EDN europe

contents
JULY-AUGUST 2014

Connected-sensor experiments made easy

The image on this month's front cover is the bare-board PCB panel for realyr's WunderBar Internet-of-Things experimenter's kit. The PCB is designed to be panelised in the normal way but the kit itself is delivered as an undivided panelised assembly. See the story on page 6 and also EDN.comment on page 4.

FEATUREARTICLES

CONTACTS

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

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

Patrick Mannion
Brand Director EDN Worldwide

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

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

ART MANAGER
Jean-Paul Speliers

ACCOUNTING
Ricardo Pinto Ferreira

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
+39-02-284 6716
info@silvera.it

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

USA & Canada

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

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

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

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

Asia

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

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

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

WHEN DID HARDWARE GET HARD?

I'm writing in this column about a comment made during a presentation introducing the internet-of-things, connected-sensor experimentation kit, WunderBar. There's coverage of the product in the following pages but, with apologies for some repetition, I'll outline again what it's about. The start-up company that designed it aims to make its living in the long term by providing, in effect, the service environment in which IoT devices and applications can function, grow and flourish. A key target customer base for such an offering is the group that has come to be known as "makers".

Makers may have software skills but likely aren't from the established embedded-software writing fraternity. They might have some hardware expertise, but very possibly, none at all – or, at least, not in the hardware skills that would normally characterise EEs. If anything does define them, it is familiarity with the possibilities of connected systems, and with the "app" model of providing future products and services. Maker Fairs (yes, I'm avoiding 'faire' and even 'fayre' but all spellings are out there) are already a major phenomenon, bringing, "hobbyists, hackers, crafters, coders, DIYers and inventors..." together with professional engineers – one of the fascinating aspects of this development is that it becomes increasingly unclear just what constitutes a professional practitioner in this space.

relayr – the start-up that produced the WunderBar kit – designed it to make things easy for that community to get ideas based around connected-sensor setups, up-and-running. So, you get one-each of frequently-used sensor types, that talk to a hub module, that hosts an ARM-core microcontroller, and a WiFi module to connect to the wider world. The production engineering is impressive, some of which you can see in the bare-board image on page 1 of this issue. The PCBs are loaded in a conventional multi-panel construction; each kit is a sub-arrangement of divisible panels that come as a single PCB in the kit. They snap apart, but prior to separating them, vias along the snap-lines route power and data to every sensor module in the kit, so you can bring them up as a single unit. Then, snap off the sensor of your choice (temperature, light, humidity, etc), give it a coin cell, and off it goes, talking to the hub over Bluetooth LE.

The line in the presentation that made me pause was to the effect that for the target market – the makers – "hardware is hard". This reveals the distinction between previous generations of experimenters, especially those at the cross-over point between professional and hobbyist, and this new category. Hardware prototyping, constructing a hand-made PCB and all the skills that that entails would previously been second nature to the experimenter fraternity: if anything, it would have been the software that would have been regarded as problematic.

To some extent the situation has reversed. The "app" phenomenon, to cite just one, has highlighted coding skills: and no promotion of STEM subjects to the young is complete without a technologically-illiterate politician pronouncing that more of our children must learn to code. With tools hosted on multiple platforms, immediate access to environments in which software can be developed gets easier, whereas the reverse is true of the hardware. With end-user markets targeting the tiny, the battery-powered, the wearable and the energy-harvesting, it is natural for semiconductor vendors to employ the smallest surface-mount packaging. Very few are offered bonded-out in any prototyping-friendly format that you might hand-solder, so constructing any remotely-realistic prototype means a full PCB design and a limited production run at an assembly house. Every MCU has its EVM, of course, but adapting those to an embryonic tiny-connected-thing notion may not be appropriate.

None of which is new; the same reasoning, in part, lies behind the runaway success of Raspberry Pi and, indeed, the many "pro" modular prototyping vehicles that have emerged in recent years. Nevertheless, and in serving this professional/experimenter cross-over market, Conrad Business Services (who will distribute the relayr product) also reports an up-tick in sales of soldering irons – so hardware construction has not gone away! If there is a slightly surprising aspect to the entire scene, it is that some form of programmable (analog/MCU/RF) hardware platform has not emerged to make a bid for dominance.

What may be new and also bound up with the maker approach, is developers who want to play in the connected-device space but who have no deep background in either the hardware or software worlds. They may come from the service side, or they may be individuals or companies who have little more than an outline idea for a product or a connected application. With ready-built and programmable prototyping vehicles, anyone can play.

Drawing on a portfolio of many designs already completed, UK-based Cambridge Consultants has created a package offering for clients who have exactly that; an idea, a concept, but no experience in taking a connected-device offering from outline to actual project. The consultancy has, in effect, parameterised the connected-device application space. In a newly-released programme, CCL can dip into its knowledge base and tell you that your concept will require this-much processing power, bracketing the processor you'll need; that-much data per second, setting its connectivity requirements... and with a given operating cycle, estimating the life you'll get from a particular battery size. With layouts already done for many practical combinations of hardware components, the consultancy can 3D-print or laser-cut a prototype housing and put "flesh on the bones" of a connected-device concept, literally within hours.

Does the emergence of a new model of product development mean that there is no place for the established electronics industry in this product space? That's unlikely, if for no other reason than (as we are repeatedly informed) there will be many billions of the things in our world before long; there should be work enough for players from all sides. What it does imply, however, is that large corporations, and start-ups in incubator units, will have to be equally nimble and able to take "concept to street" in ever-shorter time scales. That fact, in itself, might be the source of IoT problems in the not-too-distant future – but that's a thought for another day.

Keep abreast of the latest industry news with our newsletters

european business press

EDN europe

MICROWAVE engineering europe

EE Times europe LEDLighting

EE Times europe AUTOMOTIVE

EE Times europe ANALOG

EE Times europe TEST & MEASUREMENT

EE Times europe POWER MANAGEMENT

EE Times europe

EE Times europe Technical papers

# www.electronics-eetimes.com/newsletters

pulse

relayr's 'WunderBar' Internet of Things starter kit to be distributed by Conrad

Distributor Conrad Business Supplies has announced that it is working exclusively with Berlin based start-up company relayr (iThings4U GmbH) to support the development and launch of the Open Source 'Internet of Things' (IoT) starter kit known as 'WunderBar'.

relayr is a start-up that is focussed on enabling applications using connected devices, and providing the infrastructure that will allow many of the product concepts currently circulating as Internet-of-Things ideas, to get the real world and to function. Medium- and long-term, relayr's business model is to, in effect, 'commoditise' aspects such as the cloud services that will produce real and useful services from the IoT. The company is considering models similar to the "app" environment with revenue streams coming from small subscriptions or one-time payments, for IoT-based functions and services. It is developing a range of offerings such as its Open Sensor Cloud concept, which will be an environment where data from myriad smart devices can be collated and made useful by an embedded rules engine.

More immediately, relayr believes that there are many would-be IoT "application" developers for whom, "hardware is hard" - for this group, who don't have a problem in coding functions, the practicalities of actually getting sensor measurements connected to a context where they can do "connected-device" operations, is a barrier. Therefore, the company has designed

WunderBar.

The WunderBar IoT starter kit together with the relayr Open Sensor Cloud platform allows software application developers to quickly and easily begin working on wireless applications and prototype building based on data gathered from the physical world without needing to learn about hardware. The platform includes software development kits (SDKs) for iOS, Android and Node.js.
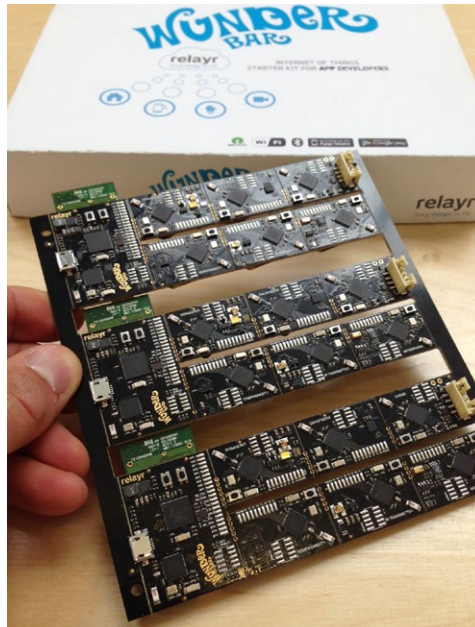
The "-Bar" part of the name alludes to a chocolate bar, in that the product comes as a single PCB that can be snapped into functional "bites". There's a host module with a Freescale ARM Cortex-M MCU ("deliberately overkill" in terms of its performance, relayr says, leaving headroom for application code) and a Gainspan WiFi module; the host and the six detachable sensor modules each carry a Nordic Semi Bluetooth LE chip. The PCB, out-of-the-box, is in effect a little panelised assembly that has yet to be separated; before separation, there is a power path to all of the modules. You power it up with a USB lead, give it Internet access and it updates and configures all the parts. Choose your module, snap it off (then, you need to give each sensor module a coin cell battery) and it connects via Bluetooth LE. More details in the full story, click right.

Complete article, here ▶

Ferro-electric memory (FRAM) MCUs enable ultra-low power operation

Texas Instruments says it is entering a "new era of ultra-low power" with its MSP430 FRAM microcontrollers; with a feature called EnergyTrace++ technology, they are designed from the outset to enable the lowest power microcontroller systems. MSP430FR59x/69x FRAM MCU families with the EnergyTrace++ real-time power profiler and debugger range from 32 to 128 kB embedded FRAM. These MSP430 MCUs are suitable for smart utility metering, wearable electronics, indus-

trial and remote sensors, energy harvesting, home automation, data acquisition systems, the Internet of Things (IoT) and many more applications that require ultra-low power consumption, flexible memory options and smart analogue integration.

TI's ultra-low-leakage (ULL) technology with embedded FRAM offers active power of 100 µA/MHz, accurate-RTC standby power of 450 nA and power performance across the tempera-

ture range from -40 to 85C. The FRAM MCUs include a variety of smart analogue peripherals, such as a differential input analogue-to-digital converter that consumes as little as 140 µA at 200 ksamples/sec and an enhanced scan interface for flow metering that can operate while the system is in standby, resulting in 10 times lower power. EnergyTrace++ technology enables developers to analyse power consumption down to 5 nA resolution in real-time for each peripheral. This new technology is now available for both MSP430FR59x and MSP430FR69x MCU families and with the new low-cost MSP430FR5969 LaunchPad development kit.

Capabilities claimed for the FRAM MCUs include; infinite endur-

MSP430™ FRAM MCUs usher in new era of ultra-low power

MSP430™ FR59x/FR69x FRAM Microcontrollers

TEXAS INSTRUMENTS

Designed from the ground up to save energy and enable a battery-free world

ance – the read/write speeds mean FRAM MCUs have more than 10 billion times more write-erase cycles than traditional non-volatile memory solutions – outlasting the product lifetime itself; and flexibility. FRAM has the unique ability to free developers from the traditional boundaries between code and data memory. Users no longer need to be confined to industry-standard flash-to-RAM ratios or pay a premium for increased RAM needs.

MSP430FR59x MCUs cost from $3.35 (1000). The MSP-EXP430FR5969 LaunchPad bundled with the Sharp 96 Memory LCD BoosterPack is available for $29.99. The MSP-FET is available for $115.

Complete article, here ▶

Raspberry Pi embedded compute module development kit

Distributor element14 has added to its development kit range; the Raspberry Pi Compute Development Kit is aimed at taking the power of the Raspberry Pi to embedded applications.

The Raspberry Pi Compute Module allows design engineers to use their own interface board that will host the compute module and deliver a smooth experience as they move from prototyping on the Pi through to shipping commercial product in volume.
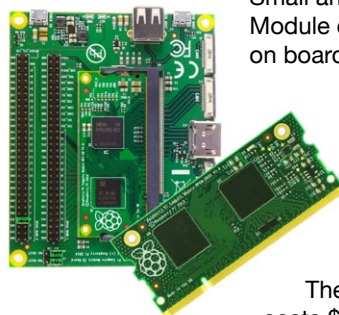
Small and powerful, SODIMM sized (6.5 x 3 cm), the Compute Module contains the BCM2835 chip with 512 MB RAM with an on board 4 GB eMMC Flash memory for booting the OS.

For flexible and rapid prototype development, the kit contains a fully functional IO board which supports an extensive GPIO and multiple connectors for rapid prototype development with access to all of the BCM2835 functionality.

The Compute Development Kit costs $200 and is available from Farnell element14 and CPC in Europe.

Complete article, here ▶

Bluetooth Smart protocol frees wearables from the smartphone

Nordic Semiconductor has a Bluetooth Smart protocol stack that enables sophisticated Bluetooth Smart wearable hub networks with wireless sensors that don't need smartphones. The S130 SoftDevice is a Bluetooth v4.1 compliant protocol stack and includes all Bluetooth Smart protocol layers up to and including GATT/GAP. The ultra low power (ULP) RF specialist's S130 SoftDevice allows the development of Nordic nRF51 Series SoC-based advanced wearable Bluetooth Smart hub network topologies such as smart watches with peripheral wireless sensors that don't always need a smartphone present to operate. The S130 SoftDevice is a concurrent multi-link Central and Peripheral Bluetooth Smart protocol stack for Nordic's nRF51 Series of System-on-Chips (SoCs), is Bluetooth 4.1 compliant, and paves the way for the next generation of Bluetooth Smart wearable hub products. A SoftDevice is Nordic's self-contained software stack for nRF51 Series SoCs that incorporates a unique separation of RF protocol and application code. Employing an architecture that cleanly separates the Bluetooth Smart, ANT or 2.4GHZ proprietary SoftDevice from the developer's application code removes

the need for the engineer to struggle with integration of their code as part of a vendor-imposed application development framework. This separation ensures the stack and application software operate independently, but communicate when necessary via event-driven APIs. The S130 SoftDevice supports multi-link central, peripheral, observer and broadcaster roles, GATT server and client, and event-driven, asynchronous and thread safe GATT/GAP and L2CAP APIs. The S130's ability to support concurrent multi-link central and peripheral roles makes it a suitable choice for Bluetooth Smarthubs that are not smartphones or tablets.

Complete article, here ▶

pulse
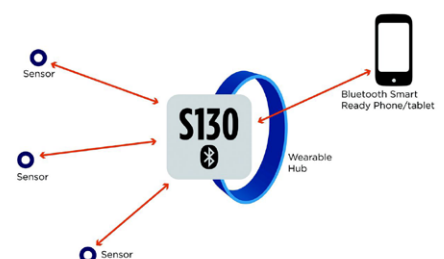
Nanopower buck-boost DC/DC converter with energy-harvesting charger

LTC3331 is a complete energy harvesting solution that delivers up to 50 mA of continuous output current to extend battery life when harvestable energy is available. A simple 10 mA shunt enables charging of a rechargeable battery with harvested energy while a low-battery disconnect function protects the battery from deep discharge. The LTC3331 requires only 200 nA of supply current from the battery when providing regulated power to the load from harvested energy and only 950 nA operating when powered from the battery under no-load conditions.

The device integrates a high voltage energy harvesting power supply, a battery charger, and a synchronous buck-boost DC/DC converter powered by a rechargeable battery, creating a single continuous regulated output for energy harvesting applications such as those in wireless sensor networks. The energy harvesting power supply, consisting of a full-wave bridge rectifier accommodating AC or DC inputs and a high efficiency buck converter, harvests energy from piezoelectric (AC), solar (DC), or magnetic (AC) sources.
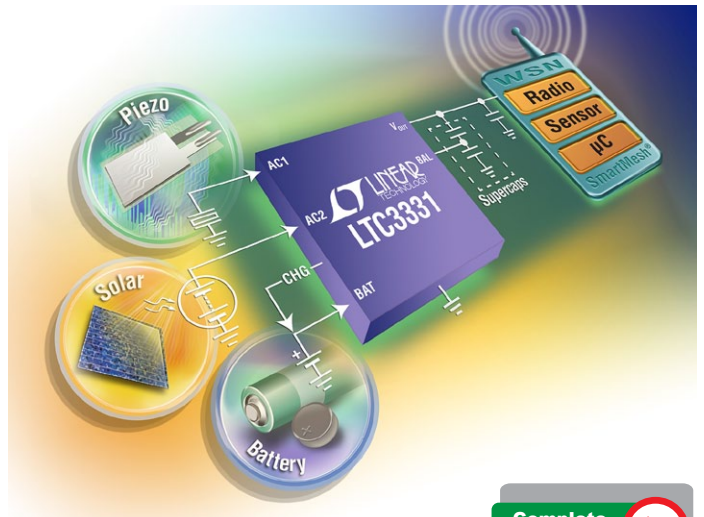
Complete article, here

TI's SimpleLink WiFi enables ubiquitous IP connection with "Internet on a chip"

Ease of use, and low power, are key attributes of Texas Instruments' SimpleLink Wi-Fi CC3100 and CC3200 devices, which comprise a single-chip, low-power Wi-Fi solution with built-in programmable microcontroller, designed for the IoT.

WiFi has been difficult to use, TI says, and offers the 3100 and 3200 to remedy that. Both contain a complete, integrated, WiFi function with the protocol entirely in ROM: the 3100 has a interface to an off-chip MCU – the 3200 has an integrated ARM Cortex-M4 MCU, fully-available for users to add their own code. In either case, you get what amounts to a standard API to the wireless IP connection, with Berkeley interface sockets and following the model that TI has established with SimpleLink. There is also an on-chip cryptography engine to establish secure connections to access points. All of this is in hardware, with very short power-up times (150 msec), enabling short wake-up-to-sleep cycles for low average power.

Features include: Lowest power consumption for battery operated devices with a low power radio and advanced low power modes; Flexibility to use any microcontroller (MCU) with the CC3100 solution or use the CC3200's integrated programmable ARM Cortex-M4 MCU, allowing customers to add their own code; Easy development for the IoT with quick connection, cloud support and on-chip Wi-Fi, Internet and robust security protocols, requiring no prior Wi-Fi experience to get a product connected; The ability to simply and securely connect their devices to Wi-Fi using a phone or tablet app or a web browser

with multiple provisioning options including SmartConfig Technology, WPS and AP mode.

Complete article, here

OpenCL software development kit targets FPGA accelerator boards

Altera and its board partners have a combined offering that they claim will accelerate development of high-performance computing, networking and SoC applications; an OpenCL development flow is capable of delivering "prototypes in minutes".

Altera says it has simplified a programmer's ability to accelerate algorithms in FPGAs, with its SDK for OpenCL version 14.0, that includes a rapid prototyping design flow that enables users to prototype designs in minutes on an FPGA accelerator board. Development of FPGA-based applications is assisted by reference designs, reference platforms and FPGA development boards that are supported by Altera's OpenCL solution. These reference platforms also streamline the development of custom FPGA accelerators to meet specific application requirements.

Altera offers a publicly available, OpenCL conformant software development kit (SDK) that allows programmers to develop algorithms with the C-based OpenCL language and harness the performance and power efficiencies of FPGAs. A rapid prototyping design flow included in the Altera SDK for OpenCL version 14.0 allows OpenCL kernel code to be emulated, debugged, optimised, profiled and re-compiled to a hardware implementation in minutes. The re-compiled kernels can be tested and run on an FPGA immediately, saving programmers weeks of development time.

Complete article, here ▶

LeCroy's WaveSurfer 3000 'scopes offer advanced user-interface

Teledyne LeCroy's WaveSurfer 3000 series of oscilloscopes features the company's MAUI user interface, previously seen on its highest-performance instruments. This range offers up to 500 MHz bandwidth with advanced measurement toolset, multi-instrument capabilities and a 10.1-in. touch-screen display.

The advanced user interface integrates a deep measurement toolset and multi-instrument capabilities with a 10.1-in. touch screen, that LeCroy says is the largest display and only touch screen in this class of oscilloscope. WaveSurfer 3000 oscilloscopes are available in bandwidths from 200 MHz to 500 MHz, with 10 Mpts/ch memory and up to 4 Gsamples/sec sample rate. Beyond traditional oscilloscope functions, the

WaveSurfer 3000 has a variety of multi-instrument capabilities including waveform generation with a built in function generator, protocol analysis with serial data trigger and decode, and logic analysis with a 16-channel mixed signal option.

MAUI is an advanced user interface that was designed for touch; all important oscilloscope controls, as well as positioning and zooming waveforms, moving cursors, configuring measurements and interacting with results are done with intuitive touch screen controls.

Complete article, here ▶

Angle sensors eliminate dynamic errors at high axial speeds

ams' latest magnetic position sensors claim to offer the best-available accuracy at high rotation speeds; AS5047D, AS5147 and AS5247 have a technology that ams terms DAEC compensation technology to reduce dynamic angle error to almost zero; these magnetic position sensors are capable of producing extremely accurate angle measurements of rotors spinning at high speed. ams developed DAEC (Dynamic Angle Error Compensation) to eliminate measurement error attributable to propagation delay.

Use the AS5047D for industrial applications including robots and encoder modules; AS5147 is an AEC-Q100 qualified part for automotive applications such as electronic power steering and pumps; the dual-die AS5247 (also AEC-Q100 qualified)

is suited for automotive applications requiring the highest level of functional safety compliance.
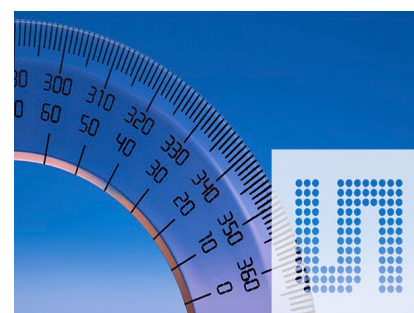
All three parts are specified with a maximum ±0.17° angle error (excluding integral non-linearity).
This precision measurement performance is the result of implementing the algorithm which performs error compensation internally and responds automatically to changes in the speed of rotation. The 47 series sensors provide angle measurements accurate to ±0.08° at 7,000rpm, to ±0.14° at 12,000rpm and to ±0.17° at 14,500rpm.

Complete article, here ▶

pulse

Breaker technologies for DC power grids: research seeks a solid-state solution

In a joint announcement from the partners in the "NEST-DC" research project – Airbus Group, E-T-A Elektrotechnische Apparate, Infineon, Siemens, and University of Bremen – the research objectives of the project are set out as developing an innovative electronic circuit breaker for renewable energy and on-board grids.

According to a statement released by Infineon, direct current offers many advantages compared to the conventional alternating current used today: For example, losses in power grids and electric devices are a total of 5 to 7% smaller than with alternating current. Direct current also makes it possible to more efficiently feed electric energy from regenerative sources into power grids and energy storage and to improve grid stability; with direct current it would be possible to build much more compact electric devices.

In the past the lack of efficient and cost-effective circuit breaker technologies has made it impossible to fully exploit the potentials of direct current, e.g. in distribution grids in data centres, photovoltaics and telecommunication systems or in on-board grids for aviation and shipping, electric vehicles and railway technology. The only electromechanical circuit breakers available today carry the risk of arcing when switching direct current and voltages; furthermore they are slow to react, heavy, unwieldy and expensive.

Funded by the German Federal Ministry of Education and Research (BMBF), the "NEST-DC" research project aims to investigate the foundations of an innovative semiconductor-based and completely electronic circuit breaker for DC power grids and applications. The new circuit breaker should be able to switch direct current on, and most importantly switch it off, as quickly and safely as possible at voltages of up to 1,500V.

The NEST-DC research project is receiving approximately €2.3 million in support from the BMBF in the context of the "Power Electronics for Increasing Energy Efficiency" funding focus area. The project began in October 2013 and will run for three years. NEST-DC is an abbreviation for (in German) " Innovative Electronic Direct Current Circuit Breakers for Renewable Energies and On-Board Power Networks". Infineon is leading the project, contributing its power semiconductor expertise and researching power semiconductors intended for use in the OCB-FETs.

Complete article, here ▶

ST Micro expands into production of long-life paper-thin batteries

STMicroelectronics has announced limited production of its EnFilm advanced rechargeable batteries that are less than 0.25 mm thick. These paper-thin batteries free designers from the constraints of standard battery sizes for personal technology and Internet-of-Things (IoT) devices.

At 220μm thick and measuring 25.7 x 25.7 mm, ST's EFL700A39 EnFilm solid-state lithium thin-film battery is suited for use in ultra-low-profile devices. Surface-mount terminals allow direct attachment to the circuit board, which simplifies assembly and eliminates wires and connectors. Optional tape-and-reel packaging allows high-speed automated placement.

With 3.9V nominal voltage and 0.7 mAh capacity, the EFL700A39 can power a wide range of applications. Its lithium technology recharges rapidly from a 4.2V charging circuit and displays low capacity loss as well as long cycle life allowing some 10 years of use if charged once per day. The EFL700A39 is RoHS compliant and UL certified, satisfies UN tests and criteria for battery transportation, meets IEC 62133 safety specifications, and meets the ISO7816/IEC10373 mechanical and flexibility standards for smart cards.

ST is ready to fulfil orders for engineering samples and small production quantities, targeting applications including wireless sensor nodes, RFID tags, smart cards, wearable technology, non-implantable medical monitors, and back-up or storage for energy-harvesting devices. The price is $30.00 per unit for orders of five units minimum.

EnFilm™: the Energy of Things
Thin-film rechargeable battery

Complete article, here ▶

Tiny amplifier drives 200-mW loads BY MARK REISIGER

Some amplifiers must deliver a moderate amount of power to a load while maintaining dc accuracy. A precision op amp can drive loads requiring less than 50 mW, and a compound amplifier can drive loads requiring Watts, but no good solution exists in the middle of the range. Either the op amp can't drive the load, or the circuit becomes large, complex, and expensive.

When driving a Wheatstone bridge, the excitation voltage directly affects the offset and span, so dc precision is required. Operating with 7-V to 15-V power supplies, the circuit must drive the bridge with unity gain from 100 mV to 5 V. A variety of different bridge resistances can be used. Strain gauges, for example, have standard impedances of 120 Ω or 350 Ω. With a 120-Ω bridge, the amplifier must source 42 mA to maintain a 5-V bridge drive. In addition, the circuit must be able to drive up to 10 nF including the cable and decoupling capacitor.

Amplifier selection

The first step in designing this circuit is selecting an amplifier capable of driving the load. Its dropout voltage must be less than the available headroom for the circuit at the required load current. For this design, the minimum power supply is 7V and the maximum output is 5V. Providing 250 mV for margin, the available headroom is 1.75V. The ADA4661-2 precision, dual op amp with rail-to-rail inputs and outputs specifies 900-mV dropout when sourcing 40 mA, so it should easily meet the headroom requirement; and it was designed to drive significant power while rejecting thermal gradients.

Feedback-loop stabilisation

Meeting the load-capacitance specification is tricky, as most op amps can't drive 10 nF without external compensation. One classic technique for driving large capacitive loads is to use multiple feedback, where an isolation resistor, R_{ISO}, shields the amplifier output from the load capacitance. The dc precision is maintained by feeding back the output signal through a resistor. The loop stability is maintained by feeding back the amplifier output through a capacitor. Unfortunately, R_{ISO} cannot be made large enough for the total load impedance to look purely resistive at the amplifier's unity-gain frequency without limiting the headroom due to the IR drop.

A second stabilisation technique is the hybrid unity follower topology. Rather than trying to move the load-capacitance pole,

Figures (a) multiple feedback (b) hybrid unity follower (c) proposed bridge driver

this approach reduces the feedback factor to force a lower frequency crossover. An easy way to think about this circuit is the superposition of an inverting gain of $-R_F/R_S$ and a noninverting gain of $(1 + R_F/R_S)$. The result is a circuit that operates with a signal gain of +1 and a noise gain of $(R_S + R_F)/R_S$. Independent control over the feedback factor and signal gain allows this circuit to stabilise any size load at the expense of circuit bandwidth. This circuit also has drawbacks, however. The noise gain is high for all frequencies, so dc errors such as offset voltage (V_{OS}) are magnified. In addition, the external feedback loop can be stable while the output stage feedback loop is unstable.

These issues can be overcome by combining the operating principles of both circuits. Multiple feedback separates low-frequency and high-frequency feedback paths, adding enough capacitive load isolation to minimise output-stage stability problems. The low-frequency feedback is driven from the bridge voltage through feedback resistor R_F. The high-frequency feedback is driven from the amplifier output through feedback capacitor C_F. The circuit behaves like the hybrid unity follower at high frequency. The high-frequency noise gain, determined by the impedance of the capacitors, is equal to $(C_S + C_F)/C_F$. This noise gain allows the feedback loop to crossover at a low enough frequency where its stability is not degraded by the load capacitance. Since the low-frequency noise gain is unity, the dc precision of the circuit is maintained.

The circuit presented here can apply 5V to resistive loads as low as 120 Ω with less than 1-mV of total error and stably drive up to 10 nF of total capacitance. The circuit meets its rated performance while operating with a wide range of power supplies from 7V to 15V and dissipating almost 400 mW. The basic circuit can be extended to drive positive and negative loads by powering the amplifier with ±7-V power supplies. All of this capability is accomplished with one tiny 3 × 3-mm amplifier and four passive components.
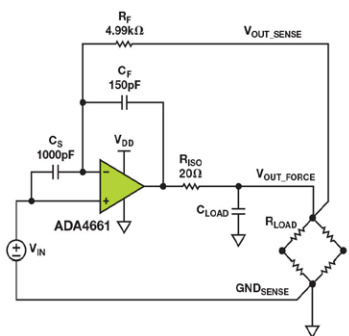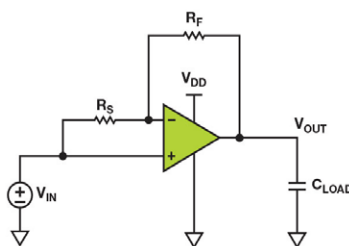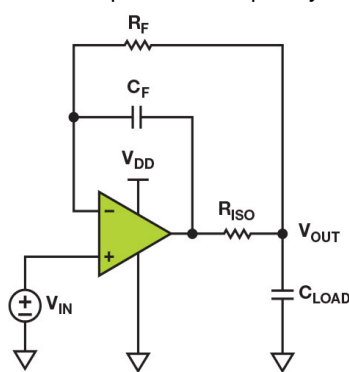
Mark Reisiger [mark.reisiger@analog.com] is a staff design engineer in the Linear Products Group at Analog Devices specialising in CMOS amplifier design. Mark has BSEE and MSEE degrees from Rochester Institute of Technology, and has worked for ADI since 2005.

BY GUSTAVO CASTRO AND SCOTT HUNT, ANALOG DEVICES

HOW TO STAY OUT OF DEEP WATER WHEN DESIGNING WITH BRIDGE SENSORS

Instrumentation amplifiers (in-amps) can condition the electrical signals generated by sensors, allowing them to be digitised, stored, or used to control processes. The signal is typically small, so the amplifier may need to be operated at high gain. In addition, the signal may sit on top of a large common-mode voltage, or it may be embedded in a substantial dc offset. Precision in-amps can provide high gain, selectively amplifying the difference between the two input voltages while rejecting signals common to both inputs.

Wheatstone bridges are classic examples of this situation, but galvanic cells such as biosensors have similar characteristics. The bridge output signal is differential, so an in-amp is the preferred device for high-precision measurements. Ideally, the unloaded bridge output is zero, but this is true only when all four resistors are exactly equal. Consider a bridge built with discrete resistors, as shown in Figure 1. The worst case differential offset, V_{OS}, is

$$V_{OS} = \pm V_{EX} \frac{TOL}{100}$$

where V_{EX} is the bridge excitation voltage and TOL is the resistor tolerance (in percent).
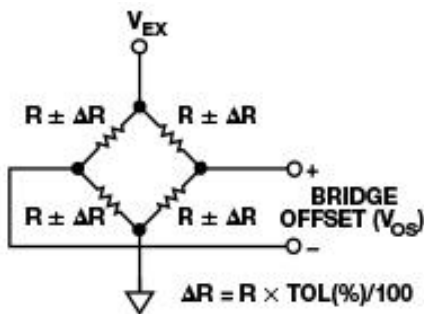

Figure 1. *Wheatstone bridge offset.*

For example, with 0.1% tolerance for each one of the individual elements and a 5-V excitation voltage, the differential offset can be as high as 5 mV. If a gain of 400 is required to achieve the desired bridge sensitivity, the offset becomes ±2V at the amplifier output.

Assuming that the amplifier is powered by the same supply, and that its output can swing rail-to-rail, more than 80% of the output swing could be consumed by the bridge offset alone. As the industry trends to smaller supply voltages, this problem only gets worse.

The traditional three-op-amp in-amp architecture, shown in Figure 2, has a differential gain stage followed by a subtractor that removes the common-mode voltage. The gain is applied on the first stage, so the offset is amplified by the same factor as the signal of interest. Thus, the only way to remove it is to apply the opposite voltage to the reference (REF) terminal. The main limitation of this method is that adjusting the voltage on REF cannot correct the offset if the first stage of the amplifier is already saturated. A few approaches to get around this limitation include:

- Shunting the bridge with an external resistor on a case-by-case basis, but this is impractical for automated production and does not allow for adjustments after leaving the factory
- Reducing the first-stage gain, removing the offset by trimming the voltage on REF, and adding a second amplifier circuit to achieve the desired gain
- Reducing the first-stage gain, digitising the output with a high-resolution ADC, and removing the offset in software
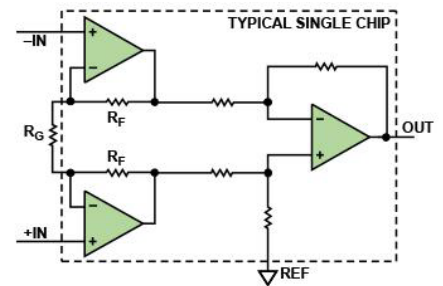
The two last options also need to account for worst-case deviations from the original offset value, further reducing the maximum gain of the first stage. These solutions are not ideal, as they require extra power, board space, or cost to obtain the high first-stage gain needed to obtain high CMRR and low noise. In addition, ac coupling is not an option for measuring dc or very slow-moving signals.

Indirect current feedback (ICF) in-amps, such as the AD8237 and AD8420, make it possible to remove the offset before it is amplified. Figure 3 shows a schematic of the ICF topology.

The transfer function for this instrumentation amplifier is of the same form as that of the classical three-op-amp topology, and is given by

$$V_{OUT} = (1 + R_2/R_1)(V_{+IN} - V_{-IN}) + V_{REF}$$

Because the feedback to the amplifier is satisfied when the voltage between the inputs is equal to the voltage between the feedback (FB) and reference (REF) terminals, we can rewrite this as

$$V_{OUT} = (1 + R_2/R_1)(V_{FB} - V_{REF}) + V_{REF}$$

This suggests that introducing a voltage equal to the offset across the feedback and reference terminals allows the output to be adjusted to zero volts even in the presence of a large input offset.

This article continues by showing how this adjustment can be accomplished; and goes on to look at a detailed design procedure, and design example

Figure 2. *3-op-amp instrumentation amplifier topology.*

Figure 3. *Indirect current feedback in-amp topology.*

Complete article, here ▶

BY BENJAMIN JORDAN, ALTIUM

WHAT ABOUT DC POWER INTEGRITY? - PART 1

When I first learned to design digital electronics and layout a PCB, I was taught to put all the 74-series chips and the microprocessor in neat rows, and the rule of thumb was to add a single 0.1 µF ceramic capacitor for decoupling to each device, and sometimes adding an additional 1 µF tantalum or electrolytic for the micros in parallel. I never worried too much about getting power to each device - using a 20 or 30 mil trace was enough for a chip that never drew more than 100 mA, along with the classic interdigitated +5V/GND "grid". Of course, power electronic designs are a whole different ball game. And I always took a lot more time, care and planning with power supply and amplifier designs - making sure to use proper (star) grounding and keeping high-current loops as tight as possible.

Some of this was more than 20 years ago now, and of course there has been a lot of development in the decoupling and power network topic since then. More elaborate and carefully placed decoupling schemes have to be designed for each new silicon process node, each new chip package generation and for each new PCB design as they become more densely packed with parts than ever. It's getting difficult to find room for all the "rule of thumb" decoupling caps! And with BGA packaged devices down to 0.4 mm pitch, that meanwhile draw several amps of current during use, it's getting really difficult to plan and design a good power network on the PCB. Whether we like it or not, Power Integrity is a challenge that all PCB designers and engineers have to address.

Power Integrity is talked about a lot these days. But a lot of the talk is really on the signal integrity side - I call it AC power integrity, which is really about the impedances of the power network at high frequencies. This deals with how the decoupling is designed as well as return paths for high-speed signals. While it is non-trivial, I don't want to simply regurgitate this already very commonly discussed topic. I want to get down to DC… why? Well, it just seems to me that learning to walk before trying to run is a good idea. So let's talk DC Power Integrity.

At face value, it seems to be a simple enough topic - you just need to make sure there's enough copper to get the necessary current to each device on the board. But that's just at face value. When you start to work with fine-pitch device packages, manufacturing constraints and power requirements of said devices are almost completely at odds. Not only is it difficult to get the current needed to all the power pins, but you are also working with multiple supply voltages. This means that unless you want a high-layer-count PCB, you are going to have to get power to your devices through various split planes, and that's just where the trouble begins.

But before I go too far down into the rabbit hole of designing the power distribution networks, how can you tell if you even have a power integrity problem? Power Integrity issues are sneaky little creatures. Like cockroaches that rapidly scamper into the crevices when the light turns on - the moment you try looking for these issues is the moment they can't easily be reproduced. But you may have a power integrity issue if any of the following symptoms occur to your assemblies:

1. The CPU is resetting unexpectedly, or when a high-utilisation thread enters execution.
2. Memory devices keep failing content retention / corruption tests.
3. Analogue front-end circuits are randomly inaccurate or out of design specs.
4. CPU or FPGA devices fail catastrophically.
5. FPGA configurations are corrupted during power up.
6. PCB vias go open-circuit after a period of use cycles or maybe even at first power on.
7. Production PCBs suffer blistering in the common locations.
8. PCBs suffer delamination in common locations.
9. Trace or polygon neckdowns are fusing.
10. Discolouration of laminate or solder mask material in some regions of the PCB.

These symptoms fall into two broad categories of DC Power Integrity problems. For example, items 1 through 5 are the more sinister misbehaviours caused by transient voltage drops across the board. Sometimes they can be fixed with better decoupling but when talking DC, really, more copper will improve the design. Items 6 through 10 are more serious power integrity issues where current density regularly exceeds the safe limits for temperature rise and the board is suffering from localised heating, or copper is outright fusing.

There are some useful tools for avoiding these sorts of problems before prototype; for example the IPC-2152 conductor sizing charts. I would say it's a must that every design begins with these charts as the basis for power network design rules for the PCB layout. However, there are designs that now approach a part density that make it necessary to design "on the edge" and work with means (that is, average values) and duty cycles to make sure the board doesn't fail.

So, DC Power Integrity is the concern over making sure that each device in the design gets the power it needs, without suffering the problems mentioned, all while ensuring a reliable power network on the PCB.

Design with PI in mind

A great question to answer though, is what is a good starting point for designing with PI in mind?

Benjamin Jordan, who is Senior Manager, Content Marketing Strategy, Altium, picks up this thread in the continuation of this article; while there are well established "best practices" in designing with power in mind, he notes that he sees a large number of designs where the power distribution on the PCB is basically copied from a chip manufacturer's reference design.

Complete article, here ▶

BY BONNIE BAKER

Can you find the missing 1?

Surprise, surprise. The bandwidth of an amplifier in a gain of +1 V/V is not the same as the bandwidth of the same amplifier in a gain of –1 V/V. After more than five years of working every day with amplifiers for an industry analogue leader, this was a total surprise to me.
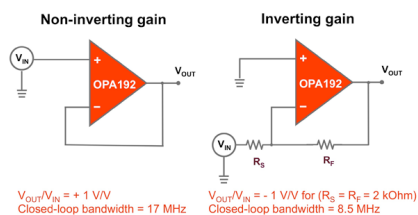

Figure 1 *The simple non-inverting and inverting gain configurations*

In Figure 1, notice that the closed-loop bandwidth of the non-inverting gain circuit in a gain of +1 V/V is 17 MHz, and the inverting gain circuit in a gain of –1 V/V has a 8.5 MHz bandwidth. The definition of noise gain (GN) is the gain of a non-inverting amplifier circuit. With that definition, the GN of the non-inverting circuit is +1 V/V, and the GN of the inverting gain circuit is +2 V/V. Figure 2 shows the open loop gain of the amplifier in Figure 1.
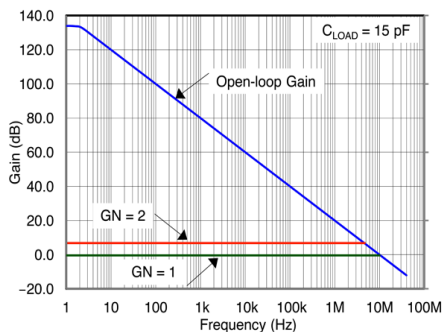

Figure 2 *Using the OPA192, we have the open loop gain curve (blue), noise gain curve for closed loop gain of 2 (red), and noise gain curve for closed loop gain of 1 (green).*

In Figure 2, the OPA192 open loop gain spans from 1 to 100 MHz. See how the 0 dB crossing (GN = 1 V/V) occurs at 10 MHz. Additionally, the 6 dB crossing (GN = 2 V/V) occurs at 5 MHz.

This is a good start if you want to understand these circuits. But let's go back and look at some control theory.

Figure 1 shows a circuit diagram representation of two gain configurations. Figure 3 shows a control theory representation for these two gain configurations.
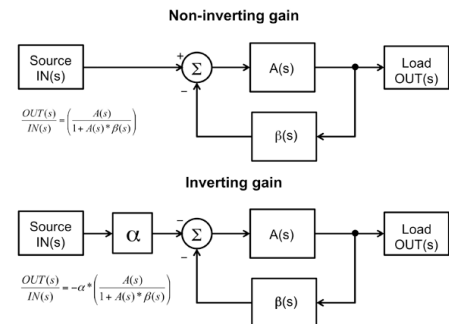

Figure 3 *Control theory representation of circuits found in Figure 1..*

In Figure 3, alpha (α) is the closed loop inverting gain. A(s) is the open loop gain of the amplifier, or the amplifier's internal gain. Beta (β(s)) is the feedback factor. In laymen's terms, it's the amount of the output that is fed back to the input. For both circuits in
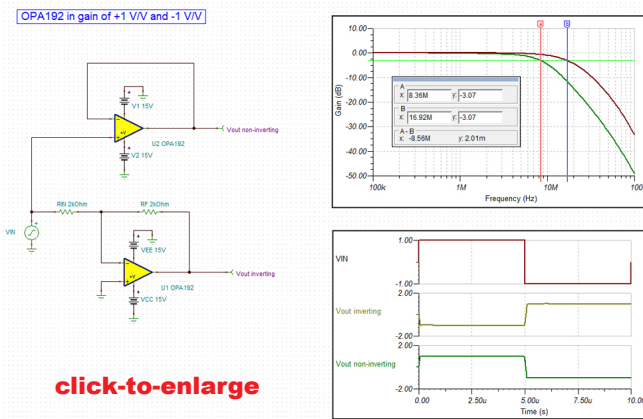

click-to-enlarge

Figure 4 *TINA-TI circuit configuration with the AC and transient response. The AC curve shows the two signal paths with the inverting bandwidth signal equal to half of the non-inverting signal.*

Figure 1, beta equals (R_S / (R_S + R_F)). Another interesting fact is that 1/beta is equal to the circuit's GN.

With the non-inverting gain representation in Figure 3, when A(s) is very big the closed loop gain is equal to 1/beta. If you examine the same circuit with A(s) very small or near the 1/beta value, the intesect frequency of the 1/beta curve and A(s) is equal to the gain bandwidth product divided by 1/beta.

The inverting amplifier circuit presents a surprise in this discussion. If A(s) is very big in the Figure 3 inverting gain block diagram, the circuit gain is equal to –R2/R1. If you examine the same circuit with A(s) very small or near the 1/beta value, this occurs at a frequency equal to the gain bandwidth product divided by 1/beta. In this situation, 1/beta is equal to 2.

Don't take my word for it – give it a try! Run some simulations with TINA-TI and this file. If you have TINA-TI on your computer, you can double click on this line to start your own SPICE simulation. Or better yet, bread board your own circuits to prove or disprove this phenomenon yourself. See the proper circuit configuration in Figure 4.

References

1. "Op amp Bandwidth," eCircuit Center
2. Download a datasheet for the OPA192
3. Download a free version of TINA-TI
4. TINA simulation– bandwidth gain differences between 1 V/V and -1 V/V amplifiers compared

BY CHRIS BURNS, DALHOUSIE UNIVERSITY
AND CHARLES CIMINO, KEITHLEY INSTRUMENTS

IMPROVING THE COULOMBIC EFFICIENCY AND LIFETIME OF LI-ION CELLS DEMANDS HIGH PRECISION, HIGH ACCURACY TESTING

Canadian battery chemistry researchers boost measurement integrity by two orders of magnitude over commercial test systems with an ultra-high-precision charger.

Lithium-ion batteries have become staples of modern electrical and electronic products; in fact, they are now used in a wide range of applications from smartphones and power tools to medical equipment and electric vehicles. Well-made Li-ion cells can easily meet the needs of smartphone, laptop and power tool applications because of their high energy density, long calendar/cycle life, and relatively low cost in comparison with other rechargeable battery technologies, such as nickel-metal hydride. However, developing cells capable of cycling with minimal capacity and energy loss for ten years as required by automotive applications or even longer for grid energy storage is far from trivial.

Designing such high quality cells can be complicated, as can the process of evaluating their performance. Cell manufacturers constantly make small changes to the cell design that can impact the cell chemistry as part of their efforts to extend the cycle life or reduce the cost without sacrificing performance. However, testing all of these different experimental cell chemistries under real-world conditions would be impractical. Consider, for example, a battery for an electric vehicle that is only cycled once per day over its ten-year lifetime; obviously, it would take far too long to determine whether the experimental changes made were beneficial and therefore should be implemented in commercial cells. No one can afford to wait years to complete the R&D feedback loop. Cell chemistry researchers need tools and techniques that allow them to perform reliable accelerated lifetime testing to determine if proposed changes extend cell lifetimes.

Accelerated lifetime testing

The most common form of accelerated lifetime testing involves high rate cycling; cells are charged and discharged at rates up to one cycle per hour (up to around 20 cycles per day, many more than would occur in actual use) in order to acquire data for hundreds or thousands of cycles over an experiment of reasonable length (several months). Over these cycles, the experimental cell's loss of capacity and energy are measured and compared with those of a control cell or a cell already in production to decide whether the experimental chemistry offers any benefits.

The degradation of Li-ion cells is not only cycle-dependent but has a strong time dependency. Smith et al. [Ref. 1] showed that cells being cycled at different low rates (either one cycle per two, four or eight days) all failed after the same amount of time, despite differing in the number of cycles completed by factors of two and four. This means that if a cell is measured to only lose 10% of its initial capacity after 1,000 cycles in two months, it does not mean that the same cell would only lose only 10% of its capacity after 1,000 cycles over three years.

Work in Dr. Jeffery Dahn's research group in the Department of Physics and Atmospheric Sciences at Dalhousie University has suggested a new method for distinguishing between the lifetimes of different experimental cells within just a few weeks: High Precision Coulometry. Given that the side reactions within cells (that is, solid electrolyte interphase or SEI growth, electrolyte oxidation, transition metal dissolution, etc.) all involve transferring charge that is not associated with the intercalation/deintercalation of lithium from the electrodes, those reactions can be detected coulometrically. If all of the lithium stored in the negative electrode on charge was returned during the subsequent discharge, then the charge (QC) and discharge (QD) capacity would be equal, so the coulombic efficiency (CE = QD/QC) would be exactly unity and the cell should be able to cycle indefinitely. However, due to these parasitic reactions occurring within a cell, the coulombic efficiency is less than the ideal value of 1.0000 and the cell degrades. This causes a typical voltage versus capacity curve to "slip" to high absolute capacities with subsequent cycles because the discharge capacity is always less than the previous charge capacity.
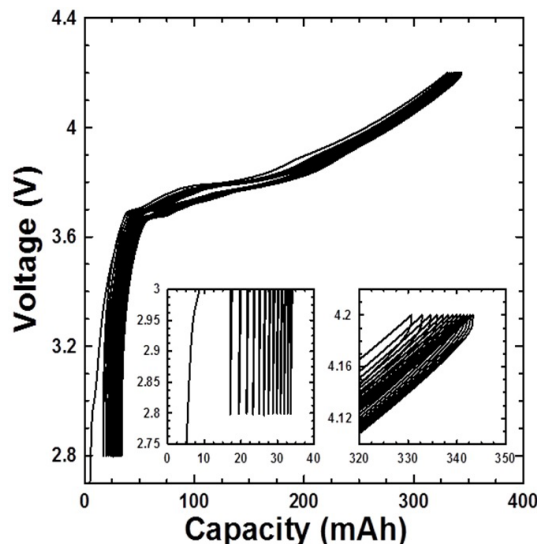
Figure 1 is a typical V-Q curve showing this type of behaviour, with the insets showing the top of charge and bottom of discharge endpoints to better illustrate the rate at which the curve slips to the right. Because the coulombic efficiency is defined as the discharge divided by previous charge capacity, it is directly related to the rate of motion of the bottom of the discharge endpoint, referred to as discharge endpoint slippage - ΔD (CE = QD/QC = 1 − ΔD/QC). The top of charge endpoint slips to higher capacity with subsequent cycling as well, referred to as charge endpoint slippage. This can be measured independently of the coulombic efficiency because it primarily relates to reactions that occur at the positive electrode [Ref. 2]. Given that parasitic reactions cause the voltage curve to slip to the right (decreasing coulombic efficiencies and increasing charge endpoint slippage), then cells with higher coulombic efficiencies and lower charge endpoint slippage rates must have lower rates of parasitic reactions and therefore should have longer cycle lives. This idea is the underlying premise for using High Precision Coulometry as a way to compare cell lifetimes in short-term experiments.

Commercially available battery testing equipment can't differentiate between the coulombic efficiencies of cells at this high level of accuracy and precision: in the continuation of this article, the authors describe how Dr. Dahn's group has configured a battery cycler using high precision current sources and multimeters called the Ultra High Precision Charger (UHPC).

BY JACOB BENINGO

7 TIPS FOR USING INTERRUPTS

Traditionally programmers and organisations have had an irrational fear of using interrupts. One might think that statement is facetious but on more than one occasion in the last year the author has worked with an organisation whose coding standard strictly said something like "Interrupts are to be avoided at all costs and manager approval must be obtained before their use". This seems absolutely silly! Interrupts are an important instrument for designing embedded software so what is all this fuss about?

It is undoubtedly true that in the past some developer disobeyed the design rules and patterns for writing good interrupt service routines and a costly and chaotic debug session debug session resulted. The result was that the company swore off using interrupts unless it was absolutely necessary and decided that polling was a better solution. There are of course many problems with this decision. Polling is very inefficient. Rather than putting a processor to sleep to conserve energy the processor must always be awake looking to see if an event has just occurred. Even worse, without interrupts, the real-time response of the system could be compromised!

Strange things can certainly happen to a system if proper interrupt implementation is not followed. It is absolutely guaranteed thanks to Murphy that this strange behaviour will happen at the most inopportune time. So what can be done to ensure that interrupts are implemented properly and utilised as they were meant to be?

Tip #1 – Create an ISR Table
An interrupt service routine table is an array that contains a list of every possible interrupt that can occur on the microcontroller. Each interrupt is populated with the function (pointer to a function) that is executed when the interrupt occurs. There are many advantages to doing this. The first is that it becomes very easy to assign a function call to an interrupt. Simply change the function name in the table, recompile and now that function is called when the interrupt occurs.

Next, the programmer is forced to include a function for every interrupt. This is a good practice because every interrupt gets initialised with code! During debugging if an errant interrupt occurs rather than jumping off and executing unknown code, a "DummyISR" could be executed instead. This allows for that errant interrupt to be trapped and debugged. Finally, using an interrupt table in this manner forces the interrupt code to take on an organised structure that is easy to understand and configure.

There are many different ways in which the table can be implemented. The most common is through the use of #pragma. This allows the table to be placed at a specific memory location in flash. The two implementations that are most common are to allow the starting address of the flash location to be specified or a linker memory label to be defined. Using #pragma is something that a developer should try to avoid but if this is the only method available to implement the interrupt table then this would be a good exception to the rule. An example of such a table can be found in Figure 1, *in the online version of this article, click the link.*

Tip #2 – Keep them short and fast
An interrupt by definition is an interruption to the normal flow of the application that is being executed. The program literally stops doing whatever it was doing in order to handle the interrupt. With this being the case it would seem obvious that an interrupt service routine should be short and to the point so that the primary application can resume execution.

The real point of an interrupt is to handle an urgent event that requires the system's attention. To keep the routine short, only do the minimum of what really needs to be done at that moment. For example, if communication data is triggering the interrupt, stuff the data into a buffer, set a flag and let the main program process the data. Don't try to process it in the interrupt!

Keeping code short and fast can sometimes be deceiving. For example, doing a simple floating point calculation that is a single line of code may appear short, but a microcontroller without a hardware floating point unit would spend a near eternity processing the math (possibly milliseconds)! There are a couple of simple rules to ensure that your interrupt service routines run fast:

- Don't call a function from within your interrupt (unless they are inline functions). The function call overhead will kill your timing.
- Any processor intensive activity such as processing a data buffer, performing a calculation, etc. should instead set a flag and let the main application do the processing.
- Wait statements should be avoided
- Loops or any time intensive logic such as for loops, division or modulus operations should be shunned as well.

Tip #3 – Double check your initialisation
For one reason or another, interrupts always seem to be a pain to get working properly. They are conceptually straightforward but some implementations require a good amount of forethought to get it right. There are a number of questions that a developer should be asking when setting up and debugging an interrupt. Some of these seem simple but they should nonetheless be checked and asked.

- Is the interrupt enabled? (Sure, I thought I enabled it but what does the MCU register say?)
- Has the priority of the interrupt been set?
- Was the ISR placed properly in the interrupt table?
- Is the interrupt mapped to the correct hardware pin? To the correct peripheral?
- Is the interrupt acknowledged ASAP?
- Is the interrupt flag cleared at the right time in the ISR?

Read tips 4 to 7 in the continuation of this article

Complete article, here
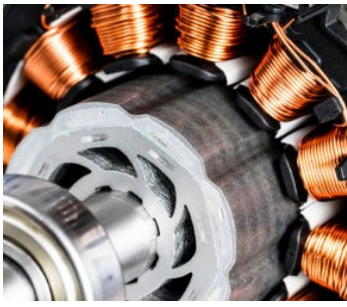
Spin Cycle

Motor control for functional safety

By Brian Fortman, Texas Instruments

The basic evolution

Over the years, electric motors have been applied in an ever increasing number of applications. Think back to the replacement of steam engines with AC induction motors in locomotives, the migration of hydraulic and pneumatic actuation systems to DC electric motors, and include today's augmentation and replacement of internal combustion motors in hybrid and electric vehicles (Figure 1.).

Electric motors can be easier to install and maintain, produce less carbon emissions, provide a more compact actuation of torque as well as provide opportunities for new system control features since their controllers are electric or electronic – think of self-parallel-parking cars enabled by electronic power steering systems.

Figure 1 *Simple evolution of electric motor use*

Over time, new techniques for commutation using electronics instead of brushes enable greater benefits from permanent magnet synchronous machine (PMSM) motors. These new techniques require more performance from the controlling electronics.

As electric motor use proliferates, these motors are found in more applications where a fault in the system, the electronics, sensors or motor, could pose a serious risk to people or the environment, so functional safety standards now apply to the electronics. Oversimplified, this means that new guidelines must be followed to ensure that risks are managed at acceptable levels. This happens through the detection and disposition of faults in the system as well as by putting best practices for product development in place.

At the same time, the number of industry-specific safety standards has grown, and many have been updated just within the last couple of years, with many more on the horizon. However, today, only a few have had specific requirements placed on the components used: IEC 61508, ISO 26262 for Automotive, and Aerospace. A growing number of embedded electronics systems have already been or soon will be subject to an international safety standard. Following the trend, I believe the component-level safety performance requirements will only become more pervasive across the standards over time (Figure 2).

In the mid-2000s, semiconductor technology allowed high levels of functional integration on chip which were previously only possible at the system level. Functional safety techniques that had previously been applied at a system level had to be rethought in order to best take advantage of functional safety diagnostics that could now be integrated on-chip, that is, at the component level. New diagnostic techniques became economically more practical, such as dual-core lockstep, which is particularly useful for applying safety at runtime in tight control loop systems. Component-level compliance, the "white box" review of internal component design, to safety standards is currently state of the art in the ISO 26262:2011 and IEC 61508:2010 standards.

Figure 2 *End-equipment-specific functional safety standards*

The trend towards application of functional safety at deeper levels of the electronics hierarchy is expected to continue. Indeed, some standards working groups are starting to talk about compliance at the IP module level.

So, what does all of this mean for a motor control developer? Simply put, you are probably a victim of your own success. You've figured out how to apply the benefits of electric motor actuation in more useful ways across many applications. But now you are or will soon be confronted by requirements to conform and give evidence that your system provides some level of functional safety.

Help is available. For example, Texas Instruments' SafeTI design components offering assists with delivering solutions for software and chipsets for use in functional safety standards that need to adhere to the ISO 26262, IEC61508 and IEC 60730 standards. SafeTI design packages can help provide documented evidence of their suitability for use in functional safety systems, addressing the greatest challenges faced by a functional safety motor control system designer.
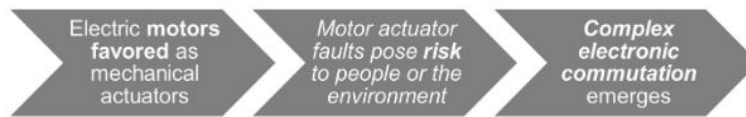
BY PAUL PARKINSON, WIND RIVER

IMPLEMENTING MULTIPLE INDEPENDENT LEVELS OF SECURITY ARCHITECTURES ON MULTI-CORE PROCESSORS

In this article, we will introduce the Multiple Independent Levels of Security (MILS) system architecture, and describe how MILS can be deployed on multi-core processor architectures to provide a high-performance, high-assurance foundation to protect critical national infrastructure systems against cyber threats.

The Cyber Security Landscape

In recent years, many European countries have recognised the growing cyber threats against both civilian infrastructure and defence systems, and have responded by developing national cyber security policies which define the objectives for the protection of critical national infrastructure against cyber attacks, and a range of strategies for achieving these objectives.

The Advent of the MILS Architecture

Commercial organisations and European national governments have long categorised information at different security classifications, based on criteria such as information value, sensitivity, and the impact of disclosure. Historically, information at different security classifications has been physically isolated in separate domains, initially in manual systems, and subsequently in computerised systems.

More recently, as organisations have become increasingly reliant on computer systems, there has been a drive towards automation of the information flow process between different security domains. This enables decision-making to be accelerated in fields as diverse as commercial business, banking, government and armed forces. Traditionally, multi-level secure (MLS) computer systems were built as bridges between these domains using multiple, physically separated computers, networks, and displays. This technique, known as "air gap" security, required expensive equipment and occupied a large footprint in terms of Size, Weight and Power (SWaP), and has limitations in the cyber era.

The Multiple Independent Levels of Security (MILS) architecture was proposed as an alternative approach for secure embedded systems many years ago. MILS uses a layered software architecture, with a separation kernel (SK) built on four fundamental security policies:

1) Data Isolation, which ensures that a partition cannot access resources in other partitions.

2) Periods Processing, which ensures that applications within partitions cannot consume more than their allocated share of CPU usage.

3) Information Flow, which defines the permitted information flows between partitions.

4) Fault Isolation, which defines that a failure in one partition does not impact any other partition within the system.

These four policies create an architecture where the separation kernel is Non-Bypassable, Evaluatable, Always Invoked and Tamper Proof, collectively known as NEAT. The small separation kernel and layered assurance approach means that the amount of code which would need to undergo rigorous scrutiny as part of a high assurance security evaluation can be very small, to reduce both evaluation time and cost. This would be an important benefit when developing a critical national **infrustrive system which needs to undergo a high-assurance security evaluation.

The implementation of MILS systems on single-core (also known as unicore) processor architectures has utilised the hardware capabilities provided by the processor to enforce the four fundamental security policies, and to also minimise unintended hidden channels of communication between applications, known as covert channels. This is an important consideration for high-assurance systems, as covert channels can be used by advanced adversaries in sophisticated cyber-attacks against critical national assets. However, as unicore processor architectures have been in use for many years, the characteristics of covert channels on these architectures are well understood. This means that on certain unicore processor architectures it may be possible to use appropriate countermeasures in the MILS implementation to reduce their bandwidth.

Advent of Multi-core Processors

For decades, gains in processor performance were achieved through increasing the processor clock frequency. This approach eventually reached the limits of viability due to dramatically increased power consumption, so semiconductor companies turned to focus on the development of multi-core processor architectures to achieve both performance gains and power reduction. However, the advent of multi-core processor architecture is a disruptive change, as applications and systems need to be designed with multi-core in mind in order to efficiently utilise the potential performance available.

Multi-core processors also present new challenges for security with respect to application isolation and core separation, because applications can truly run concurrently (compared to running sequentially on a unicore processor). Research into these areas has been undertaken in recent years in the related discipline of safety-critical systems, and although the end-goals of safety certification and high-assurance security are different, there is some overlap in their requirements. Recent research undertaken by the European Aviation Safety Agency (EASA) has highlighted issues around suitability of individual multi-core processor architectures for safety-critical avionics applications, due to contention for shared resources. This issue could have a potential impact for a security-critical application in terms of a covert channel. However, analysis also reveals that some multi-core processor designs provide hardware features which can be utilised to enforce application isolation and core separation. The result is that careful scrutiny is required when selecting multi-core processor architectures for high-assurance security applications.

Figure 1 *Cross-domain system network gateway*

In the continuation of this article, the author discusses cross-domain system requirements (Figure 1) and how a gateway might be configured
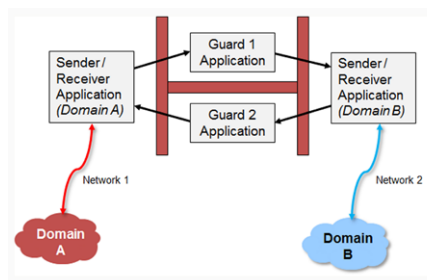
Complete article, here ▶

BY BRUNO PUTZEYS

THE G WORD:
HOW TO GET YOUR AUDIO OFF THE GROUND

My most daunting challenge is an ongoing one. I am trying to expunge my language of words that aggravate, cause hurt, misunderstandings or that are just meaningless blather. The one I find most troubling is the G word. I can't bear to imagine saying it in polite company and yet all too often I catch myself doing it unwittingly. Deep breath… I'm talking about… "GND." There. Forgive me. No more nastiness.

Audio signals are voltages. A voltage is the potential difference developed between two points. We grab a voltmeter and connect the two test leads to probe the two points, or "nodes" that we want to know the potential difference between. We don't just attach one lead and hope to get a reading.

GND-think

And yet it is not unusual for audio engineers to think of an audio signal as only one circuit node or wire next to which a voltage is written or a waveform drawn, as though this single node were magically capable of having a voltage all on its own. The second node, it seems, is too unimportant, too obvious to mention. And this is where the rub lies: what on earth is ground?

According to GND Gurus the root cause of all hum and buzz problems is current flowing through "the same ground" as that used as voltage reference. So, they suggest, we use "different grounds."

Figure 2: It's got wonderful Powerpoint appeal, though.

The hidden assumption is that a signal is just one wire. But as anyone with a voltmeter knows, the second wire is every bit as important as the first. Still we seem to think it makes sense to use as the second wire the central sewage pipe that also carries waste electrons, supply return currents, shield currents etc., back to the recycling plant. And then we're surprised to find rubbish on it.

The supposed solution is called a "star ground," a common point where "different grounds" connect.

It looks nice at first glance and its practitioners defend it as though it were a fundamental truth (Figure 1). Practically speaking though it's a nonstarter. It only works at all when it's rigorously done.

You can star a power amp. You can star a preamp. And then you connect the two. Oops. Which of the two stars guards that mythical common potential that all signals in the combined circuit are referenced to? That's where GND Gurus get into their stride. Chains of stars, stars of stars, the whole celestial menagerie (figure 3). All hinge on minimising current flow through the connections that tie the local stars together. And so the saga

Figure 1 *It's got wonderful Powerpoint appeal, though.*

Figure 2 *Advanced GND Guruship in action. Yes, I found this on the Internet.*

continues with "floating grounds," disconnecting mains safety earth and whatnot.

You heard me correctly. Most audio equipment has no safety earth connection simply because we can't seem to imagine signal connections without a common reference.

And often that doesn't even work. Suppose I have a TV, a DVD player and an amplifier. When I want to watch TV I want to hear the sound over my stereo. When I watch a DVD I'd rather run the audio straight from the player to the amp, not through the TV's rotten signal processor. So we connect the video output of a DVD player to the TV and the audio to the preamp and we also connect the TV's audio to the preamp. The dreaded "Ground Loop" scenario.

Other than the most minimalist audiophile stereos there is no way of putting a system together without creating current loops. Current loops are a fact of life. Any scheme to avoid buzz and hum had better not rely on avoiding "ground loops."

The final nail in the star's coffin is that it only works at DC. A wire has inductance and two wires have mutual inductance on top of that. Accidentally lay a "dirty" return wire next to a "clean" reference wire and bam, noise. How do we add power supply decoupling? Do we run long wires from the decoupling capacitor to the star and add exactly as much inductance as we were hoping to get rid of ?

With a star you can just about build a mildly comatose class A amplifier. Anything faster and you'll run into stability problems. Try switching circuits and all assumptions go out the window.

In short: any exposé that takes as its premise that hum, noise and distortion have something to do with "grounding" should be stamped on and ground into the ground. We need to design circuits that read voltages like voltmeters: with two wires. The result should not depend on the contents of the local electron tip (rubbish dump, for non-UK readers).

In summary; When a change in "grounding" causes hum, this is because we're naïvely thinking of a signal as one wire. Stars are a sticking-plaster to try to make this flawed assumption work.

The link on the right takes you to the continuation of this article which delves further into the topic of grounding and "GND-think." Subsequent parts consider the ideal differential input; impedance balance vs. current balance, instrumentation amps and cable shielding; and a demonstration fully-balanced pre-amp.
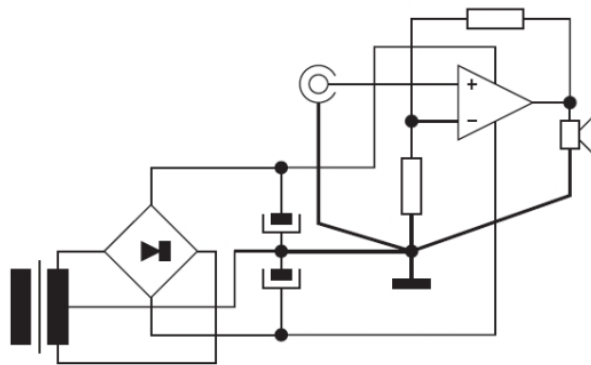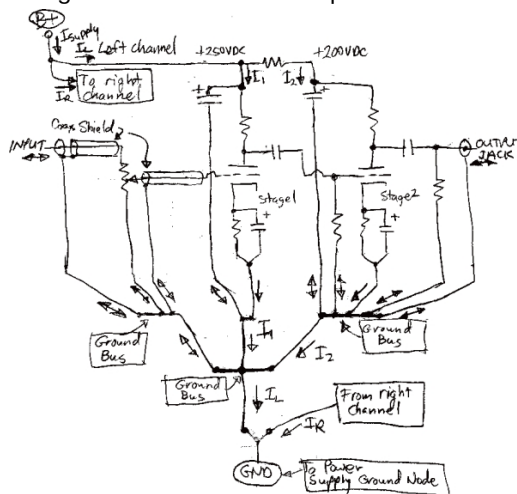
Complete article, here ▶

designideas

Milliohm Squawker: great at finding shorts and reverse engineering PCBs

Glen Chenier

Beeping continuity testers have been around for a long time, but for PCB reverse engineering purposes they leave a lot to be desired. They respond to a "short circuit" of several ohms whereas one would much prefer to discriminate PCB trace and test probe resistance of less than 1Ω to avoid false alarms. Then one would want the beep tone pitch to indicate a few milliohms of ΔR, to determine which side of a closed relay contact, transformer winding, fuse, or low-resistance current sense resistor is actually connected to the net of interest, especially when the copper traces are hidden under components.

When you are rapidly sweeping a wire broom across a PCB to find common net points, no time delay can be tolerated; the beep must sound instantly, and be extended to a noticeable duration. Secondary requirements are low current drain for long battery life, low test voltage, to avoid biasing semiconductor junctions, immunity to 50-60Hz pickup, tolerance to ESD and charged capacitors, and headphone operation to avoid sonically annoying colleagues in the lab or office environment (really, this thing sounds like a scalded cat).

A wire broom?

The fastest way to find all the points connected to a single net to which a clip has been attached at one point is to sweep a broom probe across the rest of the PCB while listening for the squawks. The probe (Figure 1a) uses very fine (3 mil) (0.076 mm, say; 0.1mm – Ed.) phosphor bronze bristles to avoid physical damage to small surface-mount components. Pogo pins zero in on the specific device pins once the general areas are located with the broom, and are useful for their gold plating and sharp points, minimising contact resistance. Their telescoping sections are soldered together to avoid adding unwanted ΔR movement to the measurement. Use multi-point pogo pins; you are less likely to accidentally skewer your hand than with a single-point one, while still making good low-resistance contact. If needed, a single-point pogo pin can be used on a separate probe attachment that lies flat on the bench for very fine-pitch

surface mount IC pins, but be careful – these are very sharp. To make the probes easy to handle while sweeping, you can use coiled cords – in this case, the four AWG26 conductors in the cord are paralleled to minimise resistance. The stationary probe can use a banana plug to attach to various sizes and types of grabber clips. Periodic alcohol cleanings will minimize the ΔR variations caused by the banana plug connection. A typical reverse engineering setup is shown below (Figure 2).
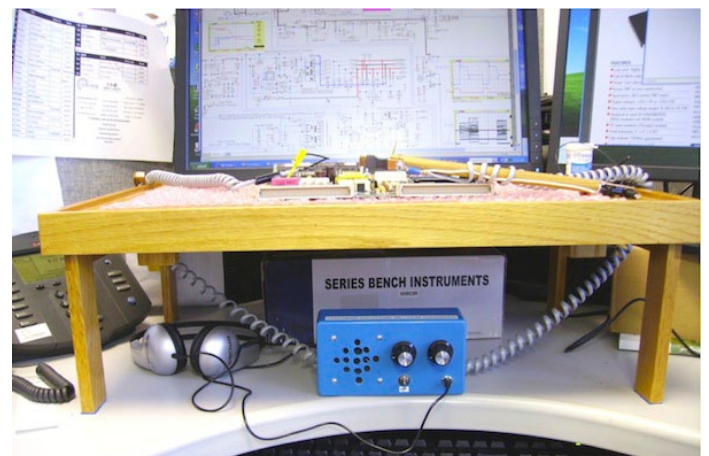

Figure 2. *A blue-box Milliohm Squawker fits nicely under an ergonomic microbench which raises the BUT (board under test) to close-up magnifying-visor eye level, with schematic capture software display in the same plane of view.*
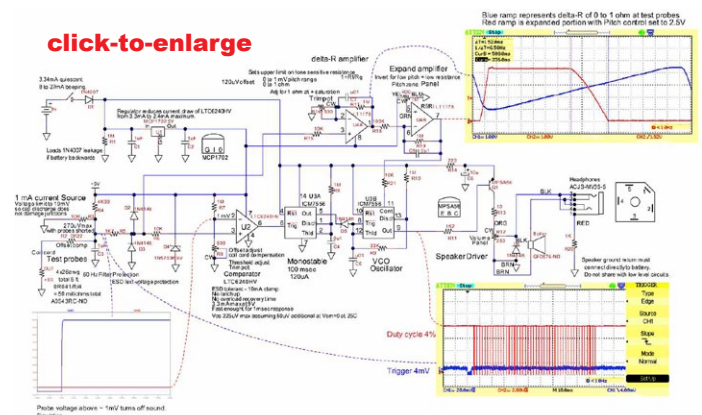

Figure 3. *The Milliohm Squawker schematic (TinyCAD drawing)*

The 9V battery is regulated to 5V for the low-level analogue circuits. R4 sets the probe test current at 1mA, and R3 limits the test voltage to 10mV. R2 adds to the test lead resistance to ensure a positive offset voltage for U2, which is compensated

Figure 1a. *Broom and radial pogo pin*

Figure 1b. *Axial pogo pin at opposite end*

by trimpot R8. This is necessary since U2 uses a single supply; its offset could be negative, and the test lead resistance is compensated later with other circuitry.

C3 removes 50-60Hz AC stray pickup, but discharges instantly into a short at the test probes for fast response. R5, D1, D2, & D4 clamp ESD and any voltages from charged BUT capacitors. U2 is a low-level comparator chosen for low current consumption and low input offset voltage, but is fast enough to respond to a broom sweep pulse of 1 msec. It is available only in surface mount, so if one builds this with leaded components, an adapter board is needed. You can experiment with other fast low-power low-offset op-amps; I chose the LTC6240 simply because it responded fast enough in the LTspice simulation. Trimpot R8 sets the beep threshold resistance: 1Ω is a good choice based on the long thin traces of a typical PCB. Neglecting U2's input offset voltage, 1 mV at U2's negative input sets the threshold for the 1 mA test current at 1Ω probe + R2 + PCB trace resistance. R8 can be set for different thresholds if desired.

A probe voltage of less than 1 mV causes U2 to trigger the 100 msec monostable U3A. This serves to extend the beep so it will be noticed during a fast broom sweep. U3A enables the U3B VCO, which drives the speaker with a 4% duty cycle. During idle (no beep) periods, U3B holds speaker driver transistor Q1 off; the low duty cycle ensures Q1 is mostly off so to minimise battery current drain. R14 isolates speaker current pulses from the battery to prevent any interaction between the speaker current and low-level analogue circuits. C6 provides the current peaks necessary to drive the speaker to a reasonable loudness. If headphones only are used, then Q1 will not be necessary; U3B can drive the headphones directly (a weakness in this speaker circuit is that even when headphones are used, there is current wasted through the volume control. Since most of the time the Squawker is quiet, I ignored this battery-wasting problem.)

The final version was built with SMT on a PCB with a solid ground plane, so I got away with sending the speaker return current through the plane. However, if you build this on vector-

board, keep the speaker return current separated from the low-level analogue ground system with its own return path directly to the battery. The initial solderless breadboard version had all sorts of problems related to this.

U4A and U4B provide the tone pitch vs. ΔR feature. Capacitors C7 and C8 were found to be unnecessary in the PCB version, and are shown here as a 'just in case of trouble', CYA move.

The 0-1 mV across the probes is amplified by U4A, whose gain is set by trimpot R16. Normally, R16 is kept fully clockwise for minimum gain; I have found this to be quite adequate for easily distinguishing pitch tone changes down to 5 mΩ ΔR. Trimpot R16 can be set for higher gain if it is ever necessary to increase the ΔR resolution; so far, I have not found this to be necessary. Do not overdo it – U4A can saturate on input offset if the gain is set too high.

U4B and front panel adjustment R20 let the user "zoom in" to the ΔR range of interest. U4B drives the VCO U3B control input to set the beep pitch. R20 sets the ΔR measurement window and adjusts out the resistance of the test probes, banana plug attachments, and BUT trace resistance. Start by shorting the probes together and tune R20 until the beep just starts to rise from its lowest pitch. A few more milliohms between the probes will cause a further increase in the beep pitch. If your circuit sniffing finds long BUT traces, readjust R20 to accommodate the increased trace resistance and lower the beep pitch back into its linear measurement range.

Eventually, you may reach a point where your net of interest ends in closed relay contacts or a transformer winding. Both sides of these will produce a beep, but the side with the lower pitch (lowest resistance) is where your net under test is connected.

You must hold the pogo pins firmly on the test points for minimum contact resistance. Note that temperature changes of the pogo pins will also result in pitch changes, so if you have just soldered the BUT or installed a new pogo pin onto the probe assembly, give them time to stabilise to room temperature. Also, do not touch the pogo pins during use. The warmth of your fingers will change the resulting beep pitch.

Differential amp has 6dB lower noise, twice the bandwidth

Roy McCammon

The traditional three op-amp differential amplifier's signal to noise ratio can be improved by 6 dB by adding a resistor and slightly changing the connections. There is a trade-off though: The traditional topology has a high input impedance, whereas the low-noise version has a lower input impedance.
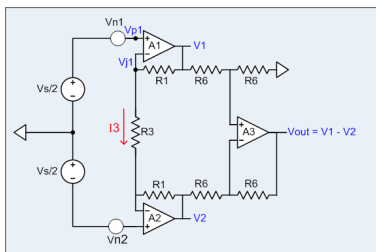

Figure 1 depicts a traditional three op-amp differential amplifier. For the sake of simplicity we will assume that A3 is an ideal noiseless opamp and the four resistors labelled R6 are exactly equal, so that the output of A3 is given exactly by V1

- V2. The A3 circuit is a perfect subtractor. We will also assume that the two resistors labelled R1 are exactly equal. The outputs are given by:

$$V1 = + (Vs/2)(1 + 2r) + Vn1 (1 + r) + Vn2 (r)$$

$$V2 = - (Vs/2)(1 + 2r) - Vn2 (1 + r) - Vn1 (r)$$

$$Vout = Vs (1 + 2r) + Vn1 (1 + 2r) + Vn2 (1 + 2r)$$

Where:

r = R1/R3
Vn1 = voltage noise from opamp A1, RMS value = Vn
Vn2 = voltage noise from opamp A2, RMS value = Vn

If R1 = 20kΩ and R3 = 1kΩ, then:
    Vout = 41Vs + 41Vn1 + 41Vn2

Vn1 and Vn2 are random, so that the polarity can be chosen for convenience. The signal and the noise sources all have the same gain, so the signal to noise power ratio is $Vs^2 / (2 Vn^2)$.
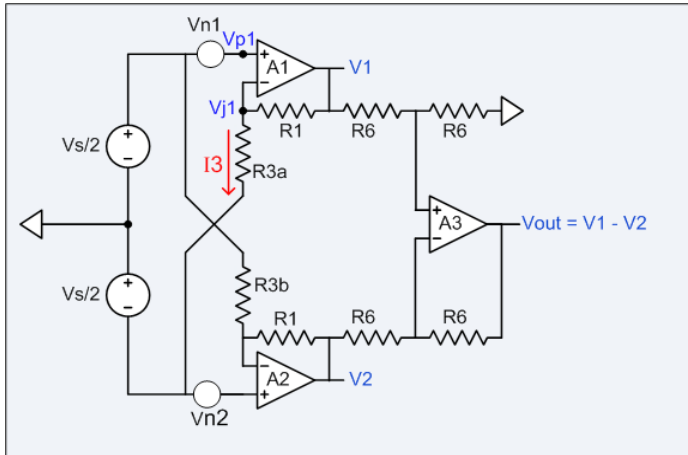

Figure 2 depicts the circuit of this Design Idea. Resistor R3 has been duplicated and its hookup modified. The outputs of this circuit are given by:

$$V1 = + (Vs/2)(1 + 2r) + Vn1 (1 + r)$$

$$V2 = - (Vs/2)(1 + 2r) - Vn2 (1 + r)$$

$$Vout = Vs (1 + 2r) + Vn1 (1 + r) + Vn2 (1 + r)$$
Where r = R1/R3

If R1 = 20kΩ, and R3a = R3b = R3 = 1kΩ then:

$$Vout = 41Vs + 21Vn1 + 21Vn2$$

The output signal to noise power ratio is $(41 Vs)^2 / ((21 Vn)^2 + (21 Vn)^2)$, which is almost equal to $2 Vs^2 / Vn^2$. That is about 6dB better than the traditional circuit. What is going on?

The answer is revealed by looking at the outputs of A1 and A2 in the traditional circuit. The output of A1 has a noise component due to its own noise, and an almost equal component due to the noise from A2, which also has noise components from both op-amps. In the new circuit, the output of each opamp has noise only from its own internal noise source.

The connection between A1 and A2 is through R3. A1's internal voltage noise source, Vn1, drives its internal non-inverting input, Vp1, which causes the output of A1 to have a noise component from its own noise source. Normal negative feedback action causes A1's summing junction to follow its non-inverting input, meaning A1's internal noise will appear on its summing junction, Vj1, which will drive a noise current into R3 that will be injected into the summing junction of A2. In this way, the noise source in A1 appears at the output of both A1 and A2. Likewise the noise source in A2 appears in the output of both A1 and A2.

The new circuit works the same way except that the noise at the summing junction drives R3a, which connects back to the low impedance source instead of the summing junction of the other amplifier.

The new circuit also has about twice the bandwidth of the traditional circuit. In the traditional circuit, some of the drive for A1 comes from A2 and some of the drive for A2 comes from A1. When the input frequency reaches the point where A1 and A2 begin to roll off, they also begin to see less drive from the other amplifier. In the new circuit, the drive comes exclusively from the input source. The net effect is that the traditional amplifier reaches its 3dB frequency about an octave lower than the new circuit.

Compared to the traditional circuit, the new circuit has all the advantages, including the same common mode rejection, less noise, and more bandwidth, but has one drawback: the input impedance is R3, which in this case would be 1kΩ. That would be adequate if the source was, for example, a balanced twisted pair transmission line (typically 100Ω).

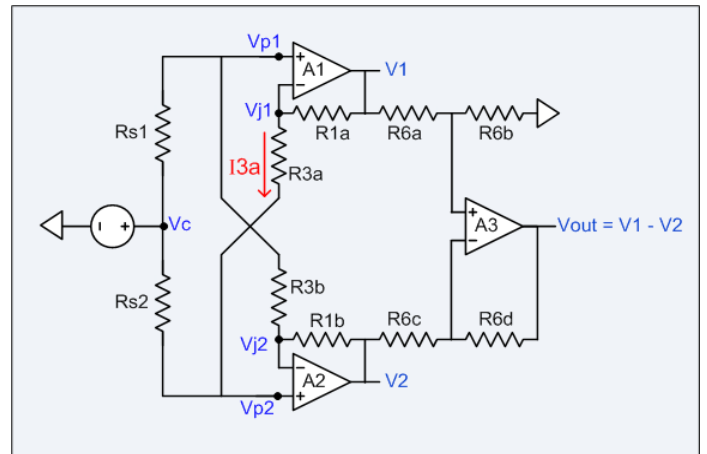Here is an update regarding the circuit's common-mode response:

Figure 3 illustrates the common mode responses of the new circuit. Vc is the common mode voltage.

If A3 and its resistors form an ideal subtractor, and the op-amps are acting ideally (common mode response is zero and differential input voltage is driven to zero by negative feedback), then overall common mode gain of the traditional circuit is zero. Surprisingly, the new circuit has the same response under the same assumptions, even if the source impedances are not matched.

Linear circuits have only one solution. If you guess the solution and all the equations are satisfied, then you have the solution. In this case, the solution is V1 = V2 = Vp1 = Vp2 = Vj1 = Vj2 = Vc, and Vout = 0. It is easy to verify that this is the solution regardless of the values of Rs1, Rs2, R1a, R1b, R3a, or R3b.

To check, I simulated a case where the source impedances were 45Ω and 55Ω, and the common mode voltage was 1V. I use ideal op-amps, except they have a finite DC gain of 107 and a gain-bandwidth of 50 MHz. To make it a little more realistic, I let A1 have a gain-bandwidth of 51 MHz, and the subtractor (A3 and resistors labeled R6) have a small common mode gain of -80 dB (equivalent to resistor matching on the order of 0.01%). Both circuits exhibited -80 dB of common mode gain from DC to about 1 kHz, at which point the common mode gain began to rise on both circuits, with the new circuit being about half a decibel better.

productroundup

Development tool sets up stepper motor parameters

A PC-based GUI control panel eases design and debug for robotics, lab automation and other motion control applications; German company Trinamic Motion Control's TMCM-1043-KIT development system is released for its integrated, NEMA 17-compatible TMCM-1043 stepDancer stepper motor module. The development kit offers designers a PC-based GUI that allows one-click modification of motor drive current, micro-stepping and other key parameters. The kit is custom designed for use with Trinamic's highly integrated TMCM-1043 electronic assembly, which is developed for direct mounting on a NEMA 17 motor. Pre-programmed and pre-configured with all operating parameters to support standard 1.5A (peak) motors, the TMCM-1043 board uses Trinamic's single-axis TMC2660 stepper motor driver IC, which integrates a motor pre-driver and power MOSFETs. Controlled by industry standard step/direction pulses, the TMCM-1043 module automatically performs all motor coil current calculations.

Complete article, here ▶

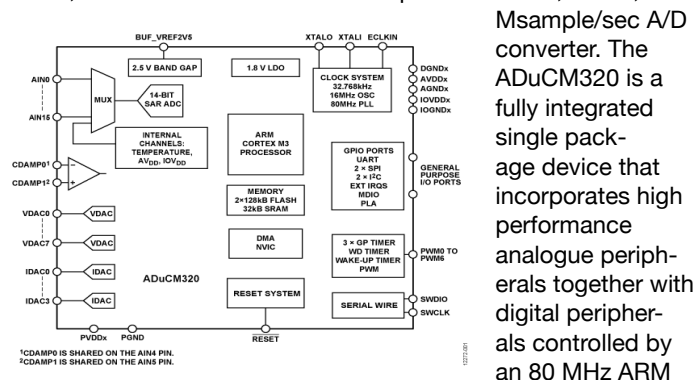## IO-Link transceiver integrates step-down regulator & LDO

LT3669 is an IO-Link PHY compatible (COM1/COM2/COM3) industrial transceiver that includes a high efficiency step-down regulator and a low dropout linear regulator. The device offers cable interface protection to ±60V. Wake-up detect functionality as well as a programmable power-on reset timer are included for system reliability. The LT3669/-2 operates from a 7.5V to 40V input voltage range, making it suitable for industrial sensor applications. The LT3669's internal switching regulator can deliver up to 100 mA of load current, whereas the LT3669-2 delivers up to 300 mA. Both versions can deliver high efficiency at output voltages from 0.8V to 16V. The integrated LDO, powered from the switcher output, provides an additional output to offer additional design flexibility while maintaining high efficiency. The LT3669EUFD and LT3669EUFD-2 are both in a 28-pin thermally enhanced 4 x 5mm QFN package from $3.55 (1000).
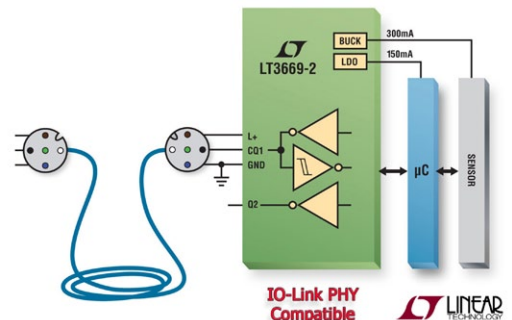
Complete article, here ▶

Precision analogue functions in a Cortex-M3 MCU

Analog Devices has posted details of an addition to its microcontrollers; for use in areas such as industrial control and automation systems, smart sensors, and precision instrumentat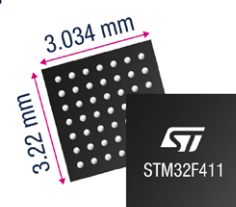ion, the ADuCM320 has an on-chip multichannel, 14-bit, 1 Msample/sec A/D converter. The ADuCM320 is a fully integrated single package device that incorporates high performance analogue peripherals together with digital peripherals controlled by an 80 MHz ARM Cortex-M3 processor and integral flash for code and data. The ADC on the ADuCM320 provides 14-bit, 1 MSPS data acquisition on up to 16 input pins that can be programmed for single-ended or differential operation. The voltage at the IDAC output pins may also be measured by the ADC, which is useful for controlling the power consumption of the current DACs. Additionally, chip temperature and supply voltages can be measured. Up to eight voltage D/A converters (VDACs) and four current D/A converter (IDAC) sources are provided.

Complete article, here ▶

ST's Dynamic Efficiency MCUs with Batch Acquisition Mode

STM32F411 STM32 Dynamic Efficiency MCUs extends the family with more advanced power-saving features including Batch Acquisition Mode (BAM): in smartphone sensor-hub applications, BAM enhances the energy savings made possible by Android KitKat operating system, and BAM extends power savings beyond smartphones into consumer, industrial and medical devices. STM32 Dynamic Efficiency MCUs improve the power-saving performance of data batching – the technique also used in Google's latest Android 4.4 (KitKat) operating system to maximise battery life - and extend the advantages to many more applications besides smartphones and tablets. Android 4.4 uses a low-power sensor hub to manage "always-on" sensors such as accelerometers or pressure sensors, allowing the main system processor to consume less battery energy. STM32F411's Batch Acquisition Mode (BAM) saves up to 50% extra power by storing sensor data directly into SRAM while its own CPU core sleeps.

Complete article, here ▶

productroundup

Microcontrollers major on capacitive sensing for HMI applications

Silicon Labs' energy-efficient capacitive sensing microcontrollers (MCUs) are aimed at human-machine interfaces (HMI). The C8051F97x MCU family combines Silicon Labs' ultra-low-power technology with the industry's fastest, most accurate capacitive sensing to provide touch control solutions for the Internet of Things, home/building automation, consumer and industrial markets. Silicon Labs' F97x MCUs offer lowest energy consumption in active, sleep and deep-sleep modes, claiming the longest battery life of any 8-bit capacitive sensing MCUs. With 200 µA/MHz active current, the F97x MCUs combine low energy consumption and system performance. The MCUs' 2 µsec wake time minimises energy consumption while transitioning from sleep to active mode. The F97x MCUs offer the lowest sleep mode energy consumption in their class: 55 nA sleep current with brownout detector enabled and 280 nA sleep current with a 16.4 kHz internal oscillator.

Complete article, here ▶

Reduced acoustic noise from ceramic capacitors

A range of MLCCs from Murata have an extra interposer substrate that reduces acoustic "squealing" noise. The ZRB series of monolithic ceramic capacitors (MLCC) is packaged on an interposer substrate designed specifically to reduce acoustic "squealing" noise typically induced by mechanical vibration of the capacitor. Available in the same size as conventional MLCCs, the ZRB is available in EIA 0402 (1.0 x 0.05 mm) and EIA 0603 (1.6 x 0.8 mm) package formats with working voltages of 6.3, 10, 16 and 25 VDC, With this approach, the ZRB becomes a replacement part to update an end-application design without the need for modification of the PCB layout. This form of acoustic noise has become of concern for the electronics industry and affects many types of consumer electronics devices such as laptops, tablets and smartphones. The ZRB series is available with capacitance values of 4.7, 10 or 22 µF with X5R temperature characteristics.
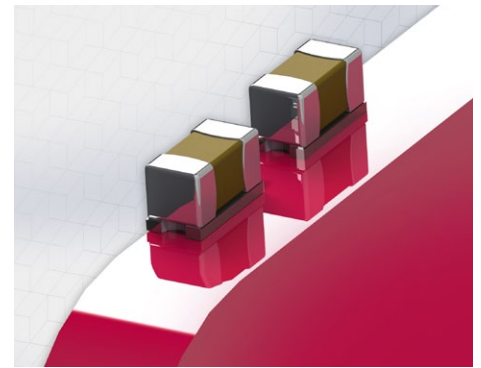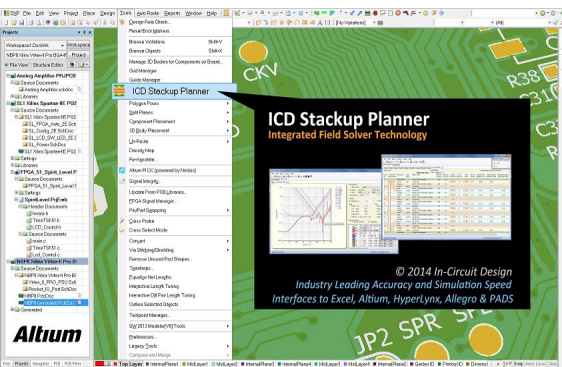
Complete article, here ▶

Field-solvers in Altium Designer for high-speed design

Altium, in cooperation with Australian based In-Circuit Design (ICD), has developed extensions for Altium Designer for advanced stackup planning and power distribution network analysis to bring comprehensive high-speed design capabilities to the mainstream market. With the increasing challenges concerning high-speed signals – not only because of high clock frequencies, but also because of faster edge rates – more and more PCB designers need to have analysis tools that allow them to successfully design with fewer iterations. The two new extensions for Altium Designer, the ICD Stackup Planner and ICD Power Distribution Network (PDN) Planner, are accessible from within the design tool to provide for seamless analysis. ICD provides a centralised, shared, impedance planning environment: ICD PDN Planner analyses the AC impedance of each on-board PDN, including capacitor selection, to ensure a broad spectrum of noise reduction.

Complete article, here ▶

Industrial-strength driver is Darlington replacement

This 7-channel, NMOS low-side driver replaces Darlington transistor arrays in high-voltage systems; the drop-in compatible relay driver reduces power, cost and board space. Texas Instruments' TPL7407L replaces half of the transistor arrays required to drive high current loads, providing a new option for high voltage systems that previously required a number of transistor arrays or a motor driver. Reducing power by 40%, this new device drives the LED matrix, relay or stepper motor in high-voltage applications, such as white goods, building automation, lighting and HVAC. When used with the SN74HC595 register, one or multiple TPL7404 can be controlled with three GPIO pins; the TPL7407L handles drain current of 600 mA per channel, reducing the number of relay drivers required per board. It supports energy-efficiency in high-voltage systems, reducing power dissipation by 40% compared to Darlington arrays, and is pin-to-pin compatible with traditional arrays.

Complete article, here ▶

Digital point-of-load regulator employs embedded dynamic loop compensation

With output uprated to 50A, this regulator offers high power density of 32.61 W/cm3 (534 W/in3) and 25% more output current compared to previous version and in the same footprint. Dynamic Loop Compensation improves stability: paralleling of seven modules with phase-spreading delivers 350A, while ripple-and-noise is reduced to the lowest possible level. Ericsson's 3E series BMR464-50A is a third-generation digital point-of-load (POL) regulator with a full set of PMBus commands enabling systems architects to fully monitor and dynamically control the energy delivered to strategic components, such as processors, FPGAs and ASICs, down to a very low and highly economical level. Embedding the latest Dynamic Loop Compensation technology, the BMR464-50A runs the DLC algorithm as default following the enabling of the output. However, three more settings are also available via the PMBus.
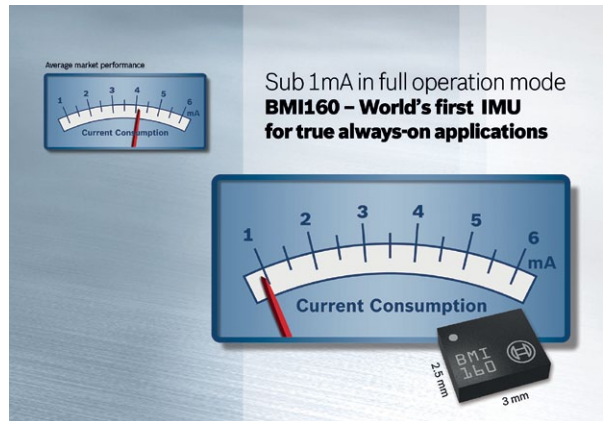
Complete article, here ▶

Inertial sensor unit runs on under 1mA

This IMU (inertial measurement unit) combines high accuracy, low noise, current consumption and footprint; the high precision 6-axis IMU occupies a small footprint, enables always-on applications for wearable devices and supports precise 9-axis sensor data fusion computation. Bosch Sensortec's BMI160 Inertial Measurement Unit (IMU) integrates a 16 bit 3-axis, low-g accelerometer and a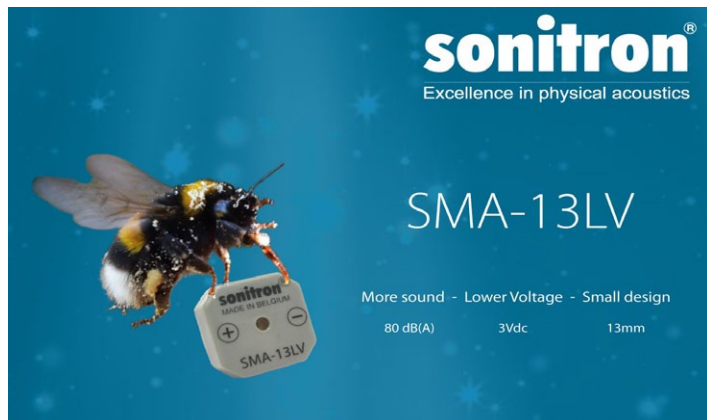n ultra-low power 3-axis gyroscope into a single package. It has been designed specifically for high precision, always-on 6-axis and 9-axis applications in smart phones, tablets, wearable devices, remote controls, game controllers, head-mounted devices and toys. The BMI160 is available in a 14-pin 2.5 × 3.0 × 0.8 mm3 LGA package. When the accelerometer and gyroscope are in full operation mode, the typical current consumption is 950 µA.

Complete article, here ▶

Loudest miniature buzzer, claims Belgian manufacturer

The SMA-13LV joins Sonitron's multi-application buzzer range. With only 3 Vdc it adds an alerts to your appliance with a sound of 80dB(A) at 30 cm. This buzzer is ideal for low voltage applications that need a small design and an alerting sound. It is available in SMD and pin version, with 7.5 mm or 10mm spacing. Adding a product option can strengthen the buzzer for use in specific circumstances. A foam patch option can be added for extra mechanical and acoustic stabilisation. The Wash tab option secures and protects the buzzer during assembly in an automatic washing process. It is, say the makers, a very loud buzzer for miniature and portable equipment.

Complete article, here ▶

Tx/Rx pair converts analogue to HD video

Intersil's HD-SDI transmitter and receiver support upgrading from analogue to HD video while using previous-generation coax cabling. The TW6872 Triple-rate SDI transmitter and TW6874 quad receiver lower system costs and extend reach for video applications; they use Dirac VC-2 mezzanine compression, an end-to-end solution enabling latency-free video over extended cable distances. The new devices are designed to conform to the SMPTE standards for SD, HD and 3G serial digital transmission, for HD video transmission at SD-SDI cable distances of 300+ metres. The TW6872 transmitter with low-jitter clocking (0.09 UIpp) and an integrated cable driver with pre-emphasis enhances output signal transitions to achieve superior cable reach. The TW6874 receiver with an adaptive equaliser compensates for frequency-dependent cable attenuation, delivering longer cable reach in combination with the TW6872 or other HD-SDI transmitters.

Complete article, here ▶

productroundup

Wireless connectivity "to the cloud" comes in pre-configured WiFi kit

Mouser Electronics has the Ayla IoT Design Kit, equipped with a Murata Wireless Wi-Fi connectivity module. This design and development kit allows you to securely connect devices to the cloud from anywhere in the world with an internet connection. The Ayla Design Kit with Murata WiFi Connectivity allows developers to easily connect their projects to Ayla's cloud service. Wireless connectivity is supported by a Murata Type-YD 2.4 GHz 802.11b/g/n radio module supporting WEP, WPA-PSK, and WPA2-PSK encryption. The Murata Type-YD module includes a TCP/IP stack, security firmware, and other network application features. Murata's Wi-Fi module mounted on Ayla's design kit allows devices to be securely controlled using OAuth-based authentication from anywhere. Developers can use these technologies to provide interactive control of industrial systems, lighting applications, HVAC, and more, all with minimal modifications to existing systems.

Complete article, here

Upgrade adds new functions to free circuit design software

Components distributor Digi-Key has added enhancements to its Scheme-it circuit design tool, available at no cost on the Digi-Key website. The tool was co-designed and built by Aspen Labs, a business-media company focused on the needs of engineers. Capable of being used as an "idea generator", Digi-Key's Scheme-it design tool provides users with a simple, free-to-use way to record their circuit design idea in a shareable, electronic form. The tool implements the entire Digi-Key catalogue, allowing users to design with actual parts available for immediate shipment. Scheme-it comes equipped with an enhanced feature set, including: The ability to diagram at the Block, Icon, System, or Schematic level; A library of over 700 generic symbols, as well as custom symbol creation; Access to over 4 million components; Freedom to keep designs private, make public, share via link, or embedded into web pages, blogs or emails; Rapid design evolution via Bill of Material (BOM) upload capability; A direct link to Digi-Key Technical Support.

Complete article, here

5 kΩ digital potentiometers for 36V operation

Microchip has expanded its 36V digital potentiometer (digipot) portfolio with two new volatile-configuration, I2C devices—the MCP45HV31 and MCP45HV51 (MCP45HV31-51). These are the first digipots to offer a 5 kΩ resistance with a specified operating voltage of 36V. They provide 10V to 36V anal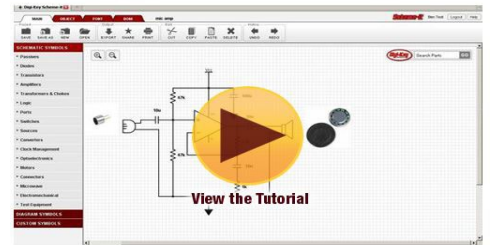ogue operation and 1.8V to 5.5V digital operation, for systems requiring wide signal swings or high power-supply voltages. The MCP45HV31-51 digipots support both 7-bit and 8-bit resistor configurations, and a high terminal/wiper current, including the ability to sink/source up to 25 mA on all terminal pins for driving larger loads. The MCP45HV31's 7-bit resistor network resolution enables 127 resistors and 128 taps, while the MCP45HV51's 8-bit configuration supports 255 resistors and 256 taps. Both digipots provide RAB resistance options of 5, 10, 50 and 100 kΩ. Both devices also feature a 1 µA typical serial-interface inactive current.

Complete article, here

Industrial automation protocols plus ARM cores

TI positions its latest processor family as offering higher performance with real-time processing; the Sitara AM437x processors, based on an ARM Cortex-A9 core, provide options for enhanced processing performance and industrial connectivity. Sitara AM437x processor family members integrate support for industrial protocols for both automation and industrial drives and include new features such as dual camera for data terminals with bar code scanning. AM437x processors enable real-time processing with available quad core programmable real-time units (PRU). The PRU offloads real-time processing from the ARM to manage deterministic tasks such as controlling motors and is robust enough to enable complex functions like multiple industrial fieldbus protocols. A quad core PRU-ICSS (industrial communications subsystem) connectivity peripheral enables dual, simultaneous industrial protocols such as EtherCAT, EtherNet/IP, Profibus, PROFINET-RT/IRT, POWERLINK, Sercos III, IEC61850, as well as motor feedback protocols such as EnDat or interfaces for sensors and actuators such as BiSS.
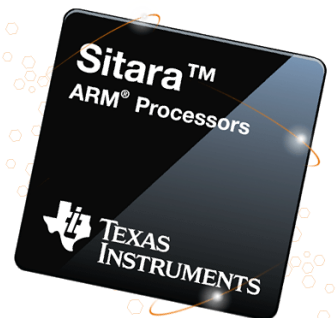
Complete article, here

The need to succeed

At the Design Automation Conference (DAC) two years ago, my company Oski Technology decided to demonstrate the value of formal verification technology (End-to-End formal in particular) by putting the technology to the test verifying a sight-unseen design in 72 hours. It was under the watchful eye of a CCTV camera publicly streaming to our DAC booth the entire time. To say this was stressful is an understatement! It took guts and trust in formal to do something like this.

While formal verification technology has been around for a couple of decades, its usage in model checking is still limited. Companies may have used automatic formal checks or formal apps, and a few advanced users might use formal to verify embedded register transfer level (RTL) assertions to find bugs.

We verify complete functionality of a designer's size blocks from inputs to outputs. While this may seem simple and obvious, the inherent complexity formal verification tools have to handle often renders formal verification results incomplete without proper techniques and methodology. The reward of End-to-End formal is formal sign-off with no bugs being left behind, and improved verification efficiency and productivity.

That's what we were promoting at DAC with the 72-Hour Verification Challenge. I was the formal verification engineer assigned to this task and worked from a hotel room near the Moscone Center where DAC was held.

I was faced with challenges on two levels. The engineering challenge was, as anyone who has used formal verification technology will know, allowing only 72 hours to deliver measurable results is an extremely compressed timeframe. Even for a known design, sometimes it takes more than that to compile the RTL code files. In this case, I was expected to understand the design, compile the design, identify and code input constraints, think of, then write, End-to-End checkers for the design, run End-to-End checkers using formal and Cadence's IEV, and find bugs in the design. I wondered if I would get any sleep.

The psychological challenge was that I was on camera throughout. As an engineer, I am not used to working under the spotlight. Lack of results would be a public setback for the company and formal verification technology. All of this weighed on me.

Going into the challenge, I knew to succeed I must have a well-defined goal to focus on. I needed to do all of the above, and generate a few interesting witnesses or counter examples for the challenge to be considered a success. Additionally, if I could find a corner case bug, identify a performance issue or even get a checker to pass — difficult under any circumstance — that would be considered a bonus.

To get there, I needed to control the panic and disregard the desire to take the first flight home, feign illness, or anything else. I needed to stick to basics. To calm myself, I went through the process, which meant spending a significant amount of time, precious under the circumstances, in formal test planning. All the while, I forced myself to resist the strong urge to dive into the RTL code and start coding the testbench.

Just before the clock started ticking, I got a 10-minute overview of the design I was to debug from the NVIDIA engineer who provided it. Then it was game on.

I started my formal test planning phase that consisted of reading the specification thoroughly to learn the design intent and developed a list of constraints and End-to-End checkers. Most importantly, I worked to understand the complexity of the design and efforts to verify these checkers to decide what could actually be done in the timeframe. I ranked the list of checkers by formal return on investment (ROI) and what was important to verify in the design to find corner case bugs versus what could be accomplished in the short time window. This took a good few of hours, but the exercise was important to having a laser-sharp focus on what to target.

Then, I started executing the plan. I compiled the design and discovered library files were missing. I worked around that. I familiarized myself with the design by writing covers. I then wrote initial constraints, and the End-to-End checkers with the highest formal ROI. By the end of day one, I had some counter examples. The next day, after some sleep, I started debugging and realized a bug was due to missing constraints. I added more constraints, and worked through complexity issues with the proper techniques. The second day ticked by as I was in the thick of the verification phase.
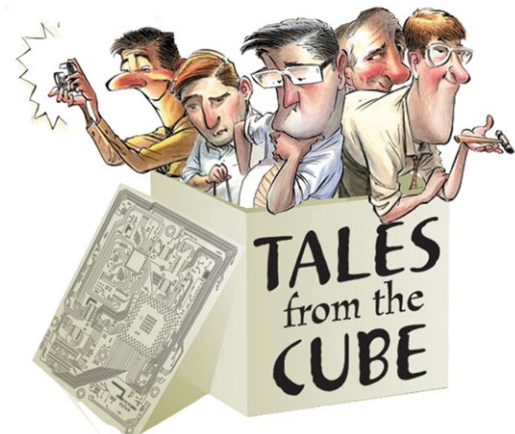
I was getting more nervous because there wasn't much time left and I understood why people lose sleep as tapeout deadlines draw near with bugs left to be found in the design. I felt the same way that night, but I knew that if I stuck to the plan and executed flawlessly, I would get to the end goal.

Sure enough, on third day, I started getting some real counter examples. Not one bug, but three all together in a design that was mature and close to tapeout. Needless to say, the engineer was happy when he confirmed the results were real bugs missed in NVIDIA's simulation-based verification environment.

Throughout the 72 hours, interested onlookers strolled in. Those who had done formal verification before would stand around for a while, pat me on the back and walk away shaking their heads. Some wouldn't believe I hadn't seen the design before, convinced that it was just a show, but I knew I had been put through a real challenge. A challenge not just of my own formal skills, but the technology itself and the promise it brings to the design and verification community to solve real verification challenges.

The 72-hour challenge ended well. I was able to report three bugs that necessitated changes in the RTL code. I've been asked what can be learned from this challenge, and I'd say don't start formal verification execution without formal test planning, no matter how little time. The time spent on planning will pay off.

Chirag Agarwal is an architect at Oski Technology with 15 years of experience verifying complex IC designs. A formal verification expert, he has led more than 20 formal verification projects at leading semiconductor companies, and his expertise ranges from processor verification and cache controller verification to interconnect verification.