EDN europe

JUNE 2016

www.edn-europe.com

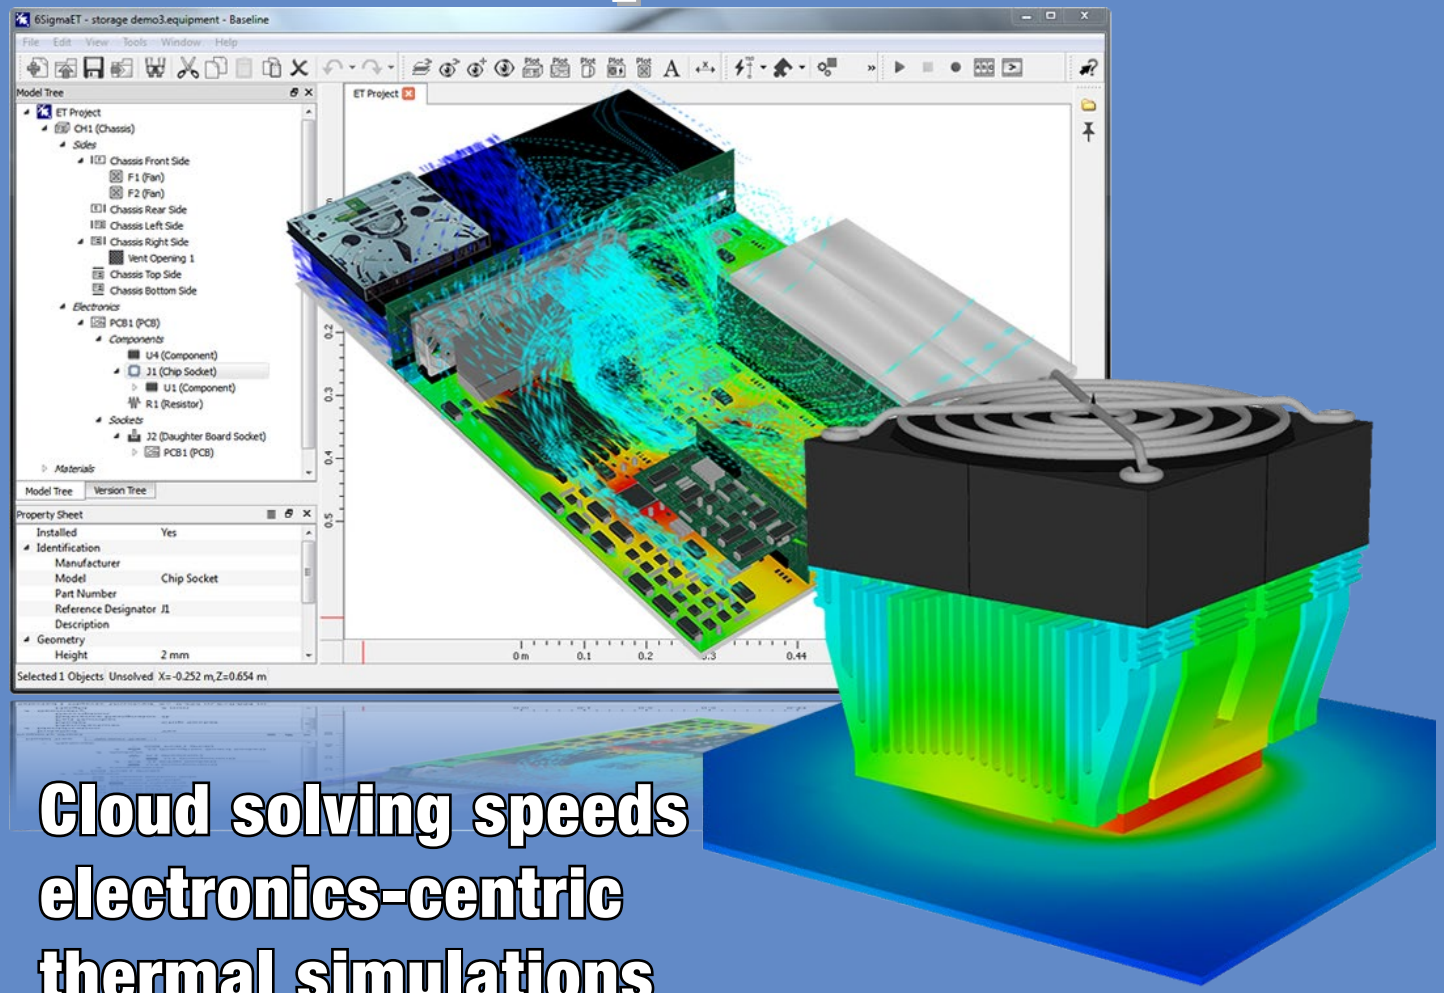

Cloud solving speeds electronics-centric thermal simulations


IMMEDIATE SHIPMENT FROM THE WORLD'S LARGEST SELECTION OF ELECTRONIC COMPONENTS™
1,300,000+ PRODUCTS IN STOCK
4.8 MILLION PARTS ONLINE
OVER 650 INDUSTRY-LEADING SUPPLIERS
100% AUTHORIZED DISTRIBUTOR
FREE SHIPPING
ON ORDERS OVER £50!*
Digi-Key
DIGIKEY.COM
Digi-Key® ELECTRONICS
FIND CONTACT AND ORDERING INFORMATION FOR YOUR REGION AT
DIGIKEY.COM/EUROPE
*A shipping charge of €18.00 (£12.00) will be billed on all orders of less than €65.00 (£50.00). All orders are shipped via UPS for delivery within 1-3 days (dependent on final destination). No handling fees. All prices are in euro and British pound sterling. If excessive weight or unique circumstances require deviation from this charge, customers will be contacted prior to shipping order. Digi-Key is an authorized distributor for all supplier partners. New product added daily. © 2016 Digi-Key Electronics, 701 Brooks Ave. South, Thief River Falls, MN 56701, USA
ecia MEMBER
ecsn member
CEDA

COVER

Thermal simulation update adds CFD solutions in the cloud

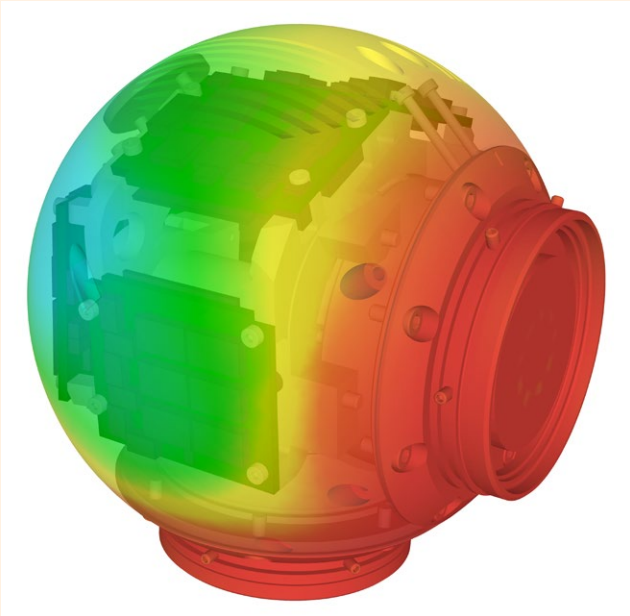

Future Facilities' latest release of its 6SigmaET thermal simulation tool for electronics system design is, the company says, differentiated by being written explicitly for electronics system analysis. The Release includes cloud solving, improved gridding (for CFD, computational fluid dynamics) and enhanced speed and accuracy in the thermal design process. *See item on page 9*

FEATUREARTICLES

ONLINE THIS MONTH

EDN's columns

IOT. AGAIN. AND NOT FOR THE LAST TIME.

Did you want to read another column about the IoT? No, I thought perhaps not, but stick with me a moment. First of all, let's give a mention to the sheer ubiquity of the expression; anyone who is making anything from sensor to cloud server farm – and, very definitely, everything in between – lays claim to making a product "for the IoT". The more so, if they can contrive a low-power angle, or a wireless aspect, or a security angle to their efforts. So "IoT" has become casual shorthand for any product that can be, or some day might be, producing or handling data that is communicated by IP (internet protocol) and sent off to do.... something. Often unspecified.

However, the feedback I get indicates that the EDN Europe reader is already adept at identifying the information he or she needs from the flood of "IoT-related" product introductions and technology announcements. The phenomenon is completely consistent with the early stages of any technology wave; we have had the grand concepts set out, the outpouring of first-round enabling technologies, and the tentative emergence of the first real products. What is likely to happen from here?

It's going to take a while. The parallel of industrial automation may be useful. The notion of the totally integrated factory has been around for a long time and in some places has even been implemented – especially in new-build, green-field projects. A useful expression, though, is, "islands of automation". Rather than having an all-encompassing system implemented in a single pass, many industrial automation installations proceeded piecemeal with individual processes instrumented and automated, with fragmented interconnection, joined up over time (or not). In the era of the IoT, the idea comes to the fore again, this time in the form of IIoT, or industrial-IoT, and this time it comes with all the associated trappings of data mining and extraction of key information from torrents of raw sensor inputs: but the basic idea has been around before. (I am not even going anywhere near the huge span of concepts that different players imply in that other over-used piece of jargon, "Industry 4.0".)

So the reality of first deliverables in the IoT space is that they will be – indeed, are – point solutions to carry our specific tasks, that happen to use IP as their medium of connection. It's worth remembering that very few of these product and service notions are, to date, in any way market tested. We can have wearables that monitor our vital signs and pronounce on our hour-to-hour or minute-to minute health. Other than those of us unfortunate enough to have pre-existing medical conditions, will we care enough about what the system might be telling us, to strap on the sensors and pay the monthly bills? Will our health-care systems have any budget to act on the data, even if we gather it, reduce, and process it? (I have seen the slides, heard the argument, that targeted medicine will be cost-effective in novel ways. I am sceptical.)

In the home-improvement and domestic technology sector, the racks are out on the shop floor with "smart" thermostats, light bulbs, controllers and sensors; it may be different where you live but my casual observation is that they are not exactly flying off the shelves; over time, will the gains be seen to be worth the extra cost to the consumer?

As I said above, it might take a while. You can form your own opinion on the credibility of the many, many reports predicting impressive multipliers in the tens-of-billions of future connected devices. Again, it may be my own scepticism at work, but when the target number is a round number and so too is the date of the prediction, I am immediately suspicious that I'm looking at little more than guesswork.

I'll leave the topic for now, with a couple of (nearly) random thoughts that I'll revisit later. Firstly; the security aspect is very far from fixed, and it's going to get bad before it gets better. And second, if there really are going to be 50 billion (round number!) devices built in a few short years, then the development environments and tools are going to have to get a lot smarter if what we create is not to be a software disaster of epic proportions.

pulse

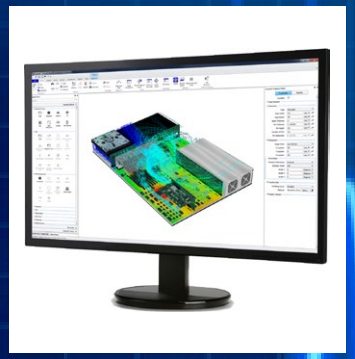


vaunix
Lab Brick
Signal Generator


USB
BLE ANTENNA
RESET
BATTERY
micro:bit

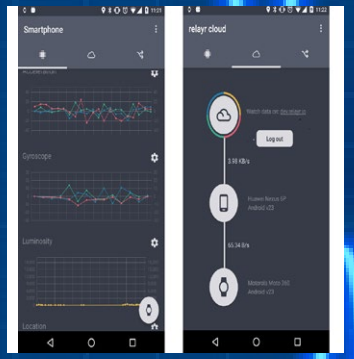
Smartphone
relayr cloud
Log out


LoRa® Technology Evaluation Kit - 800
(Part # DV164140-1)

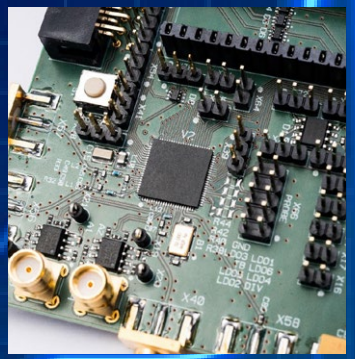

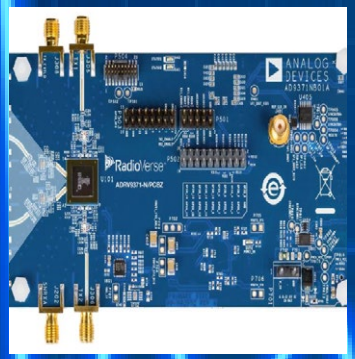
ANALOG DEVICES
Radio/erse


BGT60TR24B


RF sampling 14-bit, 3-GSPS ADC
Data Converter
ADC32RF45
Texas Instruments

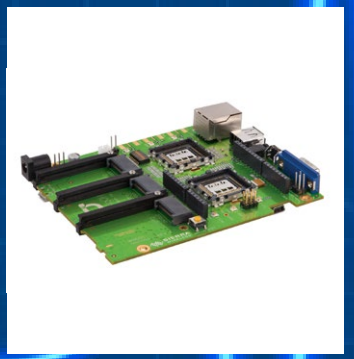


IoT Hardware Enabled Cryptogra
CEC1302

Wideband configurable RF transceiver offers 100 MHz BW to 6 GHz

Analog Devices has extended the features and specification of its AD9361 transceiver chip, with the AD9371 integrated wideband RF transceiver, placing it at the centre of a package it has called RadioVerse, to offer versatile solutions for wireless infrastructure, aerospace and defence, and instrumentation applications. The aims of the RadioVerse programme, says ADI, are to solve more complex radio problems, and to relieve non-specialists of radio problems: increasingly, the radio link may be jut one part of a complex system architecture and it is essential the radio portion is executed quickly. Accordingly, the programme comprises integrated transceiver technologies, a design environment, and market-specific technical expertise intended to take radio designs from concept to creation quickly. Transceiver technologies reduce radio size, weight and power (SWaP), while the design environment offers board support packages, software and tools to simplify and accelerate radio development across a range of applications.

AD9371 is a versatile, carrier-grade, system-on-chip radio with RF tuning range of 300 MHz to 6 GHz, 100-MHz signal bandwidth, and power consumption of less than 5W under standard operating conditions. It has been configured expressly for wideband applications (hence much higher power than the 9361); the company hints that future variants on the same theme will bring similar versatility to narrowband and low power sockets.

The 9371 contains, among other blocks, two independent receive paths with 100 MHz bandwidth, capable of FDD or TDD operation; two transmitters with max large signal BW of 100MHz and max synthesis BW of 250 MHz; and further observation and sniffer receive paths (that you might use for setting digital pre-distortion, for example) with 250 MHz BW. On the baseband side all I/O is purely digital, with JESD204B interfaces; auxiliary DACs, ADCs and GPIO are provided. A spokesman commented that most applications of the device are expected to be in fixed-function sockets; although capable of being applied to reconfigurable, software-defined-radio tasks, most such applications will likely lie in defence or T&M areas.

As well as the integrated RF transceivers, software API, design support packages, documentation, and access to ADI's EngineerZone online technical support community, RadioVerse provides integrated wideband RF transceiver evaluation boards that directly connect to an FPGA development platform, allowing engineers to perform chip-level performance evaluation and rapid prototyping of complete wireless scenarios using a single hardware platform.

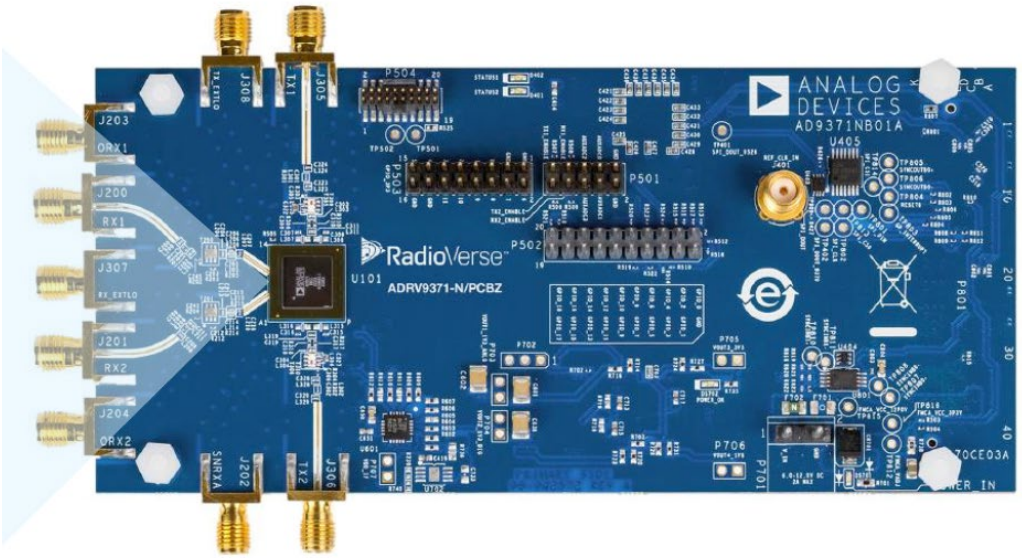

Complete article, here

BBC micro:bit 'first programmable device' now in distribution

Distributor element14 is now accepting orders for the BBC micro:bit, a pocket-sized programmable device that has been designed to introduce children to coding and computing concepts. The widespread availability follows the gifting of the BBC micro:bit to up to one million school children in Year 7 (or equivalent) across the UK, to inspire them to become creative with coding, programming and digital technology.

element14 is the manufacturing partner for the device, which has features such as;

- 25 red LEDs – children are encouraged to light them up, flash messages, create games, and invent digital stories
- Two programmable buttons – associate them with an action and interact with projects
- On-board motion detector – just like a smartphone; the BBC micro:bit can react to shakes, tilts, and even drops
- A built-in compass – sense which direction you're facing and your movement in degrees
- Bluetooth Smart Technology and Low Energy – connect to the internet and the world around
- Five Input and Output rings – connect the BBC micro:bit directly to other devices using crocodile clips or 4mm banana plugs, and send commands back and forth.

The BBC micro:bits will ship in July so children and parents can take advantage of the UK summer holidays to begin their journey into coding; following that, element14 will begin fulfilling orders from its regular sales channels.

The BBC micro:bit can be pre-ordered from element14, the Microsoft Store, Technology Will Save Us, Kitronic and Sciencescope and other resellers. Prices range from £12.99 for a single BBC micro:bit; £14.99 for the starter kit 'BBC micro:bit Go', which includes a BBC micro:bit, mini USB, battery pack and four project ideas to get users started; and £140.00 for a 'BBC micro:bit Club' pack, which includes 10 devices and everything needed to get a coding club started.

The device is supported by a website, www.microbit.co.uk, that features a range of resources and tutorials to help teachers, parents and students take advantage of the BBC micro:bit's vast potential.

The BBC micro:bit can also be coded direct from a smartphone, with apps available for Android and iOS designed by Samsung and Science-Scope respectively.

Complete article, here

6-18 GHz signal generator with 80 dB power control range

Vaunix Technology (Newburyport, Massachusetts), manufacturer of programmable, USB powered test equipment, has added a signal generator to its family of Lab Bricks. The LMS-183DX covers a frequency range of 6-18 GHz in a rugged, portable, palm-sized package.

The LMS-183DX has a power control range of 80 dB, +10 to -70 dBm, from the 6 to 13 GHz frequency range and a power control range of 65 dB, +10 to -55 dBm, from the 13 to 18 GHz frequency range. This unit typically operates with a voltage standing wave ratio (VSWR) at 1:4:1 (2.0:1 maximum) and weighs less than 1 lb (appx 0.5 kg). Maximum harmonics are -25 dBc (-38 dBc typically) when the LMS-183DX is set to +10 dBm output power.

The LMS-183DX is controlled and powered through USB with features such as phase-continuous linear-frequency sweeping, internal/external 10 MHz reference, and optional pulse modulation. This programmable signal generator can be operated through the included graphical user interface (GUI) and retails at $3,499.

Units can be used with any PC or laptop computer with USB 2.0 port (or powered USB hub) and Windows operating system. Lab Brick Signal Generators are provided with Lab Brick GUI software, 32 and 64 bit API DLL files, LabVIEW and Linux compatible drivers.

Complete article, here

Infineon adds MOSFETs to 1200V SiC device offering

Infineon has disclosed a silicon carbide (SiC) MOSFET technology that it says will enable, “unprecedented efficiency and performance” in power conversion designs for, “previously unattainable” levels of power density and performance. CoolSiC MOSFETs promise a new degree of flexibility for increasing efficiency and frequency.

The 1200V SiC MOSFETs will operate with ‘benchmark’ dynamic losses that are an order of magnitude lower than 1200V silicon (Si) IGBTs, for photovoltaic inverters, uninterruptible power supplies (UPS) or charger/storage systems, while later configurations will also extend support to industrial drives.

The MOSFETs are fully compatible with the +15 V/-5 V voltages typically used to drive IGBTs. They combine a threshold voltage rating (V_{th}) of 4V with short-circuit robustness required by the target applications and fully controllable dv/dt characteristics. Key benefits over Si IGBT alternatives include temperature-independent switching losses and threshold-voltage-free on-state characteristics.

The first discrete 1200V CoolSiC MOSFETs feature on-resistance ($R_{DS(ON)}$) ratings of 45 mΩ. They will be available in 3-pin and 4-pin TO-247 packages; both devices are ready for use in synchronous rectification schemes thanks to the integration of a commutation robust body diode operating with nearly zero reverse recovery losses. The 4-pin package incorporates an additional (Kelvin) connection to the source, which is used as a reference potential for the gate driving voltage. By eliminating the effect of voltage drops due to source inductance, this further reduces switching losses, especially at higher switching frequencies.

The MOSFETs are based on a trench semiconductor process and represent the latest evolution of Infineon’s family of CoolSiC technologies. This family includes Schottky diodes and 1200V J-FET devices, and a range of hybrid

solutions that integrate a Si IGBT and SiC diode in a module device. Infineon has also announced 1200V 'Easy1B' half-bridge and booster modules based on the SiC MOSFET technology. Combining PressFIT connections with a good thermal interface, low stray inductance and robust design, each module is available with $R_{DS(ON)}$ rating options of 11 mΩ and 23 mΩ.

Complete article, here

Thermal simulation software update increases speed & accuracy

Future Facilities' Release 10 of its 6SigmaET thermal simulation tool for electronics system design includes cloud solving, improved gridding and enhanced speed and accuracy in the thermal design process. The company points in particular to easier model creation, and to the refinement of its optimising solver.

Future Facilities presents its tool as uniquely targeted to the electronics design task. Whereas other approaches have applied a generalised computational fluid dynamics analysis to the task, 6SigmaET is totally focussed on electronics, its writers say. It works with classes of objects that represent the building blocks of an electronics system; components, PCBs, enclosures and so on. These objects are modelled with attributes and an appropriate set of parameters; for example, says a company spokesman, "a heatsink 'knows' it's a heatsink, it comes in all the standard variations and thermal properties of a heatsink, it knows it is attached to a component and should the designer move that component, it will move with it."

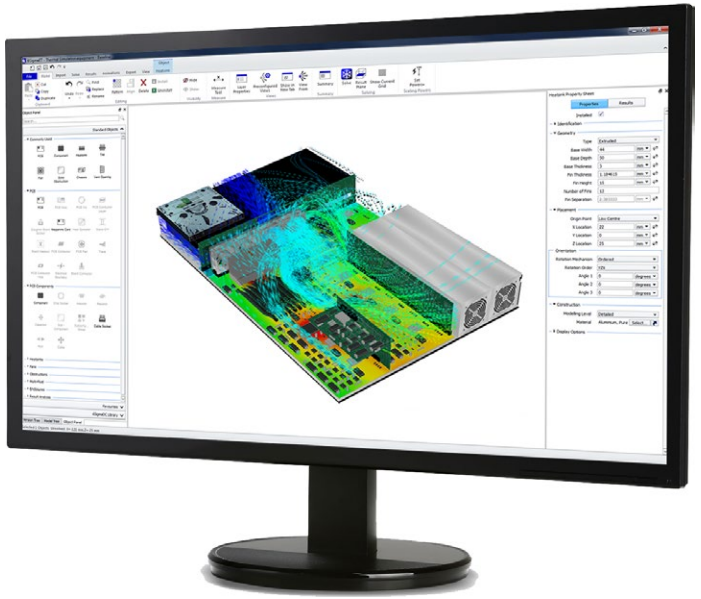

6SigmaET includes modelling down to package level, including the detailed geometry of typical components and the very variable detail of modern (especially consumer) electronics products. "Simulation tool have [in the past] struggled with complex geometry," the company adds. Release 10 introduces a host of new features designed to accelerate solve times, enhance accuracy and enable the creation of more complex designs. At the heart of Release 10 lies an improved unstructured gridding capability, which is claimed to make solving up to six times faster by altering the granularity of its grids to account for varying levels of complexity in the model.

Inputs to the tool in terms of heat dissipations of individual devices, are normally supplied by the user; either as in tabular form using expected power figures, or using worst-case maximum ratings. The engineer identifies peak-power-dissipating parts, and also thermally-sensitive items. 6SigmaET now also calculates the weight of each part so engineers can optimise their product to reduce weight. Through these enhancements, engineers can use 6SigmaET to develop designs for an even greater array of real-life and hypothetical scenarios, without the need for costly prototypes.

Complete article, here

Integrated USB Type-C Hub Controller with USB Power Delivery

Cypress has introduced the EZ-USB HX3C controller that integrates functions previously requiring up to four chips; it provides a low-component-count, space-efficient, USB 3.1 Gen 1, USB Type-C and Power Delivery Hub. Cypress's HX3C solution integrates a USB 3.1 Gen 1 hub controller, a USB Billboard controller and two USB Type-C port controllers into a 10 by 10 mm BGA package. The hub controller presents itself to connected devices as a compound hub, providing three exposed downstream ports and one non-removable port. The two USB Type-C port controllers are associated with the upstream port and one downstream port of the hub, supporting Dual Role Port (DRP) on the upstream port and Down Facing Port (DFP) on the downstream port. The non-removable port is permanently attached to the embedded USB Billboard controller, which can be enabled or disabled via configuration settings.

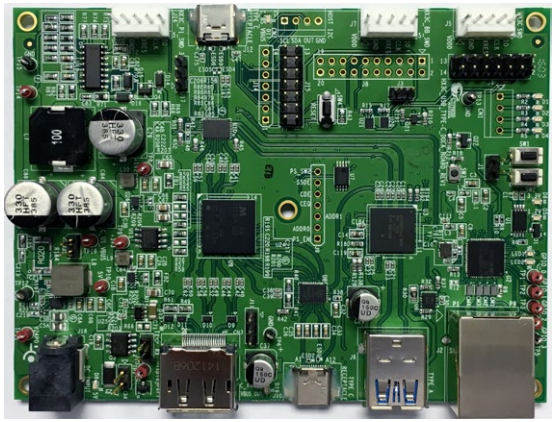

In addition to USB Power Delivery, HX3C supports USB Battery Charging v1.2 (BC 1.2) and Apple charging standards on all of its downstream ports, providing a complete USB charger solution for portable electronics. A video demonstration is available at www.cypress.com/hx3cdemo.

The EZ-USB HX3C solution's two USB Type-C port controllers contain baseband transceivers and physical-layer logic. These transceivers provide BMC and 4b/5b encoding and decoding functions as well as the 1.2-V front end. HX3C contains two ARM Cortex-M0 CPUs, each of which is optimized for low-power operation with extensive clock gating. It uses 16-bit instructions and executes a subset of the Thumb-2 instruction set. This enables fully compatible binary upward migration of the code to higher performance processors.

Complete article, here

Open-source Android app for fast IoT prototyping on smartphones

Relayr is an IoT-focussed company based in Berlin; it has previously produced quick-start IoT development tools in hardware and in software. Now, it has developed Proto-IoT, an app to run on an Android smartphone, or other Android device; Proto IoT is a mobile application which enables a phone to become the hardware behind an IoT project. Proto IoT can:

- Access sensor data easily and securely from anywhere with an internet connection.
- Add interactions and rules on top of sensors' data to trigger actions on separate devices through the mobile app or relayr dashboard.
- Build and demonstrate IoT prototype solutions quickly, with just a smartphone.

The Internet of Things, Relayr says, can be difficult to understand and execute as it requires knowledge of hardware, firmware, and client- and backend-software,

which led Relayr to develop this open-source solution that cuts down development time in the hardware/firmware prototyping stages.

It gives you access to the phone's sensors and it provides real-time data visualization, cloud connection, end-to-end delivery, and if-this-then-that rules and interactions. Additionally, a free developer dashboard monitors your IoT devices from any web browser. The company comments, "Proto IoT is the fastest way to get started in the IoT world. Stop talking about the Internet of Things and start building your solution now."

Complete article, here

LoRa evaluation kits for low-power wide-area networks

This LPWAN turnkey kit from Microchip Technology includes a LoRaWAN gateway, two sensors and a local server application: all required components to create a low-power wide-area network (LPWAN) in Europe (DV164140-1 for the 868 MHz band) or North America (DV164140-2 for the 915 MHz band).
Each kit includes two Motes (LoRaWAN sensors) based on Microchip's RN2483 or RN2903 LoRa modules, a LoRaWAN gateway and a local LoRaWAN server application. Driven by the LoRa Alliance, LoRa technology is able to achieve a range of up to 16 km and 10-year battery life. The technology targets low data rates and low-duty-cycle applications for tracking and monitoring such things as energy, location, utility infrastructure, smart city, environment, agriculture and public safety. Each kit can serve as a building block for development of a long-range LoRa network where designers can expect up to 16 km of range and 10 years of battery life using two AAA batteries. LoRa technology employs spread spectrum modulation which delivers data robustness in a noisy environment and works through physical obstructions.

LoRa® Technology Evaluation Kit - 800 (Part # DV164140-1)

Complete article, here

Lowest power, sub-GHz, multi-standard low-power wide-area transceiver IC

imec and the associated Holst Centre have designed a low-power radio chip for long-range connectivity in sensor networks and future cellular IoT applications; the design features high link budget and claims the lowest power of any such chip, for equivalent RF performance. It will be available to the companies that partnered imec in its development, and as licensable IP.
A number of stand-alone radio networks have emerged for long distance (km) transmission of IoT-class data (remote sensors, and the like), as well as narrow-band services associated with LTE cellular services. imec's radio chip can operate with a lower level of power than any other radio chip technology released to date for this role. The sub-GHz radio chip's technology can serve multiple protocols including IEEE 802.15.4g/k, W-MBUS, KNX-RF, as well as the LoRa and SIGFOX networks, and future cellular IoT.

The radio chip operates in industrial, scientific, medical (ISM) and short-range devices (SRD) bands, covering 780 MHz to 930 MHz. The design offers a high level of interference rejection and lowest bill of materials by minimizing external components as compared to of-the-shelf available chips. The radio is implemented as a complete SoC including the RF front end, power management, an ARM processor, 160 kBytes of SRAM and peripherals such as SPI, I²C and UART. It features a targeted sensitivity of -120 dBm at 0.1% BER (1 kbps) and power consumption of 8 mW (Rx) and 113 mW (Tx) for 13.5 dBm output power. The receiver supports a wide gain range to handle input signals from -120 dBm to -15 dBm, corresponding to a large dynamic range of 105 dB. The PA features automatic ramp-up and ramp-down for ARIB spectral mask compliancy. Output power is controllable from <-40dBm up to 15dBm.

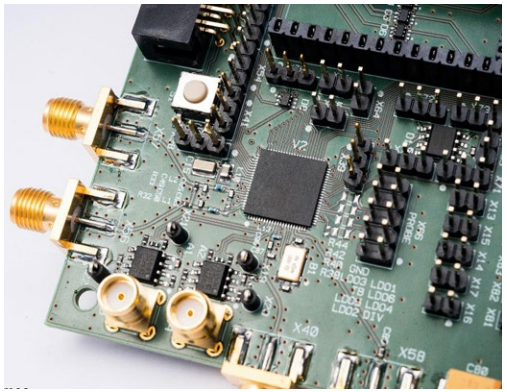

Complete article, here

Open-source IoT development platform for industrial designs

Distributor Farnell element14 has what it positions as the first all-in-one development platform for industrial IoT applications, the mangOH Green Open Hardware Platform. The “all-in-one Hardware, Software and Cloud-based solution for Industrial IoT applications” uses silicon from Sierra Wireless.

The mangOH Green Open Hardware IoT Platform is intended to allow developers to test and prototype ideas quickly and take IoT solutions to market within weeks, as a complete sensor-to-cloud platform. Suitable for data logging in a remote location, the mangOH Green Open Hardware IoT Platform can be applied to remote monitoring systems, as well as multiple sensor, IoT, and cloud-based applications.

At the heart of the board are two industrial-grade CF3 connectors to enable the use of either an AirPrime HL Series module or an AirPrime WP Series applications processor module from Sierra Wireless. The AirPrime WP Series cellular modem provides the device-to-cloud architecture enabling IoT developers to build a Linux-based product on a single module.

Features of its hardware, software and cloud services include:

- Hardware; integrated ARM-based application processor with cellular modem and GPS; Arduino and IoT Connector to add wireless, wired, and sensor technologies; battery enabled for low-power applications.
- Software; Legato open source Linux platform with robust connectivity APIs for accessing cloud and mobile network services.
- Cloud Services; pre-integrated with the IoT Acceleration Platform from Sierra Wireless

The mangOH Green IoT Platform is also compatible with other open source initiatives to allow for easy and fast prototyping; these include NXP thread module, Linear Technology Dust Networks and Texas Instruments ZigBee, Wi-Fi and Bluetooth module.

A starter mangOH kit contains Base board, WP8548 module with Cortex A5 application processor, GNSS receiver and 3G modem, socket cover and release tool, two Breakout Board IoT Connectors, a power supply unit, antennas and USB cable.

Complete article, here

Google ATAP & Infineon joint effort on gesture control with radar technology

Using semiconductor devices from Infineon, two prototypes of products controlled exclusively through gestures were demonstrated at a "Google I/O" event on May 20th 2016. The two products, a smartwatch and a wireless speaker, can both recognize gestures to replace physical switches or buttons. Both companies are preparing for the joint commercialization of the "Soli" technology. "Sophisticated haptic algorithms combined with highly integrated and miniaturized radar chips can foster a huge variety of applications," Infineon comments.
In addition to their efforts in the audio and smartwatch markets, the developers' ambitions are more comprehensive: "It is our target to create a new market standard with compelling performance and new user experience, creating a core technology for enablement of augmented reality and IoT," said Infineon division President Urschitz.
The company believes that while Virtual Reality concepts could be visualised, users could not interact with these so far. The 60 GHz radar application developed by Google and Infineon bridges the gap, as a key technology enabling Augmented Reality. This technology was developed by Google and Infineon over the last few years under the internal name "Soli". A 9 x 12.5 mm radar chip from Infineon sends and receives waves that reflect off the user's finger. Fine hand movements, such as winding a watch, can be resolved at a distance of up to 15 metres. Just a few decades ago it took a parabolic antenna with a 50 m diameter to do what the chip's technology can do today.
The chip is made by Infineon in Regensburg, Germany. It operates with the algorithm developed by Google's ATAP (short for Advanced Technology and Projects Group). The combination is used in the smartwatch from LG and the JBL wireless speaker. Beginning in mid 2017, both companies will co-market the hardware and software as a single solution.

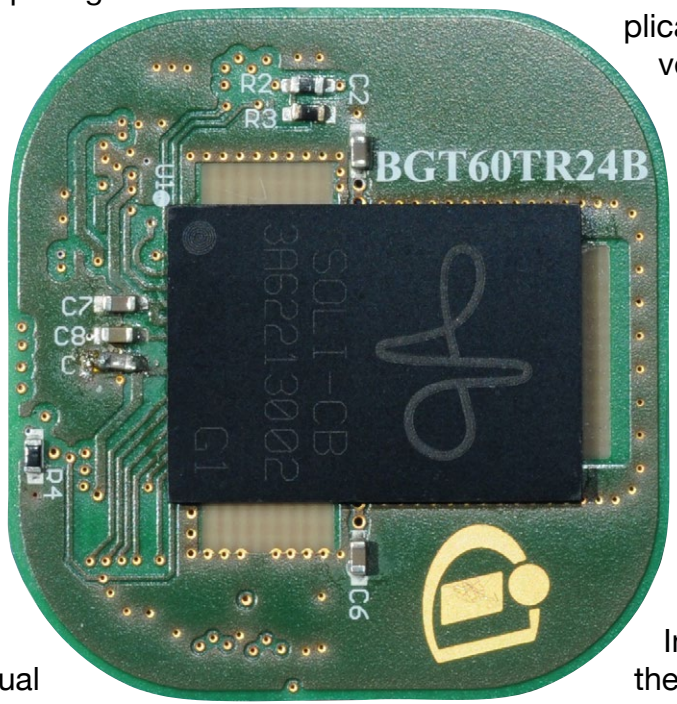

Complete article, here

ST & Arduino to expand 'maker' access to STM32 MCUs & sensors

A collaboration between the open-source board maker and the semiconductor company yields the first Arduino board based on the STM32 ARM Cortex-M MCU, which offers high-performance graphics, TFT touch display, wireless link, and connectivity for audio, MicroSD, USB OTG, and camera. The STM32 family of microcontrollers, along with ST's full portfolio of sensing, power, and connectivity technology, will be more accessible to the Arduino maker community.
The first product of the STAR (ST

and Arduino) program is the ST-M32F469-based STAR Otto base-board, allowing IoT developers to build high-performance graphics into their smart devices using accessible hardware and software to improve their applications with easy-to-use touch displays and audio for command and control as well as for media-streaming use cases.

STAR Otto Arduino Board uses STM32 MCUs and ST Sensors

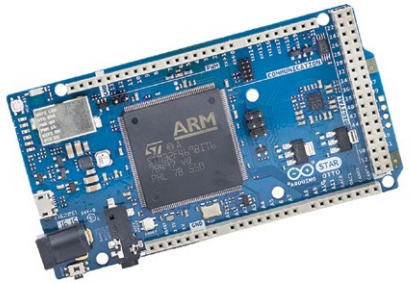

STAR Otto is built around the 32-bit ST-M32F469 MCU, which includes ST's Chrom-ART graphics accelerator and MIPI DSI display interface along with an open-source software graphics library. STAR Otto provides a pre-integrated wireless link and audio capabilities, enabled by an ST MEMS microphone together with the necessary open-source drivers. This efficient and optimized approach lets makers focus on their value-add and simplifies integration, enabling a broad range of Smart Home and Smart Industry applications.

Complete article, here

14-bit, 3-Gsps RF sampling ADC – at $2.5k each

Texas instruments has what it presents as the highest-available performance RF sampling analogue-to-digital converter, with greatest dynamic range, widest bandwidth and fastest conversion at 14-bits. The ADC32RF45, claimed as the industry's fastest 14-bit analogue-to-digital converter, is a dual-channel device accepting direct RF signal for conversion at up to 4 GHz, giving engineers access to the highest dynamic range and input bandwidth. The first in a new data converter family in TI's RF sampling portfolio, the ADC32RF45 eliminates up to four intermediate-frequency down-conversion stages in multi-band receivers, which simplifies system architecture and reduces board space.

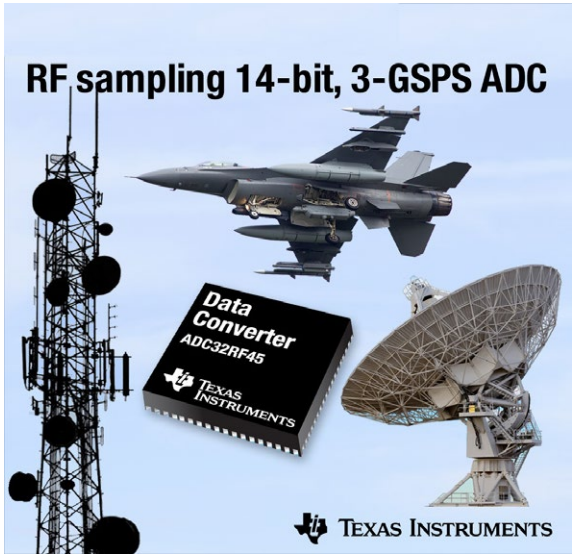

The converter has support for RF inputs up to 4 GHz and enables direct RF signal conversion in the first, second and third Nyquist zones (including all L- and S-band frequency ranges). This reduces filter complexity, saves board space and decreases component count. It detects even the weakest signals with noise-spectral density of -155 dBFS/Hz, 5 dB better than competitive devices. It offers signal-to-noise ratio of 58.5 dB at a 1.8-GHz input frequency. It is claimed as the fastest 14-bit ADC with a 3-Gsample/sec maximum sample rate that delivers an instantaneous 1.5-GHz-per-channel bandwidth, enabling engineers to implement wideband in-phase and quadrature-component receivers beyond 2.5 GHz.

An evaluation board, the AD-C32RF45EVM is priced at $2,499; the ADC32RF45 will be in volume production in the third quarter of 2016 in a 72-pin, 10 x 10-mm quad flat no-lead (QFN) package for $2,495 in (100).

Complete article, here

Crypto-enabled, ARM-based, 32-bit MCU for IoT

Claiming significant performance improvements for Internet of Things applications – over firmware-based security solutions – Microchip Technology has added the CEC1302 microcontroller with hardware-based cryptography; the MCU, says Microchip, makes it easy to add security, offering easy-to-use encryption and authentication for programming flexibility and increased levels of security.

The CEC1302 allows for pre-boot authentication of the system firmware in order to ensure that the firmware is untouched and uncorrupted, thereby preventing security attacks such as man-in-the-middle, denial-of-service and backdoor vulnerabilities. It can also be used to authenticate any firmware updates, protecting the system from malware or memory corruption.

The CEC1302 offers private key and customer programming flexibility with a full-featured microcontroller in a single-package solution in order to minimize customer risk. The device provides savings in terms of power drain and also improved execution of application performance. Since the CEC1302 is a full 32-bit microcontroller with an ARM Cortex-M4 core, adding security functionality only results in a small additional cost. The CEC1302 can be used as a standalone security coprocessor or can replace an existing microcontroller.

Complete article, here

USE BRUSHLESS MOTORS AS EASILY AS BRUSHED MOTORS

By Berko Kletzander, ON Semiconductor

Surprisingly, the brushed electric motor and the incandescent bulb have many things in common: both were invented and optimized in the 19th century, their ease of use has made them ubiquitous companions of our everyday life throughout the 20th century and they are now both falling victim to regulations emphasizing energy efficiency in the 21st.

But while the incandescent lightbulb has largely disappeared, brushed motors are only slowly being replaced by their more efficient brushless variants. This is partly for cost reasons, but also because significant technical challenges remain when implementing brushless motor systems. Semiconductor technology has helped the cheap but inefficient incandescent lightbulb disappear and will likely do the same with the brushed motor.

Electric motors are present everywhere, but unlike lightbulbs, they are less obvious. Whenever something inanimate moves, most likely it is being powered by an electric motor. These motors run our refrigerators, dishwashers, and washing machines at home. They open doors, move elevators, escalators and trains as we make our way to work. They brew our coffee, cool our computers and run the air conditioning at work. In modern cars 20 to 50 electric motors are employed to power a variety of functions from adjusting the mirrors, headlights and seats, pumping fuel, water and windshield fluid, as well as assisting with braking and steering. As the electric vehicle becomes more prevalent, electric motors will take over moving the car itself.

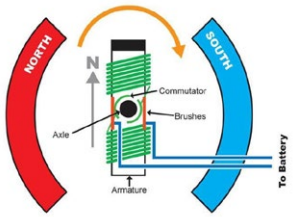

Figure 1. *Schematic of a brushed motor*

The motor of choice for most of these applications over the last century has been the brushed motor because it is beautifully simple to use. Like the lightbulb, it works as soon as it is switched on - reverse the voltage across its terminals and the motor will reverse direction. Vary that voltage and the speed changes. This simplicity is possible because a mechanical switch (the commutator) inside the motor keeps it rotating. The commutator consists of a revolving segmented assembly that connects to the rotating motor coils; and sliding contacts (brushes) that conduct the current from the fixed motor case to connecting segments. As the motor coil assembly rotates, the coil terminals slide past the brushes, directing current through the coils in such a way as to maintain rotation.

While the commutator makes brushed motors easy to use, it is also the main source of its problems. The commutator (usually, mainly copper) is subject to mechanical wear and tear, limiting the lifetime of the motor. The brushes (normally carbon) wear out generating dust which reduces motor performance. In addition, when the motor current switches from coil to coil along the slip rings, electric arcs form. Such brush arcing generates ozone, as well as causing acoustic noise and EMI, with all of its associated potential for interference. In dusty or combustible environments, brush arcing can be very hazardous as it can start fires and even cause explosions. The commutator, with its mechanical brushes, also limits the electric efficiency of the motor up to a maximum of 70%; whereas a motor without brushes can achieve efficiencies of up to 90% (as Nicola Tesla showed back in the late 19th century). If brushless motors are more efficient, more reliable and more environmentally friendly than brushed motors, why are brushed motors still so common today?

Highly integrated silicon is the answer, as the continuation of this article explains.

Download PDF of Article

Find brushless motor drivers on EETsearch

EFFICIENT PARASITIC EXTRACTION TECHNIQUES FOR FULL-CHIP VERIFICATION

By Yousry Elmaghraby, Mentor Graphics

Ever since the idea of building electronic systems using integrated circuits (ICs) was born in 1949, technology has evolved swiftly. Moving from small to very-large-scale integration, transistor count on a single chip has grown from tens to billions of transistors. As technology evolves, feature size gets smaller and the number of stack layers increases, making physical and electrical verification not only more complex, but also more critical throughout the design cycle.

Parasitic extraction and circuit simulation are major challenges in today's chip-level verification process. While most parasitic extraction tools supplied by electronic design automation (EDA) vendors offer an extraction solution for full system-on-chip (SoC) and full-chip memory designs, the huge device count in these designs, as well as their increasing routing and device complexity, make it challenging for SPICE simulators to swallow such a huge network of devices and extracted components at the chip level. Extraction becomes even more challenging for chips that include large LEF/DEF and metal fill macros.

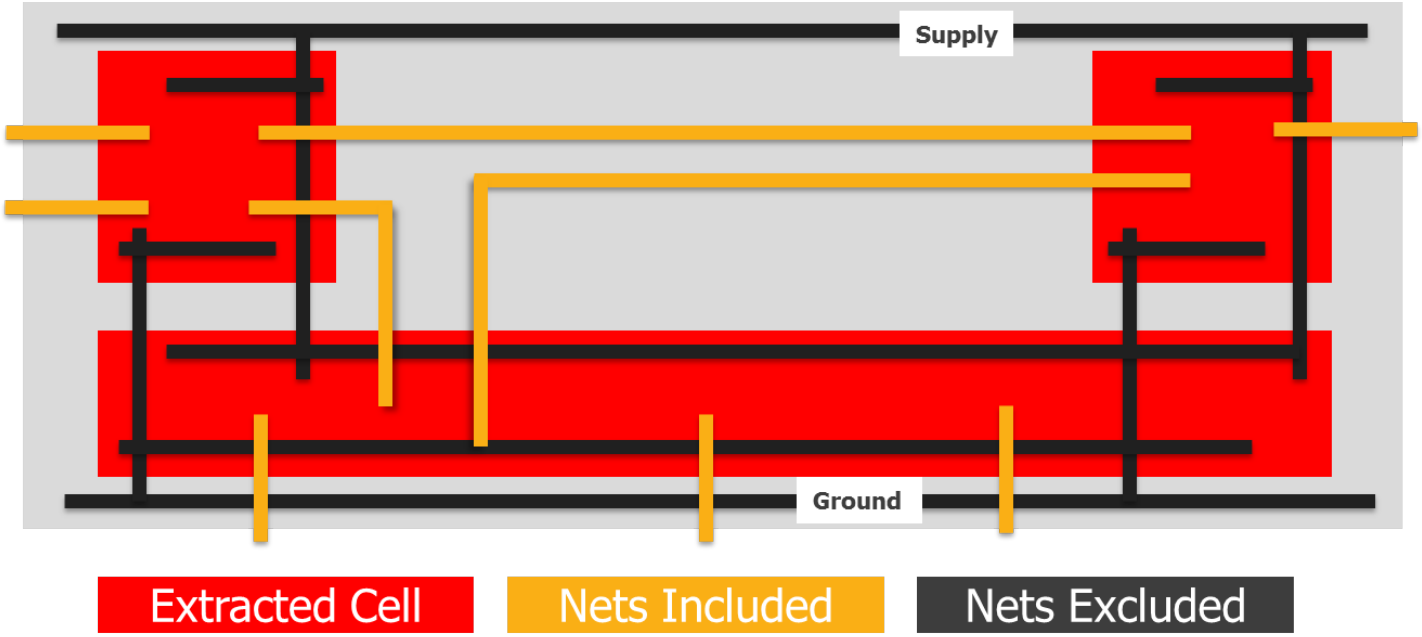

To keep up with this growth in complexity and size, some parasitic extraction engines provide multiple extraction approaches that target different design types, design stages, and circuit verification intentions. Deciding which method to use is a matter of understanding your extraction requirements in the context of the design. Let's take a look at the options, and why you might choose one technique over another.

The most accurate, yet relatively time-consuming, extraction mode is the "flat" transistor-level mode. In flat mode, extraction details are captured down to the transistor level, and all design cells are flattened to enable the extraction of all interactions between components. However, if you are creating a mixed-signal design with some pre-characterized digital modules, or you have incomplete analogue blocks in your design, using the flat mode will

not fit your extraction needs because it would result in double-counting these cells, leading to inaccurate results.

Another consideration is the size of the extracted netlist. Real designs often include modules that are repeated multiple times in the layout. For example, extracting a netlist of several gigabytes of memory would significantly increase simulation time, without providing any increase in useful parasitics information. The hierarchical extraction mode provides a way to extract any repeated modules once in isolation, but still capture their total loading effects on routing at higher levels of hierarchy. Extracting such blocks once saves both extraction and simulation run time.

To further minimize netlist size, hierarchical extraction also provides an efficient way to zoom into parasitic details for the most sensitive analogue blocks, while giving less focus to parasitic extraction in other static modules in your design. You can define specific modules/nets in the design that can be excluded during extraction, while all other blocks/nets are fully extracted.

Although full-chip verification is still essential for final signoff quality, designers can incrementally extract and simulate the connectivity of different blocks at the chip level during the design and verification flow. This approach provides a fast, yet indicative, view of the loading and coupling effects coming from top-level integration and acting on small blocks. Quick access to this knowledge reduces the turnaround time for multiple verification iterations, helping designers achieve a faster time to market.

The full version of this article illustrates some typical situations where different extraction features and techniques can help quickly capture problems in early stages of your design - click for pdf.

Download PDF of Article

Find full-chip verification on EETsearch

HANDS-ON: SEGGER SYSTEMVIEW UTILITY

BY JACOB BENINGO

Microcontrollers are becoming extraordinarily complex and feature-rich, so traditional debug techniques that have existed since the good old 8-bit days just aren't efficient or insightful enough to handle the challenges developers now face. Trace analysis and debug tools are coming to market, though, that hold the potential to revolutionize the way that embedded systems are debugged. Let's take a look at one such tool – Segger SystemView – to see how developers can get started using advanced debugging techniques.

Every time I debug an embedded system I am always uneasy. There are so many assumptions that are being made as to how the code flows, how interrupts fire, and about code timing, to name a few. A single incorrect assumption can cause the debug cycle to drag on by masking the true root cause. I believe this is where the power of trace tools has the potential to change how we debug systems. Being able to peer into the system, generate time-stamps and event logs, and track task execution is very powerful. I no longer need to assume that my system is executing in a certain way; I can visually verify that it is. If my assumption is wrong, the evidence is right there in the trace and now I know my grief's source.

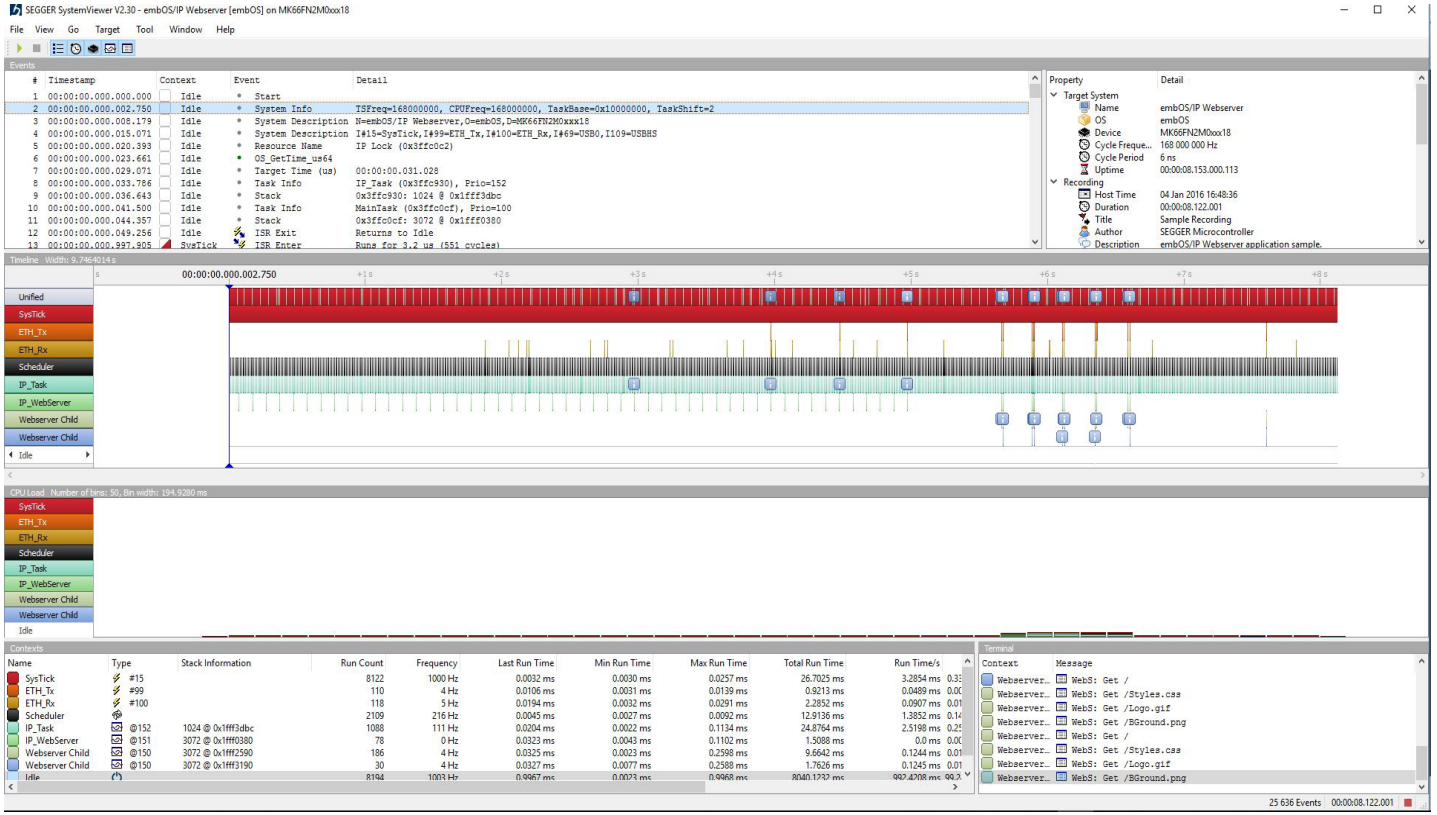

Figure 1. *A visual representation of software execution from Segger SystemView (click to expand)*

Advanced trace tools provide information in a visual form. When I'm teaching a course on real-time operating systems and discussing rate monotonic analysis, I often draw timing diagrams that demonstrate how tasks are time sliced and getting dedicated access to the CPU, show expected task completion times, and other useful design parameters. SystemView allows an embedded software developer to visually see how their software is executing in just such a manner. The SystemView window in Figure 1 demonstrates exactly what a devel-

oper would see when they start the tool. The tool traces the application execution by recording entry and exit from functions, system events, interrupt occurrences, and other parameters that are interesting to a developer. The trace information can be downloaded from the system and analyzed, for example to determine if an unexpected glitch is occurring or determine if a task is taking too long to execute. Depending on the microcontroller architecture, the trace data can be reviewed in a streaming real-time manner or downloaded in a single-shot recording.

Great debugging powers such as these do not come without some disadvantages. So far the biggest issue I have encountered is the time investment needed to get over the tool's learning curve. Segger's website has some great examples and instructions, but there is so much technology built into trace capability that wrapping one's mind around exactly how it works and how to set it up isn't trivial. I spent significant time reading the 154-page manual then setting up an embedded system and configuring it before I finally started to get trace data. Once I was through the learning curve, though, and had documented my own process, each use got easier and easier.

SystemView uses a technology known as Real Time Transfer (RTT) that must be installed on the embedded target. RTT is used to collect data on the target and store the data in a buffer. RTT then streams the data through debugger hardware to SystemView.

The documentation mentions that a typical implementation uses approximately 2 kbyte of flash space and 600 bytes of RAM. My first setup attempt used a plugin for Processor Expert, which required an up and down buffer of 1024 to get the code to compile. Eventually I discarded the plugin and manually set up the interface. I was then able to achieve numbers much closer to those in the datasheet.

The great part about the SystemView tool is that it works with any Segger J-Link debugger and can even be used with a non-Segger debugger under certain conditions. The software is a free download and doesn't have any fees, licenses, or royalties associated with it. There is even an API so that engineers can create their own tools that use the interface. If SystemView isn't to one's liking, alternative tools exist such as Percepios' Tracealyzer that can also be used with the J-Link and RTT interface.

Using a trace tool will undoubtedly become standard practice at the development cycle's start. Embedded software developers can run their system overnight and gather a baseline trace on performance, timing, and other critical system metrics. Then they can use continuous testing, a great way to catch deviations from the baseline when they first occur rather than allowing issues to stack up until an obvious problem arises. Early discovery coupled with source control should easily identify problematic code and result in a quick fix rather than a convoluted, snail's pace debug session.

In time, trace technology may become so seamless and automated that developers will no longer need to debug an embedded system. Computers will be able to automatically detect noncompliance and deviations and highlight the recently changed lines of code in the misbehaving task. The thought is a bit optimistic, but one can always hope.

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

SMALL-FORMAT FLASH MEMORY CARDS STEP UP TO INDUSTRIAL APPLICATIONS

By Ulrich Brandt, Swissbit

Removable flash memory – the familiar memory cards of personal computing products, cameras and the like – are most often thought of as consumer electronics products. They can achieve the verified performance levels needed to apply them as Industrial Technology: FTL firmware brings the enhanced performance and endurance to SD and microSD formats.

As the capacities of flash memory ICs increase, memory cards in small form factors can provide enough non-volatile storage for a growing variety of industrial applications. High-density, multi-gigabit SD and microSD cards are ideal for ruggedized embedded systems where size constraints and hazards such as shock, vibration and extreme temperatures across the industrial range prevent the use of a conventional hard disc drive.

The cards must, however, satisfy the stringent performance and reliability demands of the industrial environment, at a price that is cost-effective for the application.

One crucial metric that determines the speed and reliability of flash memory is random write performance. Industrial devices such as handheld terminals and embedded computers typically make extensive use of random writes to store data such as status updates or results from manipulating communication protocols. Random write performance is heavily influenced by the operation of the Flash Translation Layer (FTL). The FTL enables the host to interact normally with flash memory by mapping the logical sector numbers used by the host file system to the physical addresses of the data as it is stored in the memory.

Because data in flash cells cannot be overwritten with new data, any invalid, or stale, data must be erased before the new data can be written. However, flash can only be erased on a per-block basis. For this reason, after all the blocks of a given flash device have had data written at least once, the FTL also takes

care of garbage collection: moving valid data from partially filled blocks into new blocks so that the old block can be erased and re-used. If garbage collection is done while the data is being updated, random write performance can be impaired.

As an alternative, Idle Time Garbage Collection (ITGC) can avoid such time penalties, but may not be possible depending on the duty cycle of the flash. Spare blocks and Over Provisioning (OP: providing more flash capacity than is addressable by the host) are other techniques used to minimise the impact of garbage collection on memory write performance.

Address-mapping granularity

Until now, commercial FTLs have typically either used block-based address translation or page-based translation. In a block-based scheme, the logical sector referenced by the host system is assigned a logical block number that is then mapped to a physical block in the flash. Depending on the FTL, the data may be always stored at the same physical page number in the block, or the physical page number may vary. If the physical location varies, extra mapping information is needed to pinpoint the correct page in the block. Consumer flash cards often use a block-based mapping scheme for greater efficiency when handling sequential writes involving many pages of data in a single transaction.

One of the strengths of block-based mapping is that relatively little mapping data is needed to describe the physical location of the page when written. On the other hand, writing a single page of data randomly to flash involves a high number of data writes to flash cells: valid pages in the block must be saved before the block is erased, and then merged with the new data when written back into the block.

The ratio of the data written to the flash cells compared to data written by the host is known as the Write Amplification Factor (WAF); reducing WAF is essential if economical flash cards are to become acceptable in industrial applications that are characterised by intensive random-write instructions. *Read on for an explanation of how write techniques are refined for industrial use... click for pdf*

Download PDF of Article

Find industrial flash memory cards on EETsearch

CORRECT VOLTAGE REGULATOR SELECTION IS KEY TO BATTERY LIFE IN WEARABLES

By Nazzareno Rossetti and Meng He, Maxim Integrated

Wearable devices have emerged as the next big market opportunity in the electronics industry. Smart watches, as shown in Figure 1, are among the most popular wearable items today. The healthcare market, including the medical, fitness and wellness sectors, promises even broader opportunities.

The majority of wearable gadgets have a number of things in common. Wearable devices must:

- Be always ready for use
- Be small and lightweight in order to be easy to wear
- Last a sufficiently long time on a re-charge or on a disposable battery
- Support short periods of activity, spending the majority of time in idle or sleep mode
- Last a very long time in idle or sleep mode

These requirements place heightened demands on all the technologies underlying the products. Batteries must be ever-smaller and last longer. For example, smart phone batteries have a range of capacities around 2000 mAh. The battery of a smart watch, while much smaller in size and with a charge capacity about 10 times smaller, is still required to operate for the same time duration between recharges as its smart phone big brother. Correspondingly, a quantum leap in performance is required from the rest of the watch's electronic

Figure 1. *The wearable lifestyle*

components. Namely, they must be small and consume minimum power both in active mode and in passive (stand-by and shutdown) mode.

Low power consumption all around

Until recently the primary focus for voltage regulation design has been the efficiency of power delivery in active mode, from light to peak to full load. Business is routinely won and lost over a fraction of a percent of efficiency advantage. With the efficiency curve well understood and opportunities for improvement reaching saturation the focus is shifting to optimizing the power savings in passive mode. Passive mode corresponds to the system being in idle mode (still on but in standby) or in sleep mode (when the system is in shutdown).

This shift is necessary due to the popularity of wearable devices, which spend a lot of time in a passive mode with only infrequent periods of activity. It's clear that if idle and sleep mode are the dominant modes this is where the power savings need to come from. Here, indeed, every nanoampere counts since it is amplified over long periods of inactivity and ends up robbing precious charge from the battery.

As an example, a 40 mAh 1.55V silver oxide coin cell battery is a good candidate for powering a wearable device. If the current drawn by the wearable device is 4 µA the battery will have a shelf life of about one year before it runs out of charge. A reduction in current draw of a single microampere would increase the wearable shelf life by approximately three months!

With this in mind, consider the conventional product portfolio of ultra-portable voltage regulators, which contribute tens of microamps of quiescent current and several microamps of shut-down current to the overall device current consumption. We can be quickly convinced that any improvement in these parasitic currents will be beneficial. *Click for full article pdf.*

Download PDF of Article

Find voltage regulator IC on EETsearch

DON'T CONFUSE PAM4 SER AND BER

BY RANSOM STEPHENS

As we make the transition from NRZ (non-return to zero) to PAM4 it's important to distinguish BER (bit error ratio) from SER (symbol error ratio). It sounds simple enough but after saying "BER" for your entire career, it's easy to make obvious mistakes. In most cases, if you're talking about waveforms and eye diagrams you want SER, and if you're talking about anything else, stop talking, listen carefully, and double check.

As you know, PAM4 (4 level pulse amplitude modulation) encodes two bits in each of four symbol levels. The Gray coding scheme determines which pairs of bits are assigned to which symbols (see five part series links below), but Gray coding isn't relevant when you're analyzing eye diagrams. The fact that PAM4 eye diagrams come with three pupils is relevant.

We can extract bathtub plots from each of PAM4's three pupils in both the horizontal jitter direction and vertical noise direction. A PAM4 jitter bathtub plot (Figure 1) measures SER(x), where x is the time-delay position of the sampling point and the noise bathtub plot measures SER(V), where V is the vertical position of the sampling point—voltage for electrical systems and power for optical. It's easy to extract eye height, eye width, and total jitter defined with respect to SER from SER bathtub plots.

Right now, it looks like the tardy but emerging PAM4 specifications will define the system EH (eye height) and EW (eye width) as the smallest value of the three PAM4 eye diagrams:

EW = min(EWLOW, EWMID, EWUPP)

and

EH = min(EHLOW, EHMID, EHUPP)

where LOW, MID, and UPP refer to the lowest,

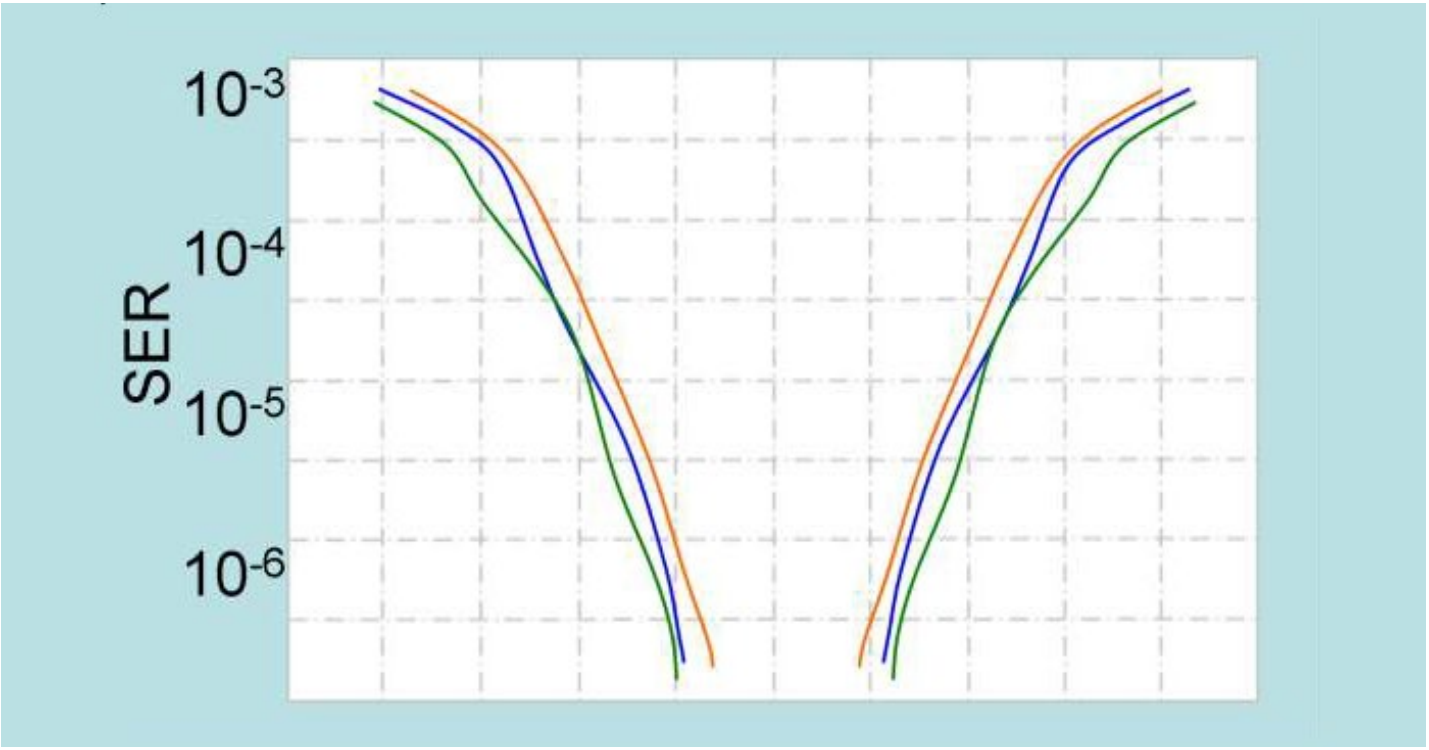

Figure 1. *Three SER bathtub plots based on a PAM4 signal.*

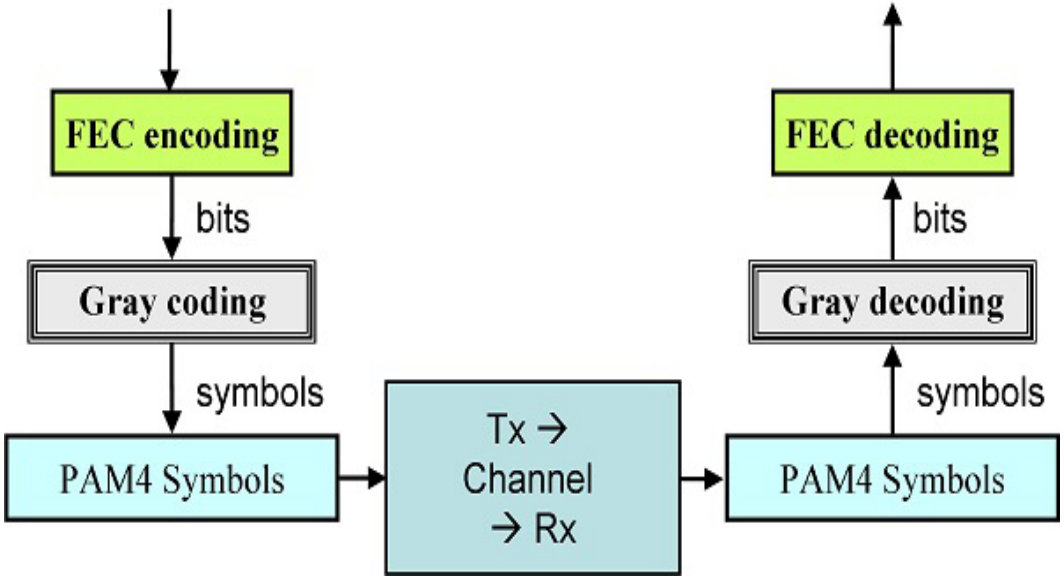

Figure 2. *Gray coding/decoding (graphic copyright Ransom Stephens)*

middle, and upper of PAM4's three eyes. Minimum allowed values for EH6 and EW6, where the 6 refers to SER = 1E-6 will be specified depending on the application.

The only PAM4 serial data standard that's been released, 100 GbE's 100GBASE-KP4, as well as (rumour has it) the tardy standards, require a common time-delay centre for analysis of each eye pupil. The time-delay centre is defined to be the midpoint of the widest horizontal opening of the middle pupil. This definition accommodates receivers with symbol identifying circuits (a.k.a., voltage slicers) that have common timing. As rates grow and PAM4 evolves, expect receivers with three independent slicers and each of the three eyes to be allowed their own centres.

Along with saving bandwidth by cramming two bits into one symbol period, PAM4 also introduces FEC (forward error correction) that relaxes the SER requirement, which is why they specify minimum requirements for SER < 1E-6 instead of NRZ's BER < 1E-12 or 1E-15.

The higher error ratio makes it possible for oscilloscopes to acquire enough data to measure SER bathtub plots and contours. In the olden days, oscilloscopes had to extrapolate to estimate eye openings at BERs of 1E-12 or 1E-15. But be careful! Some test equipment applies Gray decoding to PAM4 signals before extracting bathtub plots and contours and that means that they report BER not SER (Figure 2).

The SER-BER conundrum gets worse: The FEC required by most (but not all!) PAM4 standards is applied after Gray decoding. That is, FEC corrects BER, not SER. A symbol error can mean that both bits have been misidentified, or just one of the pair. For the Gray coding schemes we'll be using, amplitude noise is more likely to cause one bit error than jitter, but there's no easy way to tell without decoding the symbols to bits and comparing to the transmitted logic.

The number of bit errors that Reed-Solomon FEC can correct depends on the order of the bit errors. Since we can't decipher the errored bits from the errored symbols, it's almost impossible to predict the post-FEC BER from SER.

I hope to write about the SER * FEC → BER quagmire soon. If you have results that I can report, please send me a note (ransom [at] ransomsnotes dot com).

Also see

PAM4: A new measurement science
The next generation's modulation: PAM-4, NRZ, or ENRZ?
Gray Code Fundamentals – Part 1
Gray Code Fundamentals – Part 2
Gray Code Fundamentals – Part 3
Gray Code Fundamentals – Part 3
Gray Code Fundamentals – Part 4
Gray Code Fundamentals – Part 5

VERIFYING INDUSTRIAL MESH NETWORKING FOR REAL-WORLD OPERATION

By Ross Yu, Product Marketing Manager, Dust Networks

The Industrial Internet of Things (IoT) requires industrial wireless sensor networks (WSNs) with stringent reliability and security. Since such networks must operate reliably more than ten years without intervention, industrial WSNs must cope with severely changing environmental conditions over time. In addition, they must also be scalable and flexible so that the networks can support growing business needs and data traffic over a significant period of time.

This article focuses on the methods Linear Technology uses to verify data reliability through radio hardware qualification, automated network test methodology, and systematic network testing. Performance statistics from a live production network are also reviewed. The in-house technology is designated SmartMesh, and these wireless mesh network products are engineered and rigorously tested for Industrial IoT applications, delivering better than 99.999% data reliability in some of the harshest environments. Prior to each new software production release, the company's engineers accumulate over a million node-hours of network operation with a minimum of five nines (>99.999% data reliability) before it is declared production quality.

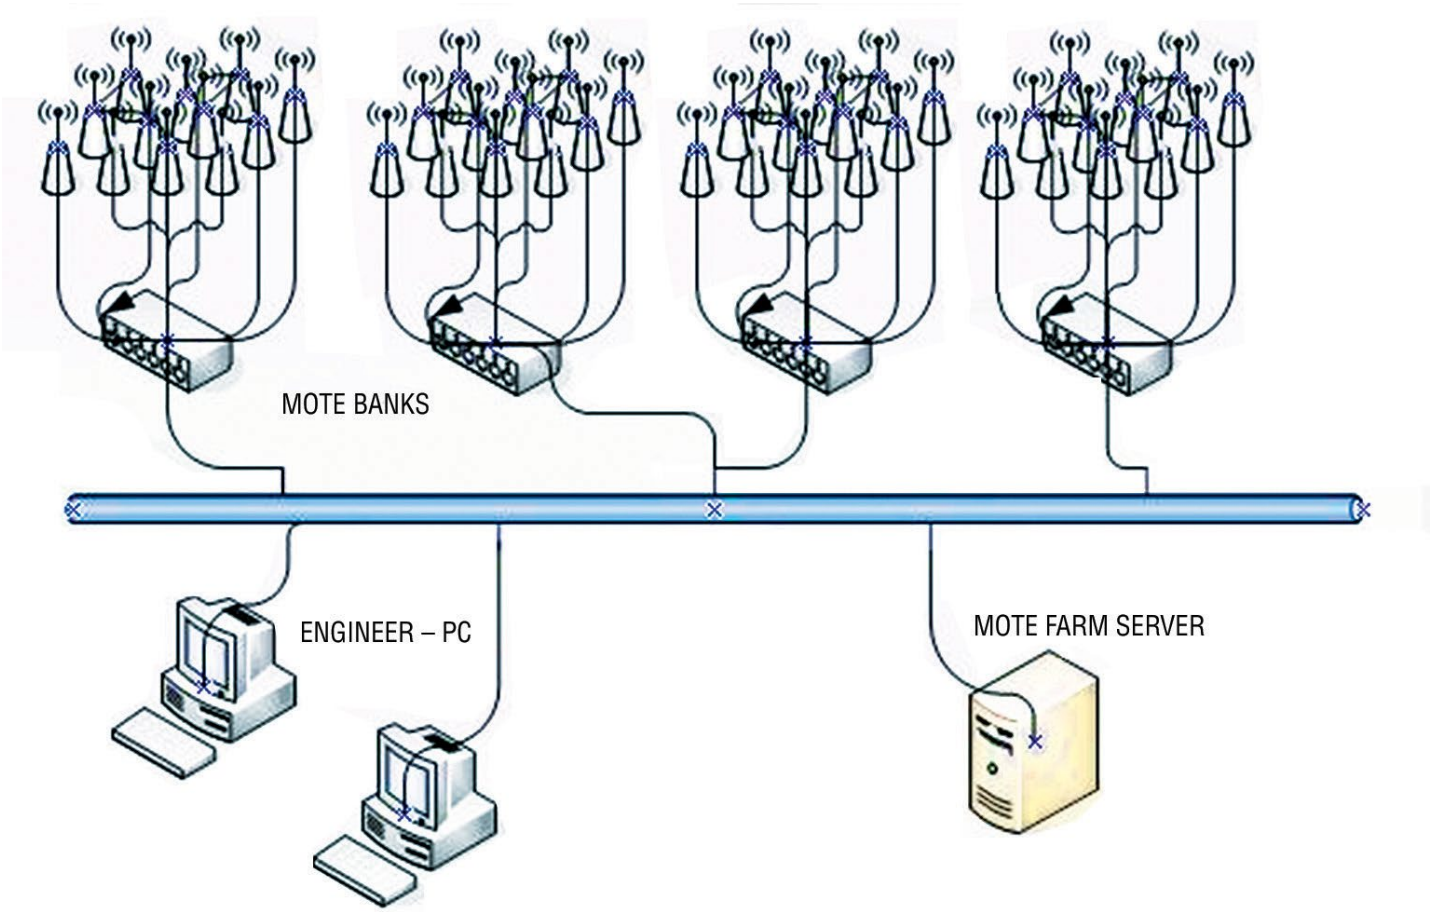

Test automation - by instrumenting hundreds of wireless nodes with an automated test fixture, a test plan of hundreds of tests can be executed in days instead of months.

Radio hardware qualification testing

The performance of a WSN is a function of both the underlying radio hardware and the protocol running on that chip. SmartMesh radios undergo rigorous testing to confirm the parameters of their operation. The results of these tests are ascertained across multiple production lots prior to publication of the data sheet, which includes all relevant specifications for the hardware.

Consistent with a detailed focus on the industrial market, the design qualification of the hardware includes operational network testing, known as highly accelerated lifetime tests (HALT), involving a live network running while subjecting the hardware to extreme conditions, including cold thermal step stress, hot thermal step stress, voltage margining, rapid thermal transitions, vibration step stress, combined thermal and vibration stress, and extended temperature tests.

Automated network test methodology

In order to assure in-service reliability, testing must comprehensively cover the situations a network will encounter throughout many years of operation. Linear Technology makes extensive use of test automation to facilitate hundreds of network tests, each verifying a unique set of test conditions. To do so, a network test bed (see Figure) consisting of banks of hundreds of wireless nodes can be readily configured into any number of test networks, large or small.

A centralized test server can quickly commission entire co-located networks, run multiple system tests and recommission nodes for the next set of tests by programming via each wireless node's application programming interface (API). Full regression tests become pragmatic with automation, ensuring that existing functions and behaviours are preserved in subsequent software releases.

The test bed has a dense, noisy RF environment because each network under test is immersed in a sea of wireless traffic from the other networks operating simultaneously. This network traffic, along with nearby Wi-Fi routers, Bluetooth and cellular radios, create an elevated RF noise floor representative of an extremely challenging RF environment. *Click for pdf version, with an expanded description of the test methodology.*

Download PDF of Article

Find industrial mesh networking on EETsearch

HDL CODING STYLES

By Adam Taylor, E2V

The major differences between coding styles relate to how the design engineer decides to handle VHDL keywords and any user-defined items (the names of signals, variables, functions, procedures, etc.). Although there are many different possibilities, in practice there are only three commonly used approaches. The first technique is to use a standard text editor and simply enter everything in lowercase (and black-and-white) as shown in Figure 1, the example of a simple synchronizer.

The second method is to follow the VHDL Language Reference Manual (LRM), which says that keywords (if, case, when, select, etc.) should be presented in lowercase. Strangely, the LRM doesn't have anything to say about how user-defined items should be presented, but the common practice (when the keywords are in lowercase) is to present user-defined items in uppercase as shown in Figure 2.

The third approach - and my personal preference - is to turn the LRM on its head; to capitalize keywords (IF, CASE, WHEN, SELECT, etc.) and to present any user-defined items in lowercase as shown in Figure 3.

The main argument for using uppercase for keywords and lowercase for user-defined

```
2   library ieee;
3   use ieee.std_logic_1164.all;
4   use ieee.numeric_std.all;
5
6   entity code_style_example is port(
7     clock : in std_logic;
8     ip    : in std_logic;
9     synch : out std_logic);
10  end entity;
11
12  architecture rtl of code_style_example is
13
14  signal synch_reg : std_logic_vector(1 downto 0) := (others => '0');
15
16  begin
17
18  synch_proc : process(clock)
19  begin
20    if rising_edge(clock) then
21      synch_reg <= synch_reg(synch_reg'low) & ip;
22    end if;
23  end process;
24
25  synch <= synch_reg(synch_reg'high);
26
27  end architecture;
```

Figure 1 *B&W: Keywords and user-defined items in lowercase.*

items, or vice versa, is that this helps design engineers and reviewers to quickly locate and identify the various syntactical elements in the design. However, this line of reasoning has been somewhat negated by the use of today's context-sensitive editors, which are language-aware and therefore able to automatically assign different colours to different items in the design.

This article next looks at our original examples through the "eyes" of a context-sensitive editor as illustrated in three images, Figures 4 – 6, showing use of colour in the code (Figures 4-6 appear in the longer version of this article,

```
2   library IEEE;
3   use IEEE.STD_LOGIC_1164.all;
4   use IEEE.NUMERIC_STD.all;
5
6   entity CODE_STYLE_EXAMPLE is port(
7     CLOCK : in STD_LOGIC;
8     IP    : in STD_LOGIC;
9     SYNCH : out STD_LOGIC);
10  end entity;
11
12  architecture RTL of CODE_STYLE_EXAMPLE is
13
14  signal SYNCH_REG : STD_LOGIC_VECTOR(1 downto 0) := (others => '0');
15
16  begin
17
18  SYNCH_PROC : process(CLOCK)
19  begin
20    if RISING_EDGE(CLOCK) then
21      SYNCH_REG <= SYNCH_REG(SYNCH_REG'LOW) & IP;
22    end if;
23  end process;
24
25  SYNCH <= SYNCH_REG(SYNCH_REG'HIGH);
26
27  end architecture;
```

Figure 2 *B&W: Keywords lowercase; user-defined items uppercase.*

click below for pdf)..

Although context-sensitive editors are extremely efficacious, it's important to remember that some users may be colour-blind. Also, even if the code is captured using a context-sensitive editor, it may be that some members of the team end up viewing it using a non-context-aware (black-and-white) editor. Furthermore, the code may well be printed out using a black-and-white printer. For all these reasons, my personal preference is to capitalize keywords and for everything else to be in lowercase.

```
2   LIBRARY ieee;
3   USE ieee.std_logic_1164.ALL;
4   USE ieee.numeric_std.ALL;
5
6   ENTITY code_style_example IS PORT(
7     clock : IN std_logic;
8     ip    : IN std_logic;
9     synch : OUT std_logic);
10  END ENTITY;
11
12  ARCHITECTURE rtl of code_style_example IS
13
14  SIGNAL synch_reg : std_logic_vector(1 DOWNTO 0) := (OTHERS => '0');
15
16  BEGIN
17
18  synch_proc : PROCESS(clock)
19  BEGIN
20    IF rising_edge(clock) THEN
21      synch_reg <= synch_reg(synch_reg'low) & ip;
22    END IF;
23  END PROCESS;
24
25  synch <= synch_reg(synch_reg'high);
26
27  END ARCHITECTURE;
```

Figure 3 *B&W: Keywords uppercase; user-defined items lowercase.*

Aside from how we handle the keywords, most companies will have their own sets of coding guidelines, which will also address other aspects of coding, such as:

- Naming conventions for clock signals; possibly requiring each clock name to include the frequency, e.g., clk_40MHz, clk_100MHz.
- Naming and handling of reset signals to the device (active-high, active-low), along with the synchronization of the signal and its assertion and de-assertion within the FPGA. The de-assertion of a reset is key, as removing this signal at the wrong time (too close to an active clock edge) could lead to metastability issues in flip-flops.
- Naming conventions for signals that are active-low (this is common with external enables). Such signals often have a "_z" or "_n" attached to the end, thereby indicating their status as active-low ("ram_cs_n" or "ram_cs_z", for example).
- The permissible and/or preferred libraries that will ensure the standard types that can be used. Commonly used libraries for any implementable code are "std_logic_1164" (this is the base library that everyone uses) and the "numeric_std" for mathematical operations (as opposed to "std_logic_arith"). Meanwhile, other libraries such as "textio" and the "math_real" and "math_complex" libraries will be used when creating test benches.
- The use of "std_logic" or the unresolved "std_ulogic" types on entities and signals ("std_ulogic" is the unresolved type of "std_logic" and is used to indicate concurrent assignments on a signal).
- Permissible port styles. Many companies prefer to use only "in," "out," and "inout" (for buses only) as opposed to "buffer," as the need to read back an output can be handled internally with a signal. Also, many companies may limit the types permissible in entities to just "std_logic" or "std_logic_vector" to ease interfacing between large design teams.
- It is also common for companies in some specialist safety-critical designs to have rules regarding the use of variables.

These coding styles can be very detailed, depending upon the end application the FPGA is being developed for. What coding style and standards do you use?

designideas

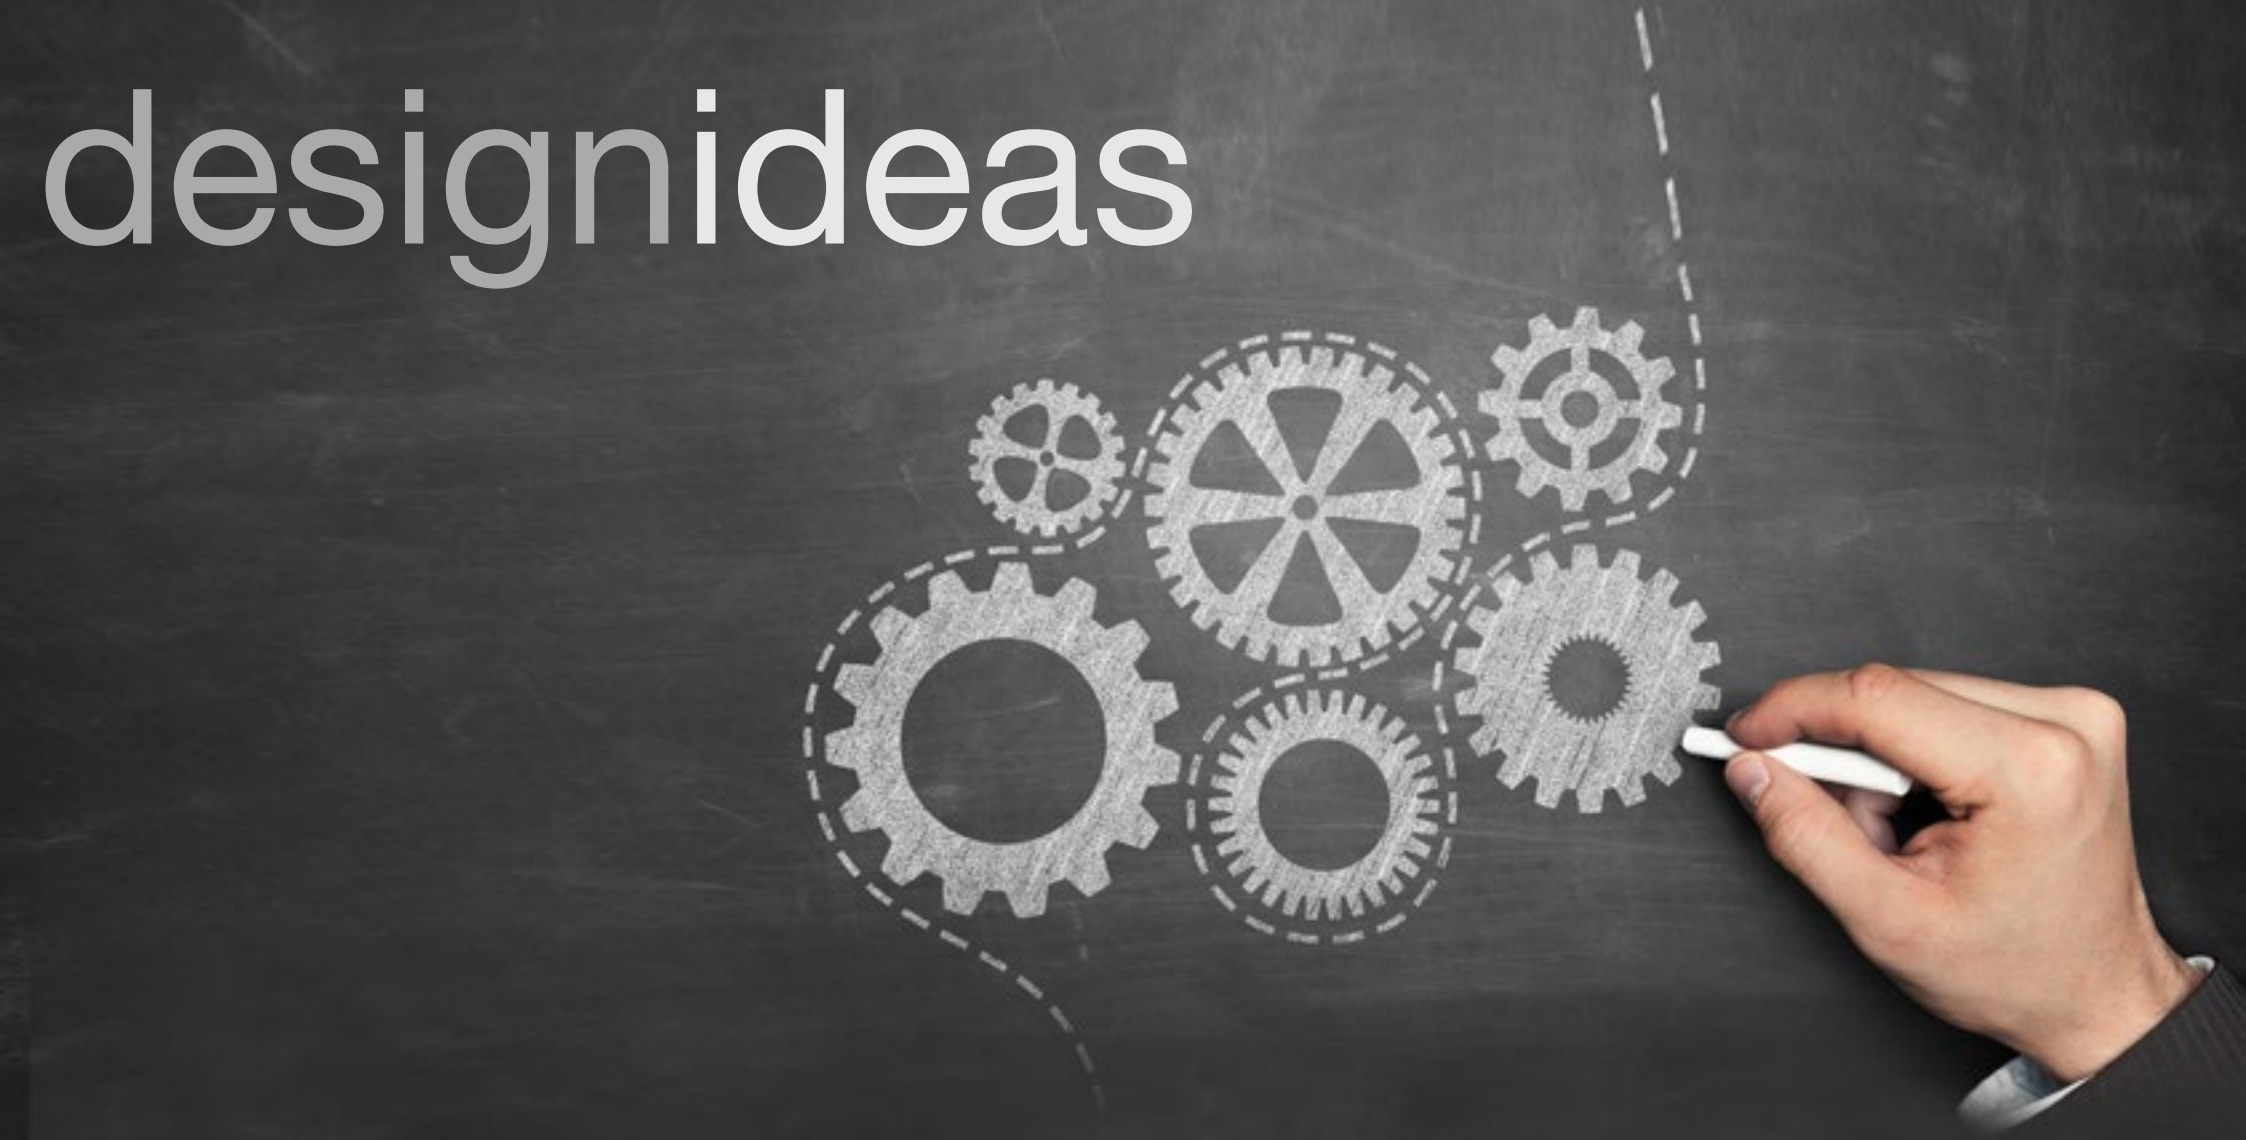

- UVP protects rechargeable batteries
- Control an FPGA bus without using the processor

UVP protects rechargeable batteries

By Peter Demchenko

Many rechargeable battery types can be damaged by deep discharge. This Design Idea circuit performs under-voltage protection (UVP) to prevent this, as well as acting as a load switch. With little or no modification, it can work for almost any battery type, with voltage from 4.5V - 19V. Standby current is under 1 µA.

High-side N-channel MOSFET Q2 reduces cost compared to a P-channel part. Turn-on and -off are inherently soft, so switching spikes are avoided.

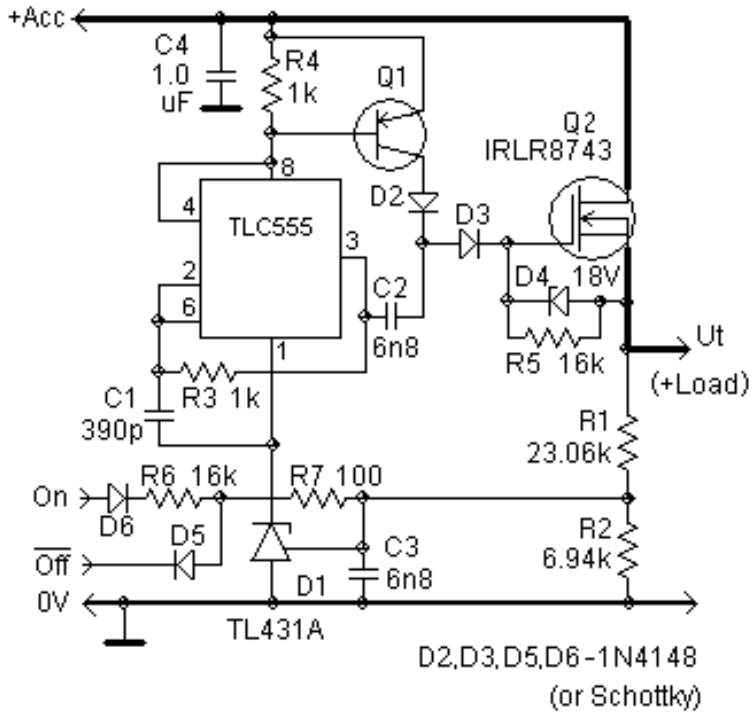

Figure 1. *The circuit is set up for a 12V lead-acid accumulator with gelled electrolyte at 20°C. Another type of accumulator may require component value changes.*

Operation

The first time the battery/accumulator is connected to **+Acc,** the circuit is off. C3 is discharged, so the programmable reference TL431A (D1) will be off, with leakage under 1 µA. Thus, all other elements of the circuit are disabled, and Q2 is off because its gate is discharged through R5.

In this state the circuit is waiting for an enabling positive pulse at the **On** input – via a button or other controller. During this pulse, the TL431A turns on, providing power to the TLC555, configured as an oscillator. Through C2 & D3, the oscillator produces a bootstrap voltage on the gate capacitance of Q2, turning it on.

After the **On** signal is removed, the circuit remains in the active state because the voltage from divider R1-R2 charges C3, maintaining the loop from D1 to Q2 back to D1.

The circuit disconnects the load and turns itself off if the **/Off** input goes low or the under-voltage trip point is reached.

The trip point expression is:

$V_t = (1 + R1/R2)\ V_{ref}$

(V_{ref} is the TL431 2.5V reference voltage)

So $R1/R2 = V_t / V_{ref} - 1$

To make the influence of I_{ref} (4µA max.) negligible, make the current through the divider at least 100 times I_{ref} :

$R1 + R2 \leq 30\ k\Omega$

So for a trip point of 10.8V, the calculated values are:

$R2 = 30\ k\Omega / (V_t / V_{ref}) = 6.94\ k\Omega$

$R1 = 30\ k\Omega - R2 = 23.06\ k\Omega$

The TL431A has a V_{ref} tolerance of 1%, so tolerances of R1 and R2 should be better to minimally degrade the trip point accuracy, or a trimpot could be added.

As a matter of good practice the current I_{ref} should be less than half of its absolute maximum rating of 10 mA. Hence:

$R6 \geq VOn / 5mA$

The upper limit of R6 is defined by:

$V_{ref} = (R1 \parallel R2)\ V_t / (R6 + (R1 \parallel R2))$

$R6 \leq (R1 \parallel R2)\ (V_t / V_{ref} - 1)$

So, $R6 \leq 17.8k\Omega$

If you select R6 near this limit, you can prohibit any attempts to switch on the load while the accumulator is depleted. The /Off input has priority over On.

Battery voltage monitoring takes place after switch Q2, so a low $R_{DS(on)}$ is essential for correct operation.

With a high load current, the switching times should be minimized to reduce power dissipation. For a quick charge of Q2's gate capacitance C_g during turn-on, oscillator frequency should be high (it's about 900 kHz here). For quick turn-off, R5 should not be too high: the time depends on $R5 \times C_g$.

Q1 prevents connection of Q2's gate to **+Acc** through D2 & D3 when the circuit is disabled. Any PNP transistor with a moderate gain (30-150) can be used (e.g., 2N2904). Transistors with higher gain (e.g., BC556-BC560) may require R4 reduction to ensure Q1 is off during the circuit's off state. Q1 though should be reliably turned on by the current through the 555; working near its upper frequency limit and using a low R3 value ensures enough R4 drop.

Low power Zener diode D4 should be spec'd for $V_{GS(max)}$ of Q2.

The choice of Q2's main parameters is system-dependent. Because C_g is used as a filter capacitor, attention should be paid to it too. Adequate values lie in the 2 nF – 10 nF range; greater capacitance may require C2's value be increased. As a rule of thumb, C2 can range from C_g to $2C_g$. The value of C2 has an influence on the turn-on time.

Dedicated to V.R., without whom this idea would not have emerged at all.

Control an FPGA bus without using the processor

By Noe Quintero

Many FPGA designs use an embedded processor for control. A typical solution involves the use of a soft processor such as a Nios, though FPGA SoCs with a built-in hard processor have become popular too. Figure 1 shows a typical Altera FPGA system that contains the processor and a mix of peripherals that are connected via Altera's Avalon Memory Mapped (MM) bus. These processors greatly simplify the end application, but require a strong programming background and knowledge of complicated toolchains. This can hinder debug, especially if a hardware engineer needs a simple way to read and write to the peripherals without pestering the software engineer.

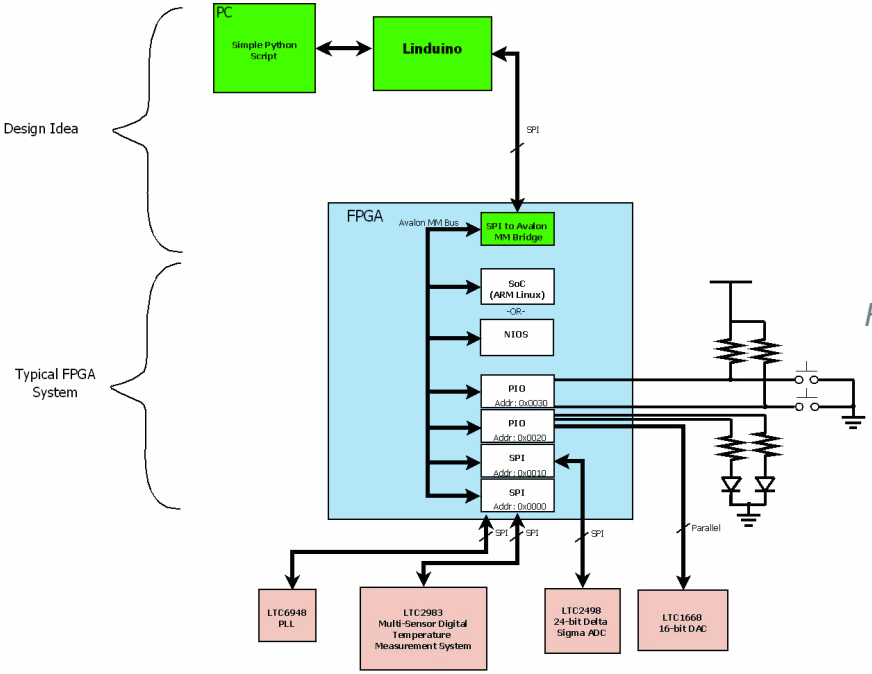

Figure 1. *Typical Altera FPGA system connected using the Avalon memory-mapped bus*

This Design Idea uses Altera's *SPI Slave to Avalon MM Bridge* to provide a simple way to hop onto the Avalon bus. There are two advantages to this technique: It does not compromise the original system design, and the bridge can co-exist with the embedded processor. For the system shown in Figure 1, the SPI bridge allows the engineer to directly control the frequency of the LTC6948 fractional-N PLL, set the LTC1668 DAC voltage, read a voltage from the LTC2498 ADC, or read temperatures from the LTC2983, just like the processor can.

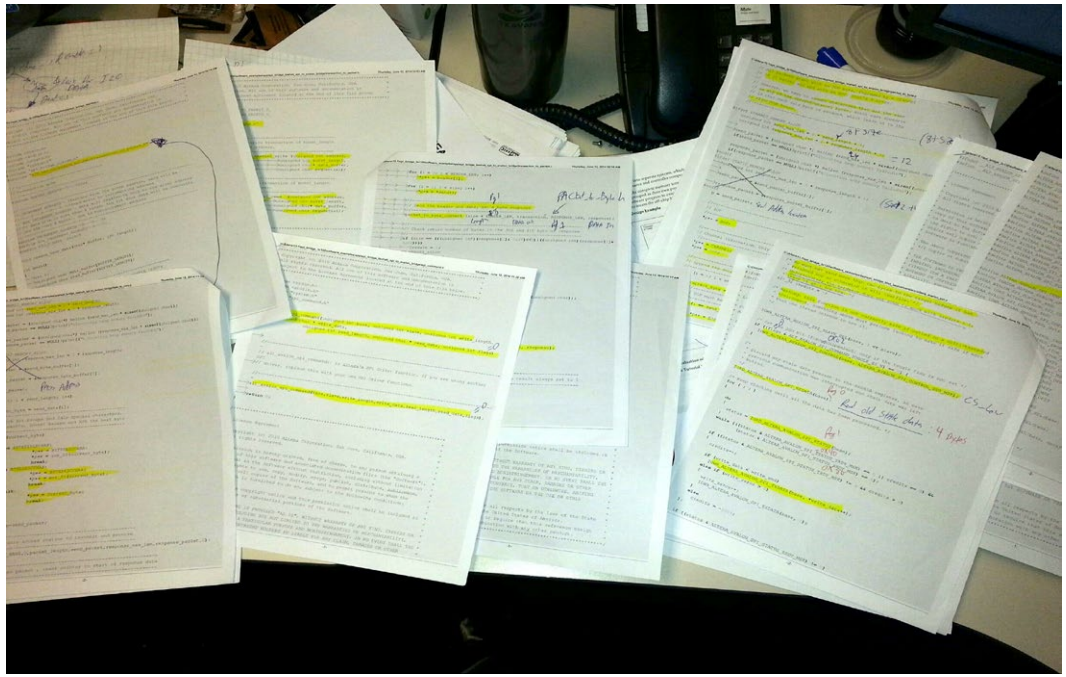

Figure 2. *Highlighter + Example code + Reverse engineering = Python script - click here for script image*

Altera provides a reference design for the SPI-Avalon MM bridge. Unfortunately, the documentation is sparse at best, and uses a Nios processor as the SPI master. This effectively defeats the purpose of the SPI bridge, as the Nios can interface directly to the Avalon MM bus. A practical SPI master is Linear Technology's *Linduino* microcontroller, which is an Arduino clone with extra features to interface to LT demo boards. One extra feature is a level-shifted SPI port. This level-shifting function is especially helpful when interfacing to FPGA I/O banks with voltages as low as 1.2V. The Linduino firmware can be used to accept commands through a virtual COM port and translate the commands to SPI transactions.

After reverse engineering the Altera example design (Figure 2), a Python library (click here for image) was developed to create packets that the bridge would accept. These packets are then translated into Linduino commands. A Python script then allows the hardware engineer to have complete control of the project without needing to reinvent the interfacing protocols. An example Python script to control the frequency of a digital pattern generator for an LTC1668 DAC is provided in the LinearLabTools Python folder. Figure 3 shows the demo setup.

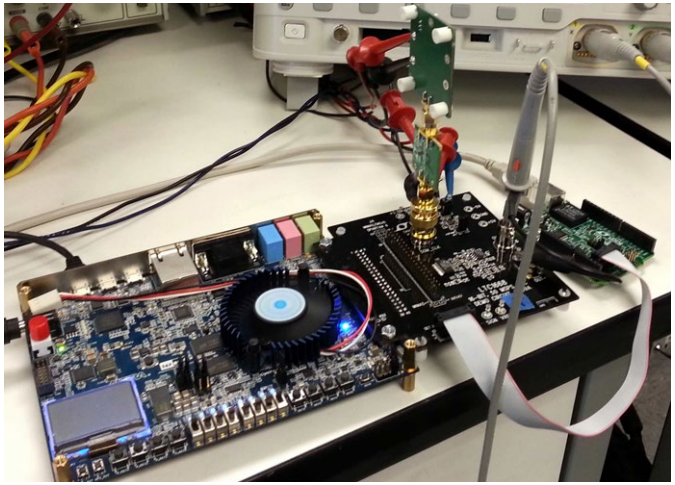

Figure 3. *DC2459 DAC demo board (R) plugged into an FPGA board (L)*

Figure 4 shows the system block diagram. Note that the numerically controlled oscillator (NCO) can be controlled by the shift register or the PIO core. The shift register is included for debug, as it allows direct control of the NCO. Setting the GPIO line high enables the SPI-Avalon bridge, which in turn controls a 32-bit PIO port over the Avalon bus. The PIO output then controls the NCO frequency.

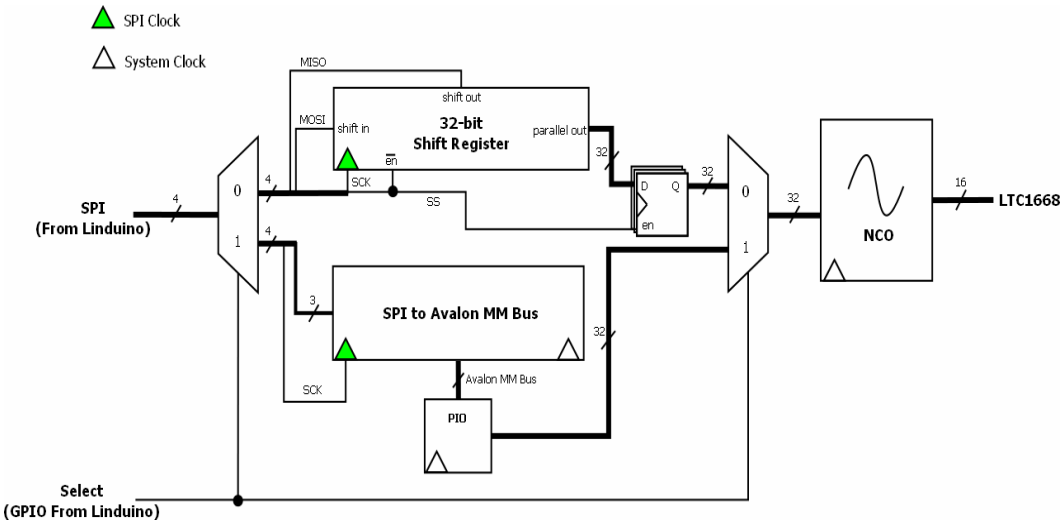

Figure 4. *FPGA system block diagram*

With the basic system operational, additional peripheral cores can be connected to the bus. To design the system, Altera provides a tool called Qsys, which provides a GUI to connect the IPs to one another. Qsys translates the GUI-designed system (Figure 5) to HDL. Peripheral addresses are fully configurable. In this case, the PIO is set to a base of 0x0.

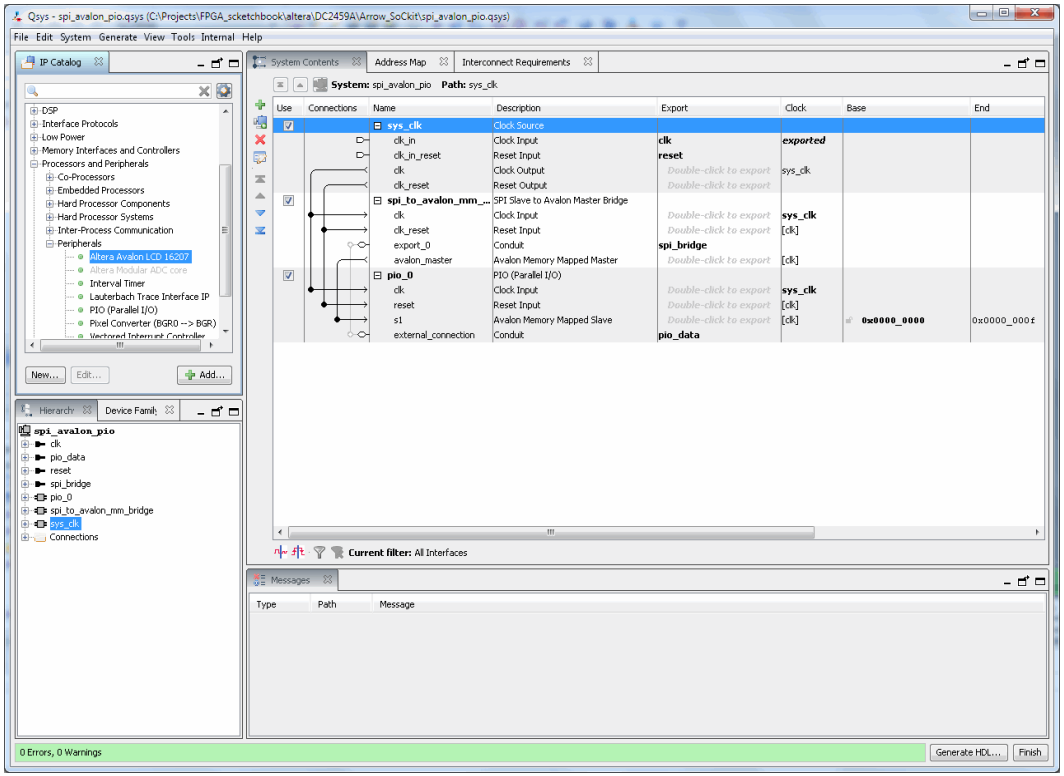

Figure 5. *Qsys GUI*

Once the design is implemented in the FPGA, the provided Python library in LinearLabTools contains two functions to interface to the design:

transaction_write(dc2026, base, write_size, data)
transaction_read(dc2026, base, read_size)

The first argument to these functions is the Linduino serial port instance. The second argument is the peripheral's address on the Avalon bus. The functions accept and return lists of bytes respectively. These two functions are written to allow flexibility when writing and reading to IP. To set the NCO for the provided example, the **transaction_write** function is all that is needed. Equation 1 is used to determine the tuning word.

$$tuning\ word = \frac{desired\ frequency}{system\ clock\ frequency} \times 2^{32} \quad (1)$$

To set the NCO to 1 kHz with a 50 MSPS sample rate, the tuning value is 85899, or 0x00014F8B, which is passed as a list of four bytes. Thus, the python code to set the DAC to 1 kHz is:

transaction_write(linduino_serial_instance, 0, 0, [0x0, 0x01, 0x4F, 0x8B])

The Python script in Figure 6 illustrates the simple text interface that configures the NCO. An important note: the bridge uses SPI mode 3. This was painfully determined to be the correct mode by trial and error, and verified by analyzing the Nios processor's SPI interface in Altera's example.
This Design Idea provides the ability to control a system without touching the embedded processor, allowing the hardware engineer to progress on a project without bothering the software engineer, and with minimal impact to the hardware design.

Python code is available at:
http://www.linear.com/solutions/linearlabtools.

HDL for the LTC1668 demo board is included in the DC2459 design files, available at http://www.linear.com/demo/DC2459

```
    The views and conclusions contained in the software and documentation are
    those of the authors and should not be interpreted as representing official
    policies, either expressed or implied, of Linear Technology Corp.

    Description:
        The purpose of this module is to use the DC2026C as a Avalon MM Bus
        interface to set the DC2459A frequency out.
"""

################################################################################
# Libraries
################################################################################

import sys
sys.path.append('../../utils')
import connect_to_linduino as duino
import ltc_spi_avalon as avalon


if __name__ == "__main__":
    linduino = duino.Linduino() # Find the Linduino
    linduino.port.write("M3")   # Set to SPI Mode 3
    linduino.port.write('G')    # Set the GPIO HIGH
    try:
        print "Command Summary"
        print "    1-Send raw code"
        print "    2-Set frequency"
        print "    3-Exit program"
        user_input = input("Enter a command: ")
        while(user_input != 3):
            if(user_input == 1):
               code = int(input( "Enter raw 32 bit code: "))

               # Send the data to the Avalon MM bus addr 0
               avalon.transaction_write(linduino, 0, 4,
                                [code & 0xFF ,(code>>8) & 0xFF ,
                                (code>>16) & 0xFF, (code>>24) & 0xFF])
            elif (user_input == 2):
                freq = float(input("Enter desired frequency(Hz): "))
                float_code = freq/50000000*(2**32-1)
                code = int(float_code)

                # Send the data to the Avalon MM bus addr 0
                avalon.transaction_write(linduino, 0, 4,
                                [code & 0xFF ,(code>>8) & 0xFF ,
                                (code>>16) & 0xFF, (code>>24) & 0xFF])
            else:
                print "**** Invalid Command ****"
            print "Command Summary"
            print "    1-Send raw code"
            print "    2-Set frequency"
            print "    3-Exit program"
            user_input = input("Enter a command: ")
    finally:
        linduino.close() # Close the port
```

```
Looking for COM ports ...
Available ports: [(23, 'COM24')]

Looking for Linduino ...
    Found Linduino!!!!
Command Summary
    1-Send raw code
    2-Set frequency
    3-Exit program
Enter a command:
```

Figure 6. *Python Avalon bus example*

About the author

Noe Quintero is an Associate Application Engineer in the Mixed Signal group at Linear Technology. He joined the company in 2012 as an intern, and started full time in 2015. Noe has a BS in Electrical Engineering from San Jose State University. He has an identical twin, who is also in the EE field. Noe's hobbies includes, cycling, building robots, and saltwater reef aquariums.

productroundup

productroundup

Digital isolators provide robust galvanic isolation

These two-channel, bi-directional digital isolators, MAX14933 and MAX14937, join Maxim's portfolio of quad-channel digital isolators, which provide speed, propagation delay, skew, and jitter for timing sensitive applications while maintaining an extremely robust isolation barrier between two power domains. They provide high voltage isolation of 2.75 kVrms or 5 kVrms, respectively. They support data rates from DC up to 1.7 MHz and can be used in isolated I²C busses with or without clock stretching. Both products feature independent 2.25V to 5.5V supplies on each side of the isolator.

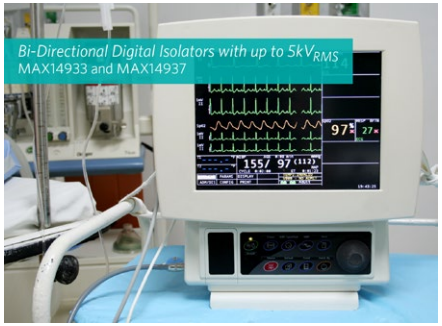

Complete article, here

40-A, 16-VIN DC/DC with differential remote sensing

Texas Instruments claims the first 40-A, 16-VIN synchronous step-down DC/DC converter with true differential remote-voltage sensing; its SWIFT TPS548D22 buck converter features a small PowerStack package and integrated MOSFETs to power ASICs and processors in space-constrained applications.The TPS548D22 DC/DC converter employs the D-CAP3 control-mode topology, features 0.5% reference-voltage accuracy over temperature, and includes true differential remote-voltage sensing to meet voltage requirements of deep submicron processors.

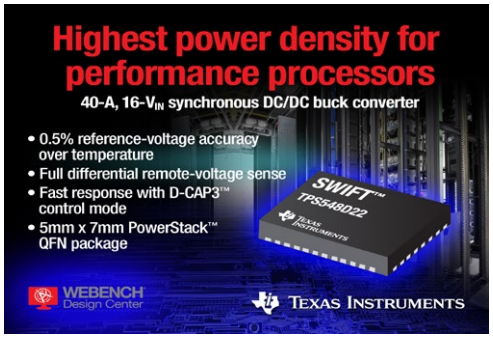

Complete article, here

SimpleLink Wi-Fi modules for IoT connectivity

Murata has added two 2.4 GHz IEEE802.11 b/g/n compliant wireless modules based around the TI CC3200 and TI CC3100 SimpleLink chipsets. These miniature modules, measuring 13.2 x 21.45 x 2.65 mm ease incorporating wireless internet connectivity to embedded applications. Both modules, Type 1JP and Type1JQ, have SPI and UART interfaces, 16 Mbit of Flash memory and a 40 MHz clock. The module is constructed on a PCB with all matching components and a PCB antenna and the wireless transceiver and associated logic enclosed in a metal case. Peripheral interfaces include SPI, UART and numerous GPIO pins.

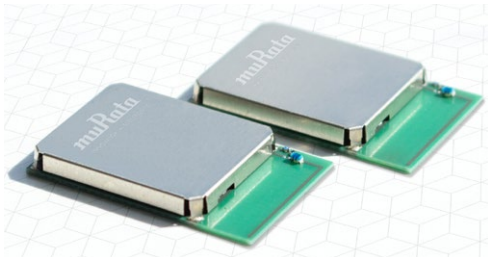

Complete article, here

Dev kit targets industrial IoT with Zynq FPGA

Avnet is to offer an IoT development kit that provides a simple path to production; the MicroZed Industrial IoT starter kit builds on Avnet's existing MicroZed development module, which is based on the Xilinx Zynq-7000 SoC. The SoC combines dual ARM Cortex-A9 processor cores with numerous programmable logic cells to provide both hardware and software programmability. Avnet has created a carrier board for the MicroZed to plug into, with an R3 Arduino-compatible expansion slot, two 2x6 peripheral module expansion slots, and a header for connection to the built-in UART, SPI, I²C, and GPIO interfaces.

Complete article, here

productroundup

Graphite heat transfer material maximises thermal conductivity

Panasonic's high-thermal-conductivity thermal interface materials, LDPGS – Low Density Pyrolytic Graphite Sheet – is a thin (available in several thicknesses) graphite sheet that you would use as a heat-conducting medium to remove heat from critical components. This material has a heat conductivity that Panasonic says is five times better than copper, and "approaches" that of diamond. The specified value for its conductivity is up to 1950 W/m K, and the principal manufactured thickness of the 2D carbon matrix is 200 µm (0.2 mm). The material is also compressible.

Complete article, here

400 mW radio modem for long range telemetry

The Radiometrix FPL3 radio modem simplifies the design of long range wireless appliances that can be marketed in any European country. This highly integrated multi-channel radio with 400mW RF output power operates in the 869.4-869.65 MHz frequency band, which is the only band permitted for license-free operation at up to 500 mW in all EU countries. The 67 x 30 x 12 mm module operates from a single 5V supply, and draws 50 mA in receive mode and 500 mA when transmitting at 400 mW (+26 dBm) RF output power. The FPL3 conforms to ETSI EN 300 220-2 (radio) and EN 301 489-3 (EMC) standards.

Complete article, here

150V synchronous step-down DC/DC

Linear Technology designed the LTC3895 to be a non-isolated DC/DC controller with a high input voltage capability that can eliminate the need for surge suppression and operate continuously with a high input voltage. The synchronous step-down switching regulator controller drives an all N-channel MOSFET power stage. Its 4V to 140V (150V abs. Max.) input voltage range is designed to operate from a high input voltage source or from an input that has high voltage surges, eliminating the need for external surge suppression devices. The LTC3895 continues to operate at up to 100% duty cycle during input voltage dips down to 4V.

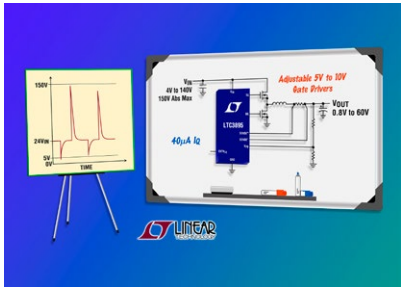

Complete article, here

Arduino Primo board hosts Nordic Bluetooth LE

The Arduino Primo integrates Bluetooth low energy and NFC Touch-to-Pair (using the Nordic Semiconductor nRF52832 SoC), Wi-Fi, and infrared connectivity, eliminating the cost and complexity of shields. IoT developers creating innovative applications now have built-in wireless connectivity all on one board. Disclosing its design win on the board, Nordic notes that to add Bluetooth low energy wireless connectivity and Near Field Communication (NFC), users previously required to add shields to the base board. The Arduino Primo can also act as a fully-functional TCP / IP Internet client and server over Wi-Fi.

Complete article, here

productroundup

1200V, 325A silicon carbide module is low-inductance

To realise the potential of its silicon carbide power devices, Wolfspeed (Cree) has designed a completely new format of half-bridge module that has an inductance of 5.5 nH; to further minimise current paths and simplify the use of the package, it comes with a physically-matching new design of driver board. This all-SiC module allows systems designers to realize lighter weight systems that are up to 67% smaller by achieving efficiencies of over 98% and improvements in power density of up to 10 times compared to systems built with silicon-based technologies.

Complete article, here

Silicon carbide Schottky diodes with lowest Vf

Rohm Semiconductor's third-generation silicon carbide Schottky Barrier Diodes(SBDs) promise optimum stability and power efficiency for high power applications. The diodes, packaged in TO-220AC outline, are rated at 650V/6, 8, 10A; they realize the lowest fF forward voltage, of 1.35V at 25C and 1.44V at 150C, for a forward current of 10A: and lowest IR (reverse leakage current) of 0.03μA at 650V (25C). Rohm says that previously, it has been difficult to combine low forward voltage with high surge current ratings; but in these devices the Ifsm rating is extended to 82A (50Hz, single pulse).

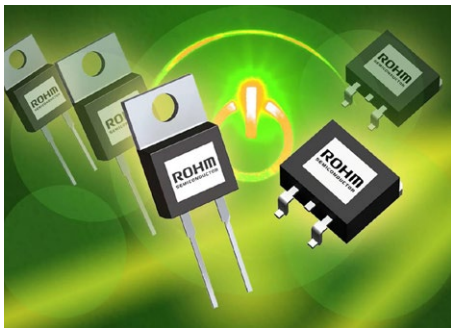

Complete article, here

Major expansion in Li-ion battery factory capacity in Europe

Labelling it as "Europe's Gigafactory" (an allusion to the Tesla battery plant being built in the western USA), BMZ has opened the first section of what will be Europe's biggest lithium-ion battery factory; BMZ is going to more than quadruple production areas in Germany by 2020. With construction taking only a little over a year, BMZ GmbH (Karlstein-Großwelzheim, Germany) European developer and producer of intelligent lithium-ion batteries, has now opened up the first two production, logistics and office buildings in addition to an existing production areas of 7,000 m² at the company's present headquarters. In the manufacturing units comprising 4,800 m² each, up to 200 million lithium-ion batteries of the most varied kind and size with an overall storage capacity of about 15 GWh can be developed, produced and tested every year.

Complete article, here

Infineon puts 650V FETs in 1nH TO packaging

Infineon's CoolMOS C7 Gold 650V MOSFETs come in a TO-Leadless package. This combination of improved superjunction (SJ) semiconductor process and advanced SMD package design is aimed at hard switching applications. The small footprint offers power density advantages for server, telecom and solar applications. C7 Gold CoolMOS technology comes with 4-pin Kelvin Source capability and improved thermal properties of the TO-Leadless package. This enables a viable SMD solution for high current topologies such as Power Factor Correction (PFC) up to 3 kW.

Complete article, here

productroundup

Power Conversion Platform SoC for renewable energy

Analog Devices has introduced a single-chip, integrated power conversion controller that is designed for next-generation solar power, energy storage, and electric vehicle infrastructure applications. The platform, which includes processing, gate driving, and sensing components, was designed to enable new faster-switching architectures and accommodate increasing safety regulations. Achieving certification by TUV, with a single-die design, the IC has two cores, implements extensive redundancy, and also has multiple specific-function-accelerators.

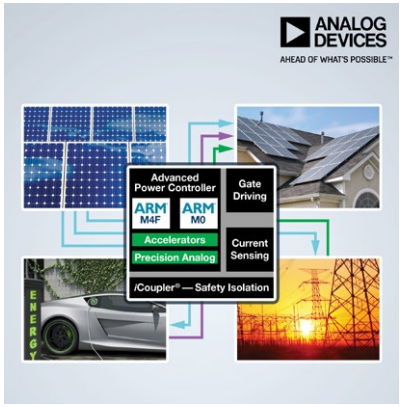

Complete article, here

Isolated driver for IGBTs & MOSFETs, to 1200V & 8A

In 2012 Power Integrations acquired Swiss company CT-Concept Technologie (Concept), a specialist in driving high-power semiconductor switches, especially IGBTs. Separately, recent introductions in its lower-power off-line converter ICs have featured FluxLink, an isolation/feedback mechanism that is fundamentally a single-turn transformer contained within a single IC package. Combining the two yields a series of driver ICs; FluxLink is extended to support two-way communication across the isolation barrier (drive waveforms in one direction/fault diagnostics in the reverse direction).

Complete article, here

1200V IGBTs with Ultra Field Stop technology

ON Semiconductor has introduced a new series of IGBTs which use its Ultra Field Stop trench technology; the 1200V devices have minimal total switching loss (Ets) characteristics; the remarkable improvement in performance is attributable in part to a very wide highly activated field-stop layer and optimised co-pack diode. The NGTB40N120FL3WG has an Ets of 2.7 millijoules (mJ), while the NGTB-25N120FL3WG has an Ets of 1.7 mJ. Both devices have a VCEsat of 1.7V at their respective rated currents. The NGTB40N120L3WG is optimized for low conduction losses and has a VCEsat of 1.55V, at rated current, with an Ets of 3 mJ.

Complete article, here

Interactive programming, using Forth, for C language

MPE (Southampton, UK) has configured Interactive C support for ARM architecture via its SockPuppet software. Interactive debugging and adding test harnesses from the start of an embedded coding exercise ensures, MPE asserts, that the code stays close to the initial specification. The introduction adds Interactive Mixed-Language Programming to the tool box of C designers. The Forth Language was designed for control, interactive programming and debugging from the start. C code is worked with as in the past. Test and Debugging is dealt with in Forth where needed, independent of the C code.

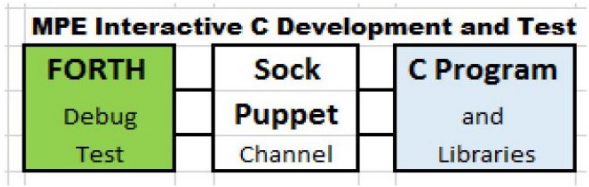

Complete article, here

MEGAVOLT/KILOAMP TESTS REVEAL EXTREME ENGINEERING CHALLENGES

BY BILL SCHWEBER

With all the attention given to low-power design and products, it's sometimes easy to forget that there's another world out there with power levels that are tens of orders of magnitude greater—and working with them is a radically different world in every respect. A recent article in IEEE Spectrum, "Inside the Lab That Pushes Supergrid Circuit Breakers to the Limit," was a dramatic reminder of the engineering challenges of anything having to do with the electric grid's high-tension power lines in the megavolt/kiloamp regime. Next time you hear engineers moan that their design has to "sip electrons," perhaps you should suggest that they read this article and they'll hopefully soon stop.

Forget everything you think you know about "electricity" when you are at these levels. The article focuses on testing of circuit breakers for these power lines, and it's truly another world. There's no need for me to reiterate the article; you can read it yourself. The author's detailed description of what a circuit breaker must do and deal with at these levels is astonishing: we're in the world of plasmas, gas quenching, switching of kA in milliseconds, mechanical contact issues that you won't be able to anticipate, and more. Every aspect of the test setup had to be custom built, as there are no off-the-shelf fixtures for this kind of work. The test facility even had to build a special generating and storage system to supply the power for the tests. Every decision and action requires careful, deliberate thinking and risk assessment.

I'm fascinated by engineers who are not only at the extremes of design along one or more performance or operational parameters, but must also devise and build ways to test their designs. Sometimes, as in the case of the power-grid circuit breakers, there is a set of operational features in their favor: the tests are reasonably close to final conditions, they can be repeated as needed under carefully controlled conditions, and changes can be made and then the devices are retested.

The limits of testing

However, not all tests involving systems at high power levels (whether electrical, chemical, or mechanical) have this characteristic. Often, the test process is so complex and difficult to set up or execute that any the test/modify/retest cycle is too expensive or time consuming. The implications of this point were clearly explained in the excellent book "Apollo: The Race to the Moon," where an "interlude chapter" steps back and provides a big-picture overview of the differences between aircraft test and big-rocket test. They paraphrase Joe Shea, Apollo Program Manager, as saying it didn't matter if they tested the Saturn V rocket six times instead of four, or eight instead of six … statistically, the extra successes (if they were successes) would be meaningless – all they would do is use up pieces of precious, costly hardware (time and dollars) that could have been used for real missions.

There are other cases when a project can't be fully tested. The recently published book "The Right Kind of Crazy" about the Mars Curiosity Rover mission discussed the implications of the obvious: that many aspects of the design of space-exploration systems must be simulated, assessed, and analyzed to an extreme, because there is no way to replicate some of the real operating conditions. For the Mars mission, one such topic was the parachute which slowed the Lander down so the "sky crane" could hover and lower the Rover to the Martian surface. You can replicate the extreme cold and vacuum of space, but how do you test deploying that parachute at hundreds of miles/hour in the Martian atmosphere and then using retrorockets to stabilize the platform in a low-g environment? The answer is that you can't.

What's your experience with higher voltages and currents? What's the highest power level for which you have had to design for? What was your biggest surprise or memorable example of culture shock?

EDN europe

EDN-EUROPE is published 11 times in 2016 by
European Business Press SA,
533 Chaussée de Louvain, 1380 Lasne, Belgium
Tel: +32-2-740 00 50 Fax: +32-2-740 00 59
email: info@eetimes.be.
VAT Registration: BE 461.357.437.
RPM: Nivelles.
It is is free to qualified engineers and managers involved in engineering decisions – see:
http://www.edn-europe.com/subscribe

CONTACTS

PUBLISHER

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

EDITOR-IN-CHIEF

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

Patrick Mannion
Brand Director EDN Worldwide

european business press

CIRCULATION & FINANCE

Luc Desimpel
luc.desimpel@eetimes.be

ADVERTISING PRODUCTION & REPRINTS

Lydia Gijsegom
lydia.gijsegom@eetimes.be

ART MANAGER

Jean-Paul Speliers

ACCOUNTING

Ricardo Pinto Ferreira

SALES CONTACTS

Europe

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

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

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

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

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

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

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

USA & Canada

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

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

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

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

Asia

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

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

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