EDN europe

APRIL 2014
www.edn-europe.com
Issue 4

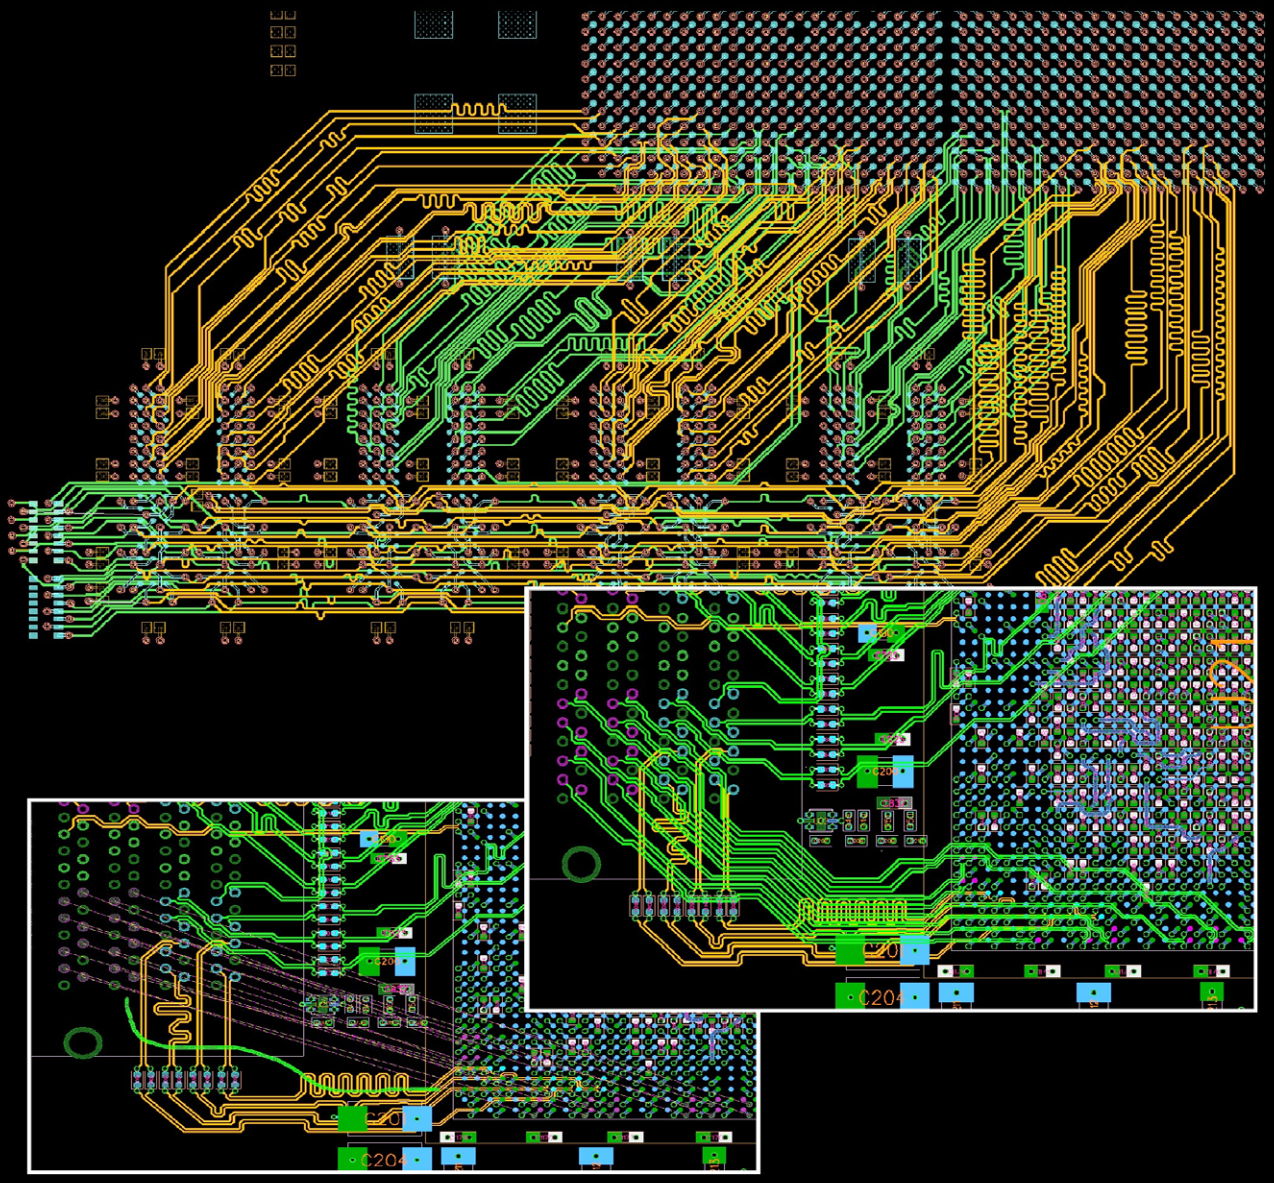

EDN europe

contents

APRIL 2014

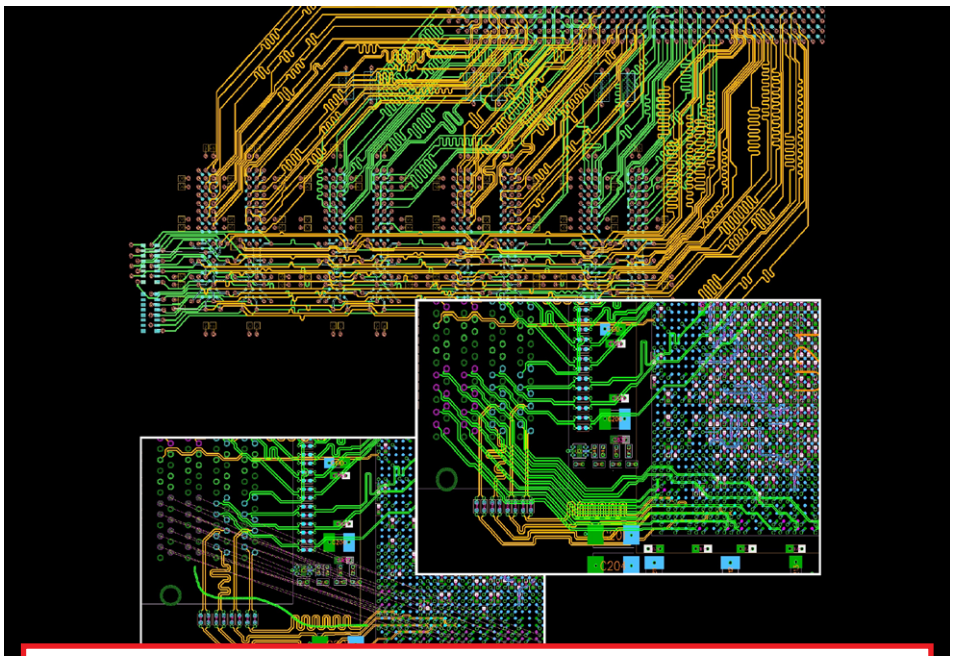

Cover Image
this month's cover is a composite of screen shots issued by Mentor Graphics to display some of the routing capabilities of its Xpedition PCB software release (see page 8). The background picture shows some of the high-frequency routing features such as serpentine traces to equalise signal paths. The lower pictures show "Sketch routing"; below, left, the user has indicated with a "swipe" (green) the approximate area where the traces in a group should be placed: lower right, the software automatically completes according to that guidance.

FEATUREARTICLES

EDN. COMMENT

BAKER'S BEST

Spin Cycle

TALES FROM THE CUBE

productroundup

DESIGNIDEAS

pulse

Easy. Powerful.
The R&S®RTE:
Simply more scope.

Quickly master everyday T&M challenges.
More confidence in your measurements, more tools and fast results, more fun to use – that's the R&S®RTE oscilloscope from Rohde&Schwarz. From embedded design development to power electronics analysis to general debugging, the R&S®RTE offers quick solutions for everyday T&M tasks.

See for yourself!
Click here for more information:
www.scope-of-the-art.com/ad/rte

ROHDE&SCHWARZ

EDN europe

APRIL 2014
www.edn-europe.com
Issue 4

CONTACTS

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

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

Patrick Mannion
Brand Director EDN Worldwide

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

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

ART MANAGER
Jean-Paul Speliers

ACCOUNTING
Ricardo Pinto Ferreira

european business press

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

EDN-EUROPE is published 11 times in 2014 by European Business Press SA,
7 Avenue Reine Astrid, 1310 La Hulpe, Belgium
Tel: +32-2-740 00 50 Fax: +32-2-740 00 59
email: info@eetimes.be. VAT Registration: BE 461.357.437.
RPM: Nivelles.
It is is free to qualified engineers and managers involved in engineering decisions – see:
http://www.edn-europe.com/subscribe

SALES CONTACTS

Europe

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

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

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

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

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

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

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

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

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

ALL TOGETHER... NOW?

For as long as we have been designing embedded systems with software running on microprocessors or microcontrollers, managing the parallel flows of hardware and software development has been a challenge; a key part of the challenge being, naturally, to have those two flows converge to completion at about the same time. When this fails to happen, it's almost always the case that software lags hardware; either software development awaits the availability of hardware, or it proceeds and when hardware does show up, the code fails to run on it as intended.
This is equally true of board-level designs as it is with ASIC or SoC exercises; it has attracted more attention in the case of IC designs as the costs of hardware re-work (to get around "issues") are so crippling. Thus, for all that time – and I'll let you choose your own starting point – the desirability of having some early – but realistic – representation of your nascent hardware on which to run your embryonic software, has been all too apparent.

The EDA industry has responded to this need with several generations of products, and there has been a certain cyclical aspect to its offerings. "Hardware/software co-design"; "concurrent design"; virtual prototyping"; all proposed, and in many cases used, as a solution. At some stages, advances in simulation held out the possibility of hosting the early hardware design entirely in software and running application software on top of it. The classic fault-line in that approach is the trade-off between accuracy and speed – in general, pure-software simulations for this purpose ranged from not-quite-fast-enough to way-too-slow. And, every time simulators looked like they might get fast enough to be usable, design complexity bounded ahead. So we have seen part-hardware solutions, and fully-hardware-based offerings (emulators and their companions) in which your hardware is represented on an array of programmable logic so that it can run at some manageable percentage of real-world speeds. The emulation business has long been a welcome customer for the biggest and fastest chips, in successive generations, that Altera and Xilinx have had to offer.

You might buy complete system from an EDA vendor; or a board from a third-party with FPGAs all ready to go; or you might build your own emulation of your intended hardware, from scratch. I've seen, or heard of, all three. Their utility tends to revolve around the project manager's dilemma; how much effort must I divert to creating this facsimile of my eventual product? Will that facsimile be accurate enough to see a return on that effort, or would my team's resources be better used pushing ahead with the main product development? Will I really short-circuit a nightmarish debug episode and have v1.0 of the software run sweetly on v1.0 of the hardware?

As I said earlier, this issue has come to the forefront many times, not unconnected to the coming to market of assorted generations of emulator hardware – the main EDA vendors have always been happy to sell you a box of logic about the size of a washing machine (sometimes, they also looked like one) plus accompanying methodologies for transferring your design intent, into the box. And to the accompaniment, at least as far as the associated presentations to technical editors have been concerned, of a remarkably similar set of Powerpoint slides. I have sat through many talks delivered to slide sets that depict the twin development flows, and their failure to converge until the later stages of a project, with an accompanying narrative that emphasises the escalating cost of not identifying the fault-lines in your design until late in the cycle. Sometimes, they would include the stock Powerpoint "brick wall" slide, to depict the two disciplines operating in isolation. With a degree of irony, it's notable that today's presenters lack the historical perspective on how many times the message has been aired before.

It's apparent that no one vendor has captured this market, indeed, that no one strategy has managed to do so. Generations of product have come and gone, some have sold well, all of them have eventually been outpaced by design complexity. When you factor in the number and variety of "home-grown" solutions that have been deployed, the range of solutions to tightening the connections between hardware and software flows is wide indeed. Developments such as multicore designs has only added to the problem.

Right now, the message is out there again, with the EDA companies adding an extra push, saying that this time there's a degree of maturity to the offering that increases its usability and gives greater confidence that it will deliver substantial returns. Cadence, for example, reports healthy business in its Palladium product line – aimed primarily at those the big-chip developer, it promises to emulate "HW/SW designs at up to MHz speeds – months before silicon tapeout". Most recently, comes a book from Synopsys (available as a free download, details here); "Better Software. Faster!" has input from more than 10 contributing companies on the topic of advanced software development with virtual prototypes, and presenting a practical guide for using virtual prototypes to develop, test and debug software up to 12 months before hardware availability. The text includes case studies from mobile, consumer, industrial and automotive companies that are highly experienced using virtual prototypes for software development and hardware/software integration. It promises that you will better understand software development methods that will result in products with higher quality and reliability, and reduce costs.

That has long been the promised outcome of adopting this strategy of product design; maybe the most remarkable aspect of the subject is that the outline – even, the detail – of the promotional message has changed so little over time. Can a standard offering ever serve a big fraction of this market, or will there always be a diversity of approaches, with individual teams blending their own solutions?

pulse

Precision 50A sensor in 7-mm sq. IC-style packaging

TLI4970 from Infineon is a high-precision, AC/DC, ±50A current sensor that claims to resolve many compromises in today's current measurement circuitry, and which requires only a sixth of the board space taken up by existing sensors on the market today.

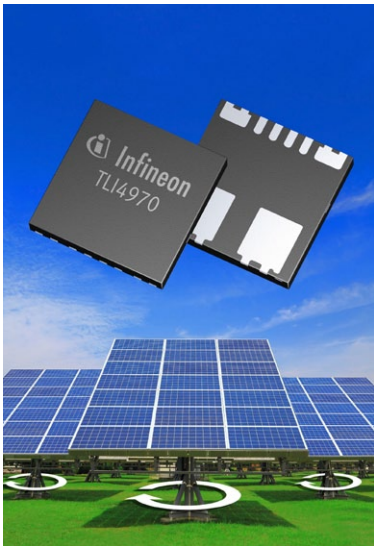

The fully digital sensor does not require external calibration. Thanks to the stray-field suppression it incorporates, the sensor is extremely robust against external magnetic fields. It continues to measure with absolute precision even after years of continuous operation. It also provides additional functions such as fast overcurrent detection at a pre-configurable level.

The TLI4970 is suited for use in solar inverters, charging devices and power supplies as well as for electric drives and for controlling energy-saving LED lighting units.

The measuring principle of the TLI4970 is based on Infineon's implementation of Hall effect technology. This involves using integral Hall probes to detect the magnetic field of a current-carrying conductor. The TLI4970 dispenses with the field concentrators that other designs require and, as a result, avoids hysteresis effects, which can lead to measuring inaccuracies. The differential measuring principle integrated in the TLI4970 suppresses interference caused by external magnetic fields, achieving an extremely low offset of 25 mA.

Besides high-precision current measurement, it is also possible to implement highly efficient protection of a power output stage. In the event of external short circuits, critical overcurrents can occur. To keep the latency extremely short, the TLI4970 provides a parallel signal path. This allows the error condition detected by the sensor to be processed with an extremely short delay of less than 3 µsec directly in the driver stage or in the microcontroller. The TLI4970 is one of the first current sensors to transmit the measured values via the digital SPI interface. It integrates differential amplifiers, filters and signal processing and can carry out galvanically isolated measurements up to 600V operating voltage and up to 3,600V test voltage.

Complete article, here

Add touchscreen display module starter kits to Raspberry Pi

Distributor RS Components has two touchscreen display starter kits for the Raspberry Pi; the kits for the 2.8-in. uLCD-28PTU-PI and 4.3-in. uLCD-43PTU-PI modules have been developed by 4D Systems and provide all the components needed for users to connect their Raspberry Pi to the touchscreen displays and get their project up and running and ready for programming within a few minutes. The kits feature the 4D Pi Adapter Shield, which connects to the serial GPIO port on the Raspberry Pi and to the touchscreen display module via a 5-way cable. The Adapter Shield also replicates the Raspberry Pi GPIO lines and allows users to gain access to other Raspberry Pi functions.

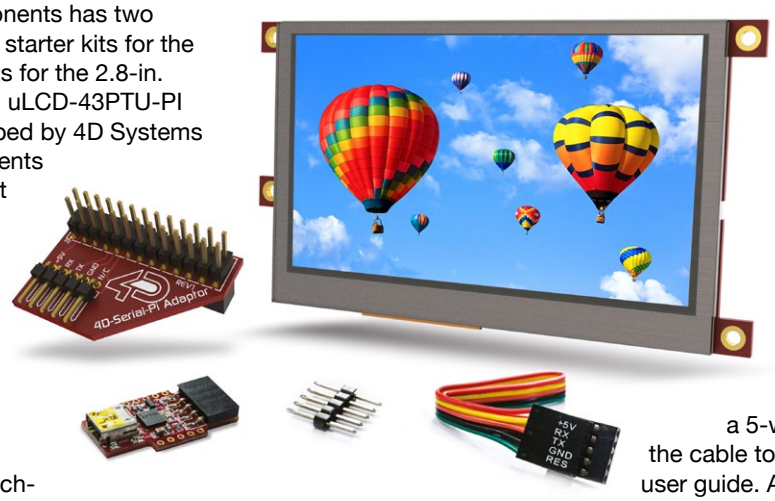

These displays deliver a powerful HMI (Human Machine Interface) solution for Raspberry Pi applications featuring resistive-touch capability, on-board audio amplifier and speaker, and resolutions of 480 x 272 and 240 x 320 pixels for the 4.3- and 2.8-in. screens respectively. Each starter kit contains: a 4D serial Pi adaptor shield; a 2 GB microSD card; a uUSB-PA5 programming adaptor; a 150 mm 5-way female-female jumper cable for quick connection to another device or a breadboard; a 5-way male-male adaptor to convert the cable to male-female; and a quick-start user guide. A series of ViSi-Genie application software libraries are also available to communicate with the Raspberry Pi board and user code. The two display options are priced at around £60 and £95 (UK Pounds, or local equivalent)

Complete article, here

Agilent Technologies

Agilent's Electronic Measurement Group, including its 9,500 employees and 12,000 products, is becoming **Keysight Technologies**.

Learn more at **www.keysight.com**

Mentor Graphics renews high-end PCB design tools: starting with layout

Mentor has embarked on a revision of its printed-board design software at the high-end, with the re-launch of its Xpedition platform. Mentor says it is addressing today's printed circuit board (PCB) systems challenges of increased complexity, changing team demographics, and systems-aware design requirements

Over time, the Xpedition platform will see introductions in a number of different technology domains – signal integrity, power-awareness, and thermal design, for example – but this step focusses on layout and routing. Director of Engineering at Mentor's System Design Division, Charles Pfeil, emphasises that although this introduction brings to market work that has been under way as long as six years, it is based on prior Mentor products and in that sense is a next-generation continuation, for existing users.

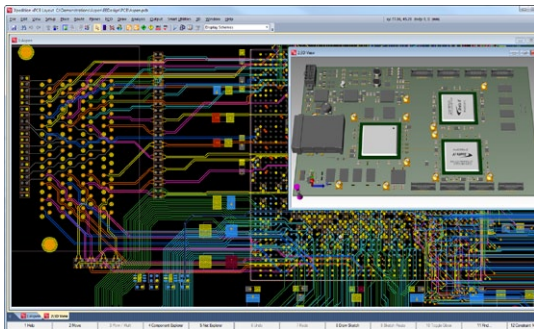

The Mentor team says that the methodologies embodied in Xpedition recognise the changing way that complex PCBs are designed; the traditional "PCB engineer" has largely disappeared and the specialist knowledge of the circuit designer has to be applied directly to the board design process in a "systems-aware" manner that also connects to manufacturing. An intuitive design environment makes use of hierarchical groupings of components and functions, with full awareness of connectivity. Many of the features of Xpedition look familiar – initial placement of functional groups, display of un-routed connection density to effect placement, floorplanning and routing automation – but Mentor says that each has new depths of features, effectiveness and usability; the company terms it as "designer-driven automation". A key feature in this platform is Sketch Router which provides the designer with extensive interactive control of the auto-assisted routing process by delivering hand-routed-quality results in dramatically reduced time. With a simple "swipe" you can indicate to the router the general path you want it to take – it will then take over and complete the routing. Detail such as route-length adjustment to minimise skew (adding serpentines where needed); creation of balanced differential-pair tracks, with phase-match tuning, and overall strategies of minimising vias, are all implemented and updated in real time. Route edits by push-aside are all implemented, even when moving many, complex tracks, with all other constraints and design rules active at all times. The illustration on the cover of this issue shows typical routing patterns.

The system includes a 3D mechanical kernel, to impart 3D spatial awareness, that Mentor sourced from CAD provider Catia. This draws on libraries of millions of components and their outlines; "the most important part of 3D checking is in placement," Pfeil says, noting that the system uses the same information regarding groups of components – with their connectivity – that the native Xpedition modules employ. At all times the software monitors designs rules and constraints and flags them as either optimal, or "minimal" - meaning that the resutl is acceptable but possibly at the cost of reduced yield. You can move items in three dimensions and the system will re-route them: Mentor is, Pfeil adds, working on a full "3D routing" solution.

Complete article, here

DC/DC IC provides 1A of carry-over current

LTC3355 is a complete ride-through DC/DC system IC that features a main buck regulator with a built-in boost converter for temporary backup, or ride-through, of V_{OUT} during a sudden loss of V_{IN} power, from a single supercapacitor energy source. The device includes all functions necessary to provide seamless charging of a supercapacitor (or other storage element), including the monitoring of V_{IN}, V_{OUT} and V_{CAP} and automatic switchover to backup power. The LTC3355's nonsynchronous constant frequency current mode monolithic 1A buck switching regulator provides a 2.7V to 5V regulated output voltage from an input supply up to 20V. The device is suitable for ride-through, "dying gasp" or data backup supplies commonly used in power meters, industrial alarms and solid state drives.

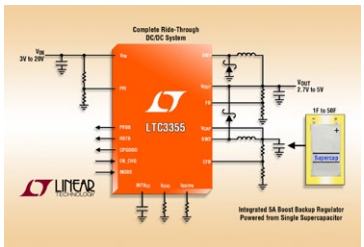

Complete article, here

ST schedules 200°C SiC MOSFET production

STMicroelectronics has announced that it will soon begin production of the first device in a family of silicon-carbide high-voltage MOSFETS, and asserts that his makes it "among the first" companies to commercialise SiC power MOSFETs, with a 200°C rating for more efficient, simplified designs. Major features of SCT30N120 include; On-state resistance (RDS(ON)): 80 mΩ typ. at 25°C, and ≤100mΩ typical over the entire temperature range to 200°C; low turn-off energy and gate charge (ensures efficient, high-speed switching); and leakage current lower than 10 µA typ (enhances system energy efficiency and reliability, compared to other structures based on the same material). The transistors have a very fast intrinsic and robust body diode (saves an external freewheeling diode for cost/size reduction); and use simplified gate drive circuitry (reduces costs of network driving).

Complete article, here

Synopsys updates IC place-&-route flagship, with IC Compiler II

For a new generation of the company's physical-design place-and-route tool for complex SoC designs in leading-edge silicon processes, the EDA vendor claims as much as a 10-fold increase in physical design throughput.

IC Compiler, Synopsys says, is the biggest single product in its overall offering: in looking to renew it, the question the company asked was, "what more [in P&R] can we do?". Built from the ground up on a new, multi-threaded infrastructure, IC Compiler II introduces ultra-high-capacity design planning, new clock-building technology and advanced global-analytical closure techniques. As a major step up from the existing IC Compiler, the new tool will be offered in parallel; the existing suite will continue to be supported and developed. IC Compiler II will be most attractive to those working at the leading-edge nodes; Synopsys says that its collaborators who have helped develop it have employed it at 28 nm and below, "...and one at 45 nm." Users will migrate from IC compiler, "according to design need and at a time of their choosing." However, Synopsys sees many of its customers opting to stay at more mature process nodes, for longer, as the increased complexity of moving to a new node can result in a diminishing-returns effect. If market advantage cannot be so easily gained by technology progression, then the need for effective tools to "get chips out" becomes more prominent. The claim of 10-times faster throughput is based on separate speed-ups of 10x in design planning, 5x in implementation, and doubled capacity. This enables, Synopsys says, a more speculative approach to IC layout, with more freedom for designers and architects to try out design variants.

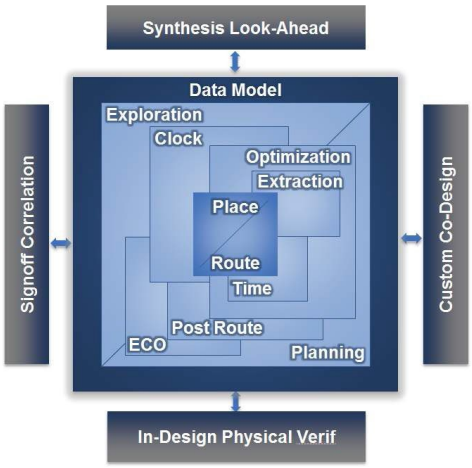

The new tool has been "several years" in the making, the company says, and according to a spokesman, "We ended up rebuilding almost everything," although some modules, for example the ZRoute component, have been re-used. IC Compiler II is able to handle designs with more than 500 million instances. Implementing a "rethink, rebuild and reuse" development strategy, IC Compiler II relies on industry standard input and output formats, as well as familiar interfaces and process technology files, while introducing innovative design storage capability. It was architected from the outset with a full chip-level focus, deploying design planning capabilities that provide a 10-fold performance boost while consuming 5-times less (host system) memory. This enables designers to quickly evaluate many floor-planning alternatives to arrive at the right starting point for implementation. The "rethink" part of the process involved looking at how the tool handles timing, design hierarchy and clock generation, among other factors. A key component is an optimisation element called APS, that carries out "global analytical optimisation". Placement and clock optimisation are melded into a single step. IC Compiler II also incorporates leading technologies used in IC Compiler, such as the conjugate-gradient placer as well as the ZRoute router.

Complete article, here

STM32 Nucleo boards, free from distributor's "Board Club"

Broadline technical distributor Future Electronics recently announced that it intended to expand the scope of the scheme it call its "Board Club". Under Board Club, Future itself funds the a stock of development boards relating to microcontrollers, and other products, that it stocks as silicon. Those boards are then available to be issued free of charge to engineers who can make a case that the project they have in development is likely to have a certain on-going value (to Future, in subsequent orders). Board Club was initially available to established customers of Future: the company recently announced that it would accept applications, on the same basis, from engineers from the wider design community.

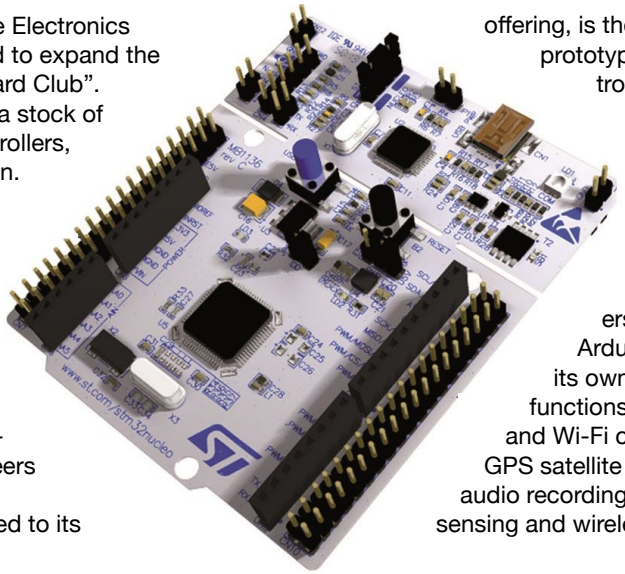

The first board that the company has added to its offering, is the new range of STM32 Nucleo prototyping boards from STMicroelectronics; which are are compatible with ARM's mbed application development platform. They also include ST Morpho extension headers to allow access to all of the microcontroller's on-chip peripherals, and Arduino headers which accept shields from the Arduino ecosystem. ST will offer its own dedicated shields supporting functions such as Bluetooth Low Energy and Wi-Fi connectivity, GPS satellite positioning, audio recording, proximity sensing and wireless control.

Complete article, here

Appliance PSU reference design with “zero standby” power rating

Power Integrations has a reference design for an 8W, universal-input auxiliary power supply that achieves zero standby power consumption, using the definition in the standard IEC 62301 Clause 4.5 that describes zero standby as ‘power dissipation below 5 mW’. Based on a member of Power Integrations’ LinkZero-LP family of ICs, DER-417 describes a universal-input, 5V, 1600 mA flyback power supply that consumes less than 4 mW at 230 VAC and provides 1 mW of power in standby mode. PI notes that with recent advances in low-power sensors and processors driven by the mobile-phone market, 1 mW of standby power is more than enough to operate functions such as a digital clock, environmental monitoring, wireless signal reception and infra-red activation; “DER-417 will inspire designers to implement zero-standby auxiliary power supplies for their products – leading to an overall reduction in wasted standby energy.”

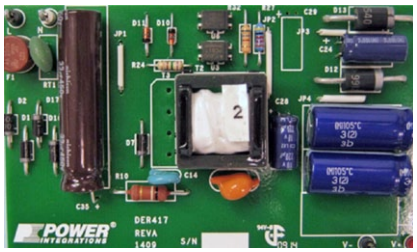

Complete article, here

Power e-bikes with Li-ion battery reference design

ams has produced a cost-saving system management reference design for lithium pedelec/e-bike batteries; implementing controller-less balancing and voltage monitoring, it shows how to reduce component count and BoM cost while providing essential safety functions. The example design for lithium pedelec/e-bike batteries implements accurate cell monitoring and balancing without the need for a microcontroller in the Battery Management System (BMS). The ams design is for a 48V pedelec battery consisting of up to 14 lithium-ion cells. It uses two AS8506 smart cell monitoring ICs, with some supporting components, to monitor the temperature and voltage of up to seven cells each and to implement passive balancing of the cells when charging. Conventional BMS designs in pedelecs use simple voltage monitoring ICs to measure the voltage and temperature of cells, reporting the values to a dedicated battery management microcontroller via a serial communications link. The MCU is required to control safety and protection functions (over- and under-voltage and over-temperature shut-down) and cell balancing.

The AS8506 by contrast, includes built-in logic functions for controlling cell safety, protection and balancing. These functions can be configured by the user, with the settings saved in an on-board OTP memory. The device also features integrated MOSFETs for use in passive cell balancing operations. During charging, each cell’s voltage is compared to a user-programmable reference voltage threshold. Up to 100 mA may be discharged through the MOSFET from any cell exceeding the threshold, until all cells have reached the threshold and the battery module is fully charged.

Complete article, here

Digital power controller matches analogue bandwidth, transient response

ADP1055 is an advanced digital power controller with PMBus interface for high-density, isolated DC-to-DC power supply system applications; it uses Analog Devices’ high-resolution, high-speed A/D converter sensing technology combined with proprietary non-linear transfer functionality to achieve high-bandwidth performance and transient response equivalent to traditional analogue switching controllers. The controller has six PWM (pulse-width modulation) logic outputs that can be programmed using a graphic user interface via the PMBus interface. The device enables high-energy-efficiency topologies including full bridge with precision drive timing and control of secondary synchronous rectifiers. The controller’s GPIO (general-purpose I/O) can be configured for active clamp secondary energy efficient snubbing. Energy efficiency is further improved using adaptive dead-time compensation to improve the efficiency over the load range. Programmable light-load-mode operation combined with low device power consumption (<150 mW) reduces system low-power standby losses.

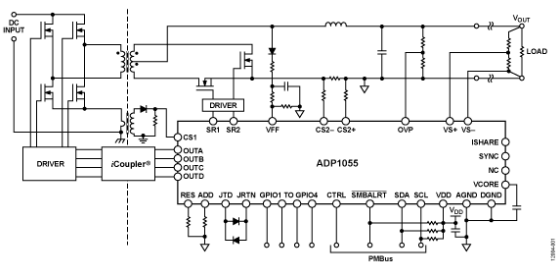

Complete article, here

Intersil adds 25/33A PoL integrated regulator modules

The ISL8270M/71M family of digital modules meets the power needs of high-current 25/33A applications in cloud computing and wired and wireless infrastructure; the company has based these digital power modules around the controller technology it has developed since acquiring Zilker Labs.

The underlying power efficiency advantages enabled by the ISL8270M/71M digital power modules are based on Intersil's ChargeMode Control modulation technology. This built-in, compensation-free architecture enables very fast transient

response and reduced output capacitance resulting in dynamic performance that minimises the output voltage variation. The prior generation of controller technology (Intersil termed it "third generation") used auto-compensation in which the controller injected known perturbations into the control loop, measured the response, and adjusted loop parameters to optimise stable operation. These devices do not need to do that, yet maintain stable operation through changing circuit conditions and component ageing. The same technology is employed in the company's ZL8800 controller IC – which will, Intersil says, be the last product to be distinguished by the "ZL" prefix.

These 25/33 Amp digital power modules are fully PMBus compatible, and report from an extensive menu of parameters as well as allowing external configuration. They are supported by the company's latest release of its PowerNavigator graphical user interface, which allows you to use these digital power products without writing code.

Complete article, here

BY CHRISTIAN LEGARE

DESIGNING IOT—PART II—THE THING

The IoT Devices

In the first article in this series, I presented the communication technologies used in Wireless Sensor Network (WSN) nodes. WSN nodes are not the only types of IoT device, but to provide the hardware and software architectures for each type of IoT node, I needed to establish the kind of software load that is typical on such equipment.

In this segment, I will describe what an IoT device actually is – the "Thing" in the Internet of Things. How is it structured, what processor should be used for each case, and which software languages should be used for the framework and the application?

What is a Thing?

The definition of a "Thing" in the Internet of Things varies a lot. I would define a Thing as an embedded computing device (or embedded system) that transmits and receives information over a network for the purpose of controlling another device or interacting with a user. A Thing is also a microcontroller- or microprocessor-based device. The capabilities of these embedded devices have been expanding at the speed of Moore's law.

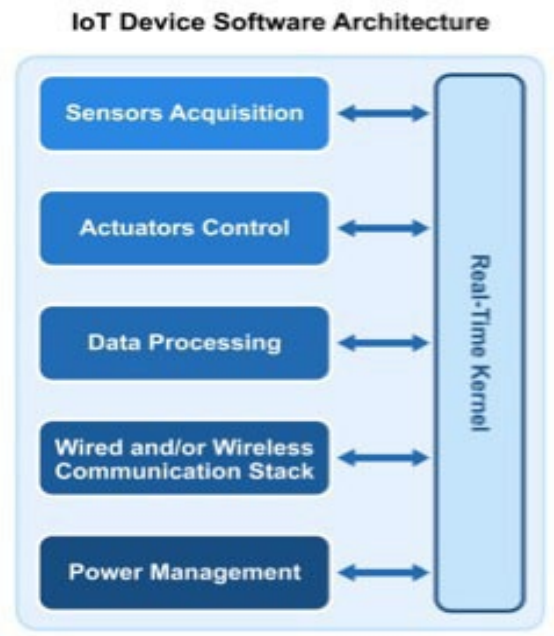

Figure 1. *IoT device software architecture*

What is an Embedded System?

When I am asked what an embedded system is, I like to use the following definition: "An embedded system is a microcontroller/microprocessor-based system encapsulating one single-purpose process for the lifetime of this system."

As the CTO of Micrium, I consider an embedded system to be a system based on a microcontroller (MCU). Some Linux and Android-based systems can also be defined as embedded systems, but usually, these general-purpose operating systems require an application processor, and have additional capabilities such as dynamic application loading (which does not match my definition of an embedded system). This is why MCU-based embedded systems are often described as deeply embedded systems, versus the more general definition of embedded systems. For us at Micrium, these deeply embedded systems are the Things in the Internet of Things.

MCUs featuring 32-bit architectures have dropped in price dramatically in the last several years, and are becoming increasingly common in embedded systems. The greater capabilities of 32-bit MCUs present new choices for embedded systems developers. For 8 and 16-bit MCUs, software has often been written using a foreground/background approach (that is, a super-loop). With 32-bit CPUs becoming more and more cost effective, using a real-time operating system (RTOS) for embedded devices has become the preferred option, allowing for more flexible and extensible software to run on these systems. A complete RTOS – with kernel, GUI, file system, USB Host and Device, TCP/IP (Ethernet, Wi-Fi), and more – can fit in a memory space of less than 1 MB. With an RTOS, the software architecture of an embedded system can be more flexible, and the device can perform its processing and communication tasks more easily. Troubleshooting and adding new features becomes drastically simplified. It is also simpler to perform firmware upgrades. In summary, it just makes sense to use an RTOS with a 32-bit processor.

Figure 1 illustrates how all activities within an IoT device can be mediated by a real-time operating system. Of course, the same can be done with Linux and Android, but those platforms require larger CPUs with many megabytes of RAM, and much greater power requirements. This is why the majority of the Things in IoT are – and will be – based on 32-bit MCUs.

Networking choices

An IoT device can be very simple and low cost (for example, the nodes in a WSN), or can be quite sophisticated (as, for example, the NEST Thermostat, which runs on a Cortex-A8 and uses Wi-Fi). In a WSN node, the networking technology used is usually a short-range access link. But there are a variety of other networking technologies used in devices where a WSN is inappropriate. These include:

Bluetooth
Zigbee, Zigbee IP, and other 802.15.4-based protocols such as 6LoWPAN
ANT, Z-Wave
Wi-Fi
DASH7
ISA100: Wireless Systems for Industrial Automation: Process Control and Related Applications
HART, WirelessHART
EnOcean
Wireless M-Bus
Ethernet, EtherCAT, EtherIP
Modbus, Profinet
HomePlug, GreenPhy, G.hnn (HomeGrid)

And more…

In the complete article Christian Legare goes on to cover networking technologies, processors for the Things, and programming languages.

Complete article, here

BY BONNIE BAKER

Will the right voltage reference stand up?

This is the scenario. You have a good idea about your application needs, and you have finally zeroed in on the correct analogue-to-digital converter (ADC). The next step is to feed the ADC with the proper voltage reference chip and driving amplifier. In this article, we will talk about how to choose your voltage reference (Figure 1).

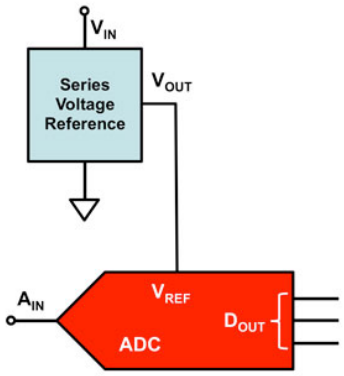

Figure 1 *Block diagram of an ADC/Voltage reference configuration. The voltage reference output (V_{OUT}) determines the full-scale range of the ADC.*

For some converters the full-scale range is equal to the voltage reference's output voltage, or V_{OUT}. For other converters the full-scale range is equal to twice V_{OUT}. You can find this relationship between the ADC's full-scale range (FSR) of the voltage reference chip's (V_{OUT}) output value in the ADC's product data sheet.

In a perfect world, that is all you have to do as you find the IC chip with the correct output voltage for your ADC. However, in our non-perfect world the voltage reference chip has initial output errors, and an inherent inability to drive the ADC's reference pin directly. If you want to learn about voltage reference driving circuitry, refer to references 1 and 2. But first let's come to terms with the reference's initial output errors.

When it comes to the voltage reference output errors, what you should be concerned about is initial accuracy, temperature drift, and noise. Part of this evaluation is to convert all of these errors to units of volts.

Usually the specification unit for the voltage reference's initial accuracy is a percentage. This unit can be converted to millivolts with one of these two equations:

$$\text{Initial Accuracy} = V_{OUT} * \%\ \text{error} * 10 \quad (\text{for } V_{REF} = \text{FSR of ADC}) \quad (1)$$
$$\text{Initial Accuracy} = 2 * V_{OUT} * \%\ \text{error} * 10 \quad (\text{for } 2 * V_{REF} = \text{FSR of ADC}) \quad (2)$$

Usually the specification units for the reference's temperature drift is ppm/°C, or parts per million per degree Celsius. You convert this unit to millivolts with the following equations:

$$\text{Reference drift error} = \text{ppm} / °C * \text{max app temp} / 1000 \quad (\text{for } V_{REF} = \text{FSR of ADC}) \quad (3)$$
$$\text{Reference drift error} = 2 * \text{ppm} / °C * \text{max app temp} / 1000 \quad (\text{for } 2 * V_{REF} = \text{FSR of ADC}) \quad (4)$$

Where ppm/°C is the reference maximum temperature drift value, max app temp is the maximum temperature to which you will expose your application circuit board.

The noise specification unit for the reference is µVPP, or microvolts peak-to-peak. This is an easy conversion. Simply divide the specification by 1000 to change it from microvolts to millivolts. Also you need to use a multiplier of 2 if the ADC's FSR is equal to twice the voltage at V_{REF}.

$$\text{Reference noise} = \mu V_{PP} / 1000 \quad (\text{for } V_{REF} = \text{FSR of ADC}) \quad (5)$$
$$\text{Reference noise} = 2 * \mu V_{PP} / 1000 \quad (\text{for } 2 * V_{REF} = \text{FSR of ADC}) \quad (6)$$

Now take these three values and add them together using the root-sum-square equation:

$$\text{Total Accuracy} = \sqrt{((\text{Initial accuracy})^2 + (\text{Reference drift error})^2 + (\text{Reference noise})^2)} \quad (7)$$

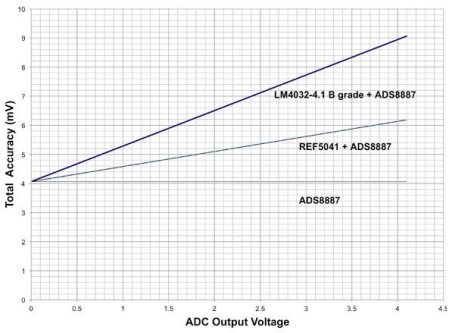

Figure 2 shows these results using the ADS8887: 18-bit ADC, REF5041 – V_{OUT} = 4.096V, and LM4032-4.1 B grade – VOUT = 4.096V. Table 1 has all of the specifications for these devices.

Figure 2 *The maximum application temperature is 85°C.*

By the way, you convert the ADC's SNR to converter noise with this formula:

$$\text{ADC noise } (V_{PP}) = 6.6 * \text{FSR} * (10^{-\text{SNR}/20}) \quad (8)$$

How do you minimise these errors? The dominant error in this evaluation is the offset error for the ADC and the initial accuracy error of the reference chips. You can digitally calibrate these errors at the ADC output. One can reduce reference noise with analogue filtering prior to the ADC's V_{REF} pin. The drift error is by far the hardest, if not impossible, to calibrate.

	ADS8887	REF5040	LM4032-4.1 B grade
Initial accuracy	4 mV	0.05%	0.10%
Noise	712.8 µVpp	12 µVpp	350 µVpp
Max drift	12.75 ppm/C	3 ppm/C	20 ppm/C

Table 1 *Maximum chip specifications for offset, noise, and drift. In this table maximum offset is analogous to the offset error for the ADC and initial accuracy for the voltage reference chips.*

	ADS8887	REF5040	LM4032-4.1 B grade
Initial accuracy	4 mV	2.048 mV	4.096 mV
Noise	0.715 mVpp	0.012 mVpp	0.35 mVpp
Max drift	0.0128 mV	0.255 mV	1.7 mV

Table 2 *Calculated offset, noise, and max drift values that appear in Figure 2. The max drift value is obtained with an 85°C temperature.*

Complete article, here

BY ERIC BOGATIN

THE BLURRING LINE BETWEEN OSCILLOSCOPES AND ON-DIE INSTRUMENTATION

"At 28 Gbits/sec everything matters," says Dave Dunham, Director of signal integrity at Molex. Even above 5 Gbits/sec, interconnects are no longer transparent. The beautiful, clean, pristine signals from the pads of the die get distorted, attenuated and turned to mush by IC packages, circuit board traces, connectors and cables.

Limitations from the fundamental physics of interconnect losses can't be eliminated. The only practical way around this limitation in the interconnects is to implement significant signal processing on-die at the transmitter in the form of FFE (feed forward equalisation) and on-die at the receiver in the form of CTLE (continuous time linear equalisation) and DFE (decision feedback equalisation).

Interconnects can distort a signal so much that at the receiver's pads, an eye diagram can be completely closed. On-die signal processing can open the eye. Figure 1 shows an example of the measured real time signal and resulting eye as measured at the pads of the receiver package, before the on-die equalisation has a chance to clean it up.

Upon looking at the measured real-time signals, you can't determine if on-die equalisation circuitry can clean the eye, recover the embedded clock and interpret the signal at an acceptable bit error ratio. Both the test and measurement industry and the semiconductor transceiver providers have responded to this challenge with innovative solutions, borrowing a little from each other.

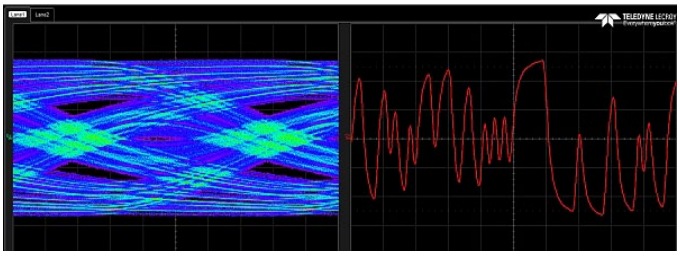

Figure 1- *A eye diagram at the pads of a receiver package before any equalisation shows how interconnects can distort the signal and close the eye.*

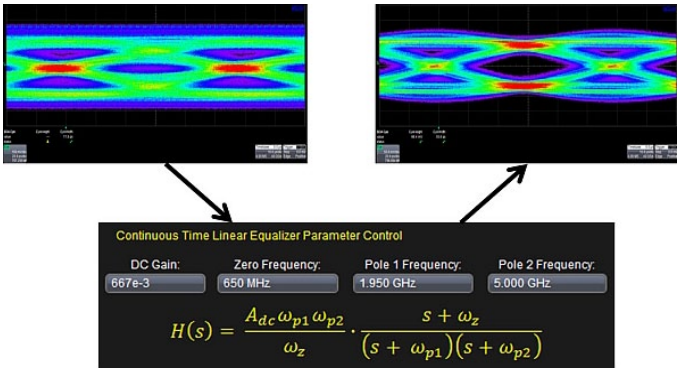

Figure 2- *Measured eye at the input to the CTLE filter implemented in the scope and the resulting eye.*

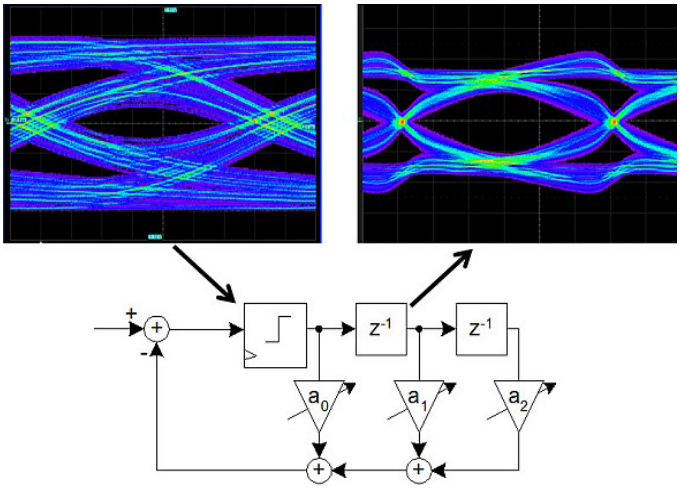

Figure 3- *Final signal emulated by the scope based on the CTLE, CDR and DFE features that would be implemented on-die at the RX.*

As a result, oscilloscopes from the top test and measurement instrument providers—Agilent Technologies, Rhode Schwartz, Tektronix, and Teledyne LeCroy—have implemented the same equalisation techniques and CDR (clock data recovery) algorithms as found in typical receivers, to emulate what the receiver might see.

"The actual signal on the pads of the receiver chip can be so distorted above 5 Gbits/sec that we have to clean it up before we can even say if it is good or bad," said Alan Blankman, product manager for serial data analyser products at Teledyne LeCroy. "We're careful to use the same algorithms in processing the signals measured by the scope as used in the receiver circuitry itself. This gives the engineer a realistic view of what the receiver would actually see."

An on-die CTLE filter is defined by a few simple pole-zero parameters. These can be selected based on the actual values used by the receiver, or optimised with software resident in the scope to maximise the eye opening. Figure 2 shows an example of the impact from the CTLE filter in the scope on the measured eye before and after the filter is applied.

After the CDR algorithm, the same DFE filter on-die is implemented in the scope to replicated the final eye opening that would appear at the receiver. This filter will take a fraction of a few previous bits and add them to a later bit. The fraction of each prior bit is the "tap coefficient." The coefficient values used in the scope can be selected as the same ones used on-die, or optimised based on the maximum eye opening. Figure 3 shows an example of the impact on the measured eye after the CDR and DFE filter are applied in the scope.

In the continuation of this article Eric Bogatin looks at the approach by semiconductor vendors – including FPGA vendors Xilinx and Altera – who have implemented "embedded instruments" on their latest 28 Gbit/sec designs.

Complete article, here

BY NICK NI, ALTERA

BEST PRACTICES FOR DESIGNING HIGH-THROUGHPUT, REAL-TIME SOC SYSTEMS

Modern SoC software often consists of multiple applications ranging from hard real-time, such as automotive motor control, to high-throughput, such as HD video streaming. Hybrid system design is becoming challenging as the modern SoC is evolving rapidly into a high-throughput system with an increasing number of processor cores and high-bandwidth interconnects.

Achieving hard real-time – less than 1 µsec jitter with µsec-level response – on such a system requires careful tradeoff analysis and system partitioning. It is also essential to consider future-proofing strategies for ever-increasing SoC complexity. There are mainly three approaches to such a system design – Asymmetric Multi-Processing (AMP), hypervisors, and Symmetric Multi-Processing (SMP) with core isolation (Figure 1) – from which system designers can choose to optimize hybrid SoC systems.

Asymmetric multi-processing

AMP is fundamentally a port of multiple Operating Systems (OSs) on physically different processor cores. An example would be to run a bare metal OS dedicated to handle real-time tasks on the first core and to execute a full-featured OS such as embedded Linux on the other cores. Most of the time, the initial porting of the OSs onto the cores is straightforward. However, the start-up code and resource managements, such as memory, cache and peripherals, are very error-prone. When multiple OSs access the same peripheral, their behaviours become non-deterministic and the system could become extremely time-consuming to debug. Hence, it often requires careful protection using an architecture such as ARM TrustZone to be in place.

	Scalability	Initial Porting	Message Passing	Throughput	Real-time jitter
AMP	Low	High	Slow	Reduced by core assignment	Low if no message passing
Hypervisor	High	Low	Slow	Reduced by the hypervisor decision	High
SMP w/ core isolation	High	Low /High*	Fast	Reduced by core assignment	Low

* Low if starting from a SMP OS but high if starting from multiple OS'es

Figure 1- *Comparison of AMP, hypervisor and SMP with core isolation.*

To add more complexity, message passing between the OSs requires memory sharing and needs to be managed together with the other protection measures. Because the cache is usually not shareable between different OSs, message passing needs to happen through non-cache memory regions, which adds latency and jitter to the overall performance. It is also poor software architecture from the scalability viewpoint as it requires significant re-porting when the number of cores increases.

Hypervisors

A hypervisor is a low-level software layer that runs directly on the hardware and manages multiple independent OSs on top of it. Though the initial porting is similar to AMP, the benefit is that the hypervisor hides the non-trivial details of the resource management and message passing. One drawback is that it incurs a performance overhead due to the extra software layer degrading the throughput and real-time performance.

Symmetric multi-processing

SMP with core isolation runs a single OS on multiple cores with internal core partitioning. An example is to instruct an SMP OS to assign a real-time application on the first core and the rest of the non-real-time applications on the remaining cores. This approach is very scalable as the SMP OS is designed to port seamlessly to an increasing number of cores. Because all cores are managed by a single OS, message passing between cores can happen at the L1 data cache level, resulting in faster communication with less jitter.

Core isolation can reserve a core for the hard real-time application to shield effects from other high-throughput cores, preserving the low-jitter, real-time data response. This is generally a good software architecture decision because it allows the designers to consider which OSs to use instead of re-inventing error-prone, low-level software to manage multiple OSs. The initial porting may require some effort if starting from multiple OSs. However, starting from an SMP architecture would be much lower effort.

Optimising a high-throughput, real-time SoC with SMP

Based on analysing the alternatives, SMP with core isolation offers the best architecture to optimise high-throughput, real-time SoC systems. The architecture we consider is a system similar to Figure 3 where an I/O data stream comes into a SoC, undergoes some form of processing in the cores, is returned to the I/O with a low-jitter and low-latency real-time response. In addition, the SoC consists of multiple cores that run other throughput-intensive applications simultaneously.

First, it is essential to understand what a real-time response (loop time) consists of:

- Transfer new data to the system memory from an I/O (DMA)
- Processor detects the new data in the system memory (Core Isolation)
- Copy the data to a private memory (memcpy)
- Compute on the data
- Copy the result back to the system memory (memcpy)
- Transfer the result back to an I/O (DMA)

Because the jitter and the latency are accumulation of the six steps, it is essential to optimise each step. With an RTOS such as VxWorks with core isolation, the polling/interrupt response can be bounded in the nanosecond range (Step 2). Data computation is also application specific and is fairly predictable (Step 4). Therefore, we focus on the trade-off of the Direct Memory Access (DMA) and the memcpy (Steps 1, 3, 5, and 6). There are two major means to transfer data: transfer with or without cache coherency. The two methods have very different consequences in the DMA and memcpy.

The author goes on to emphasise the benefits of cache coherent transfer, and to present a case study in best-practice SoC design.

Complete article, here

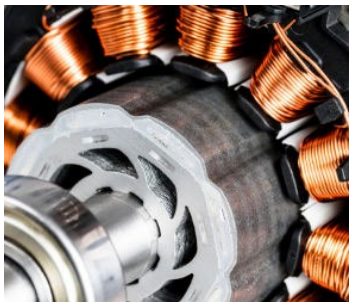

Spin Cycle

Induction Motors: Don't write them off just yet!

By Chris Clearman, Texas Instruments

I've been doing some light reading recently on the European Council for an Energy Efficient Economy website, specifically their recommendations for electric motors and controls from their recent "Ecodesign Seminars." There is some excellent information included about the efficiency and lifetime costs of various motors, including the workhorse of industry – the induction motor. There was also information on the more natural-resource-consuming, permanent-magnet synchronous machine, as well as the new kid on the block – the synchronous-reluctance machine. Some very intelligent decisions were drawn regarding where best to focus efforts in target applications with specific technology. But (you knew there had to be a downside), I feel they neglected a very important consideration.

Check out this image (Figure 1) from the overview presentation that communicates one of their key points. I took the liberty of summarising their conclusions:

- **Top graph:** Induction motors are wildly inefficient when they don't need to produce torque (the load needed to drive is not near the maximum), even when they are superiorly designed for excellent efficiency at load, as the motors described have been.
- **Middle graph:** Variable speed drives (VSDs), used typically in a digitally controlled three-phase inverter, are reasonably efficient components on their own.
- **Comparing the bottom graph to the top:** VSDs do very little to help efficiency of induction machines, especially in trying to overcome their inefficiencies in lower/partial load operation.

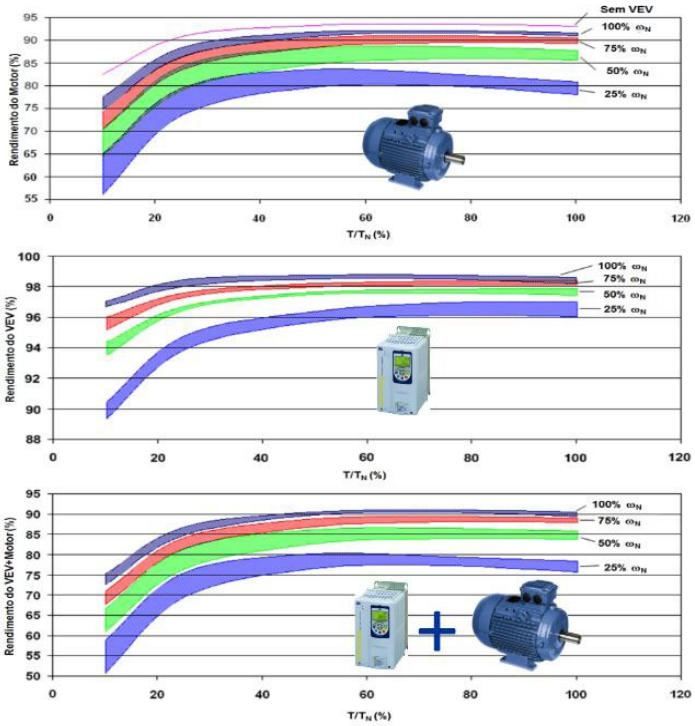

Figure 1 *is drawn from Page 22 of "Consideration of system approaches for motors and motor-driven units: rationale and implementation challenges" presented by Anibal de Almeida, Coimbra University. (online PDF)*

Induction motors are not much different than they were when invented about 130 years ago. Their efficiency is poor, primarily due to the namesake: Induction. Because there is not a permanent magnet to use for the rotor, you must induce a magnetic field on the rotor through the magnetic field you are producing in the stator. In a synchronous machine, you simply create a rotating magnetic field on the stator, and the existing magnetic field of the magnetic rotor will be attracted and follow. The simple way to think about it is that for an induction motor, we are always using a near constant amount of current (called the magnetising current), even if it's not required, to induce this rotor field to a specific level to produce a specific shaft torque. But this is also done if, at times, no torque is required! Power dissipated with no work accomplished is the definition of poor efficiency.

Permanent magnet machines

By the way, permanent magnet machines are by no means perfect either. In their case you are pre-designing a motor – and spending substantial money on rare earth magnetics – that will be able to produce the maximum torque you will require. If you often need this torque, it is money well spent. If not, for the most part that is a sunk cost in your rotor…but at least the overall efficiency – even when lower torque is required - is usually very good.

So that's it, these ancient induction motors are just what they are, right?

To quote one of my favourite movies: "I got information man. New stuff has come to light!"

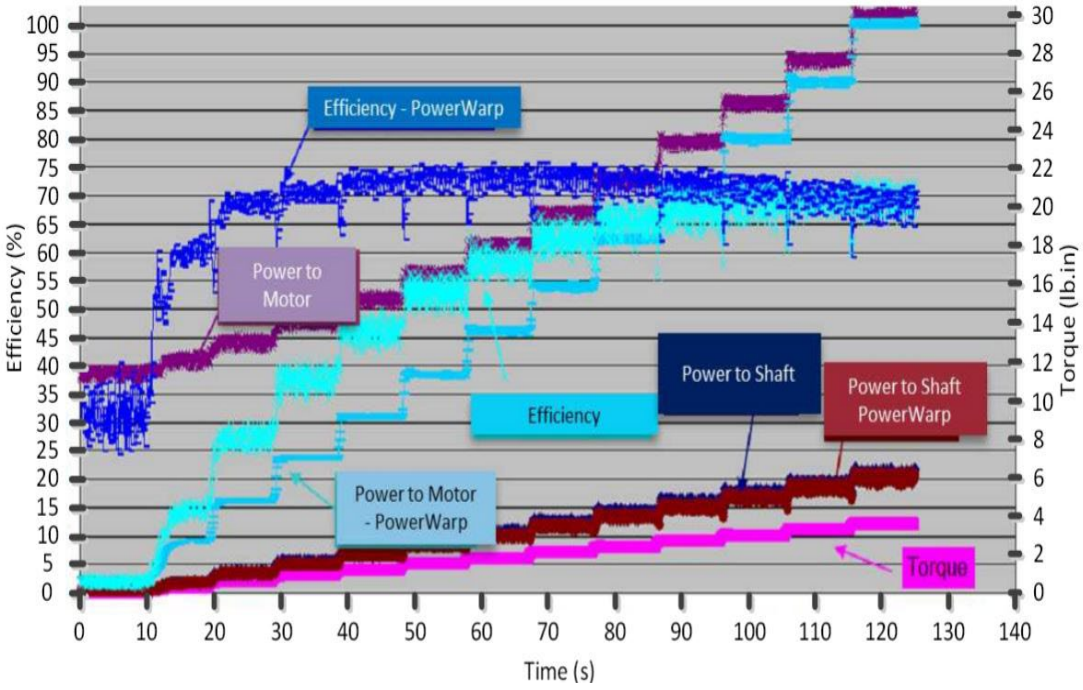

Figure 2 *PowerWarp improves motor efficiency.*

Why can't we vary the magnetising current to create a machine with a different rotor flux and hence a different rotor torque? We can! In fact, there are some applications that do this today to make their induction motor create more torque than rated for initial starting of engines, generators, compressors and large pumps. It is called field boosting, and many times 150-200% of torque can be created for short periods of time. But it certainly doesn't address the efficiency concern at loads.

So why not do the opposite? Reduce the magnetising current to always induce just the right-sized magnetic field for the torque you want to create? Well, at Texas Instruments (TI), we have created this technology. We call it PowerWarp and offer it with our FAST software observer-based TI InstaSPIN motor control technology solutions, which work with certain members of our C2000 Piccolo microcontrollers (MCUs). PowerWarp adaptively reduces current consumption to minimise the **combined** (rotor and stator) copper losses to the lowest possible, without compromising the motor output power levels.

That means we can create a small torque machine when that's all that's needed, but if the load changes, we can dynamically increase to create a large torque machine. And, due to the stability and tracking performance of our FAST software observer – which detects the dynamic changes in torque demand and accurately estimates rotor angle (even during over-torque stall recovery) - PowerWarp can be used in low-cost, "rotor sensorless" implementations for even the lowest cost, variable speed or variable load applications. Even those where simple, single-phase induction motors have typically been used.

Figure 2 is a real-world test of a very standard one-fourth horsepower (HP) induction motor used widely in industry, which struggles to achieve higher than 70-75% efficiency at rated speed, rated load. When the load is lower, notice how the efficiency quickly degrades. However, with PowerWarp enabled the efficiency curve is nearly flat with just a nominal load.

The contination of this column quotes some practical application examples, and goes on to derive cumulative energy – and

Complete article, here

BY MICHAEL DUNN

A PRACTICAL TAKE ON FFTS

Virtually all scopes today provide the FFT (Fast Fourier Transform) as a standard feature, but in my experience, few users know how, when, or why to use it. And if they do, they rarely use it well.

Why use an FFT? The list of answers is almost as long as the one for "Why use a scope?" Measuring distortion is a classic example, whether it's at 50/60 Hz, audio, or RF frequency. Perhaps you want to see the spectral characteristics of some noise, or maybe the occupied spectrum of a serial data signal.

As an aside, the first real FFT equipment I got my own hands on was a lovely HP dedicated audio signal analyser in the mid-1980s.

Unfortunately, this spoiled me in terms of UI and functionality, and I've yet to meet a scope that comes close. Factoid: This instrument used bubble memory to store setups and data (think hard disc on a chip).

Reviewing the utter basics: An FFT takes a signal in the time domain and converts it (transforms it, you might say) to the frequency domain. A 60 Hz sinewave (time domain) transforms to a spike at 60 Hz (frequency domain). If there are spikes at 120 Hz, 180 Hz, and so on, then you know there's distortion.

And reviewing our basic waveshapes: A square wave has only odd harmonics, falling off at 1/N. A triangle has only odd harmonics, falling off at $1/N^2$. A ramp has all harmonics, at 1/N. A pulse-train can have missing harmonics. For example, a 10% duty-cycle pulse will lack the 10th, 20th, 30th harmonics, and so on.

OK, we want to look at the FFT of a signal on our scope. What parameters will sculpt the view? There's:

- Record size
- Ratio of signal period to record time
- Windowing function

Record size is limited by the scope's buffer size – and your patience. Well, patience is becoming less of an issue, because a modern scope can churn out million-point FFTs in the blink of an eye. The main effect of record size is on frequency resolution. For example, an 8k record FFT will generate 4k frequency bins (it's always half). FFTs always start at 0 Hz, so if the frequency span is 20 kHz, our resolution is 20k/4k = 5 Hz. That's pretty good on a linear scale, but if you prefer to view your frequency axis logarithmically, and you don't want the low end to be too coarse, more points may be called for.

Figure 1- *The HP 3561A Dynamic Signal Analyser.*

The ratio of signal period to record time determines where on the X-axis frequency scale the signal components appear. If you time things such that exactly one cycle fills the FFT's input record, its fundamental spectral line will be in the second bin (the first being DC). Any harmonics in the signal will fill up successive bins. It might seem neat and tidy, but it is not usually possible.

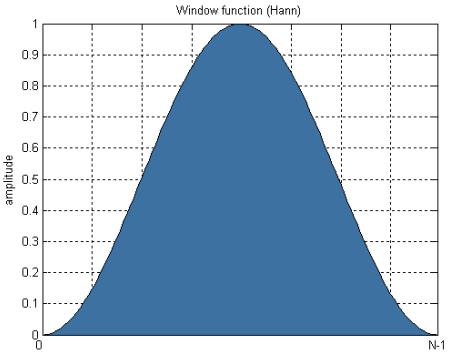

Figure 2- *A von Hann, or Hanning, window.*

First off, in the real world, it's virtually impossible to get exactly one cycle in the buffer. Let's say you get 1.01 cycles. This will cause every spectral line to be smeared into its neighbours. It's usable but messy. You're much better off capturing perhaps 10 cycles or 100 (depending on the FFT size and the signal's harmonic content). All the spectral lines will be moved up and spread out. That's much better.

I mentioned frequency span, but not in the parameter list. That's because it's a secondary value. Span simply equals half the sampling rate. Therefore, we can redefine frequency resolution as the sampling rate divided by FFT size.

In normal scope operation, there is usually some interaction between sampling rate, sweep speed, and record size. If the FFT UI isn't well thought out, these interactions can drive one to drink (or worse), as you adjust controls and find your careful FFT settings ruined by a well-meaning turn of a knob. If the FFT is something you'll be using a lot on a new scope, make sure to try before you buy.

The last item under our control is the window function. "What nonsense is this?" you ask. It's an unfortunate necessity in dealing with real-world signals, and it's actually pretty ingenious.

The Fourier transform, regardless of flavour, assumes a continuous signal over infinite time. Discrete versions of the transform still adhere to this rule – it's assumed the finite waveform record repeats infinitely. So, imagine multiple records spliced together: The end of the record must splice smoothly to the beginning of the next – identical – record (recall the previous talk about fitting exactly one cycle in your FFT input). If we could do this – feed a perfectly integral number of cycles to the FFT – all would be simple (and this also assumes we're only interested in analysing harmonically related signals).

But we both know it ain't gonna work out that way. The inevitable glitch at the join point between the end and start of the record will cause a whole slew of nasty faux harmonics to appear. Try it. That's where the window function comes in.

The window above is a commonly used one, and its basic shape is typical. This curve is aligned with the captured record, and is used to scale each point by the curve value at that X-position. The resulting record looks like a tone burst, and as the record is assumed to repeat infinitely, the FFT performs its magic on, in effect, this repeating tone burst. "And that's better!?" you exclaim.

In the continuation of this article, and in answer to this last question, Michael says, "Yes it is!" - and explains why.

Complete article, here

BY STEVE SANDLER, PICOTEST

MEASURE SMALL IMPEDANCES WITH ROGOWSKI CURRENT PROBES

You're probably familiar with clamp on current probes, but chances are you don't know much about the Rogowski current probe. A Rogowski coil produces a voltage that is proportional to the rate of change (derivative) of current enclosed by the coil-loop. The coil voltage must be integrated for the probe to provide an output voltage that is proportional to the current signal. This also means that the Rogowski can't perform at DC, but can operate to a frequency lower than most AC clamp-on current probes.

	Rogowski	AC Current probe	Hall effect probe
Limited to one equipment manufacturer	yes	yes	No
Works with VNA and Spectrum Analyzer	yes	yes	no
Requires degauss and offset	no	no	yes
Suffers from magnetic saturation	no	yes	yes
Noise	high	low	low
Bandwidth	≈100Hz-30MHz	≈1kHz-120MHz	DC-100MHz
Insertion impedance	zero	moderate	highest
Head size	smallest	medium	largest
Relative cost	lowest	moderate	highest

Table 1- *A comparison of current probes.*

Table 1 highlights the key performance attributes of three different clamp on current probes. The comparison results are based on the PEM CWT015, Tektronix P6022 and Teledyne Lecroy CP031 current probes. The three probe heads are shown in Figure 1.

Both the AC current probe and the Rogowski current probe can be used with test equipment from any manufacturer. This includes oscilloscopes, VNAs (vector network analysers), and spectrum analysers. Hall-effect probes, on the other hand, are generally keyed to a particular equipment manufacturer. Comparing the AC and Rogowski probes, the Rogowski current probe excels in all performance characteristics except noise and the ability to measure DC. You can, though, manage the noise and make the Rogowski probe one of your most-used probes.

Figure 1- *CP031 (top), P6022 (middle), and Rogowski (bottom).*

Consider the noise

The low frequency noise from the Rogowski current probe is a result of the high gain integrating amplifier needed by the probe. The CWT015 noise is predominantly below 1.5kHz as shown in the noise density measurement in Figure 2. (Note; the noise is mostly below 1 kHz because of the opamp 1/f noise which is greatly amplified at low frequencies due to the integrator configuration: all Rogowski probes require a high gain integrator to create an output that varies linearly with current. Most linear regulators also exhibit 1/f noise at low frequencies (typically below 1kHz))

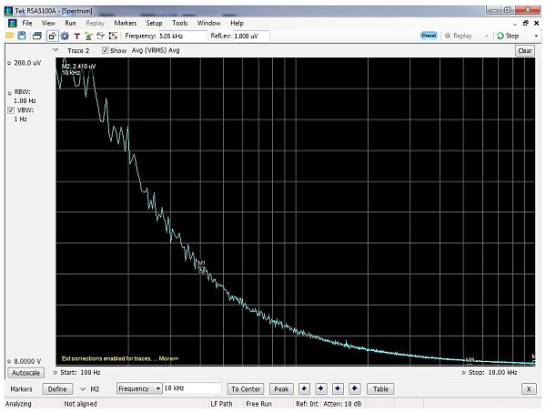

Figure 2- *The noise density is 200 µV/√Hz at 100 Hz, 20 µV/√Hz at 1 kHz, and 2.4 µV/√Hz at 10 kHz. When used with a spectrum analyser or VNA the resolution bandwidth can be set to 30 Hz or even lower. At 1 kHz with a 10 Hz resolution bandwidth the noise is approximately 100 µV and at 10 kHz and above only 10 µV.*

The frequency domain

When using the Rogowski current probe with a VNA or spectrum analyser, the resolution bandwidth can be set to 30 Hz or even lower. This narrow bandwidth minimises the measurement noise. The sensitivity of the CWT-015 probe is 100 mV/A. At 1 kHz, the noise density indicates 20 µV/√Hz noise. Converting through the sensitivity this equates to a noise current density of 200 µA/√Hz.

The peak-to-peak noise can be calculated from the noise density and the resolution bandwidth as:

$$Noise_{pp} = 2 \cdot Noise\ Density \cdot \sqrt{resolution\ bandwidth}$$

Using this relationship, a 30 Hz resolution bandwidth results in a peak-to-peak noise of 2 mAP-P at 1 kHz and 260 μA_{P-P} at 10kHz. If the current being measured is maintained above 20 mA_{P-P} at 1 kHz or 2 mA_{P-P} at 10 kHz, then the signal-to-noise ratio will result in measurements with acceptable fidelity.

In the continuation of this article, the author looks at measurement of impedance with a VNA and the Rogowski probe; and sets out the optimum instrument setups for a variety of measurement tasks.

Complete article, here

BY ALBERT WU, LINEAR TECHNOLOGY

FOUR-QUADRANT DC/DC REGULATOR SMOOTHLY TRANSITIONS FROM POSITIVE TO NEGATIVE OUTPUT

In many electronic systems, it is necessary to provide a bipolar (positive and negative) voltage or current into a particular type of load. Loads that require bipolar voltages/currents include FPGA body biasing applications, thermoelectric coolers, DC motors, and many others.

There are many traditional methods to provide a bipolar voltage/current to a load. H-bridge designs are frequently used, but require that neither of the load's terminals is tied directly to ground. Each of the load's two terminals has to swing between the positive supply rail and ground and usually an inductor is placed in series with the load to filter out this chopped waveform. This lack of a direct ground connection to the load can complicate the mechanical and electrical design of the overall system. The H-bridge method also requires four switching elements and a more complex control scheme. There are loads that have a negative terminal that cannot be biased high relative to ground, such as an FPGA back-biasing application.

Another traditional method is to build two power rails, a positive one and a negative one. Various circuits are used to "swap" in the positive or negative rail with regulation to achieve bipolar voltage operation that can go below ground. This results in a very complex system, generally with poor efficiency, and a non-linear response at the point at which the output voltage crosses the ground potential.

A new DC/DC switching architecture is described here that has the ability to generate true four-quadrant operation, meaning the output voltage can be positive or negative and the current flow can be in either direction as well. Additionally, this new architecture can generate an output voltage that transitions from one polarity to another, through the ground potential, smoothly and without any non-linearity from mode transitions.

Four-quadrant DC/DC converter

Figure 1 shows the basic connections and elements of the four-quadrant converter. NFET, MN and PFET, MP are operated out of phase from one another and at a constant switching frequency. Current mode control is used (not shown) to modulate the duty cycle of MN as needed.

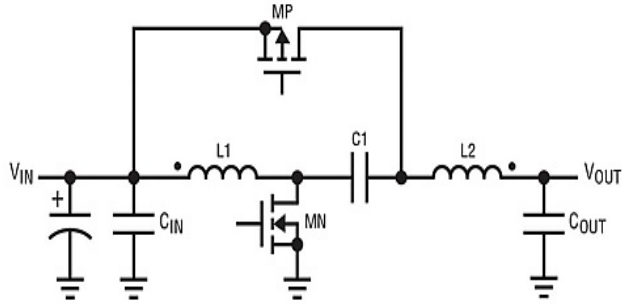

Figure 1. *Four-quadrant DC/DC converter topology*

If we assume fixed frequency operation, the duty cycle for the ON time of MN can be calculated as

$$D.C. = \frac{V_{IN} - V_{OUT}}{2V_{IN} - V_{OUT}}$$

From this equation, it is clear that with a positive V_{IN} voltage, the output voltage V_{OUT} can be positive (up to V_{IN}) or negative (limited only by practical DC considerations) and can go to 0V as well. In fact, there is nothing special about the 0V output level since the DC of the converter is 50% at that operating point.

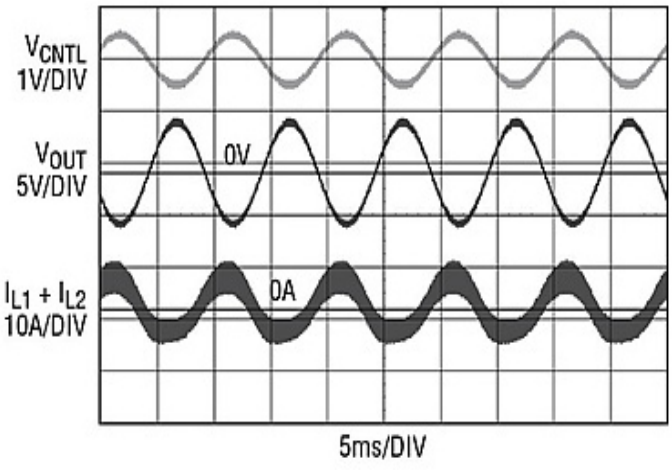

Figure 3. *Sine-wave output voltage passes through 0V*

The output of this converter can sink or source current regardless of the polarity of the output voltage, making this a true four-quadrant operating topology. The maximum drain to source voltage stress on MN and MP are both $2V_{IN} - V_{OUT}$. For example, if V_{IN} is +12V and V_{OUT} is -12V, then the BV_{DSS} ratings for both FETs must be greater then 36V.

The four-quadrant operating capability of the converter is shown in Figure 3. Here, a sinusoidal control signal is used to generate a sinusoidal output voltage centred on 0V. The inductor currents can go positive or negative; whatever is necessary to have the output voltage go to the commanded level. The operating waveforms show clean and smooth operation through the ground potential. The choice of using a sine wave control signal is arbitrary; a DC signal, square wave signal or any other type of signal could also be used.

The complete version of this article (click the link) provides a fully-tested circuit, and continues to describe typical applications.

Complete article, here

Position sensing of liquid metal antennas for beam scanning antenna arrays

by Ahmad Gheethan

This Idea was the 1st prize winner in the TI LDC1000 inductive sensor design contest, a challenge posed by Texas Instruments following the launch of its inductive sensor in 2013.

High gain beam scanning arrays are usually implemented using phased antenna arrays or mechanically controlled parabolic dishes. However, the aforementioned technologies are widely known as costly ones. Recently, a new design of beam scanning arrays has been introduced using the liquid metal technology. Specifically, the antenna is made of a liquid metal volume that is encapsulated in channels and moved mechanically using a single mechanical pump.

To implement this technology successfully, the liquid metal antenna should be exactly located in the a reservoir with minimum physical misalignment. Therefore, it is essential to sense the antenna location with respect to the desired reservoir. Moreover, the sensing system should be cost effective and simple, yet accurate and highly sensitive. Having that mentioned, the LDC1000 appears to be a potential candidate to develop a position sensing system for the proposed antenna array.

The LDC1000 can be used to carry out the sensing operation by designing and constructing a feedback system. By using the LDC1000, an inductance vs. misalignment design curve can be designed to determine and sense the antenna position and send feedback commands to the pump in order to adjust and move the antenna. If successful, this proposed solution will significantly contribute in implementing low cost beam scanning arrays. In addition, the solution can be extended to modify the LDC1000 design to allow sensing different positions at the same time using only one LDC1000 piece.

In order to test the possibility of using the LDC1000 in detecting the antenna in a microfluidic channel, a prototype microfluidic channel with an antenna encapsulated in it was first fabricated. Later, the LDC1000 inductance was measured using the GUI software when the inductive sensor was located on the top of an empty reservoir location and the inductance was recorded. Next, the sensor was moved to a reservoir where the antenna was located and the sensor inductance change was monitored. The sensor inductance dropped as the sensor overlapped with the antenna.

As a result, the inductive sensor can be utilised for detecting the antenna location, and thus, the beam direction by scanning the sensor among the channel and noticing the sensor inductance reduction. Scanning the sensor can be performed using a step motor and the whole system can be programmed using a compact micro-controller with low cost and very sensitive monitoring.

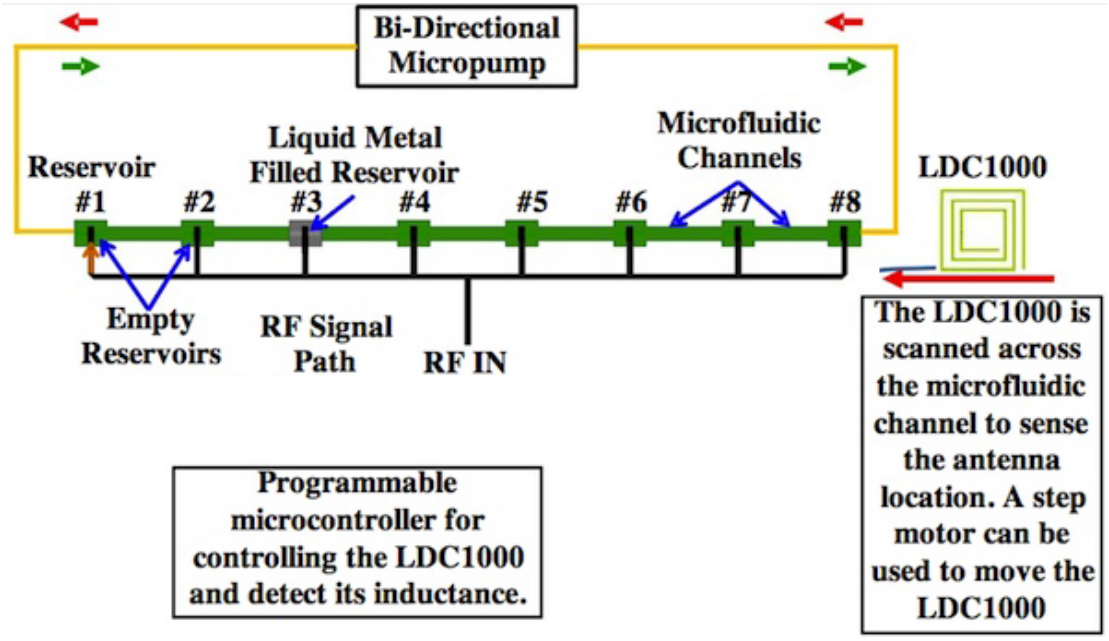

Filtered noise generator using a voltage regulator source

by Vladimir Rentyuk

Noise generators are very useful devices. They are used in test equipment, random number generators, etc. Usually, designers use special Zener noise diodes (at very low current) and transistors (with reverse-biased BE junction) as a source of noise. All of these sources of noise have low level noise voltage and need additional amplifiers.

But voltage regulators can also be used as source of noise signals. Output noise voltage of a low dropout regulator, for example, the LM2931-5.0, is 500 µV (typ., 10 Hz – 100 kHz, C_{OUT} = 100µF). It is a few times higher than the noise level of commonly used noise diodes. A circuit which shows an LM2931-5.0 as a noise generator has been published. It needs an additional amplifier, because the noise signal is about 0.5 mV.

This Design Idea shows another way of using a voltage regulator which I accidentally discovered. It is a proven circuit, but this startling effect does not show in computer simulations, for example, using Multisim. The effect is inherent to all voltage regulator ICs, and such reference ICs as the ADR425ARZ. I checked this effect with: LM2931-5.0, L78L15ABU, SPX1117M3-L, and others. The highest noise is generated by the LM2931-5.0; the lowest, by the ADR425ARZ.

The simple circuit (Figure 1) generates hundreds of millivolts of noise. The regulator supplies power to filter IC2 and works as a noise source at the same time.

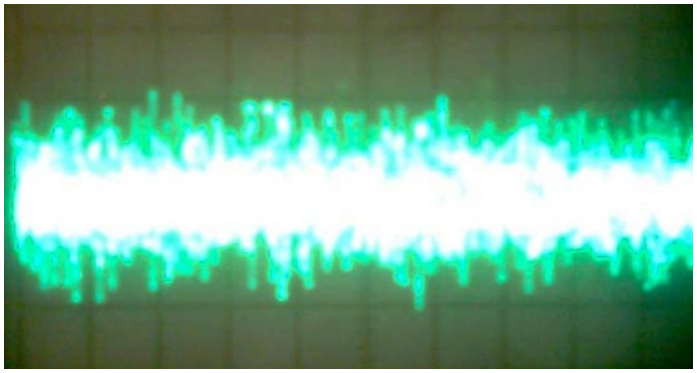

As opposed to the LM2931 circuit linked above, the source of noise here is the input of the voltage regulator IC1. R2 is used to convert the noise current to a voltage, resulting in about 800 mV P-P for R2=1.5 kΩ. A 4th-order active RC lowpass filter in IC2 selects the desired noise bandwidth (about 20 Hz-5 kHz), and presents a small load to IC1.

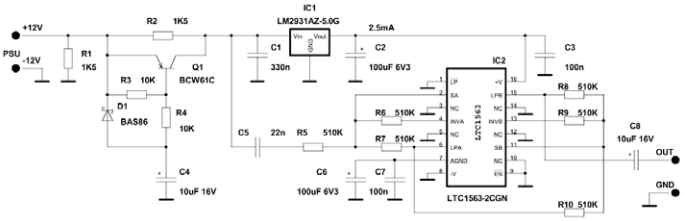

R2 is calculated as R2 < (V_{IN} – (V_s+V_{DV}) – V_{PN}) / (I_{out}+I_q); where V_{IN} = min. input voltage (+12V in this case); V_s = max. output voltage of IC1 (5.19V); V_{DV} = max. dropout voltage of IC1 (0.2V); I_{out} = max. output current of IC1 (power supply current in LP mode of IC2 is 2.5 mA max.); I_q = quiescent current of IC1 (1 mA max.); and V_{PN} = estimated max. amplitude of noise signal (0.6V). So, R2 < (12 – (5.19+0.2) – 0.6) / (2.5 + 1) = 1.7 kΩ. Thus a resistor 1.5 kΩ was used.

The Q1 circuit enables startup of IC1. The LM2931-5.0 needs an initial current of 26 mA for 40 msec during startup. The Q1 circuit may not be necessary when using other types of voltage regulators (L78Lxx, SPX1117, etc). This circuit can be used with another filter, or without a filter. In this case it will be necessary to provide at least 1.5 mA load on IC1. Figure 2 shows the output noise.

About the Author;
Vladimir Rentyuk is a development engineer at Modul-98, a developer of embedded electronic systems, robotic equipment, image recognition systems, and other products for European and world markets. He holds a master's degree in radio engineering from Zaporizhzhya Machine Construction Institute (now Zaporizhzhya National Technical University) in Ukraine.

Linearise thermistors with new formula

by Mark Biegert

I often have to get a reasonably accurate temperature measurement for a laser, transistor, or printed circuit board. I usually use a thermistor for my temperature sensor because they are inexpensive. However, the thermistor's nonlinear resistance characteristic makes accurate temperature conversion complex. To minimise this difficulty, I often linearise the thermistor's resistance characteristic by placing the thermistor within a resistor divider. This linearised response is simpler to convert to a temperature value than the thermistor's raw response. To obtain an optimum level of linearity, I have derived a pair of formulae that are useful in determining component values (R_S and R_P) for this common linearisation circuit.

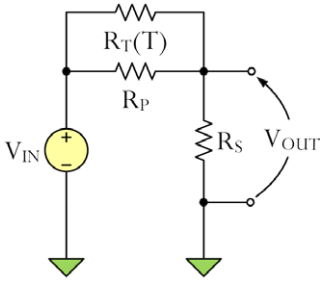

Figure 1. *Two-resistor thermistor linearisation circuitO*

Figure 1 shows the thermistor linearisation circuit that I am addressing in this Design Idea. The temperature of the thermistor is linearly related to the output voltage (approximately).

The design process begins by picking a temperature, which I call the inflection temperature TI, at which we want the flattest possible transfer function (V_{OUT}/V_{IN}). My design task is to compute values for R_S and R_P given T_I and the resistor ratio $\mu=V_{OUT}/V_{IN}$ at T_I.

Three parameters, R_0, β, and T_{REF}, are often used with Eq. 1 to model the thermistor's resistance versus temperature characteristic.

$$R_T(T) = R_0 \cdot e^{\frac{\beta}{T}-\frac{\beta}{T_{Ref}}}$$

I determine R_S and R_P by setting the second derivative of V_{OUT}/V_{IN} in Figure 1 to 0 at temperature T_I. After much algebra, I obtained the two formulae shown in Eq. 2.

$$R_S = R_T(T_I)\cdot\left(\frac{\beta-2\cdot T}{2\cdot\beta\cdot(1-\mu)}\right)$$

$$R_P = \frac{R_S\cdot R_T(T_I)\cdot(\beta-2\cdot T_I)}{R_S\cdot(\beta+2\cdot T_I)-R_T(T_I)\cdot(\beta-2\cdot T_I)}$$

Not all ratios are possible with passive components – ratios less than;

$$\mu_{Critical} = \frac{(\beta-2\cdot T_I)}{2\cdot\beta}$$

result in negative RP values.

Figure 2 shows a graph of the linearised thermistor voltage transfer function for a common thermistor.
I have put together an Excel spreadsheet and a Mathcad worksheet that models this data, downloadable here.

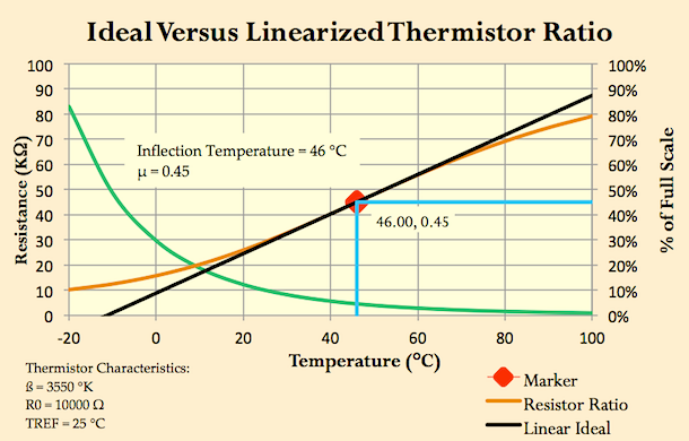

Figure 2. *Graph of the linearised thermistor transfer function*

productroundup

DC brushed motor driver

A single-channel DC brush motor driver IC for industrial equipment, the TB-67H303HG is the first monolithic IC to deliver 10A output current with 50V withstand voltage. The high output current and voltage are required for motor driver circuits have been difficult to design into a single driver IC, for thermal reasons; multiple discrete semiconductors are generally used in order to reduce heat generation in each component. Housed in an HZIP25 package that enables high power dissipation, TB67H303HG can sustain a 10A output current without overheating. TB67H303HG uses the latest high voltage analogue process, reducing ON resistance to 0.2Ω or less. The motor driver can be operated in two driving modes: standard direct PWM mode, or constant current PWM mode. Both driving modes support Forward, Reverse, Short brake, and Stop functions when used with a single brush motor.

Complete article, here

700W L-band transistor with highest output power

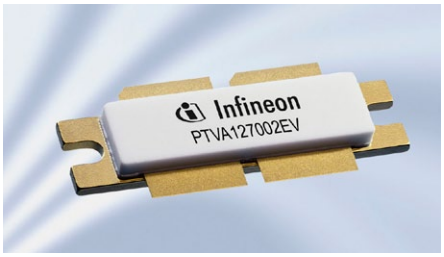

Infineon's 700W L-Band RF power transistor offers the highest-available L-Band output power for radar systems operating in the 1200 – 1400 MHz frequency range. By lowering part counts, the new device can reduce system cost and improve reliability while maintaining high ruggedness. The power transistor, PTVA127002EV, is suited for L-Band radar systems used in air traffic control and weather observation applications. High efficiency corresponds with low heat output, and high ruggedness (ability to withstand 10:1 VSWR load mismatch) further contributes to the advantages of low component count made possible by the 700W output. Based on the company's 50V LDMOS power transistor technology, the PTVA127002EV exhibits high efficiency; typically 55% across the 1200-1400 MHz band, with a P1dB output power of 700W, 16 dB gain and low thermal resistance characteristics when measured with a 300 µsec 10% duty cycle pulse.

Complete article, here

ARM Cortex A9 MCU development board runs Linux

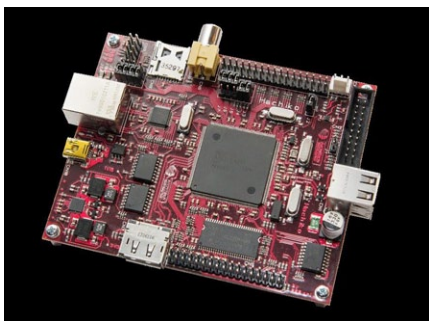

The ArchiTech development board offers a low cost platform for Linux based designs. From distributor Silica, the ArchiTech Hachiko board is supplied with a Linux kernel optimised for the Renesas RZ/A1H MCU, to work with a small memory footprint together with a board support package for the on-board peripherals. The Linux distribution with Hachiko is Yocto compatible, enabling you to create custom Linux based systems quickly whilst avoiding the complexity of standard MPU boards. The developers at Silica have ported GUI libraries to the distribution with a demo application and source code is available in open source. The Renesas RZ/A1H processor used in the Hachiko board is based on the ARM Cortex A9 MCU core and offers a large capacity, 10 MB, RAM. Other features include an HDMI output, a USB host, 128 MByte SPIFI, an Ethernet port and an expansion connector for all unused pins.

Complete article, here

IGBT driver cores implement soft shutdown

IGBT driver manufacturer CT-Concept Technologie GmbH, a Power Integrations company, has announced products that include its SCALE-2+ gate driver chip set. SCALE-2+ technology enables soft shutdown (SSD) to be implemented in the event of a short circuit without requiring additional components. This is beneficial in applications with low stray-inductance where full Advanced Active Clamping – a method invented by Concept for shutting down IGBTs or MOSFETs in a controlled manner under any circumstances - may not be necessary. The first product to ship with the SCALE-2+ chip set is Concept's 2SC0106T2A0-12. This is a high-performance two-channel IGBT/MOSFET gate driver core for 1200V IGBTs in the 37 kW to 110 kW power range. The 2SC0106T2A0-12 is suited for servo drives used in industrial applications, UPS and solar inverters. It meets IEC 61800-5-1 and IEC 60664-1 standards for reinforced insulation.

Complete article, here

productroundup

2-W DC-DC cuts footprint 50%

Murata's MTU2 series of ultra miniature surface mounted 2W DC-DC converters measures 8.2 x 8.4 x 8.5 mm with a 0.69 cm2 footprint: it is 50% smaller than the current 1.67 cm^2 industry standard. With a typical conversion efficiency of 85% across the full load range and a power density of 3.403 W/cm^3, the MTU2 series is

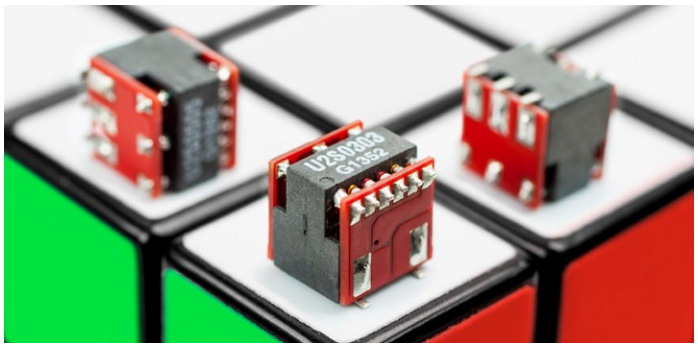

available with either a single or dual output voltage. Input voltages cover the common nominal inputs from 3.3 to 24 VDC. Output voltage options include 5, 12, or 24 VDC. Dual output models provide ±5 or ±12 VDC. Load regulation is typically 5% better than other products, reducing the need for any additional regulation components. The MTU2 has a 1 kVDC input/output galvanic isolation that helps to reduce switching noise and allows the converter to be configured to provide an isolated negative rail in systems where only positive rails exist.

Complete article, here

Basic scopes gain learning support package

Tektronix' TBS1000B-EDU, TBS1000B 2-channel oscilloscopes offer an improved learning experience for students, and gain expanded capabilities for basic R&D. The TBS1000B-EDU and TBS1000B series of 2-channel oscilloscopes gets enhancements such as a high-resolution 7-in. display, dual channel frequency meter and 34 automated measurements; aimed at students and teachers, they also support a wide range of general commercial applications including basic research and development. The TBS1000B-EDU series have an integrated courseware system designed to help students learn more efficiently and effectively while minimising the time instructors need to spend preparing and conducting lab sessions. Instructors can make their lab exercise content viewable on the oscilloscope and students can conveniently capture lab results straight from the instrument.

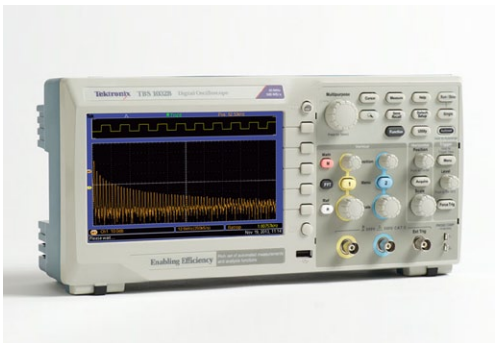

Complete article, here

Range extension up to 7x for low-power 2.4 GHz

Texas Instruments' SimpleLink CC2592 range extender delivers a low cost, high temperature, small form factor design for worldwide operation. The SimpleLink CC2592 range

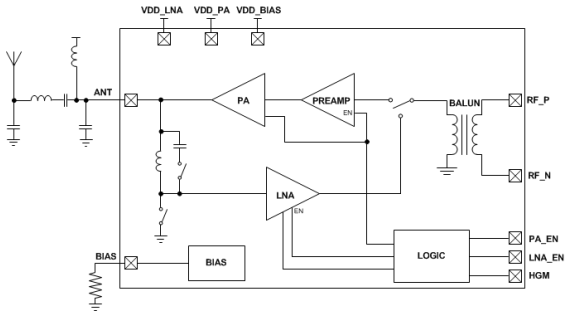

extender provides up to a 7x improvement in range when paired with TI's 2.4 GHz low power RF solutions for ZigBee, 802.15.4, 6LoWPAN and Bluetooth low energy networks. The first such pairing to be made available is the CC2592 with the SimpleLink ZigBee CC2538 wireless microcontroller (MCU). The combined solution has -101 dBm sensitivity and +20 dBm output power, which raises the link budget by 17 dB and equates to up to a 7x range improvement. The CC2592 range extender operates over -40 to +125C; uses fewer passive devices for a complete configuration; and occupies a 4 x 4 mm QFN package. It costs $1.55 (100).

Complete article, here

SEPIC-fed buck DC/DC in 720W I-bus converter

Calling it a second-generation high density isolated intermediate bus converter, CUI has used its recently introduced Solus technology, which combines SEPIC and buck circuit configurations, to build the NQBS series; housed in a quarter brick package, these DC/DC units deliver a power density of 27.2 W/cm^3 (445 W/in^3) as well as

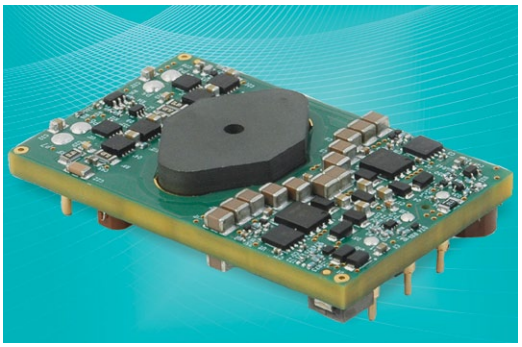

efficiencies in excess of 96%. The Solus Power Topology, CUI claims, delivers a greater power density, higher efficiency, faster transient response, and lower EMI in both isolated and non-isolated DC/DC converter designs. The module incorporates full regulation across the entire 42~60 Vdc input range and a droop load share feature with 10% current share accuracy for higher power boards. As part of CUI's Novum Advanced Power portfolio, the 720W NQBS series outputs 12 Vdc at 60A in an industry standard DOSA footprint (58.4 x 36.8 x 12.3 mm) and provides input to output isolation of 2250 Vdc.

Complete article, here

Cat5e data and 250W power in a single cable

Available at distributor TTI, TE Connectivity's CloudSplitter Connector System combines eight signal and two power positions: applications include intelligent building manage-

ment, video surveillance over IP networks, factory automation, and wireless network transceivers. The CloudSplitter System includes surface mount jacks, cable assemblies, plug kits, and tools. Cable is available in shielded and unshielded configuration and as pre-terminated cable assemblies in standard lengths. Cable is UL444 compliant and CM rated for safety and components are certified to UL 1863. Electrical characteristics for the CloudSplitter Connector System are 1.5A/150 VAC/DC for signal data and it provides DC power at 5A up to 50 VDC over two integrated power contacts. Operating temperature range is -40C to +85C.

Complete article, here

Power system slave controller for up to 240A

LTC3870 is a dual phase synchronous step-down slave controller that can generate currents up to 240A - when paired with a companion master controller - by extending the phase count in multiphase applications. The LTC3870 is an alternative to Linear Technology's full-featured controllers and provides the essential functions

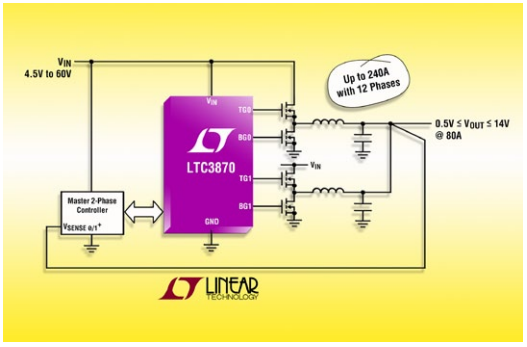

for multiphase slave designs. It operates over an input voltage range of 4.5V to 60V and produces a fixed output voltage up to 14V. Up to 12 phases can be paralleled and clocked out-of-phase to minimise filtering. Its peak current-mode architecture allows for accurate phase-to-phase current sharing even for dynamic loads. Applications include power distribution, redundant (n+1) supplies, industrial systems, DSP and ASIC power. The LTC3870 has a fixed operating frequency from 250 kHz to 1 MHz, or it can be synchronised to an external clock. The 1.1 Ohm onboard all N-channel gate drivers minimise MOSFET switching losses.

Complete article, here

Automotive angular sensor sends with SENT protocol

NXP's KMA215 is an angular sensor with digital output, designed to meet changing needs in the automotive industry; it is compliant with the most recent Single Edge Nibble Transmission (SAE J2716 JAN2010 SENT) standard. The KMA215 is also ready for future designs, already supporting 12-bit High-Speed transmission with doubled frame rate, which is under consideration for future SENT releases. Suitable for all automotive applications where a mechanical angle needs to be measured, the KMA215 offers superior EMC performance and significantly reduces system costs by eliminating the need for external components. Completely calibrated and ready to use, the KMA215 combines full sensor systems including capacitors in a single package. It meets the most recent SENT protocol, the standard interface in motor management control units.

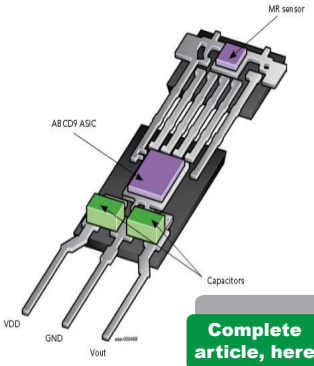

Complete article, here

Infra-red imaging modules for security markets

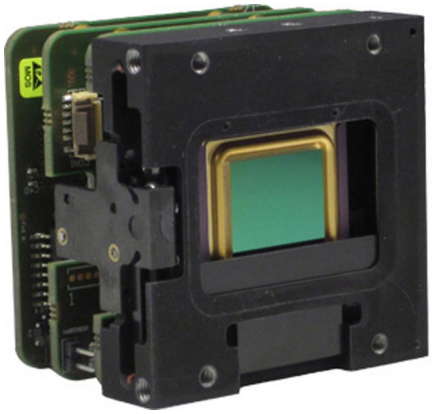

Xenics, European developer and manufacturer of InGaAs infrared detectors and customised IR imaging solutions, is expanding its reach to government and security systems applications for ShortWave InfraRed and LongWave InfraRed OEM modules. From its focus on the R&D and industrial automation markets for its infrared detectors and camera systems, Xenics is putting a stronger focus on applications, such as active and passive day/night vision for enhanced vision systems, situational awareness, UAV (defence, surveillance and fire-fighting), homeland security and SAR (search and rescue) missions. The XenicsCores product family of high-resolution infrared OEM modules offers integration and image quality, with the short-wave module XSW-640 and the thermal imaging XTM-640 module. These uncooled OEM modules are compact and light-weight, and offer low power consumption.

Complete article, here

Distributor expands 3D printer range; from under 1200euro

RS Components (RS), has introduced an array of 3D printers and print technologies to its product offering. The new introductions are from 3D Systems, provider of 3D content-to-print solutions including 3D printers and print materials designed for professional, educational, and personal use.

The 3D Systems products now carried by RS include a range of 3D printers, 3D scanners and cartridges; Second generation Cube - a simple plug-and-play printer offering 3D printing capability with fast print speed and high accuracy for printed parts up to 5.5 in. cubed; CubeX - a desktop 3D printer, offering the largest print volume in its category and mono, dual or triple-colour printing. CubeX provides professional printability with 1,070 cubic inches of space, or 10.8 x 10.45 x 9.5in. There are several options for print modes, including a choice of accuracy and a choice of print fill density: CubeX prints in both PLA and ABS plastics.

Complete article, here

Voice-activation with always-on power levels

Dialog's DA7212 audio codec has low always-on power consumption (650 µW) which can extend the battery life of portable devices – including new 'wearables' such as smart watches and connected fitness devices. Voice-trigger activated applications are feasible with the 24 bit audio codec that supports digital or analogue MEMS microphones and offers a number of DSP features including ALC, 5 band EQ and noise gate. Dialog designed the DA7212 audio codec designed with a novel footprint for cost effective PCB manufacturing; it meets needs of 'wearables' markets with Hi-Fi quality audio at low power levels. The 650µW always-on mode enables constant tracking and instant recognition of voice commands. It also integrates an enhanced hybrid Automatic Level Control (ALC) that offers faster reaction times and finer gain resolution to maintain a consistent audio quality level when recording.

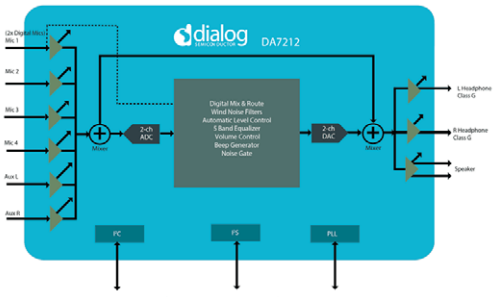

Figure 1: Component values shown for a typical application

Complete article, here

Surface mount ultrasonic transducer

Murata has developed a surface-mount device (SMD)-type ultrasonic transducer; in an ultra low profile package measuring 5.2 x 5.2 mm and 1.15 mm thick, the MA40H1S-R series is capable of delivering up to 100 dB sound pressure at 40 kHz. Operating voltage is 6 Vp-p and the transducer has a -6dB full angle beam of 80 degrees. Through the use of a proprietary structural design, including the ceramic element, Murata Manufacturing has designed what it believes to be the world's first surface-mount device (SMD)-type ultrasonic transducer. Since it is a surface-mount device (SMD), the new device requires minimal mounting space, and it enables improved functionality for distance measurement and position detection through the use of ultrasonic waves. Typical applications include distance measurement, position detection and 3D gesture detection in portable equipment.

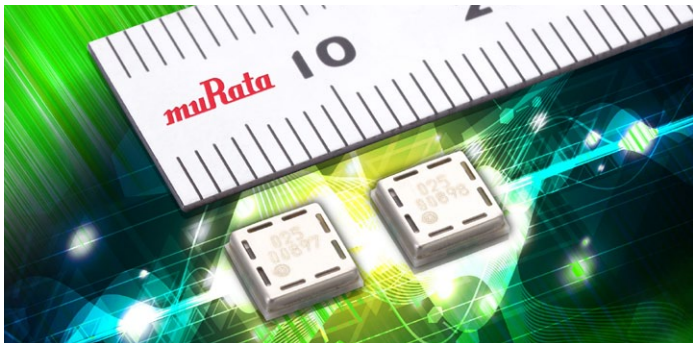

Complete article, here

Smallest transistors reduce footprint by 50%

Rohm has announced what it believes are the smallest discrete transistors available. The VML0604 package (0.6 x 0.4mm, t=0.36mm) cuts size by 50% compared with conventional models; low ON resistance contributes to greater performance and increased miniaturisation in smartphones, and wearable devices. For transistor packages, in addition to technical challenges such as bonding stability and package processing accuracy, it has been difficult to increase element size, resulting in higher ON resistances and a maximum withstand voltage of only 20V. Rohm has been able to successfully reduce package size and ON resistance while improving voltage resistance up to 60V, providing breakthrough performance and applicability. This new package will be expanded to small-signal MOSFETs, contributing to increased space-savings in a variety of sets.

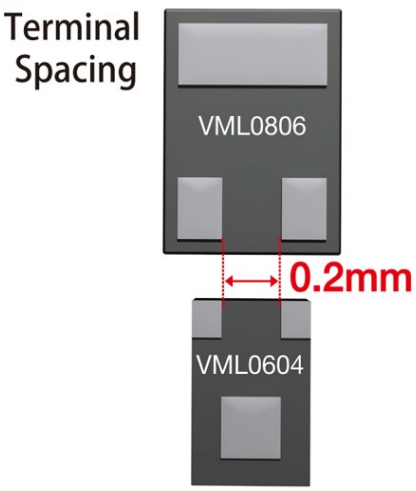

Complete article, here

Computer-on-module hosts ARM Cortex-A CPUs

Distributor Rutronik has boards in a new COM form factor for Cortex-A CPUs, by F&S Elektronik Systeme. efus measures only 47 x 62 mm: efus stands for easy, functional, universal, small: It comes with various interfaces, is expandable with wireless modules and suitable for universal use, e.g. for visualisation, communication, and control in industrial and medical applications. The first module of this new product family, efusA9, will be available by April 2014, mass production starts by the end of quarter 2/14. Further modules will follow later in 2014. The efus product family uses a common 230-pin MXM2 edge connector. The signals on the base board leading to the plug connector were routed by EasyLayout standard, on the principles of no crossing lines or avoidable through-holes. A 4-layer circuit board is sufficient for a base board; the schematic data for developing a base board under EAGLE is available free.

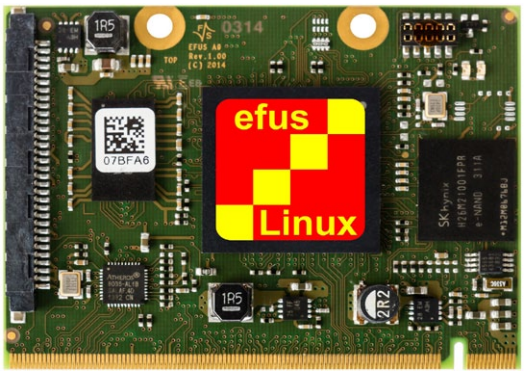

Complete article, here

Competition breeds oscillating amplifiers

I once worked for a company that provided and installed various kinds of communications systems. We got a contract to provide several racks full of 300-W power amplifiers to drive the paging speakers at a large commercial site–seven buildings connected together in all.

The amplifiers had 70-V line outputs to allow connection to 70-V ceiling speaker circuits installed throughout the complex. The amplifiers were commercial grade and were made by a prominent manufacturer of industrial, commercial, and school sound equipment. They had a bandwidth of 20 Hz to 70 kHz, which was pretty common at that time (the early 1980s).

The speakers were interconnected with twisted pair shielded cable, the kind with the foil shield and drain wire. The main feeder cables from the amplifier room to junction boxes in each building were 12 AWG twisted pair shielded cable, which were landed on terminal strips so that 16 AWG twisted pair shielded cable could be connected at each junction box for the speaker branch circuits. When everything was connected up and ready for test, I sent a technician out to commission the system.

The technician came back and reported that the system worked, but he had a suspicion that the amplifiers were oscillating. I told him to go back with an oscilloscope and run tests to see if his suspicion was correct. He did, and the measurements showed a high frequency oscillation up near 70 kHz.

I called the manufacturer, who agreed to send engineers to look at the situation. When they came, I sent my technician with them to the site.

The engineers came back later in the day and reported that they had both observed the oscillation and found a fix for it. I asked them to explain.

They told me that the use of the long runs of shielded cable caused the amplifiers to see a lot of capacitance, more than they were designed to drive. The high capacitance affected the amplifiers in such a way as to cause them to oscillate. The solution was to install an RC network across the output to reduce the amplifier bandwidth high end from 70 kHz to about 12 kHz, which was good enough for voice and background music use.

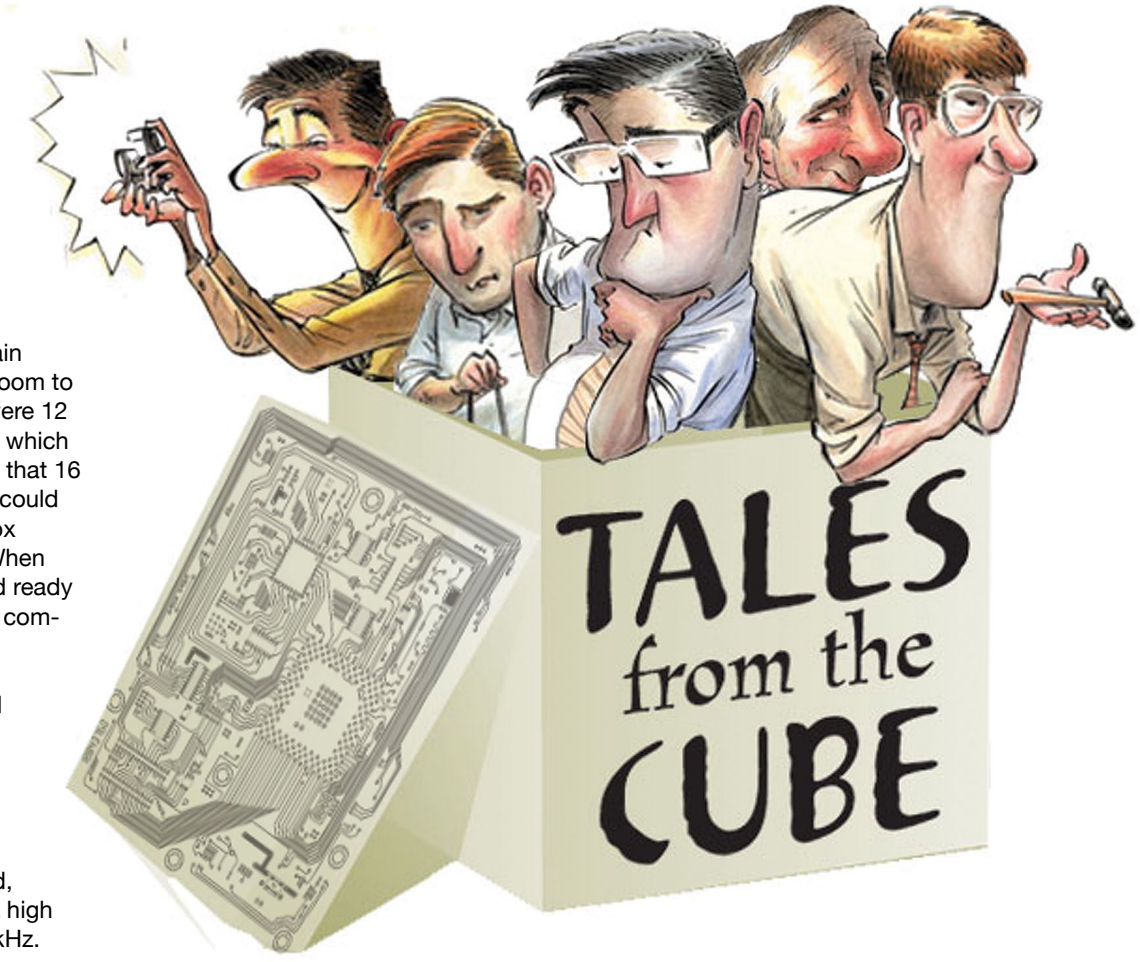

I asked the engineers why they designed amplifiers with such a wide bandwidth if they couldn't drive shielded cable, since the use of shielded cable for 70-V speaker lines was a fairly common practice. The reply was that the manufacturer's marketing department thought it was competing with big name stereo amplifier manufacturers like Crown and McIntosh, and believed that it needed an amplifier with specifications matching amplifiers made by those companies so they could specify and sell them successfully.

I had my technician go back and install the RC networks. The oscillations went away, and I marvelled at how a marketing department could cause such a problem!

Kenneth Ciszewski is a system designer and project manager in St. Louis, Missouri. He has more than 30 years experience designing electronic systems at all levels.

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