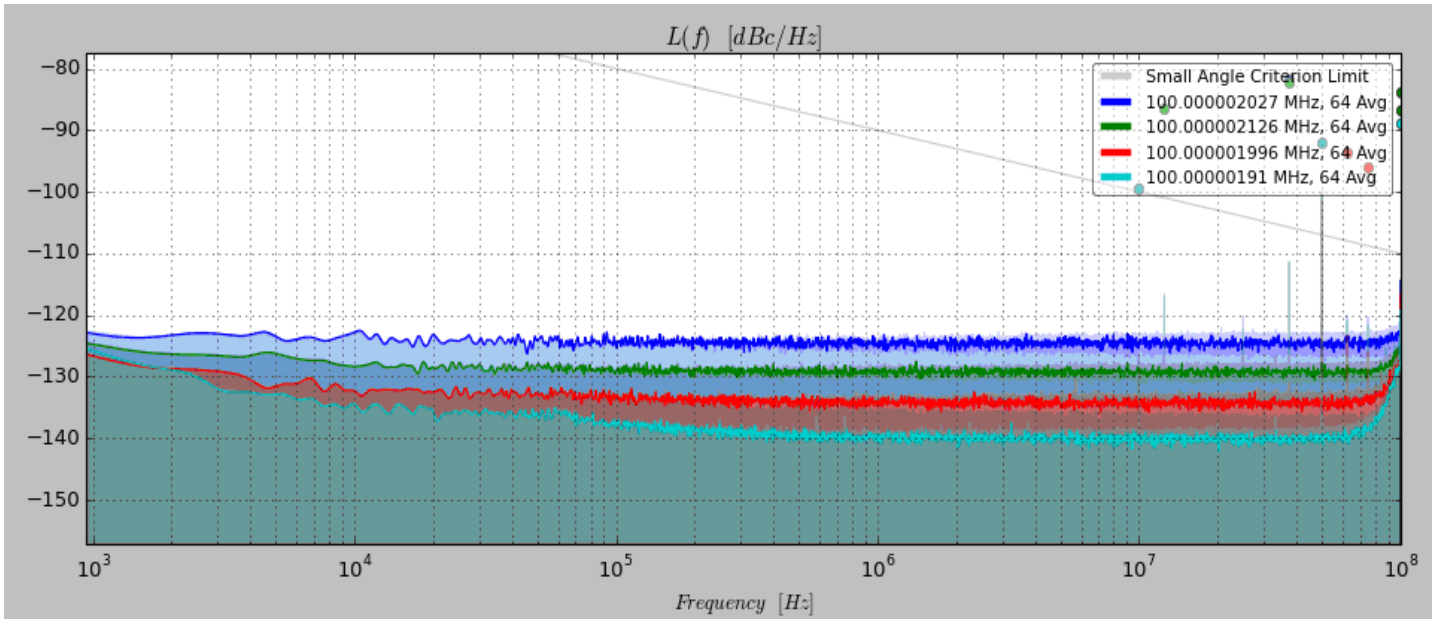

Figure 2.

ie: jitter. Thus it is crucial to choose an oscilloscope with as low a vertical noise bandwidth density as possible. A further improvement in jitter measurement floor can be achieved if the scope also has the ability to limit the bandwidth to an arbitrary frequency. Since the phase noise information is contained within a bandwidth $2*f_c$ we can drastically limit the measurement noise in many cases.

Figure 2 shows a set of phase noise measurements made using a Keysight 8 GHz S-Series oscilloscope. The signal source was a 100 MHz sine wave from an ultra-low phase noise Performance Signal Generator, E8267D. The true phase noise of the E8267D (as verified with an SSA or other suitably low phase noise instrument) is well below the scope measurements so this enables us to see the measurement floor of the scope.

The scope bandwidth was adjusted for each measurement as follows:

Blue = 8 GHz, Green = 4 GHz, Red = 1 GHz, Cyan = 200 MHz.

The phase noise floor at 1-10 MHz offsets drops from ~-124 dBc/Hz to -140 dBc/Hz when going from bandwidth of 8 GHz to 200 MHz. This can be explained by the fact that we're reducing the bandwidth by a factor of 200 MHz / 8 GHz. If the noise of the scope is fairly flat with bandwidth we should expect a drop of about 10*log10(0.2/8) = -16 dB. This is not the case at all frequencies. At low frequencies the phase noise of the scope's internal reference starts to dominate. At higher frequencies we see the limit of the ability to produce a perfect brick-wall bandwidth limit filter at 200 MHz. This means we are still getting some scope noise beyond 200 MHz included in our measurement.

The benefit gained in limiting the scope's bandwidth is highly dependent on the slew rate of the signal to be measured and the ratio of the signal frequency to the full scope bandwidth.

The article continues with a discussion of the noise floor, and the role of the scope's internal sources, and the effects of sampling rate: click for pdf

ADDING UP-FRONT PCB DESIGN INTELLIGENCE TO AVOID LAST-MINUTE MISTAKES

By Mike Griesbach, Altium

PCB designers today are dealing with an array of new challenges including greater electronic and mechanical design complexity, tighter constraints and rules, and the need to stay on top of a fast-changing supply chain. These and other emerging issues can cause extra design spins that drive up engineering costs and delay product introductions, ultimately costing you money.

Leading-edge PCB design tools are helping to address these concerns by providing up-front design intelligence that addresses these and other areas of the design process. Rules and constraints can be established earlier in the design process, reducing the need for potentially expensive changes late in the design process. Electrical design is becoming more integrated with mechanical design so that, for example, changes can be pushed from the mechanical computer aided design (MCAD) environment to the electrical computer aided design (ECAD) environment and vice versa. Finally, embedded component management systems are keeping designers up to date on availability, delivery and pricing.

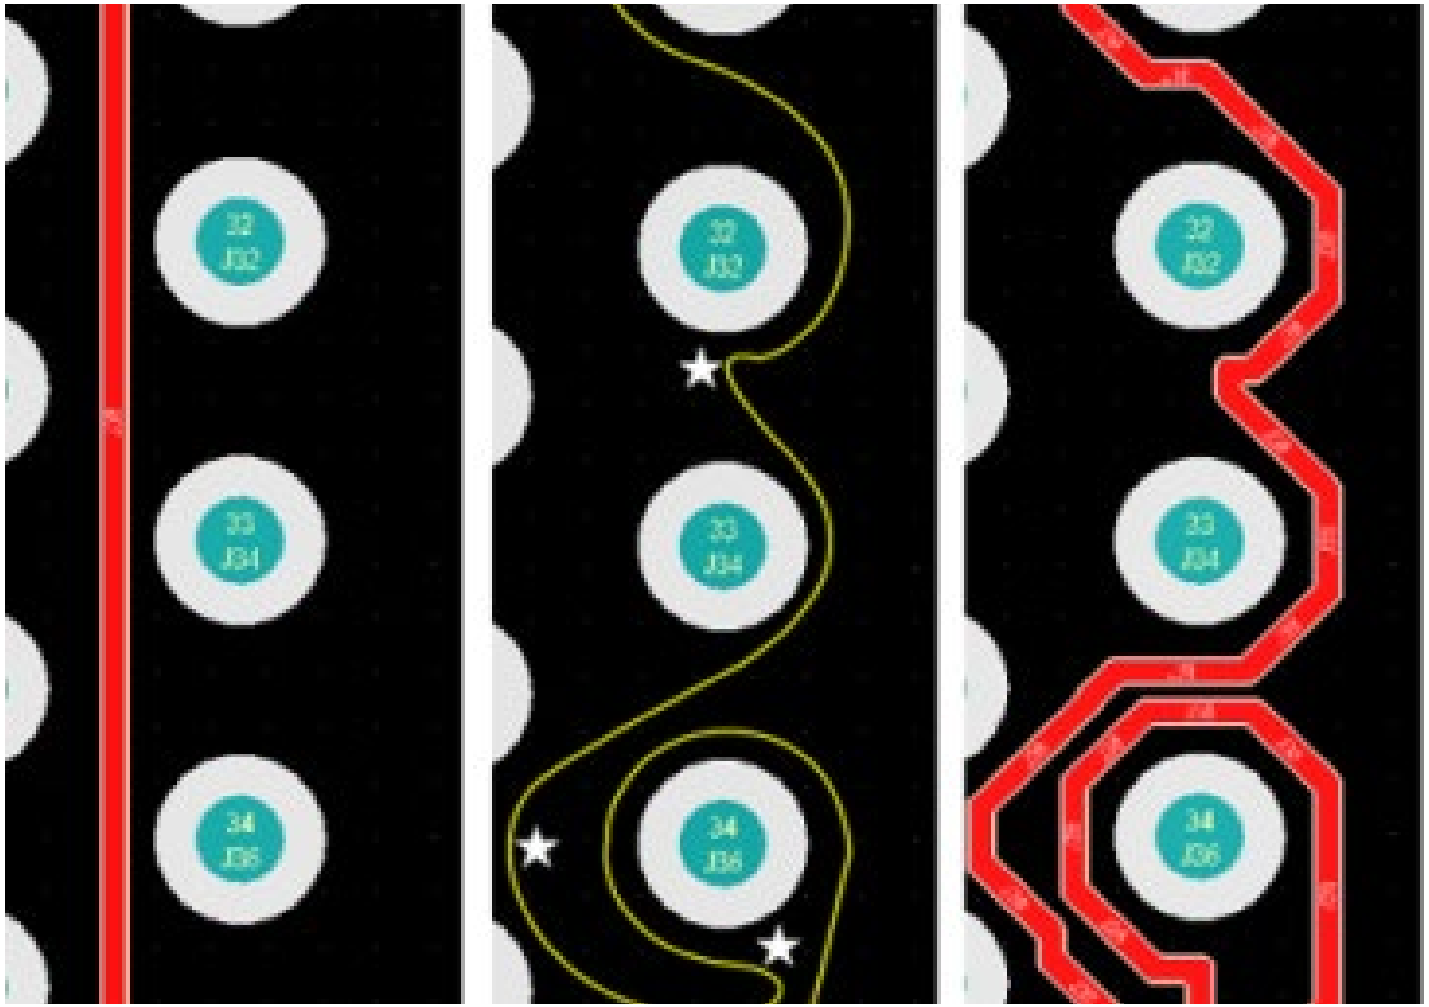

Figure 1. *Example of an interactive routing path*

PCB DESIGN TOOLS

PCB design challenges

Today's electronic products deliver more functionality than ever before so engineers are tasked with designing products with more PCBs, more layers and more components. Enclosures are shrinking and taking on more stylised outlines which makes it more critical than ever for PCB design to take into account the shape of the enclosure. Likewise, the physical aspects of the internal electronics are becoming more important than ever to the mechanical design. The increasing speed and functionality of the latest cutting edge products combined with their smaller size makes it more difficult than ever to avoid power integrity, thermal integrity and signal integrity problems. Dealing with the supply chain has also become more complicated than ever before. It is no longer possible to specify a particular component at the beginning of the design process and assume that it will still be available at the manufacturing stage. Designers need to continuously monitor changes in the supply chain and be ready to react to changes by substituting different components.

Tools reach beyond automatic routing

Component placement and routing were the original capabilities offered by PCB design tools and these capabilities are still vital to today's designers. Today's leading edge design tools help maximise routing efficiency and flexibility with interactive routing that lays route sections along the cursor path, enables pushing or walking around obstacles and automatically follows existing connections. Cursor-guided routing makes complex manual routing around obstacles fast, easy and intuitive. You create the path of the route with your mouse and the interactive router attempts to place the tracks according to that path. Figure 1 is an illustration of how complex routes can be created quickly with the routing path guided by cursor movement. The left image shows a normal, minimal length route, the centre image indicates the cursor path, with stars indicating where clicks were made to commit sections of the route. The right image is the resultant routing. Although an extreme example, it shows how few routing commits are required to place many tracks.

The article continues with consideration of component placement, and the significance of establishing design rules at the earliest stage... click for pdf

Download PDF of Article

Find PCB Design Tools on EETsearch

HIGH BOOST FACTOR STEP-UP-REGULATOR

By Florian Mueller, Texas Instruments

This article presents a high boost factor step-up regulator. The boost topology is non-isolated, the input and output voltages share a common ground and the output voltage is always higher than the input voltage. The input current is continuous because the inductor is connected to the input capacitor. However, the output current is pulsating because the high side FET on the output conducts only during the off state. Figure 1 shows the simplified schematic of a boost converter.

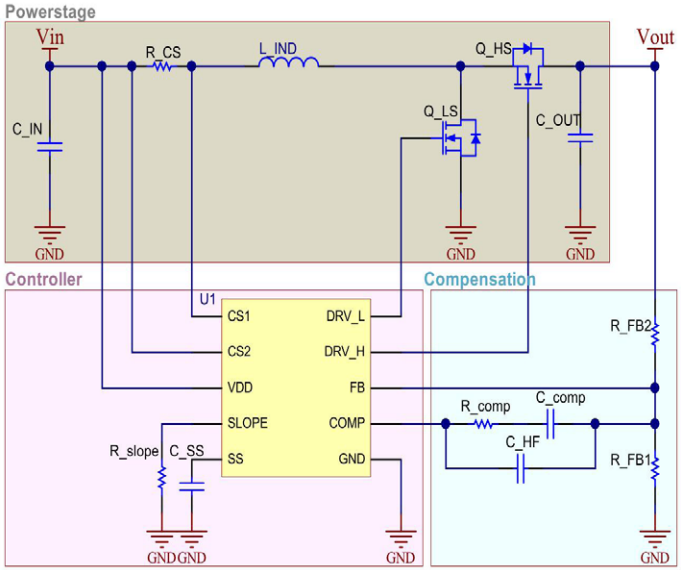

Figure 1. *Boost converter*

Powerstage high boost factor

The AC part of the input current for a boost is quite low. That is not usually the case for the output current ripple. The discontinuous output current causes a high AC current ripple, which can overstress the output capacitors. To prevent this, it is highly recommended that one calculates the maximum AC output current ripple and the current sharing between the output capacitors. Unlike the buck converter, the average of the inductor current is not equal to the output current because the inductor only delivers current to the output during the off state. The worst case condition occurs for the highest boost factor and, therefore, respectively for the minimum input voltage. At this condition the RMS (root mean square) and the peak current of the inductor is at the highest level. If the losses are neglected, the maximum peak current of the inductor is equal to the peak current through the lowside FET Q_LS during the on state, and equal to the peak current through the highside FET Q_HS during the off state. Therefore, the peak current of all powerstage components reaches its maximum value for the minimum input voltage. Two low side MOSFETs in parallel may be beneficial in a high boost factor application, because for large duty cycles the losses of the low side MOSFET may be very high.

Feedback compensation

The open loop response of the regulator is the product of the feedback transfer function, the powerstage transfer function and the modulator gain. The power stage, consisting of the output inductor and the output capacitor, must be compensated. This 2nd order system degenerates to a first order system if the controller is working in peak current mode. This is because the inductor behaves like a current source, making the transfer function of the output filter independent of the inductance. This makes the compensation simpler because the output capacitor and the equivalent load resistor remain as output filter. At the same time a Type II compensation network is sufficient to get a stable system with enough phase and gain margin. But there is a very important concern called the right half plane zero (RHPZ).... *the article continues to outline stability and compensation issues in the complete version, click for pdf.*

designideas

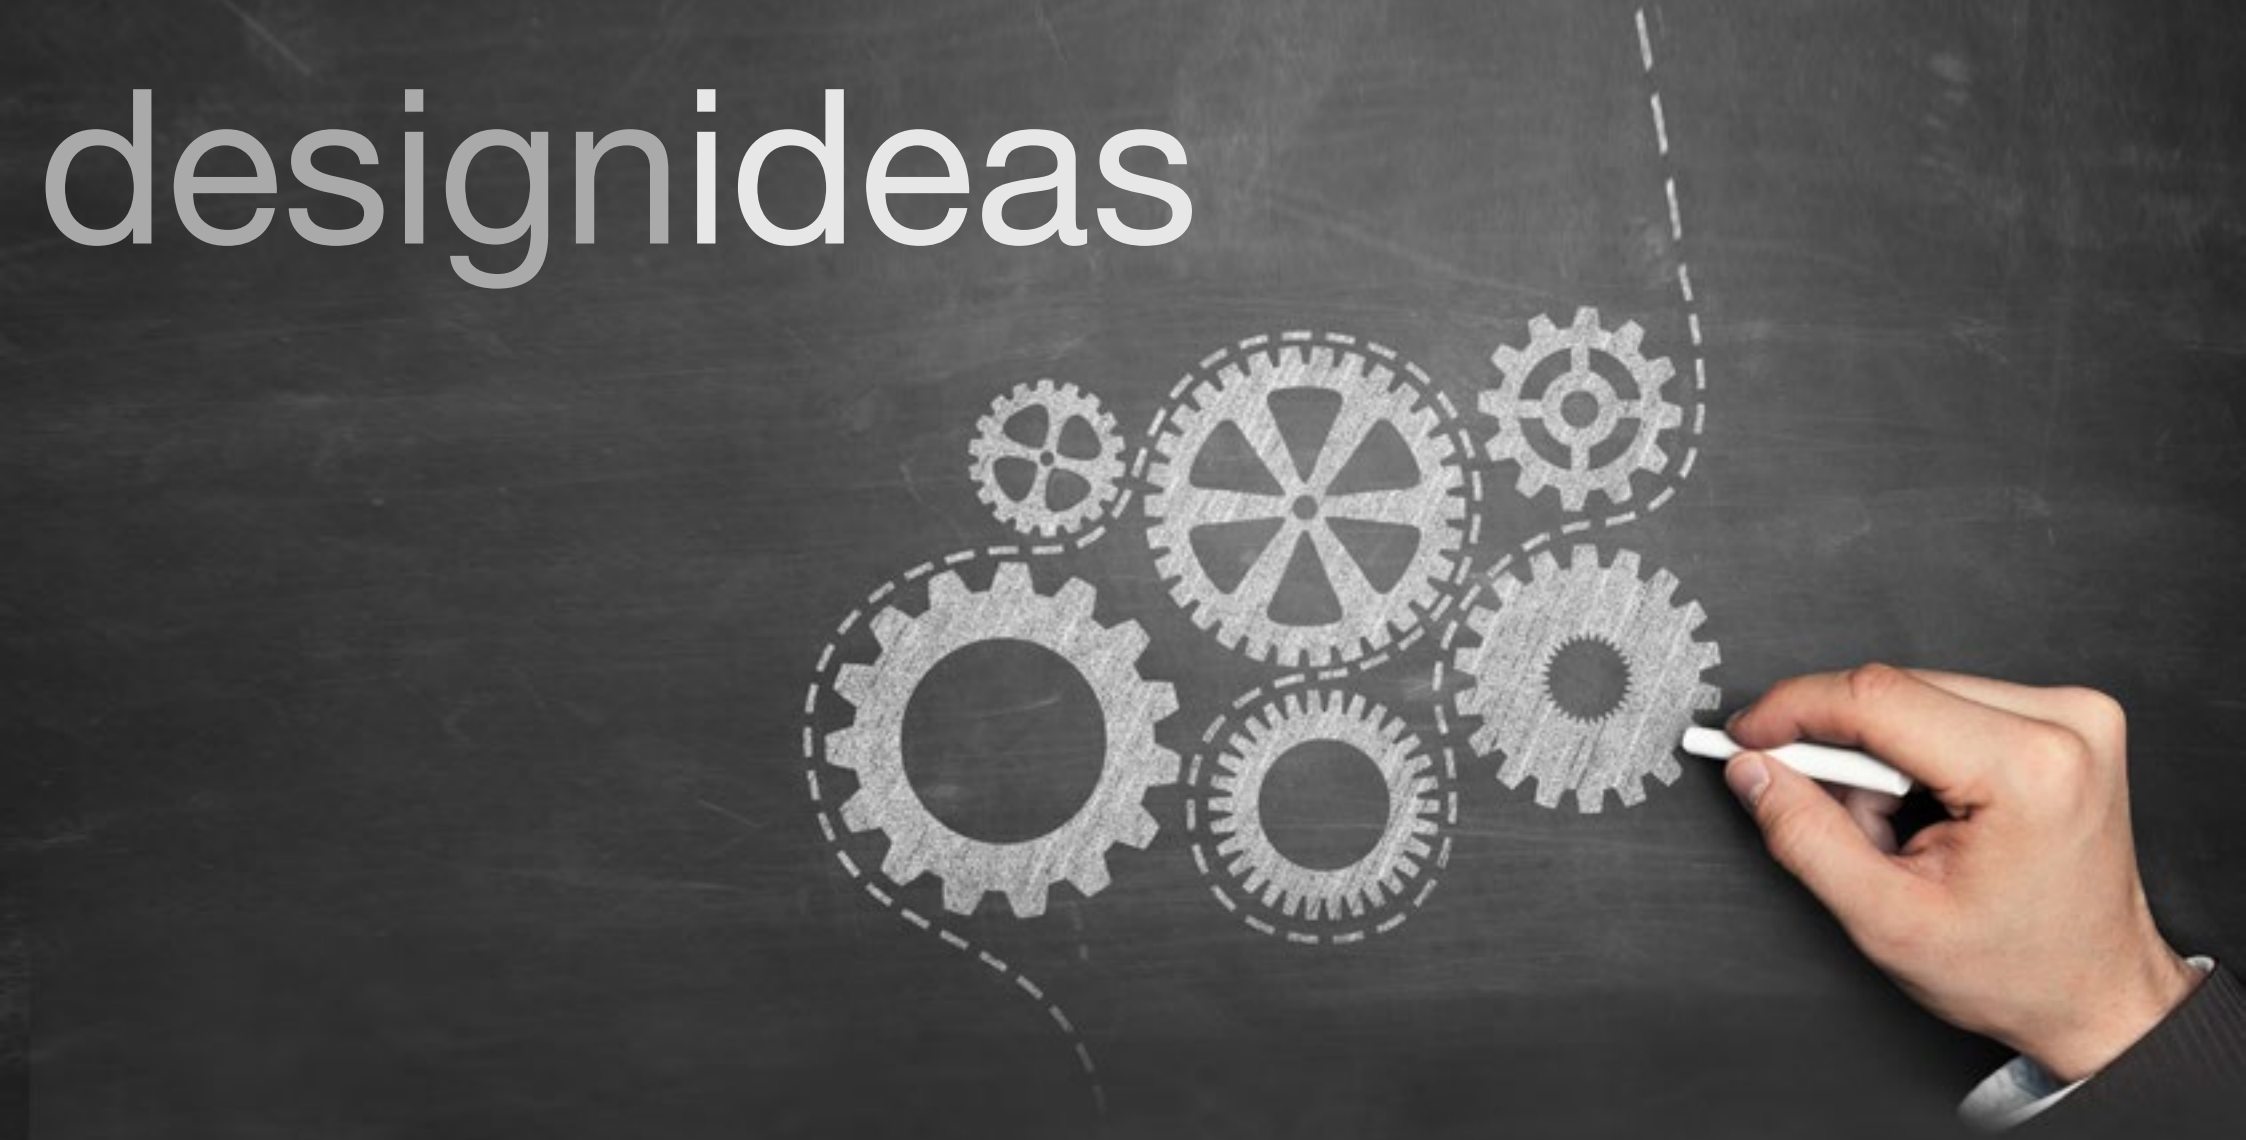

- New slope compensation method stabilizes switchers

New slope compensation method stabilizes switchers

By Francesc Casanellas

In switch-mode power converters, peak current control is very popular because of its inherent current limitation and ease of control. However, if the duty cycle is higher than 50%, there is an instability problem.

Some background: The up-slope of the current is $dI/dt = V_{CC} / L_P$, where V_{CC} is the supply voltage and L_P the inductance of the transformer or output inductor.

The down slope is $dI/dt = V_R / L_P$, where V_R is the reflected secondary voltage to the primary = $(V_O + V_F) \times N_P / N_S$. So the up slope depends on the input voltage and the down slope is constant.

The duty cycle D is: $t_{ON} / T = 1 / (1 + V_{CC}/V_R)$.

The following examples assume a flyback converter, but a buck or forward converter has the same problem.

In Figure 1, $D < 0.5$, that is, $V_{CC} > V_R$. The black waveform is the theoretical current in the inductance (inductor or transformer primary). If there is a small

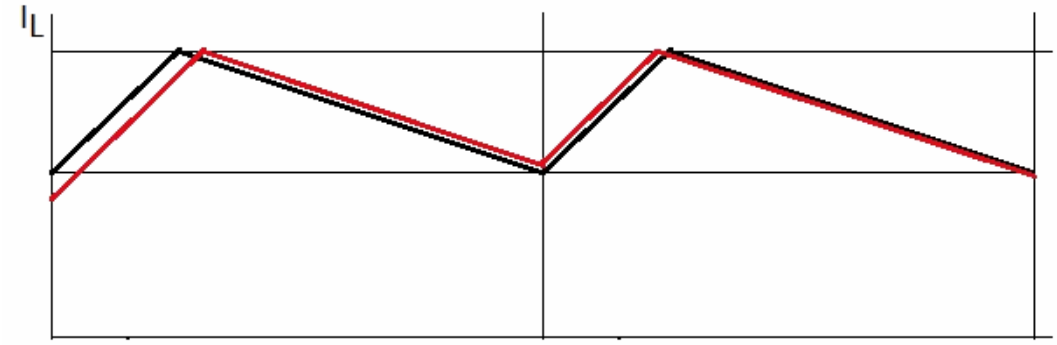

Figure 1. *Stable operation*

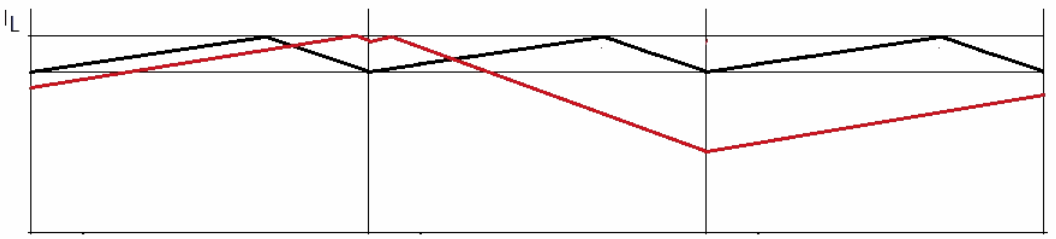

Figure 2. *Unstable operation*

perturbation of the current, as shown in the red waveform, the peak current limit corrects the error, as seen in the figure. The system is inherently stable.

In Figure 2, the same waveforms are shown when $V_{CC} < V_R$, or $D > 0.5$. Now the perturbation in the current (in red) causes a dramatic change in the duty cycle and the average current. The system is absolutely unstable. If we draw the waveforms for $D = 0.5$, it is easily seen that the current error remains the same in the following cycles; we are at the boundary of the instability.

To remedy the problem, instead of comparing the peak current with a fixed value, we compare it with a ramp, as shown in Figure 3. As we can see, there is a dramatic improvement: the stability is now as good as when the duty cycle is under 0.5. Figure 4 shows that if the reference ramp has the same slope as the down current ramp, the recovery occurs in a single cycle.

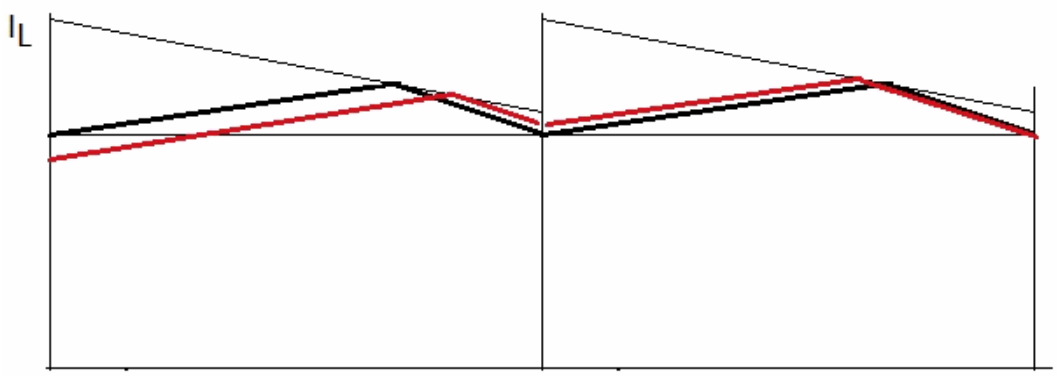

Figure 3. *Ramping current limit*

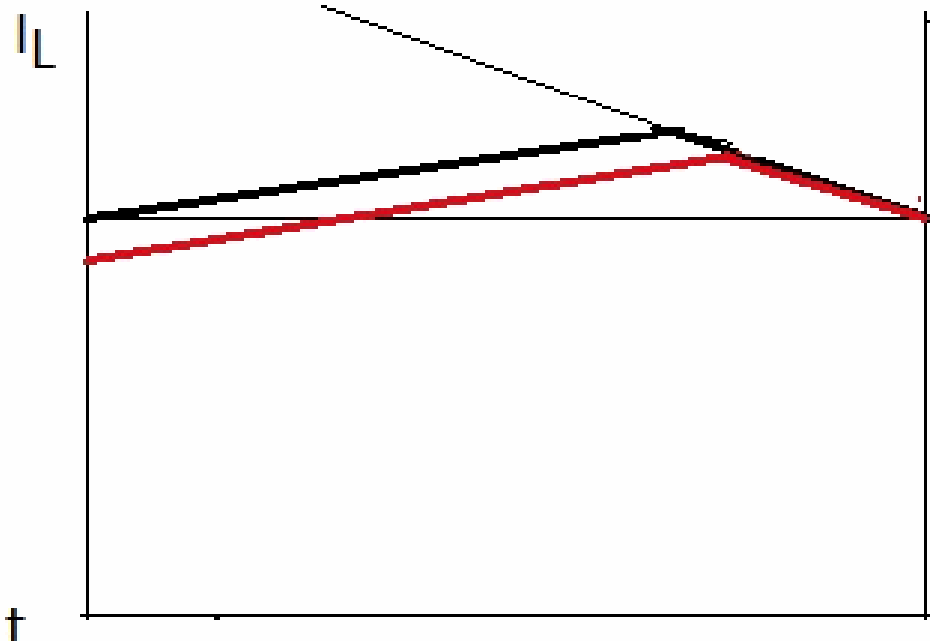

Figure 4. *Ramping current limit with the same slope as the inductive down current*

However, too much slope compensation makes the converter behave more like a voltage mode converter than a current mode one. If the slope of the reference ramp is 50% of the current slope, we are at the limit of instability. Thus, a practical slope for the reference ramp is between 50% and 100% of the current ramp; 75% is a good choice. This method of adding the reference ramp is called “slope compensation”.

The added ramp has benefits even with buck, forward, or flyback converters working at duty cycles lower that 50%. If the inductance is high and the current ripple is low, noise may cause false turn-offs. The added ramp stabilizes the converter, and a small amount may be enough.

A problem with the peak current limit is that the average current changes with the duty cycle. If the slope compensation is 50% it can be shown that the average current does not change with the duty cycle, and the current control loop is improved. However sub-harmonic oscillations can occur when the duty cycle approaches 100%.

Normally it is not possible to access the reference voltage of the IC controller. The simpler way is to add a ramp to the input current signal: a positive ramp there will have the same effect as a negative ramp in the reference voltage. The standard method is to use the ramp of the oscillator of the PWM controller, as shown in Figure 5.

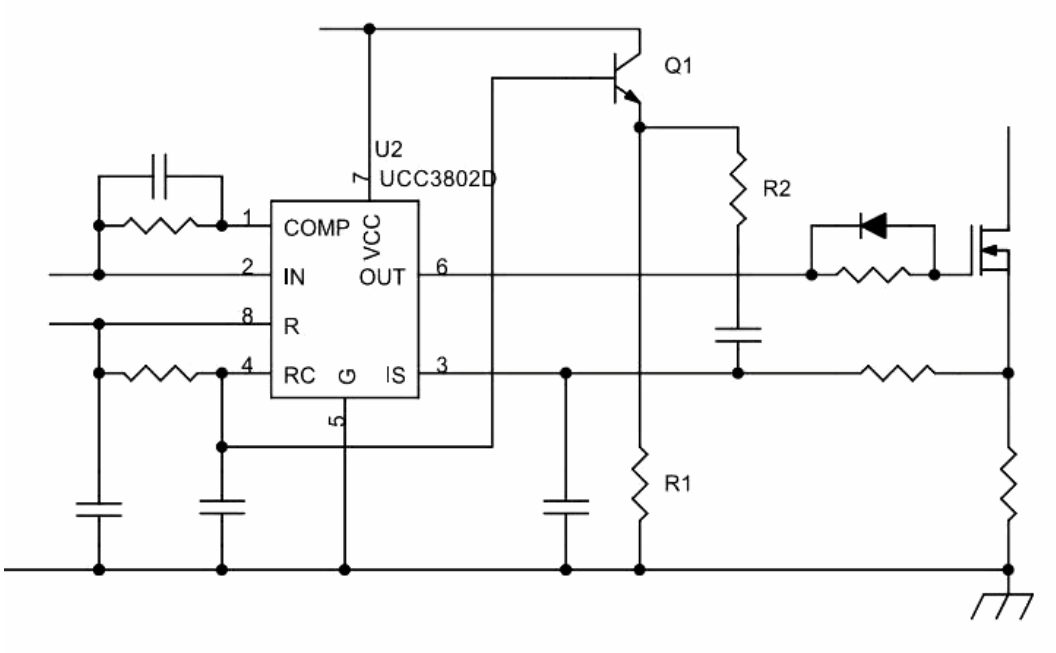

Figure 5. *Ramping current limit*

This system has two drawbacks:

- Not all the controllers have the oscillator ramp accessible.
- The value of R1 has to be quite low (R1 « R2, e.g., R1 = 0.1 × R2), so in spite of Q1 buffering the resistor, the oscillator circuit is loaded and the frequency can be impaired.

The Design Idea in Figure 6 is free of these problems. It works with any controller without relying on its oscillator circuit.

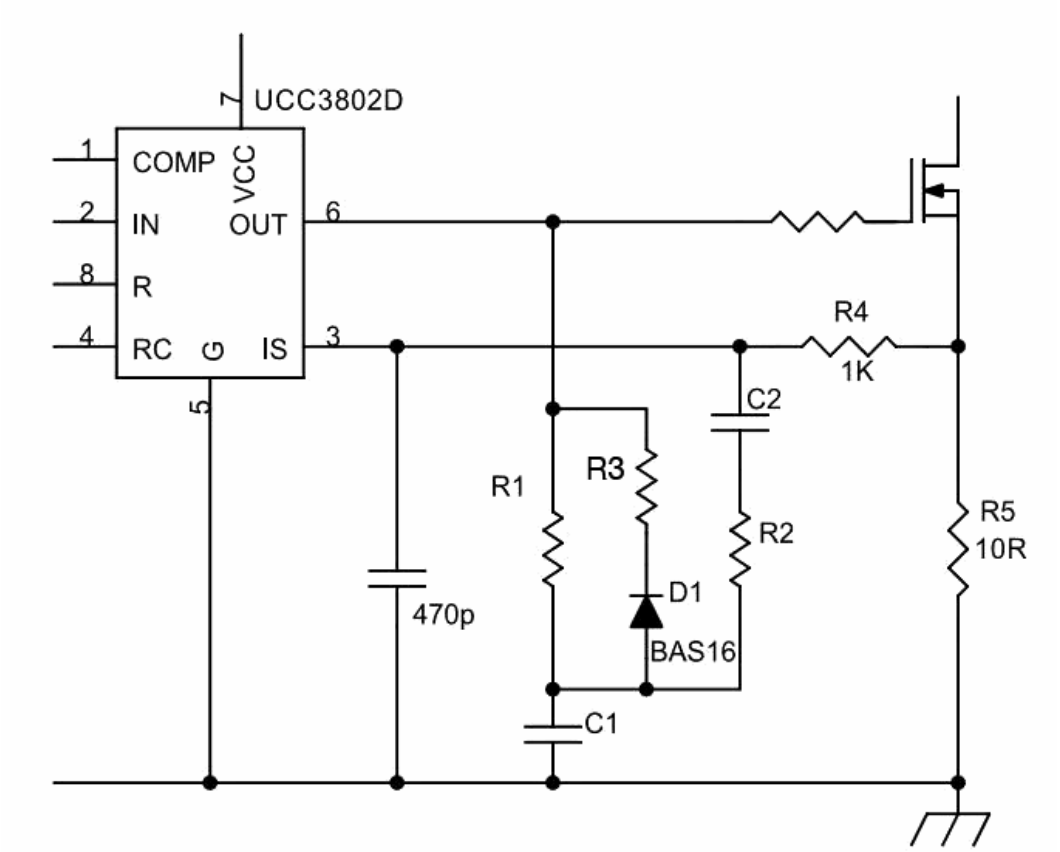

Figure 6. *Slope compensation that works for any controller*

When the output to the gate is high, the ramp goes up as R1 charges C1. When the gate output goes down, C1 is discharged through D1 & R3. The amount of slope compensation is set by R2.

A practical example will help to understand the circuit and calculate the values. The example is a 10W 12V continuous mode flyback converter. It has to work from 135 to 390 VDC input.

Primary inductance is 33 mH, I_{max} = 0.1A, so R5 = 10Ω for a 1V IS threshold.

The reflected secondary voltage to the primary is $V_R = (V_O + V_F) \times N_P / N_S$ = (12V + 0.6V) × N_P / N_S = 200V (turns ratio is 16:1). Switching frequency = 100 kHz (T = 10 µsec).

To get a fairly linear ramp, its maximum voltage can be selected to be 1/3 V_{CC}; that is, if V_{CC} = 12V, a reasonable peak voltage is 4V. Then the ramp amplitude is 4V - 0.6V = 3.4V.

Calculate the maximum duty cycle:

$D_{max} = 1 / (1 + V_{CC(min)} / V_R) = 1 / (1+140V/200V) = 0.6$
$t_{ON(max)}$ = 10 µsec × 0.6 = 6 µsec

The slope of the ramp:
$(dV/dt)_{ramp}$ = 3.4V / 6 µsec = 567×10³ V/sec

The down slope of the primary current:
dI/dt = 200V / 33 mH = 6×10³ A/sec

The slope of the voltage in R5:
$(dV/dt)_{shunt}$ = dI/dt × R5 = 60×10³ V/s

Calculate R2 for 75% slope compensation:

$R2 = R4 \times (dV/dt)_{ramp} / ((dV/dt)_{shunt} \times 0.75)$
= 1 kΩ × 567×10³ V/s / (60×10³ V/sec × 0.75) = 12.6 kΩ

The next step is to find the values for R1 and C1, with R1 « R2. We have to use the charging capacitor equation:

$t = RC \ln((V_{CC} - V_1)/(V_{CC} - V_2))$ to get suitable values for R1 and C1. In the example, t = 6 µsec, V_{CC} = 12V, V_1 = 0.6V, V_2 = 4V, so the result is RC = 17 µsec.

A good choice is C1 = 22 nF and R1 = 750Ω.

The discharge resistor R3 can be as small as possible while keeping the peak D1 current within its limits; R3 × C1 « t_{OFF}. In our example D1 is a BAS16 and R3 = 47Ω, t_{OFF} = 4 µsec, R3 × C1 = 1 µsec.
C2's reactance has to be much lower than R2; C2 = C1 is a convenient choice.

About the Author

Francesc Casanellas is a Chartered Engineer, Senior member IEEE, member IET. In a career that began in 1968 he has worked for a number of concerns in several European countries, prior to operating as Consulting Engineer since 2010; he has designed a very wide variety of power-oriented circuits from LED lighting to high-power motor drives.

productroundup

productroundup

Dual 375 nA op amp with independent shutdowns

Despite the very many op amps available on the market, configurable-mixed-signal-IC maker Silego configured one more; an amplifier that among other uses is specifically suited as a high performance front end for Silego's GPAK Mixed signal array products. Achieving low power applications specifications is possible in the "always on" mode, with 375 nA typical current consumption per channel when active. SLG88101 achieves a gain-bandwidth product of 18 kHz typical, with an open loop gain over 100 dB and voltage offset of 400 µV typical.

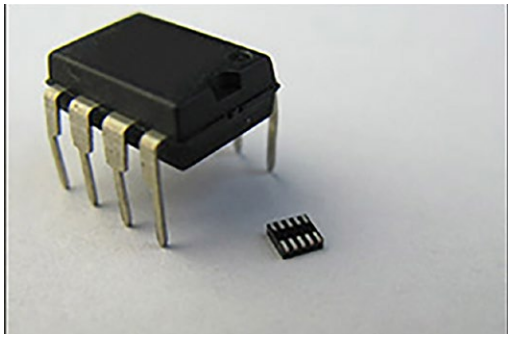

Complete article, here

GaN microwave transistor rated for 30W & 28V

Wolfspeed (part of Cree), maker of GaN-on-SiC high electron mobility transistors (HEMTs) and monolithic microwave integrated circuits (MMICs) has added a 28V, 30W GaN HEMT bare die device. Designed for up to 8 GHz operation, the HEMT die exhibits 12 dB typical small signal gain at 8 GHz, 17 dB typical small signal gain at 4 GHz, and 30W typical PSAT. This device also delivers higher breakdown voltage, higher temperature operation, higher efficiency, higher thermal conductivity, higher power density, and wider bandwidths than Si and GaAs transistors.

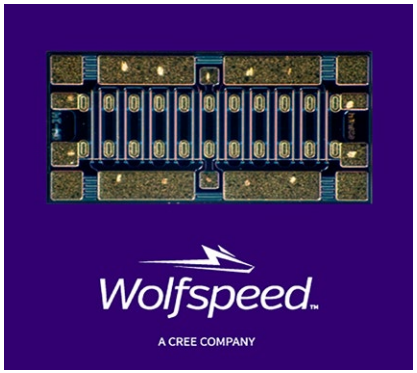

Complete article, here

Most accurate 7½-digit DMM in PXIe format, from NI

National Instruments' PXIe-4081 is a 7½-digit high-performance DMM and 1.8 Msample/sec isolated digitiser. NI positions it as the first PXI Express DMM available. The NI PXIe-4081 is claimed as the most accurate 7½-digit DMM, with 15 ppm accuracy for DC voltage measurements up to two years after calibration. It is capable of voltage measurements from nanovolts to one kilovolt and resistance measurements from microOhms to GigaOhms. A solid-state current shunt configuration offers eight DC current ranges from 1 µA to 3A and six AC rms current ranges from 100 µA to 3A.

Complete article, here

Fractional rate resampler IP core for FPGA implementation

RFEL's Fractional Rate Resampler IP core is now offered as a stand-alone IP core for FPGAs. The core enables input data rates to be modified to manage systems with multiple clock domains, unifying all signal paths onto a single clock domain. Arbitrary input data rates can be matched to support the data rates required for following algorithms such as DEMODS or CODECs. Equally, the core can be used in post-filtering applications so that data rates can be optimally set to match the data rate to the output bandwidth. This suits the core to Digital Signal Process systems development, harmonising over multiple clock domains, clock domain crossing, and algorithm integration that are particularly useful for COMINT, SIGINT, Electronic Warfare (EW), radar, sonar and similar security and surveillance applications. The architecture uses an Interpolator followed by a Low Pass Filter and a final Decimator stage, leading to the output.

Complete article, here

productroundup

Raspberry Pi distributors add uprated camera modules

Competing distributors RS Components and Farnell element 14 have both announced the addition of high-resolution camera modules to their offerings of Raspberry Pi peripherals. The module integrates the high-quality and high-sensitivity eight-megapixel Sony IMX219PQ CMOS image sensor, which offers high-speed video imaging including support for 1080p, 720p60 and VGA90 video modes. The device integrates with the image sensor pipeline on Raspberry Pi, which provides advanced de-noising, distortion and lens-shading correction, automatic gain control (AGC) and white-balance (AWB).

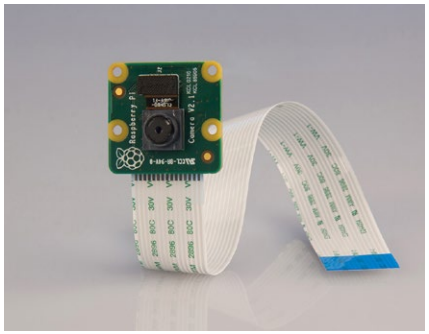

Complete article, here

GaN CMOS power switch delivers 40A to RF PAs

Specialist distributor RFMW has the 365CT000, 40A pulsed switch from Xsystor, Inc., that offers easy integration of GaN amplifiers in high speed, pulsed power systems. These CMOS, complementary MOSFET switches have clocked speeds of under 200nsec for rise and fall times. Allowing source and drain voltages from 28 to 80V, the Xsystor switch is compatible with standard and high voltage GaN amplifiers and transistors. Operating temperature is specified for -40 to 85C but will withstand 175C with derated voltage and current capacity.

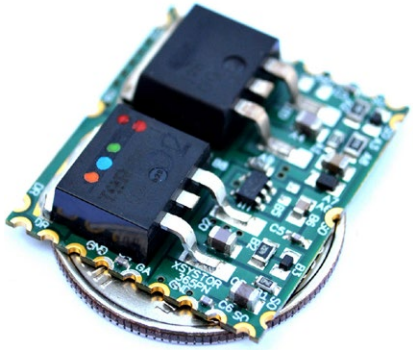

Complete article, here

iPhone/Android app for Tek/Keithley source measure units

Promising the fastest, easiest way to perform current-voltage (I-V) characterization, Tektronix has written the Keithley IVy app for iPhone and iPad, to work with the 2600B source measure unit(SMU) instruments. An updated version of the app with support for wireless instrument connection has also been released for Android devices on Google Play. IVy gives simultaneous control over source levels and monitoring of test results. This makes it easy to visually analyze I-V characteristics, DUT stability, response time, or drift in a circuit using normal touch zooming and scrolling.

Complete article, here

Multi-sensor IoT board hosts u-blox' GPS/GNSS ICs

u-blox' (Thalwil, Switzerland) EVA-7M GPS/GNSS receiver module has been embedded in the miniature IoT development board LoRa-ONE by Sodaq, now running a campaign on Kickstarter. The product is a 32-bit Arduino compatible board equipped with Low-Power WAN. The u-blox EVA-7M on the board enables a quick positioning fix. The board is equipped with an accelerometer allowing the GPS to safely switch off until the device detects movement, in order to minimize energy usage. A magnetometer identifies magnetic fields, while 14 I/O pins are available.

Complete article, here

productroundup

Software dev. kit for latest Google Eddystone open beacon

An open beacon format from Google, Eddystone-EID is designed to protect user privacy and security in sensitive applications by preventing unauthorized access and guarding against various potential types of malicious attack. The Nordic Semiconductor 'nRF5 SDK for Eddystone' allows the immediate development of Eddystone-EID beacons to provide real-world context to users in proximity-based beacon applications. The Nordic nRF5 SDK for Eddystone also features a GATT configuration that allows Eddystone beacons to be configured from a smartphone. The SDK is designed for Nordic's latest nRF52832 / nRF52 Series Bluetooth Smart Systems-on-Chips (SoCs). A future upgrade will also allow existing Nordic nRF51 Series or nRF52832-based Eddystone beacons to be updated to the latest Eddystone-EID secure capabilities via a straightforward over-the-air firmware update.

Complete article, here

SoC programmable hosted on DSP system-on-module

A collaboration between Microsemi and Solectrix (Nuremberg, Germany) aims to exploit the capabilities of Microsemi's FPGA-based SoC to enable faster time to market for imaging and video processing systems. Solectrix' SxoM-SF2 is a high performance, low power, secure and compact digital signal processing (DSP) system-on-module (SOM) solution based on Microsemi's SmartFusion2 system-on-chip (SoC) field programmable gate array (FPGA) devices. The device functions as a modular system core for high performance DSP applications and uses a carrier board with a PCIe interface.

Complete article, here

Wide input voltage buck converter for USB type-C transceivers

Maxim Integrated's MAX77596 allows always-on operation; designers of USB Type-C small form factor devices can now create always-on rails with the ultra-low quiescent current converter. Operating from a supply as low as 3.5V and as high as 24V, the MAX77596 step-down converter can regulate from traditional 5V USB power as well as the 20V upper-end of the PD range. The converter offers fixed 3.3V and 5V-output versions, and an adjustable version that allows users to program the output voltage between 1V and 10V with a resistor-divider. The converter can support up to 300 mA DC loads.

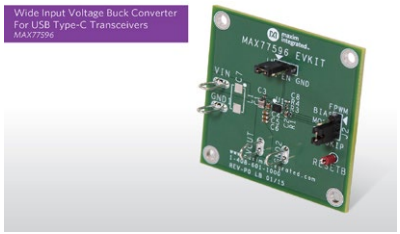

Complete article, here

MCU has low-leakage transimpedance amplifier + FRAM

This MCU is positioned by TI as the only MCU with an integrated low-leakage transimpedance amplifier (TIA) consuming 50 pA current. The microcontroller offers 20 times lower leakage than alternative voltage and current sensing solutions and provides the configurability of analogue and memory technology without sacrificing battery life or board space. Analogue integration includes ADC, op amp, comparators and TIA. This solution also integrates Ferroelectric Random Access Memory (FRAM) technology and eliminates the need for an onboard crystal in a single 3.5 x 4 mm package.

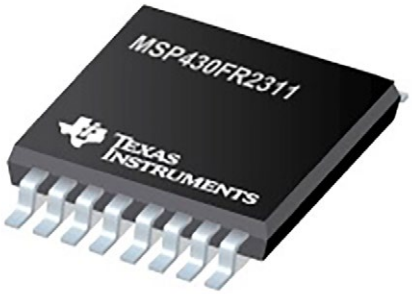

Complete article, here

productroundup

Multi-channel temp. sensor IC for diode-based readings

Measurements of lower-temperature, outdoor and industrial applications achieve greater accuracy with Microchip's MCP990X multi-channel temp sensor family; remote diode measurement with resistance error correction enables accurate readings of up to 0.5m away from the IC. Designed for monitoring of up to four channels in cold, outdoor and industrial low-temperature applications ranging from freezers and refrigerators to base stations and remote radio units, with ±1°C maximum accuracy for both external and internal diode temperatures from -40°C to +65°C, these ICs provide a low-cost, highly flexible solution where precision is needed at lower temperatures.

Complete article, here

Small-footprint Bluetooth LE module outputs +3dBm

Silicon Labs' pre-certified Blue Gecko BGM113 module, with embedded software stack, promises developers an optimal combination of small footprint, ease of use and low-energy technology for short-range wireless applications. Blue Gecko BGM113 module provides a small-form-factor Bluetooth 4.1-compliant connectivity solution with 3 dBm output power for applications typically requiring up to a 50 metre range, typical applications including smartphone accessories, wearable sports and fitness products, wireless locks and point-of-sale devices. The module combines a 2.4 GHz Blue Gecko wireless SoC and a high-efficiency chip antenna.

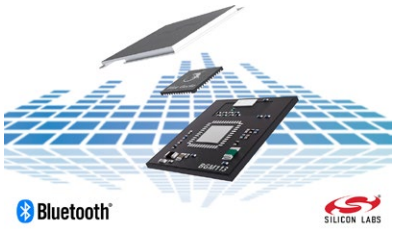

Complete article, here

Spectral and power measurements to 70 GHz on IEEE802.11ad

Tektronix claims the most accurate solution for characterising and debugging IEEE 802.11ad transmitter PHY performance. The new test package, based on the 70 GHz DPO70000SX Series ATI Oscilloscope and SignalVu Option SV30, delivers best residual error vector magnitude (EVM) measurements along with easy setup for complete characterisation and robust debugging capabilities. Using the DPO70000SX oscilloscope's low-noise ATI acquisition technology, characterisation can be performed with 20% greater accuracy than alternative offerings with an EVM that is specified at 2.5% (-32.0 dB); live triggering is supported on signals up to 70 GHz.

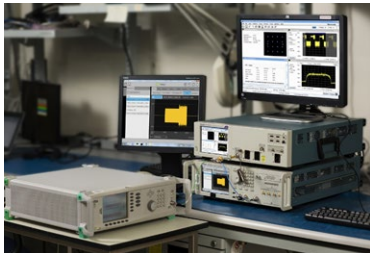

Complete article, here

60V, 4A step-down inverting output regulator

LTC7149 is a high efficiency, 60V, 4A synchronous buck regulator for negative output voltage supplies. It incorporates a current mode architecture with user-selectable internal or external loop compensation and an output voltage that is programmed with a single external resistor. The regulator operates from an input voltage range of 3.4V to 60V and provides an adjustable output voltage from 0V to -28V, while delivering up to 4A of continuous output current. The LTC7149's design includes an accurate internally generated 50 µA current source on the ISET pin that enables it to deliver outputs as low as -28V.

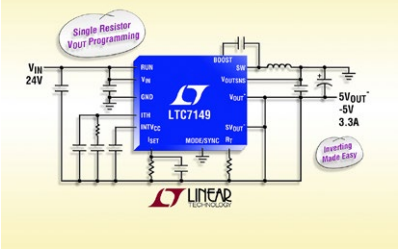

Complete article, here

DEBUGGING THE CORTEX-M MCU

BY JACOB BENINGO

Debugging embedded software is my least favourite activity, but unfortunately it's a necessary evil. Thankfully, advances in technology and tool chain innovation have resulted in a plethora of techniques to drastically speed up the debugging process. Let's examine a few of them, starting from traditional debugging by break-point through the more advanced instrument trace techniques.

Technique no. 1 – Traditional break-point debugging

Every developer is familiar with the traditional debugging technique of setting breakpoints, executing code, and then stepping through the code while monitoring registers and variable values.

Break-point debugging is a technique that I see used more than any other technique. This is disheartening, however, because break-point debugging is inefficient and generally yields sub-optimal results.

So, then, why are breakpoints used so often? The main reason seems to be because break-points are easy to use, are readily accessible, and developers are optimistic that breakpoints are the right tool for the job. That optimism needs tempering. Breakpoints have the potential to disrupt the systems real-time performance and can suck a developer into a black hole of endlessly single stepping through code blindly looking for a solution to their problem.

Technique no. 2 – IDE value graphing

Nearly all modern debuggers and IDEs now allow a developer to monitor the value stored inside a memory location such as a variable. A developer selects the memory location and the value refresh rate, and then starts the debug session. Some IDEs have the ability to monitor the value built into the IDE while others rely on using an external software.

Value monitoring can be very useful, but it is far more valuable if the data being monitored is tied to a graphical representation. Graphing the data values in real-time can be extremely useful for discovering unexpected changes or to verify that a particular waveform is generated. Take, for example, a three phase BLDC motor. A developer may want to monitor the current and voltage on each motor leg, which need very specific waveforms to be generated in order to drive the motor. Plotting the voltage and current on each motor leg can allow a developer to visualize what is happening in real-time.

Technique no. 3 – Reroute printf to SWO

On higher-end ARM Cortex-M parts, such as the M3/M4, developers are provided with additional debugging capability known as the Serial Wire Viewer (SWV). The SWV comprises the standard serial wire debugger in addition to a Serial Wire Output (SWO). The SWO can be used to do really cool things such as retrieving the program counter, event counter, and trace data, to name a few. Developers can customize the information that they would like to have transmitted over the SWO.

Many developers set up printf in order to get debugging information from their embedded system. Rather than using serial pins on the microcontroller, a developer can use SWO to reroute printf information through the debugger. Using the debugger in this way can save the need for a dedicated serial interface, eliminate the development time for the UART or USB device, and is far more efficient. Overhead that would have originally been used in the application is now off-loaded through the SWO and debugging hardware, which saves precious clock cycles that would have otherwise been used by the application code.

Figure 1. *A screenshot from Percepio's Tracealyzer.*

Technique no.4 – RTOS tracing

Trying to peer through the veil at what an RTOS is doing can be quite challenging. Developers don't want to disturb the real-time system performance but still need some method to understand the system's behaviour. Blinky LEDs have often been the go-to trick, but more recently trace tools have added an amazing technique to the developer's toolbox. Free and commercially available RTOS trace tools exist such as TraceX (Express Logic), SystemView (Segger) and Tracealyzer (Percepio) to name a few.

Trace tools allow a developer to analyze when the RTOS is idle and when each task is entered and exited. Developers can monitor for system exceptions, response time, execution time, and many other critical details needed to properly develop an embedded system. The coolest feature available in RTOS trace tools is their ability to graph what is happening in the system. Reviewing and monitoring timing diagrams in real-time or in the recorded log can help give developers a confidence level that the system is working as expected or help them to discover small issues that would have otherwise taken a lot of time to discover.

Technique no.5 – Use instruction trace technology (ETM/ETB/ETM)

Sometimes developers face debugging problems that are just at the lowest levels imaginable within the processor. Trace technologies exist that can monitor individual instructions the processor executes. Such low level tracing can be useful for monitoring branch coverage when testing and validating software. The debugger tools used for instruction trace are different from those a developer would use for Serial Wire Viewing and generally cost a little bit more.

Closing thoughts

Debugging tools and technologies have been rapidly evolving over the past few years, especially for higher end microcontrollers. Engineers in general are visual creatures and tool vendors are finding ways to reveal in visually stimulating ways what exactly is happening within a real-time system. Setting up debugging tools can require some upfront time but the potential to spend less time debugging and more time designing is well worth the time investment. At a minimum, developers should become familiar with the different debugging tools and capabilities available to them so that when problems arise and the system needs to be debugged, they can select the right tools exist to get the job done. What other techniques do you use to help engineers debug their systems faster and more efficiently?

Jacob Beningo *is principal consultant at Beningo Engineering, an embedded software consulting company. Jacob has experience developing, reviewing and critiquing drivers, frameworks and application code for companies requiring robust and scalable firmware. Jacob is actively involved in improving the general understanding of embedded software development through workshops, webinars and blogging. Feel free to contact him at jacob@beningo.com, at his website www.beningo.com, and sign-up for his monthly Embedded Bytes Newsletter.*

EDN europe

EDN-EUROPE is published 11 times in 2016 by
European Business Press SA,
533 Chaussée de Louvain, 1380 Lasne, 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

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

european business press

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

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

Andrea Rancati
Italy
+02 70 30 00 88
arancati@rancatinet.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

Keita Sato
Japan
+81-3-6824-9386
ksato@mx.itmedia.co.jp

Laura Chen
Taiwan
+886 2 2759 1366 #305
laura.chen@ubm.com

Brandon Smith
China
+86755 33248168
Brandon.Smith@ubm.com