EDN europe

DECEMBER 2014
www.edn-europe.com
Issue 11

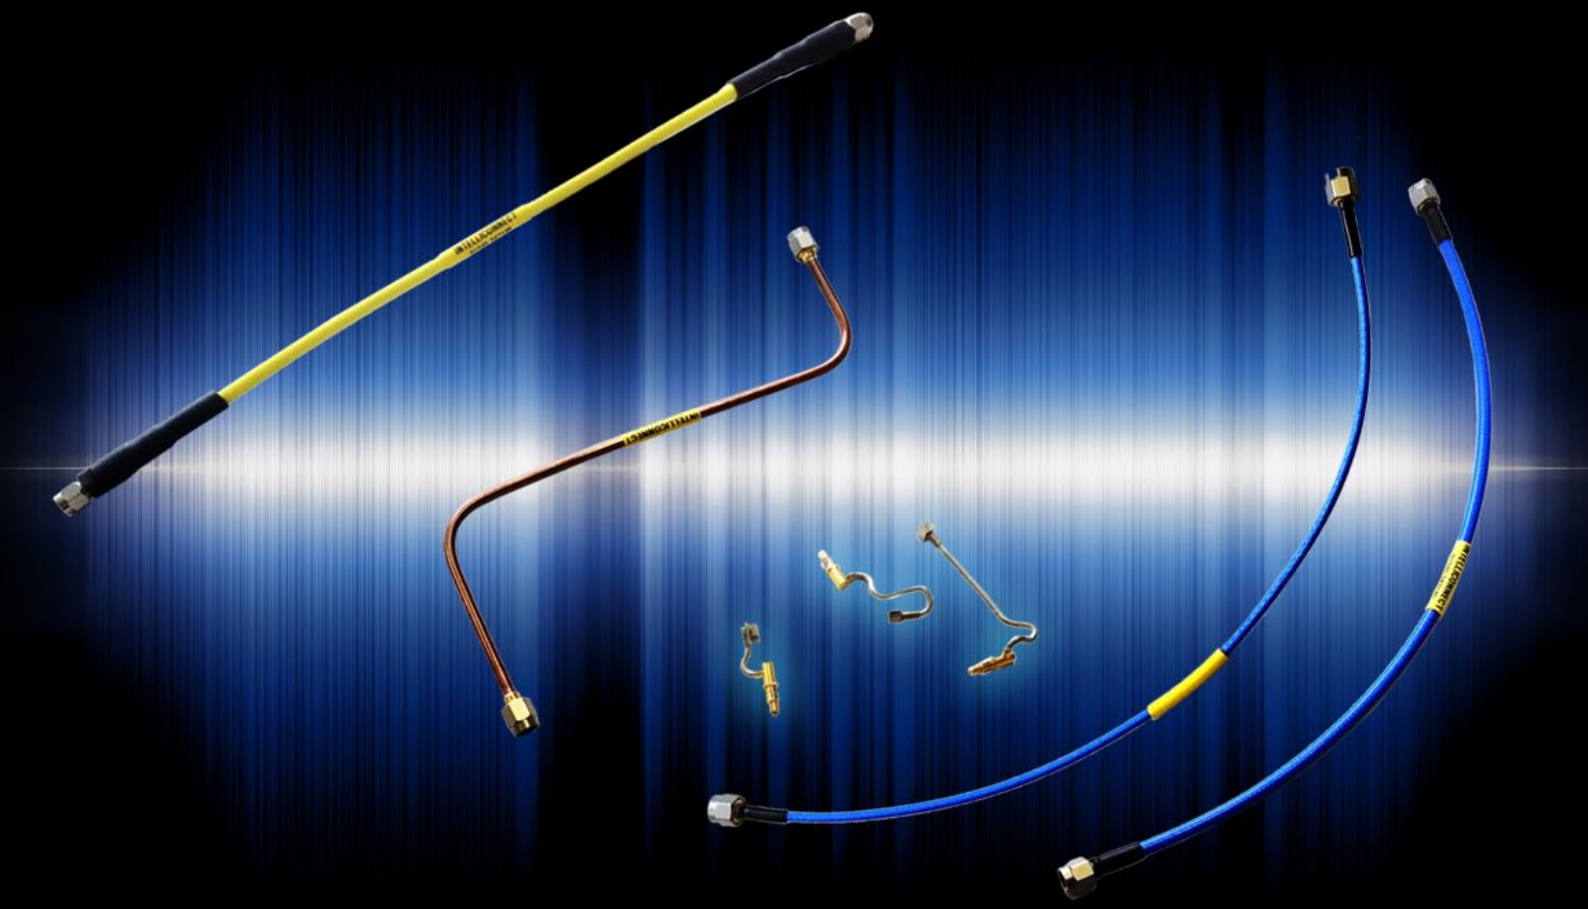

EDN europe

contents

DECEMBER 2014

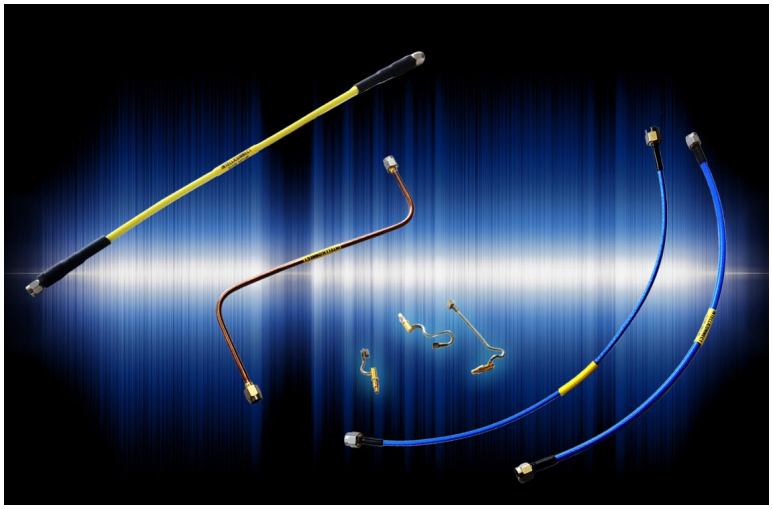

Cover image; Phase-matched cable assemblies to 50 GHz

Rated for operation from DC to 50 GHz, a range of phase matched cable assemblies allows engineers to specify electrical length in time or degrees at the frequency of operation, for a specific cable assembly. Cables can be matched in "sets" or pairs. One of two methods may be used - electrical length (matched to a golden standard), or phase matched to another cable. Typical applications include high power amplifiers, RF combiners and filters, multi-beam antenna arrays, phased array radars. Intelliconnect can provide bespoke cable assemblies for both high-volume or specialist small batch requirements operating from DC to 50GHz for RF applications including test and measurement and defence. Coaxial, triaxial and multi-way connectors may be specified and fully immersible waterproof versions are available; phase matching is available to ensure signal integrity in individual, or batches of, assemblies. Read the full product news item here.

FEATUREARTICLES

Departments and Columns

EDN. COMMENT

Analog Tips

Eye on Standards

INTERNET OF THINGS

productroundup

DESIGNIDEAS

pulse

EDN europe

DECEMBER 2014
www.edn-europe.com
Issue 11

CONTACTS

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

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

Patrick Mannion
Brand Director EDN Worldwide

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

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

ART MANAGER
Jean-Paul Speliers

ACCOUNTING
Ricardo Pinto Ferreira

european business press

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

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

SALES CONTACTS

Europe

Daniel Cardon
France, Spain, Portugal
+33 688 27 06 35
cardon.dan@orange.fr

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

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

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

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

Andrea Rancati
Italy
+39-02-284 6716
info@silvera.it

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

USA & Canada

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

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

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

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

Asia

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

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

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

SECURITY IN THE IOT ERA

Are you weary of hearing about the Internet of Things? If so, I fear have bad news; it's not going to ease up any time soon. If 2014 has been the year in which the IoT emerged as a major topic, then 2015 will see yet more frenetic activity as vendors at all stages in the product pipeline seek to establish their presence, and to carve out a foothold in whatever market eventually emerges. I could say; if it's becoming tiresome from where you sit, you should see the inbox at edn-editor@eetimes.be – but that would be quite rightly dismissed as special pleading.

There is a fairly well-understood curve associated with such phenomena; the anticipation of a lucrative market drives an initial phase of over-enthusiastic promotion; after a time, some sense of reality emerges about what types of products will actually work in the market; some first-generation products fall by the wayside, some begin the long ascent to volume, to market share and to profitability. Right now, we're firmly in phase one. There are some general areas in which a consensus is emerging, however.

There will be very large numbers of connected things; this has become a statement of the obvious. Some will seek to exploit their connectedness by acquiring, consolidating and 'mining' data culled from activity across many nodes. This we can class as the, "it's all about big-data" faction. For other devices the convenience of connection is an end in itself – "we could have built our own links but communicating over the Internet has become the default, so we'll use that." Or in other words, "it's not about big data." There's a third sub-division, possibly the most interesting; "it's not all about big data." With the emphasis on "all".

One area in which there is general agreement is that security is key; over and over again we are told, "if we can't get the security right, the IoT is not going to happen." At the higher end of IoT system complexity, there will be data streams of great intrinsic value that will have to be protected; but with universal connectivity, every connected device is a potential portal, not only at its own level, but to every system that accesses it, and that uses data abstracted from it. Security means not only protecting a device's intrinsic function, but preventing it being misappropriated to gain access elsewhere.

It is well understood that there is no such thing as absolute security. At one level, the guideline applies that if you make the cost of breaking your security much greater than the benefit that will accrue to anyone who does so, that will be sufficient. The IoT will offer targets to the potential hacker that in terms of technical challenge and peer-group prestige will render that rule-of-thumb irrelevant.

You may, or may not, consider the self-driving car as an IoT device. It will have to communicate with others like itself, and with infrastructure: and it will be Internet-connected. To that extent it qualifies and it is also a handy high-profile example of a target that would irresistible to hackers. So it was that in recent weeks, and at more or less the same time, we had Elon Musk (of Tesla electric cars) telling us that the self-driving car was [only] five or six years away: and a report from the UK's Institution of Engineering and Technology opining that the same prospect might be many more years away, exactly because of the security threat.

At the component level, every technology that has ever been applied to ensuring security is being offered as a route to building-in safe operation to IoT devices from the outset. Encryption and authentication engines that until recently would have required a PCB to themselves are now a basic requirement on the lowliest microcontroller offered as a IoT-capable part. And it scales upwards from that point.

As I noted above, I hear presentation after presentation that declares, "If we don't get security right, then; no IoT bonanza." The subtext is, of course, "...but we have the tools, the products and the concepts to ensure that it will all be OK." Which prompts me to ask of those doing the presentations; what if the technology at our disposal is not, in fact, up to the task? If we can't make the IoT both intrinsically reliable (robust against internal failure) and proof against intrusion? Most answers take the general form, "It will be OK, we can do this." With, possibly, some nervous shuffling of feet and crossing of fingers behind backs.

I received a more considered response from Freescale's Tim Summers; he is primarily a network and datacomms guru, as well as speaking for a company that is offering silicon at every level, for the expected IoT tsunami. He too, let me make it clear, thinks it (security) can be done, that the technology resources that are, and will be, at our disposal will guard out interests. But he offers this further thought, "As a society, we have already made that choice; we all carry [for example] bank cards knowing that there is a certain failure rate of the systems that support them." We know, he argues, that the technology is imperfect but we choose to use it nevertheless, for the convenience and benefits it brings. And our approach to IoT devices will be no different.

I appreciate the candour of the answer; but when you set it against the scale and pervasiveness of the systems we are proposing to construct in the name of the Internet of Things, I'm not sure I'm greatly reassured. What do you think?

Dreamteam for success.

Signal generation and analysis for demanding requirements
When working at the cutting edge of technology, you shouldn't waste your time with inferior tools. Rely on measuring instruments evolved in the spirit of innovation and based on industry-leading expertise. Instruments like the R&S®SMW200A vector signal generator and the R&S®FSW signal and spectrum analyzer. Each is at the crest of today's possibilities. As a team, they open up new horizons.

ROHDE&SCHWARZ

pulse

FTDI Chip USB 3.0 interface IC retains ease-of-use features

FT600Q and FT601Q are FTDI Chip's first generation USB 3.0 products that function as SuperSpeed USB 3.0 to FIFO bridges, providing data bursting rates of up to 3.2 Gbps.

The company says it has carried over the ease-of-use features from its previous USB devices; Windows 8 marked a "take-off point" for USB 3, and although this generation has a completely new driver, the user sees no difference, and still sees a simple read-write transfer. The FT600Q comes in 56-pin QFN package and has a 16-bit wide FIFO bus interface, while FT601Q comes in 76-pin QFN package and has a 32-bit wide FIFO bus interface. Both these chips support up to eight endpoints, other than the management endpoints. The endpoints are linked to a configurable endpoint buffers of 16 kByte length for IN and 16 kByte for OUT.

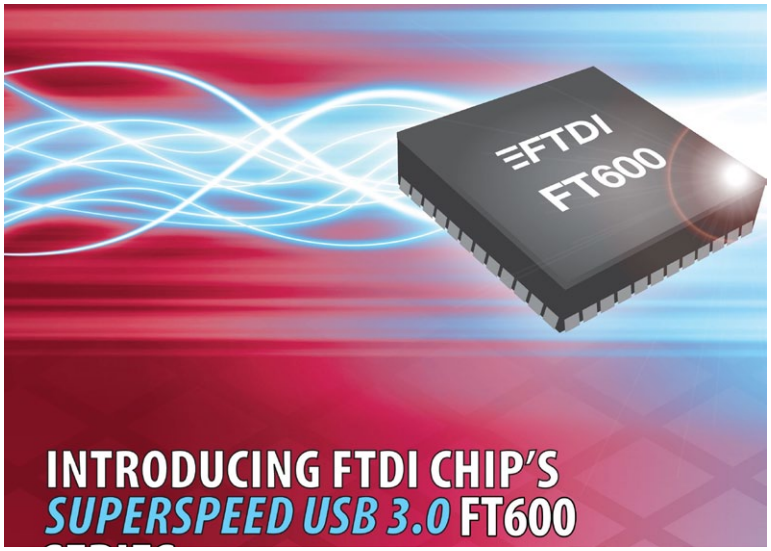

Both FT600Q and FT601Q support two interfacing modes; the 245 FIFO mode and the multi-channel FIFO mode, and thus provide more flexibility for system designers. The 245 FIFO mode has a simpler protocol, but for more sophisticated customers, the multi-channel FIFO mode supports up to 4 logical FIFO channels and data structures optimised for higher throughputs. The FIFO is provided with a 16 kByte configurable buffer.

The remote wake up function on these chips can be used to rapidly bring the USB host controller out of suspend mode. The USB battery charger detection function enables USB peripheral devices to detect the presence of a higher current power source in order to boost charging capabilities. It means that the FT600 can detect connection to a USB-compliant dedicated charging port (DCP) and transmit a signal allowing external logic to switch to charging mode. The IC can also benefit from the higher power delivery capabilities (up to 900 mA) that the USB 3.0 standard supports while still being able to transfer data.

Complete article, here

Digi-Key and Mentor cooperate on low-cost PCB design tools

Aiming to "meet the needs of today's engineers" Mentor Graphics – as the software provider – has joined forces with distributor Digi-Key – with access to a worldwide base of designers – to present an "entry-level" range of three software tools that are intended to get designs started more easily and, in many cases, take them through to production.

The software is in the Mentor product family, at a new lower entry point than either its (high-end) Xpedition or its PADS offering. Mentor and Digi-key are emphatic that this is not free software ("We've got free software as well," says a Digi-Key spokesman, "this is different") but a fully supported and maintained software offering. It may lack some of the higher functions of PADS or Xpedition, but it is only capacity-limited to 1500 connections, six layers and 50 square inches.

The three tools are Designer Schematic, Designer Layout, which are the chargeable elements, and both give access to the Mentor PartQuest website. Digi-Key points out that while the tools will directly access its own inventory and parts database, there is no restriction to use parts or devices from its own catalogue. The jointly developed Designer Schematic and Designer Layout tools are available exclusively from Digi-Key.

Designer Schematic is available starting at $199 ($299 after the introductory period). Designer Layout starts at $449 ($649 per year going forward). A free 15 day trial is offered.

Complete article, here

Power Integrations' InnoSwitch cuts optocouplers from switching supply circuit

Power Integrations has introduced several generations of off-line switching power supply ICs; the latest, InnoSwitch, returns to deriving feedback information to control output regulation direct from the output, but feeds that data back to the primary switch controller via an air-core magnetic coupler.

Calling it FluxLink, PI says that its safety-isolated communication technology combines primary- and secondary-switcher circuitry to reduce component count, eliminate slow and unreliable optocouplers, outperform primary-side controllers and reduce manufacturing costs

A single package with all necessary creepage and clearance spacings includes primary switching controller with integrated power FET on one die, secondary synchronous rectifier and associated circuitry on a second die, and the secondary-to-primary feedback path. This is implemented using the copper leadframe of the part. Part of the leadframe and chip bond wires form the 'primary' of a single-turn, air-core transformer (the 'primary' of the transformer is on the secondary or output side of the power supply layout, as it is passing output voltage level information back to the input switching controller). This loop couples to a further loop in the leadframe, connected to the controller IC; the physical layout is configured to cancel stray magnetic field effects. Clearances between the transformer loops are sufficient to give safety and regulatory isolation levels. On the PCB, the single package spans the gap between input and output sections.

The InnoSwitch family of highly integrated switcher ICs, PI says, combines primary, secondary and feedback circuits into a single, worldwide safety-rated, surface-mount package that exceeds all global regulatory standards for efficiency and no-load consumption, while minimising component count and providing highly accurate constant voltage and constant current, up to 25W. The InnoSwitch family is suitable for smart mobile device chargers and adapters for a wide range of applications such as set-top boxes, networking equipment and computer peripherals.

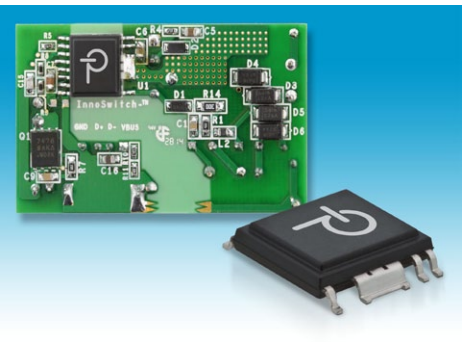

Unlike primary-side regulated switchers, InnoSwitch-based secondary-side regulated (SSR) designs are inherently less sensitive to the tolerance of external components such as transformers, diodes, resistors and capacitors. This increases manufacturing yield and reduces total power supply cost. Now, mobile device chargers up to 5A can have a total component count as low as PSR designs, with accurate CV and CC control (±3% and ±5% respectively) and low voltage ripple.

Complete article, here

Raspberry Pi Model A+; latest design for embedding – in distribution

With a smaller, lower-power board and HAT-compatible 40-pin GPIO as featured on Model B+, this board is more expandable, more embeddable and more mobile; RS Components now has the Raspberry Pi Model A+, which combines enhanced value and ease of use with smaller size and lower power consumption for applications such as industrial controls, remote monitors, and multimedia devices.

The Raspberry Pi Model A+ has the 40-pin GPIO connector introduced on the Model B+, with the same pin-out and mounting holes for standard Hardware Attached on Top (HAT) accessories allowing users to add extra functions. The Model A+ utilises the Model B+ power architecture to achieve lower power consumption than the earlier Model A. With smaller board dimensions of 65 x 56mm, compared to 86 x 56 mm for the Model B+, the new Model A+ is more easily embeddable and better suited to mobile or battery-powered applications. Click here for the full version of this item.

Prototyping accessories

Also new on the Raspberry Pi scene, distributor element14 has launched the Microstack range, a series of accessories for the Raspberry Pi that will make it quicker and simpler for all levels of user to create and prototype physical devices; the first Microstack products are GPS positioning, accelerometer and prototyping modules.

Microstack modules are positioned as the “building blocks for The Internet of Things for All”. Microstack was created by the designers of the original PiFace Digital and PiFace Control & Display accessories for the Raspberry Pi and the recent PiFace revisions for the Raspberry Pi model B+. Microstack builds on PiFace, offering many more connected-device possibilities.

The Microstack GPS module provides easy plug and play solution for designs requiring GPS positioning to create geo-location aware projects. The module features standalone data logging and also brings globally synchronised, highly accurate time keeping to the Raspberry Pi. The Microstack Accelerometer module provides plug and play integration and development for designs requiring accelerometer such as robotics, tilt sensors, gaming, tracking and motion. It is based around a MMA8491Q, a low power, 3-axis low-g accelerometer offering multi-range 14-bit ± 8g resolution. The Microstack Protoboard module is a prototyping module allowing users to create their own circuit designs to be used with the Microstack form factor.

Complete article, here

Vicor extends "component power" concepts with VIA package

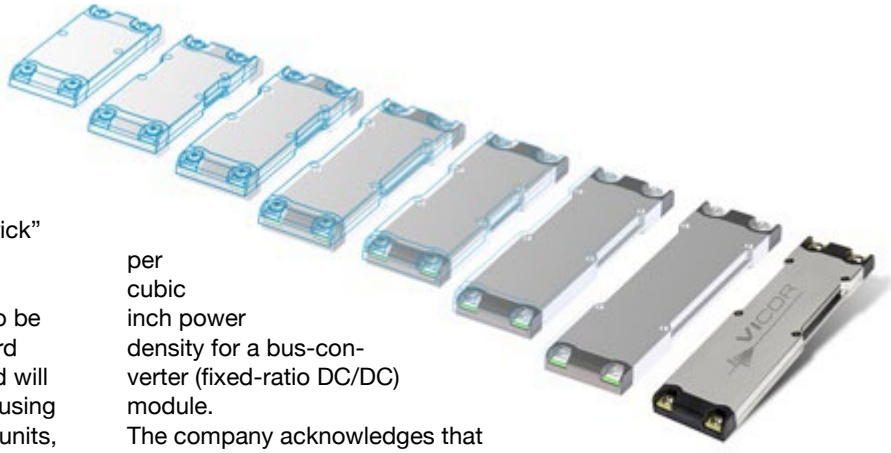

Vicor has said it will formally introduce its first product in its VIA – Vicor Integrated Adaptor – package by the end of 2014, for availability in early 2015. The package extends Vicor's component power concepts and will initially be used for front-end modules.

The VIA package is superficially not unlike a traditional "brick" format, a module for direct mounting to a heatsink or cold wall. Vicor intends to use it to house input modules for its component-power architecture; the first product is likely to be a 400W, AC/DC input block. VIA modules will be a standard width and profile of 35.3 mm and 9.3 mm respectively, and will vary in length from 72 to 141 mm. The machined metal housing encloses and provides distribution for one of Vicor's ChiP units, plus other circuitry.

The ChiP is a recent introduction that hosts a variety of power conversion functions built in an almost totally-automated process. Multiple instances of converters are assembled together on a single large panel, with structural features such as in-board magnetic components that, in the final form, present thermally-conducting faces to both top and bottom of the package. ChiP panels are over-moulded and then sawn – in a similar fashion to a scaled-up silicon wafer – into individual blocks. When used discretely, Vicor adds pins or terminations to the device's connections – that are exposed on the sawn edges – to allow conventional connections to be made. With appropriate heatsinking to top and bottom of a ChiP package, Vicor quotes over 1 kW per cubic inch power density for a bus-converter (fixed-ratio DC/DC) module.

The company acknowledges that "ultimate" (top and bottom contact, low thermal resistance) heatsinking is not always easy to achieve, which is partly the rationale behind the VIA package. The VIA houses a ChiP module with sufficient space at either end for additional PCB-mounted functions such as EMI filtering, transient protection, or digital bus control/reporting. In the VIA package, the ChiP modules are used without their added pins/surface-mount terminals, direct connections being made to the edge contacts exposed when they are "sawn". A machined metal housing surrounds the ChiP block, and provides a good thermal path from its top, as well as bottom, surface, to the base plate.

Complete article, here

16-bit, high-resolution oscilloscopes from Rohde & Schwarz

A 16-bit, high definition (HD) mode increases the vertical resolution of the R&S RTO and RTE oscilloscopes to up to 16 bits, for waveforms that are sharper and show signal details that would otherwise be masked by noise. Users benefit from even more precise analysis results.

The increased resolution is achieved by signal processing, using the same A/D converter already in use, with the trade-off of reduced bandwidth while the HD mode is engaged.

To be able to analyse waveforms in detail, the input sensitivity of the oscilloscopes has been increased to 500 µV/div. To achieve the higher resolution, the signal is low-pass filtered directly after the A/D converter. The filter reduces the noise, thereby increasing the signal-to-noise ratio. Users can adjust the bandwidth of the lowpass filter from 10 kHz to 500 MHz as needed to match the characteristics of the applied signal.

Typical applications include examining in detail low-voltage components on a signal that also exhibits high-voltage components; one example is the characterisation of switch mode power supplies. The voltages across the switching device must be determined during the off and on times within the same acquisition. Because the voltage variations during these switching cycles can be several hundred volts, a high resolution is essential for the precise measurement of small voltage components.

Switching on the high definition mode does not compromise measurement speed or functions. Since the digital filtering, which improves resolution and noise suppression, is implemented in realtime in the oscilloscope's ASIC, the acquisition and processing rates remain high.

Complete article, here

Universal temperature sensor IC linearises data to 0.1°C accuracy

LTC2983 is a high performance digital temperature measurement IC which directly digitises RTDs, thermocouples, thermistors and external diodes with 0.1°C accuracy and 0.001°C resolution.

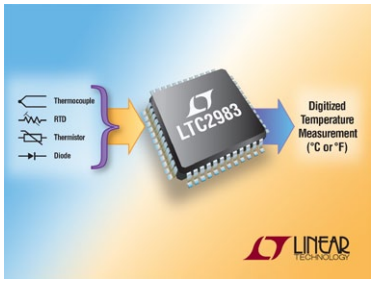

A high performance analogue front end combines low noise and low offset buffered ADCs with all the necessary excitation and control circuits for each sensor. Measurements are performed under the control of a digital engine combining all the algorithms and linearisation required for each. The LTC2983 provides a multiplexed high precision interface to virtually any sensor. It precisely measures absolute microvolt level signals from thermocouples and ratiometric resistance measurements from RTDs and thermistors, linearises the results and outputs the results in °C or °F. Up to twenty analogue inputs are available. The SPI interface works with virtually any digital system and a comprehensive software support system with drop-down menus allows easy customising of the IC.

LTC2983 interfaces with a wide variety of temperature sensors, including type B, E, J, K, N, S, R, T thermocouples, 2,3, or 4-wire RTDs, 2.25 kΩ to 30 kΩ thermistors and temperature-sensing diodes. It works with ground-referenced sensors without the need for amplifiers, negative supplies, or level shift circuitry. Signals are simultaneously digitised with three, high accuracy, 24-bit ΔΣ ADCs using an internal 10ppm/°C reference. Automatic thermocouple cold junction compensation can be done using any type of external sensor. Included on the chip are linearisation algorithms for all common sensor types. Custom sensors can be linearised with custom coefficients programmed into the chip.

Complete article, here

10-W wireless power delivery promises faster, cooler charging without plugging in

TI's receiver and transmitter system efficiently charges one- and two-cell battery-powered applications, and supports any Qi-compatible 5-W wireless charging system. Claimed as the first fully integrated 10-W wireless charging receiver and corresponding transmitter, the bq51025 and bq500215 enable waterproof and dustproof portable designs and provide a faster, cooler charge to one- and two-cell (1S and 2S) Li-Ion batteries. The charging solution is also compatible with any 5-W Qi-compliant product in the market – allowing consumers the flexibility to charge in more places.

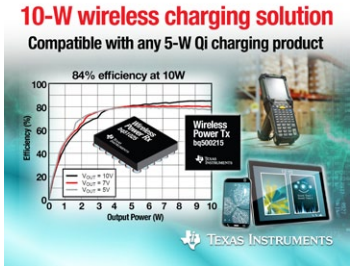

The bq51025 receiver supports a programmable output voltage of 4.5 to 10 V and achieves up to 84% charging efficiency at 10-W when paired with TI's bq500215 wireless power transmitter, significantly improving thermal performance. The fully-contained wireless power receiver solution measures 3.60 by 2.89 mm, and can be designed into many portable industrial designs.

The bq500215 is a dedicated, fixed-frequency 10-W wireless power digital controller transmitter compatible with 5-W Qi receivers. The transmitter features an enhanced foreign object detection (FOD) method that detects objects before any power is transferred, and actively reduces power if excessive loss is detected. Qi-compliant communications and control ensures compatibility with any Wireless Power Consortium Qi-certified transmitter or receiver up to 5W.

Complete article, here

Microchip's 5V "EV" dsPIC33 MCUs offer enhanced noise immunity, ECC flash

Microchip has added a family of 16-bit dsPIC33 Digital Signal Controllers (DSC), the dsPIC33 "EV" family, that has 5V operation for improved noise immunity and robustness, for devices operating in harsh environments such as appliance and automotive applications.

The dsPIC33EV family is the first dsPIC DSC with Error Correcting Code (ECC) flash for increased reliability and safety. For safety-critical applications, the dsPIC33EV devices also include CRC, Deadman Timer (DMT), and Windowed Watchdog Timer (WWDT) peripherals as well as a backup system oscillator and certified Class B software.

Other features of this family include up to six advanced motor control PWMs, 12-bit ADC, and operational amplifiers, a combination configured for motor control applications. The dsPIC33EV devices provide interface to 5V automotive sensors such as level or flow sensing, with improved noise immunity and enhanced reliability, and provide performance to execute smart sensor filter algorithms and integrate CAN communication software.

Complete article, here

Intel Edison embedded/IoT development platform, now in distribution

RS Components now has the complete Intel Edison development platform, which includes: the ultra-small Intel Edison compute module; the Intel Edison breakout board that helps rapid prototyping; and the Intel Edison board for Arduino.

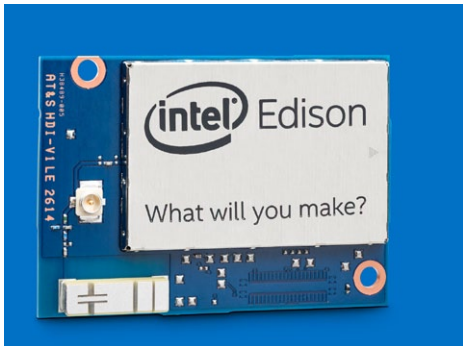

The Intel Edison module is a highly miniaturised product-ready compute module built using an advanced 22nm System-on-Chip (SoC) that includes a dual-core, dual-threaded 500 MHz Intel Atom CPU and a 100 MHz 32-bit Intel Quark microcontroller. Connectivity is provided, and built-in Wi-Fi and Bluetooth Low Energy suit this module for Internet of Things (IoT) applications. A device-to-device and device-to-cloud connectivity framework is also integrated.

Plugging the 35.5mm x 21.0mm x 3.9mm Intel Edison module into the breakout board exposes the module's 1.8V I/Os, and allows connection via the board's USB Micro Type-AB connector. The board also provides a DC power-supply jack, battery-charger circuit, and USB OTG power switch.

The Intel Edison expansion board for Arduino is hardware and software pin-compatible with Arduino shields designed for the Uno R3. The board provides important connections in the same locations as on the Arduino Uno R3, and also has a Micro SD card connector as well as micro and standard-size USB connectors.

Complete article, here

Programmable RF dual transceiver implements MIMO-on-a-chip

Lime Microsystems is shipping its second generation Field Programmable Radio Frequency (FPRF) devices. The FPRF contains dual transceivers that cover the frequency range 0.1 to 3800 MHz.

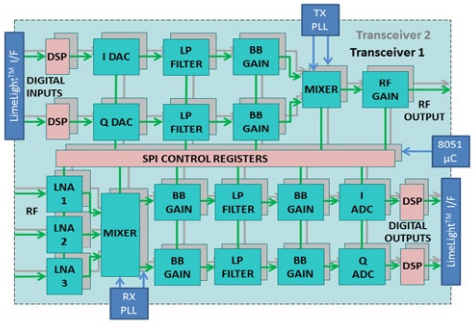

The fully programmable, low power chips cover all the cellular bands used in 2G, 3G and 4G networks, as well as many commercial and military bands. Parameters such as gain and filter bandwidth are also fully in-system programmable, making the device a highly flexible RF solution for a wide range of wireless applications.

The dual transceivers are can implement Multiple-In, Multiple-Out (MIMO) systems with a digital bandwidth up to 60 MHz in each channel. The LMS7002M includes user-programmable DSP that can equalise the gain and phase of a MIMO system, or enhance the analogue filtering to provide lower power consumption. The low power chip uses a 65 nm CMOS process, and consumes 550 mW in single-in single-out (SISO) mode or 800 mW when configured for MIMO operation. The device can operate using a single 1.8V supply rail, and requires a minimum of external components.

Complete article, here

Instruction smart trace on 8051-derivative cores – smart & non-intrusive

Polish silicon IP company Digital Core Design has enhanced its 8051 portfolio with the functionality of an Instruction Smart Trace. Starting from the most popular DP8051, through the DP80390, to the fastest 8051-architecture core, the DQ8051 – can all efficiently reduce trace memory size and increase traced program history. Instruction Smart Trace is an inherent part of the DoCD Hardware Debugger, which provides debugging capability for System-on-Chip (SoC) designs. Unlike other on-chip debuggers, the tool provides non-intrusive debugging of a running application and saves designer's time, thanks to the hardware trace, called Instructions Smart Trace buffer (IST).

The DoCD-IST captures instructions in a smart and non-intrusive way, so it doesn't capture addresses of all executed instructions, but only these related to the start of tracing, conditional jumps and interrupts. This method not only saves time, but also allows improving the size of the IST buffer and extending the trace history. "For example, by using 256 Bytes of trace memory, we can store 128 program branches and decode much more program history, since the executed program is composed of normal opcodes (mov, add, mul, anl, etc.) and branches," explains Tomek Krzyżak, VCEO of Digital Core Design.

Complete article, here

Software test libraries for STM32 Microcontrollers boost functional safety

Yogitech (Pisa, Italy), provider of functional safety solutions, has announced that its fRSTL libraries for STM32 microcontrollers from STMicroelectronics have been positively assessed by TÜV Rheinland for, among other attributes, functional safety certification.
Heinz Gall, Global Business Field Manager, Functional Safety of TÜV Rheinland, said, "Yogitech's fRSTL libraries have been positively assessed according to the IEC61508:2010 functional safety standard for the systematic capability (SC3) and the Diagnostic Coverage, making them suitable to be integrated into SIL2, HFT=0 and SIL3 HFT=1 applications. The final report and certificate will be issued by first week of December 2014".
fRSTL for ST's STM32 MCUs is a set of software libraries addressing the STM32 microcontroller series, and intended to detect hardware random faults. They are an application-independent, off-the-shelf product designed to be easily integrated into safety-relevant systems. fRSTL libraries build on fRMethodology, Yogitech's techniques that apply functional safety to integrated circuits using a white-box approach, and are duly verified by means of intensive fault injection simulations on the RTL and netlist level, to which Yogitech has access.

Complete article, here

CSR's accurate indoor location available to Android device developers

A software library brings, says the provider of SiRF location products, plug and play simplicity to Android app developers: the SiRFusion Software Development Kit (SDK) enables indoor positioning for Android developers looking to create next-generation apps.
Developers can now use the SiRFusion library to rapidly integrate new location-based capabilities and services such as indoor location tagging and analytics for social networking applications, indoor navigation, lone worker efficiency and safety capabilities, as well as indoor asset tracking and targeted e-commerce services.

Mobile applications with integrated SiRFusion can – CSR says – deliver the ubiquity of outdoor navigation to indoor environments without costly surveys or infrastructure upgrades. SiRFusion combines real-time Wi-Fi signals, satellite positioning information, pedestrian dead reckoning, and the company's cloud-based CSR Positioning Center to calculate accurate indoor location. SiRFusion automatically crowd-sources a venue's indoor Wi-Fi signatures as consumers walk through the location, and it has also been architected to accommodate future proximity and location technologies such as Bluetooth Smart beacons, Wi-Fi Round Trip Time (RTT), and Indoor Messaging System (IMES).

Complete article, here

USB-scope drivers for PicoScopes now run on BeagleBone and Raspberry Pi boards

Pico Technology has released beta drivers for its oscilloscope and data logging devices to run on the ARM-based BeagleBone Black and Raspberry Pi development boards. The drivers give programmers access to a wide range of compact, economical USB oscilloscopes and data loggers. These include high-speed oscilloscopes with bandwidths up to 500 MHz, high-resolution 12-bit and 16-bit oscilloscopes, and deep-memory oscilloscopes with buffer sizes up to 512 Msamples. Most PicoScope oscilloscopes also offer advanced digital triggering, fast block-mode and streaming-mode data capture, and a built-in signal generator or arbitrary waveform generator. The data logger range includes multichannel voltage loggers, 8-channel thermocouple loggers and the educational DrDAQ multifunction logger. Separate driver packages and snippets of C code are available for Raspbian and Debian systems. Similar code could be developed in C-compatible programming languages such as C++, Java and JavaScript.

Complete article, here

BY STEVE SANDLER, PICOTEST

HOW TO MEASURE THE FASTEST POWER SWITCHES

If you're designing power circuits with GaN devices, you need a grasp of the device's switching speed. To measure that, your oscilloscope, probes, and interconnects must be fast enough to minimise their effect(s) on the measurements.

Gallium Nitride (GaN) FETS are poised to replace silicon power devices in voltage regulators and DC-DC power supplies. Their switching speeds are significantly faster and their $R_{DS(on)}$ is lower than silicon MOSFETS. That can lead to higher power efficiency power sources, which is good for all of us.

One of the most frequent questions I receive on the subject of device performance is "how fast are they, really?" My general response is that they are blazingly fast but that we just don't know quantitatively how fast. To find out, I made some measurements using a 33-GHz real-time oscilloscope and a high-speed transmission-line probe. I'll discuss the design limitations that mask the device's speed, and what's in store for the future. With these measurements, I believe we'll be designing power supplies switching at 250 MHz before long.

Figure 1 shows two evaluation boards used to perform the measurements. Both boards include a gate-voltage regulator, driver, pulse conditioner, and two eGaN switches switches (from Efficient Power Conversion Corp.). The board on the right is a complete DC-DC converter, which includes a Gen4 monolithic half-bridge (both switched on one die) and includes an L-C output filter. The board on the left uses individual Gen3 eGaN devices in a half-bridge configuration, lacking the L-C output filter. In both cases, an external pulse generator provides a PWM (pulse-width modulated) signal through a BNC connector soldered to the test board's PWM input. The switch rise time is measured on each board at input voltages of 5 V and 12 V.

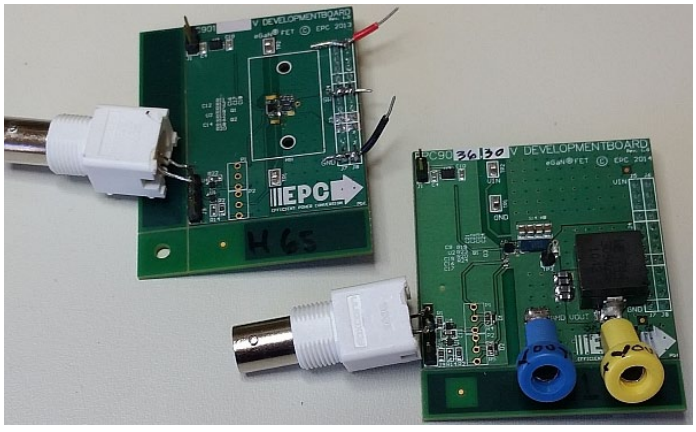

Figure 1 *The test boards are shown with the half bridge configuration only on the left and the complete DC/DC converter on the right. The banana sockets allow connection of the board to an electronic load. BNC connectors provide access to an external pulse generator.*

Instrument and probe requirements

To ensure that the instrument and probe don't significantly impact the measurement, we can assume that the rise times of the probe, oscilloscope, and the half-bridge can be added using root sum squares. This isn't always true, but for our initial estimates we'll assume this relationship holds.

The measured rise time of the half-bridge including the rise time of the oscilloscope and the rise time of the probe is:

$$t_{measured} = \sqrt{t_{scope}^2 + t_{probe}^2 + t_{half-bridge}^2}$$

The actual rise time of the half-bridge is determined as follows:

$$t_{half-bridge} = \sqrt{t_{scope}^2 + t_{probe}^2 - t_{measured}^2}$$

To restrict the measurement error to some percentage, K, the rise time of the instrumentation can be related to the actual rise time:

$$\sqrt{t_{scope}^2 + t_{probe}^2 + t_{half-bridge}^2} = (1+K)\cdot t_{half-bridge}$$

Solving for K, the ratio of the instrument rise time to the actual half-bridge rise time is:

$$Ratio(K) = \sqrt{K} \bullet \sqrt{K+2}$$

So for the two examples, if we wish the measured result to be less than 5% or less than 10%, then the rise time of the oscilloscope and the probe needs to be less than 32% or 46% of the FET rise time, respectively. Stating it differently, the instrumentation should be 3.1 or 2.2 times faster than the FET rise time, respectively.

Measuring switching performance

The oscilloscope used here is a Keysight 90000-X Series 33 GHz oscilloscope with a Teledyne LeCroy PP066 transmission line probe. The connection between the oscilloscope and the probe is made through a 50 GHz Huber+Suhner Sucoflex-100 cable. The oscilloscope and probe used for these measurements represent "overkill" because they're much faster than the minimums noted above in order to assure that the measurement is valid.

In the continuation of this article Steve Sandler measures the risetime of the instrumentation setup, setting a conservative reference against which to measure the risetime of the power switch configurations. He presents measured results and draws some real-world comparisons against state-of-the-art silicon switches.

Complete article, here

Your Global Link to the Electronics World

www.automotive-eetimes.com

www.eetsearch.com

www.electronics-eetimes.com

www.analog-eetimes.com

www.ledlighting-eetimes.com

www.edn-europe.com

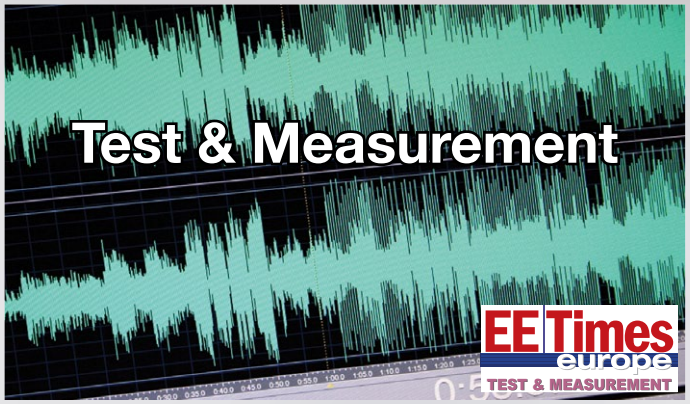

www.tm-eetimes.com

www.power-eetimes.com

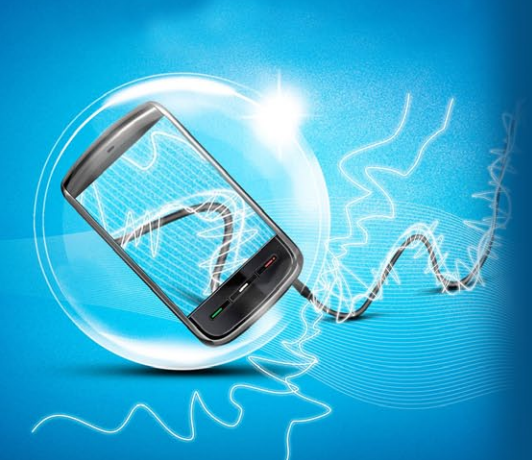

Analog Tips

AD8495 INTERFACE TO TYPE T THERMOCOUPLES

BY SCOTT HUNT, ANALOG DEVICES

Thermocouples are a common, inexpensive way to make precision temperature measurements over wide temperature ranges. Type T thermocouples (Copper-Constantan) are sensitive, stable, easy to manufacture, and moisture tolerant—and the copper-to-copper connection to the printed circuit board eliminates the need for an isothermal block. These are some of the main reasons that type-T thermocouples are chosen for many applications – from catheters to food processing – that don't need the full temperature range available from other thermocouples such as type K.

ADI makes thermocouple amplifiers with built-in cold-junction compensation, but not for type T thermocouples. Fortunately, type T and type K amplifiers have similar voltage characteristics up to about 100°C. Thermocouple conditioners intended for use with Type K thermocouples such as AD8495 or AD8497 can take advantage of this similarity to measure type T thermocouples with very low errors. The output voltage for higher temperatures will deviate from the ideal 5 mV/°C transfer function of the AD8495, so it is necessary to account for the difference with a look-up table or polynomial. The figure was generated using the type T thermocouple table from NIST and the transfer function of the AD8495 in order to show the AD8495 output at various reference junction temperatures.

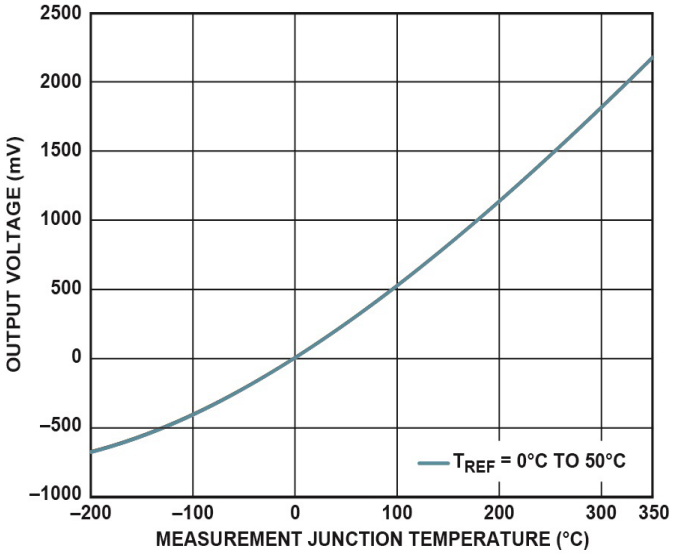

Figure 1. *Amplifier calculated output, for a Type T thermocouple and the AD8495*

The graph or this 6th order polynomial can be used to map the output voltage to the measurement temperature. V_{OUT} (mV) = $a_0 + a_1 \times T + a_2 \times T^2 + \ldots + a_6 \times T^6$, where the coefficients are shown in Table 1.

a_6	a_5	a_4	a_3	a_2	a_1	a_0
2.115E–16	–2.116E–12	2.602E–09	–4.238E–06	5.548E–03	4.728	1.099

Table 1.

Over the full –200°C to 350°C measurement temperature range and 0°C to 50°C reference junction temperature range, the calculated error of this polynomial is approximately –2.50 mV to 2.28 mV, which is less than about ±0.5°C. Thermocouple tolerance and AD8495 amplifier errors must be added.

Notice that the output voltage is negative for temperatures below 0°C. The output voltage is measured with respect to the reference pin (REF), so the voltage at REF can be raised to accommodate the negative output voltage with a single supply. See the AD8495 data sheet for important considerations when designing with thermocouples, including grounding, filtering interference, and layout practices in order to keep the device temperature equal to the reference junction temperature.

About the Author

Scott Hunt *[scott.hunt@analog.com] is a product applications engineer in the Linear Products Group in Wilmington, Massachusetts. Scott joined Analog Devices in 2011 after receiving a bachelor's degree in electrical engineering from Rensselaer Polytechnic Institute. Scott specialises in integrated precision amplifiers including instrumentation amplifiers, differential amplifiers, and thermocouple amplifiers.*

BY ARTHUR PINI , GREG TATE & OLIVER ROVINI

DIGITISER FRONT ENDS NEED THE RIGHT INPUTS

Digitisers used for capturing both low-speed and high-speed signals need to match their inputs to the fixed input range of their ADC (analogue to digital converter). To best use these instruments, you need to understand the tradeoffs.

Digitisers must minimise loading of the device under test and provide appropriate coupling. Additionally, filtering may be needed to reduce the impact of broadband noise. All of these features are provided by the instrument's "front end," which includes all the circuitry between the input and the ADC. Figure 1 shows a block diagram of a modular digitiser. Each input channel has its own front end, shaded in green.

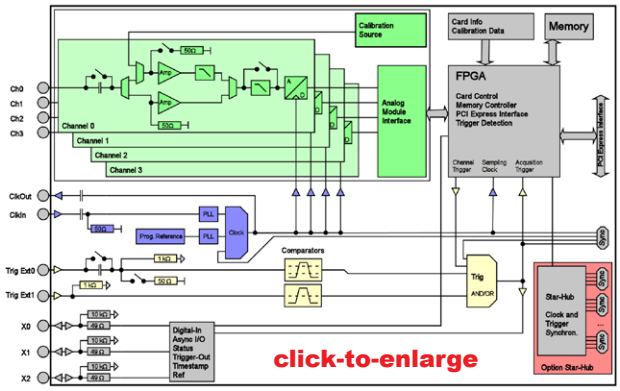

Figure 1. *The block diagram of a Spectrum M4i.44xx PCI Express 14/16 bit modular digitiser where the front ends for each channel are shown in green. The front end provides appropriate input coupling and termination along with range selection and bandwidth limit filtering.*

Maximising the versatility of a modular digitiser requires that the front end circuits have the following capabilities:

- Multiple input ranges offering the ability to capture a wide variation in input signal levels and at the same time minimising noise and distortion to maintain signal integrity.
- A selection of input termination to offer matching impedances or minimised loading with a high impedance input.
- A choice of coupling modes to offer either AC or DC coupling as needed.
- Filtering to minimise noise and reduce harmonic components if present.
- Internal calibration to maximise accuracy.

Input termination

A measuring instrument should properly terminate the signal source. For most RF measurements, this is generally a 50-Ω termination. A matching termination minimises signal losses due to reflections. The figures of merit for the 50Ω matching can be return loss or VSWR (voltage standing wave ratio). Both indicate the quality of the impedance match.

If the source device has a high output impedance, then it is more properly matched with a 1 MΩ high-impedance termination that minimises circuit loading. The 1 MΩ termination also lets you use high-impedance oscilloscope probes. The probe would increase the input resistance of the digitiser further, decreasing the loading on the circuit. Keep in mind that the probe will also decrease the signal level into the digitiser.

Because there is a tradeoff between convenience and signal integrity in designing with selectable input impedance, some modular digitiser suppliers only offer 50-Ω termination. Thus, if you need a high-impedance termination or both high impedance and 50 Ω you should verify that the manufacturer does offer both.

Input coupling

Input coupling in a measurement instrument offers the ability to AC couple or DC couple the measuring instrument to the source. DC coupling shows the entire signal, including any DC offset (non-zero mean signals). AC coupling eliminates any steady-state DC mean value. AC coupling is useful for measurements such as ripple measurements on the output of a DC power supply. Without the AC coupling, the DC output would require a large signal attenuation that would make the ripple harder to accurately measure. With AC coupling, a higher input sensitivity can be used, which results in a better ripple measurement.

The key specification for AC coupling is its low frequency cutoff (lower -3 dB point) of the AC coupled frequency response. This specifies how much a low-frequency signal will be attenuated by the AC coupling. AC coupling is related to the recovery time, the time needed for the input level to settle after a change in the DC level applied to the instrument. Generally, the lower the cutoff frequency, the larger the coupling capacitor and the longer the settling time.

Some modular digitisers offer only AC or DC coupling, but not both. Again, this is an engineering tradeoff to reduce complexity because a digitiser with fixed coupling doesn't need relays or switches. Again, your application will determine if a fixed or selectable coupling is acceptable.

Input voltage ranges

The digitiser's ADC generally has a fixed input range. The simplest interface is to have a single input with a fixed input range matching that of the ADC. While simple, this is not very practical in a measuring instrument unless the single range is exactly the one you need. To bring the input signal swing into the range of the ADC requires either an attenuator or an amplifier.

An attenuator is a simple voltage divider, generally resistive, which reduces the input signal's amplitude. When designed with quality components, it generally won't significantly affect signal integrity. One issue that appears with attenuators in the signal path is that the instrument's internal noise amplitude scales (relative to the input of the attenuator) with the front end attenuation. Thus, your digitiser's internal noise level is 58 µV rms and you add a 10:1 attenuator, then the noise level, referenced to the input, becomes 580 µV because you've reduced signal amplitude but not the digitiser's internal noise level.

In the continuation of this article, the authors note the contribution of input amplifiers; consider the configuration of overall input paths in the instrument; and show a number of side-by-side measured waveforms to illustrate the effects of the user's choices, optimising signal integrity.

Complete article, here

BY PINKESH SACHDEV, LINEAR TECHNOLOGY

SOFT MANAGEMENT OF POWER SYSTEM HARDWARE

Power system architects and designers of digital ASIC/FPGA/microprocessor boards may be justified in being a little envious of their colleagues in software engineering if you were to consider the following advantages that the software team has over their hardware colleagues.

The time delay between writing software code and observing its effect is much shorter than the days or weeks that a hardware board spin takes. The time to market is mostly limited by their coding and testing productivity and less by extraneous factors. Software updates are pushed to the customer as-needed to fix bugs and improve field performance. Hardware updates require boards to be recalled for rework. Software engineers easily track performance of their code through logged data from the comfort of their cubicles. Performance bottlenecks are quickly identified leading to rapid future improvements.

On the other floor, hardware engineers spend days in the lab, hunched over boards with voltmeter and oscilloscope probes. Software engineers write one core set of modular code and then adapt it for different customer and market needs. Customised hardware requires component and bill of materials (BOM) changes, risking designs diverging from each other.

Increasing challenges facing power system architects & designers

Exacerbating the situation, tougher challenges face the power system team on modern digital boards as nanometer-scale processor (ASIC, FPGA, microprocessor, DSP) supply voltages continue their downward march below 1V. The tolerance requirements on point of load (POL) supplies are tightening up, approaching 2% to 3%; the error budget includes DC accuracy, ripple and transient excursions during load steps. Note that 3% of a 0.9V supply is just 27 mV. As supply voltages drop and more cores are packed into processors, current levels rise, even exceeding 100A. Maintaining accuracy to a few tens of millivolts at the processor input with one hundred Amperes flowing through the power and ground planes is a daunting PDN (power distribution network) design task.

Simultaneously, there is a push for more efficient use of processing energy to lower data centre utility bills and cooling costs. Server chassis are running hotter with board temperatures approaching 100C. Design cycles are getting shorter but designs need to be refined at the last minute depending on margin test results and also for the unique needs of different markets and customers. Sequencing has been a common requirement on boards with multiple supplies, but those requirements are getting more complex as the number of supplies ranges from 20 to 50, spanning multiple power domains.

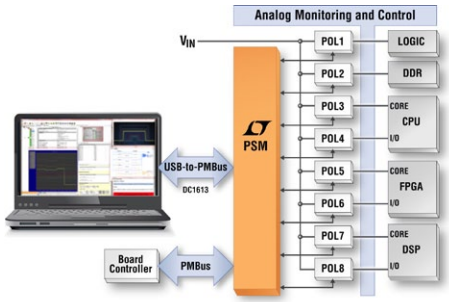

Figure 1. *The DPSM architecture. POL supplies stay analogue with digital communication and control added in.*

Solutions so far

Power management tasks such as sequencing, supervision, monitoring and margining have been handled by an assortment of devices including supervisors, sequencers, ADCs, DACs, amplifiers and microcontrollers. Coordinating these disparate devices to work together takes up most of the design effort. Integrated solutions have generally evolved or descended from supervisors and sequencers with capabilities added for margining, ADC monitoring and EEPROM fault logging.

However, these devices have poor voltage accuracy on trimming, margining and monitoring. There are also system-on-chip (SoC) devices that integrate an array of uncommitted digital gates and logic with ADCs, DACs, comparators and PWM outputs. Lacking any power management architecture, these devices require a lot of programming to perform even the most basic tasks, taking up months of design and validation effort.

The push towards digital management of power systems has led to digital power solutions where the DC/DC converter loop employs an ADC, digital compensator and digital PWM. Due to the inherent quantisation of this sampled system, digital loops generate more noise and ripple in the supply output voltage. They also tend to have slower transient response, poorer accuracy and, at worst, even erratic, unpredictable behaviour. On the other hand, analogue loops are faster, cleaner and much more predictable. Management of multiple supplies requires digital configuration and communication with the POL supply but the supply loop itself can stay analogue to obtain the best of both the analogue and digital worlds.

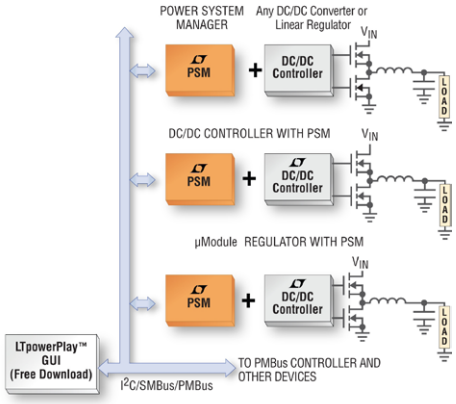

Figure 2. *DPSM device types, each offering multiple devices, depending on the number of supply rails to be controlled.*

The complete solution

Keeping POL supply trends in mind was a major part of the rationale that lay behind a complete Digital Power System Management (DPSM) solution, architected starting from first principles. The core philosophy is that the supply loops stay analogue, with digital interface and control added in. This is illustrated in Figure 1.

The DPSM family includes a broad array of interoperable devices with and without built-in DC/DC conversion, as shown in Figure 2. All of these devices communicate with a board controller via the industry standard PMBus interface. The choice of PMBus helps reduce design time by enabling firmware re-use. For those preferring autonomous operation without the need for code development, an engineering-level development software is provided to customise device configuration.

This article concludes with a description of the various digital management block functions; how fault logging becomes available; and the applicability to power environments with many power rails.

Complete article, here

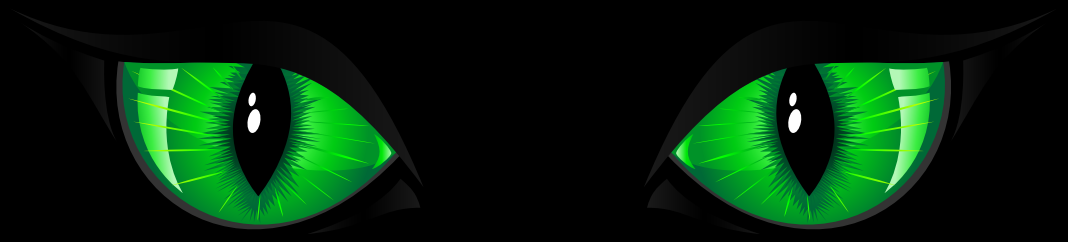

Eye on Standards

100G and things to COM

By Ransom Stephens

FEC (forward error correction) changes the way we think about BER (bit error ratio/rate) in serial data links. If we can afford seven errors in every set of 528 bits (i.e., for the Reed-Solomon coding used in 100 GbE, RS(528,514)), then we need a new way to characterise those network elements that cause errors.

In June (2014), IEEE released the 100 GbE (gigabit Ethernet) specification, 802.11bj. 100GbE> accomplishes 100 Gbits/sec by combining four 25.78125 Gbit/sec lanes; the excess data rate accommodates overhead—we'll call it 4x25 Gbits/sec. Previous releases had neglected specification details for electrical signalling on cables and backplanes. That is, fibre optic transmission had been wholly spec'ed years ago but we just got specs for these blistering data rates on PCB (printed circuit board) and cables.

The backplane spec (called 100GBASE-KR4 in the spec) accommodates 4x25 Gbits/sec over backplanes with total insertion loss up to 35 dB at 12.9 GHz, about half the data rate, and for cables (100GBASE-CR4), 4x25 Gbits/sec over shielded balanced copper cabling with reach of at least 5m. Notice that backplanes are specified in terms of their loss and cables in terms of distance of propagation.

The spec has some new concepts that open up design options but can be confusing; for example, channel operating margin. COM (channel operating margin) provides flexible choices instead of straight-forward pass/fail criteria or masks for s-parameters. The idea is to let engineers trade off between loss, reflections, and crosstalk. That last signal impairment, crosstalk, is at the root of most of this evil.

COM is the ratio of the signal amplitude to the noise amplitude given in decibels:

$$\mathrm{COM} = 20\log\frac{A_{signal}}{A_{Noise}}$$

and the spec requires COM >3 dB. Well, "require" is too strong a word. The COM spec is "informative" rather than "normative," which translates to "you don't have to meet this spec, but if you don't, you had better have all your excuses ready ahead of time."

It looks innocent enough, right? Measure COM, require it to be larger than a specified threshold and move along. But there's a devil in the denominator. The noise term has to include crosstalk, random noise and jitter (RJ), inter-symbol interference (ISI), and all the other acronymed culprits

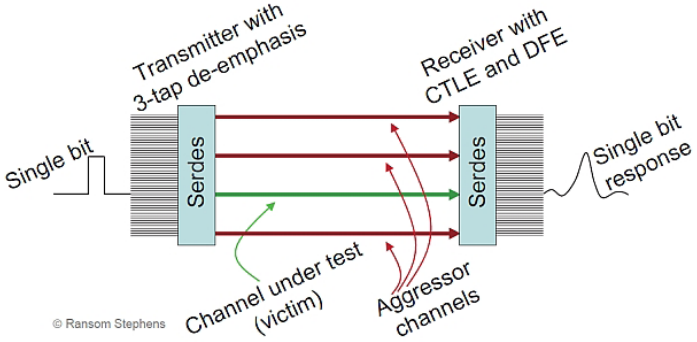

Figure 1. *COM test setup (Ransom Stephens).*

Specific, though hardly restrictive, models are used to extract COM. Three tap de-emphasis at the transmitter, called feed forward equalisation in the literature, and both CTLE (continuous time linear equalisation) and DFE (decision feedback equalisation) at the receiver. To specify the intrinsic noise of a channel, it makes sense to use s-parameters or, equivalently, the impulse response of the channel, including its response to all three other crosstalk aggressors. Instead of worrying over the complication of constructing an actual impulse with infinite amplitude and infinitesimal width, a single bit-wide pulse is used; that is, instead of impulse response, COM uses the response of the channel to a single bit, the SBR (single bit response).

The noise term is extracted by using a channel's SBR, including the aggressor SBR, combined with the best available transmitter and CTLE receiver equalisation. COM is then the ratio of the signal voltage at eye centre to the RMS (root mean square) noise after the CTLE.

Perhaps a more intuitive description is to derive the noise by combining the RMS noise of the transmitter, ISI, RJ, crosstalk, and amplitude noise something like this:

$$\mathrm{COM} = 10\log\frac{A_{signal}^2}{\sigma_{Tx}^2+\sigma_{ISI}^2+\sigma_{Jitter}^2+\sigma_{Xtalk}^2+\sigma_{AmpNoise}^2}$$

COM >3 dB is required; the design flexibility lies in how you divide up the contributions from ISI, crosstalk, and transmitter noise.

BY IAN BEAVERS & JEFFREY UGALDE, ANALOG DEVICES

DESIGNING JESD204B CONVERTER SYSTEMS FOR LOW BER, PART 1

Many real world sampling systems, such as test and measurement equipment, cannot tolerate a high rate of analogue to digital (ADC) or digital to analogue (DAC) processing errors. These systems essentially require a converter with a low bit error rate (BER). Historically, the conversion error rate of the core converter alone is what has dominated the overall BER. However, with the adoption of the new high speed serial digital interface link, known as JESD204B, between converters and FPGAs, the error rate within the digital transmission line cannot be ignored as a potential contributor to the overall BER.

As a third generation serial interface standard, JESD204B data rates top out at 12.5 Gbps per lane and can span multiple lanes per link, with the next generation of the specification pushing beyond 12.5 Gbps. The speed and quantity of data that is sent per link requires careful attention to several key design criteria in order to minimise the overall BER and prevent the digital data communication from being the dominant error contributor of the converter system.

There are several active and passive elements to the converter JESD204B link that, when considered proactively, can effectively mitigate the digital serial data path as the main contributor to the overall BER. The dielectric material used within the PCB, the inter-symbol interference (ISI) created by the lane layout and the transmission line impedance are all passive elements of the system design that can impact the BER of the link. Conversely, active on-chip transmitter and receiver channel compensation techniques as well as automatic adaptation elements can also substantially improve the BER. Several robust measurement techniques exist to isolate and measure the BER of just the JESD204B link, outside the converter core.

To understand, improve and reduce any impact of the JESD204B link BER on the converter, we answer some common questions from JESD204B system design engineers. Firstly, how is JESD204B link BER determined in the first place?

Determining BER for JESD204B links

ADCs and DACs for instrumentation systems often need to meet a strict error rate standard. The quality of many instrumentation systems can be determined by the error rate occurrence. While a typical conversion error rate can be sought within a converter's analogue core, the digital data link needs to have a better error rate in order to prevent being the dominant contributor. Therefore, not only does the analogue conversion error rate need to be measured, but also the BER of the digital link transmission.

BER in a serial or parallel digital data transmission is the ratio of the number of detected errors at the receiver divided by the total number of bits sampled. BER testing in a digital data stream implements a long pseudo random sequence that is started within a transmitter using a common seed value at both ends of the transmission. The pseudo random pattern should ideally have a long non-repetitive sequence to test as many digital combinations as possible. However, real pseudo random patterns have a repetitive sequence that can change with different seed values. The receiver will also know the seed value of the sequence and have the expectation of an ideal transmission. The BER is precisely calculated by observing the difference in the received data compared to the ideal pattern. Mis-matches in the pseudo random sequence data, based on the seed value, between both ends are counted as bit errors.

A bit error in a high speed serial interface such as JESD204B can occur due to one of the following scenarios:

- Timing jitter causing an edge transition to occur inside the sampling region due to conductor loss and dielectric loss attenuating low and high frequency signal components differently.
- Amplitude noise causing an incorrect voltage level Vhigh or Vlow capture due to conductor loss and dielectric loss attenuating low and high frequency signal components differently

In addition to the key analogue performance metrics of an ADC or DAC, key qualitative plots are also used in a converter datasheet to denote the performance of the SERDES link. Some of the most commonly used graphs are the data eye diagram, the 'bathtub' curve of error rate vs. unit interval, and a time interval error histogram.

The data eye diagram for the output of an ADC shows the cumulative persistence magnitude of the output signal swing into a defined transmission line load. It is often compared to the mask of a specified 'keep-out' area that defines the remainder of the cumulative margin available for the channel loss, inter-symbol interference and receiver clock and data recovery.

A bit error rate plot vs. the unit interval comparison size of the bit stream is commonly referred to as a 'bathtub curve'. This is because the slope of the plot resembles that of a cross-section of an actual bathtub: *an example follows in the online continuation of this article. It defines the available portion of the unit interval (UI) that can* reliably be sampled by the receiver in order to achieve a given bit error rate.

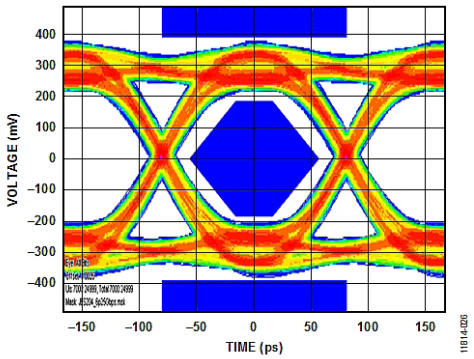

Figure 1. *A Typical JESD data eye diagram.*

This article continues with further observations on the "bathtub" curve; poses the question, "Can the JESD204B link BER be measured mutually exclusive of the normal converter operation?"; and investigates detailed effects such as the contribution of different PCB substrate dielectric materials to BER.

Complete article, here

Monolithic PWM generator runs fast, minimises silicon

Jindrich Windels , Ann Monte & Jan Doutreloigne

In IC design, the typical approach for generating a pulse width-modulated (PWM) signal from an analogue value (such as in the control loop for a switch mode power supply) uses a ramp or triangle wave fed to one input of a comparator, and the analogue control voltage connected to the second input of the comparator. This circuit, shown in Figure 1, is time-tested and can operate at high frequencies, but you need an amplifier and comparator (or two amplifiers) to generate the PWM signal. Alternatively, a pure digital approach is possible, where a free-running counter is compared with a register value, although the maximum PWM frequency is limited by the need for a clock signal several orders of magnitude over the desired PWM frequency – for acceptable resolution.

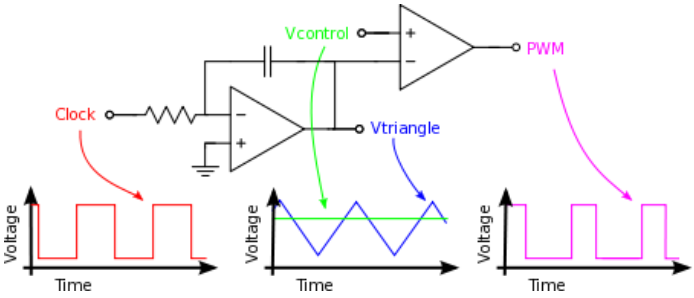

Figure 1. *Analogue PWM circuit using triangle wave generator and comparator*

In this Design Idea, we show an alternative circuit. It uses very little power, can be implemented in a small silicon area, and can achieve high PWM frequencies. In the design, shown in Figure 2, the PWM signal is generated based on an analogue control voltage and a clock signal at the PWM frequency. The analogue control voltage is used to modulate the capacitive load in a delay line – using a FET as a control element – which alters the delay, and hence the pulse width.

This simple yet powerful circuit is extremely compact to integrate on a chip. For example, a 10 MHz version of the circuit, including a driver for a 10 pF capacitive load in the ON Semiconductor C035 0.35 µm technology, can occupy as little as 120 µm × 70 µm.

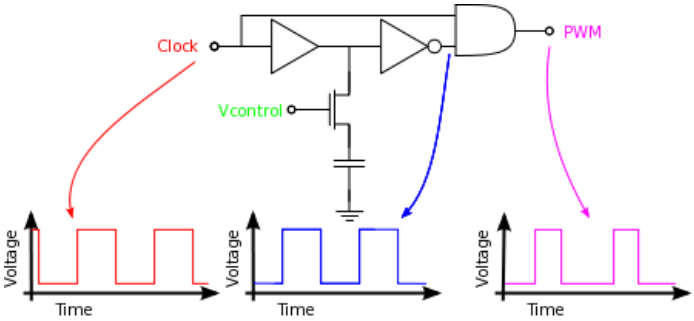

Figure 2. *Proposed PWM generator circuit*

The power consumption of the IC block is low, and mostly dependent on the load. For example, when operating at 3V, power consumption in simulation is approximately 1.057 mW at 10 MHz with a 10 pF load, and 248 µW with a 1 pF load, so most of the energy is used for charging the input of the next stage or the measurement probe.

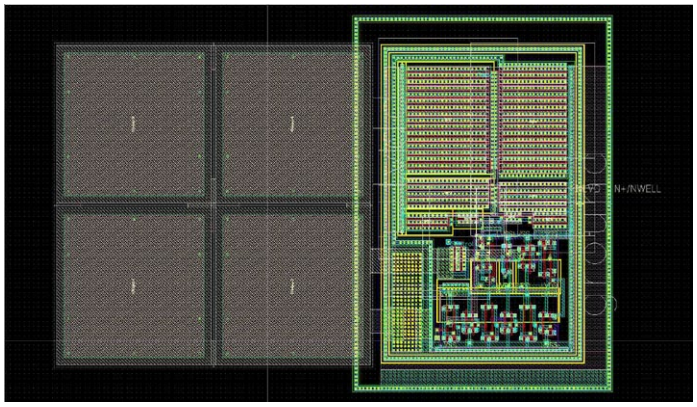

Figure 3. *Layout plot of the PWM block. In the actual chip, the block is covered with metal layers for interconnections, so a picture shows little of interest.*

The circuit was manufactured as part of a larger design through the Europractice Multi-Project Wafer (MPW) service, and the functionality was verified through measurements. In Figure 4, the pulse widths in simulation and measurement are plotted. Unfortunately, the power consumption of the PWM circuit could not be verified experimentally because other circuits use the same power rail.

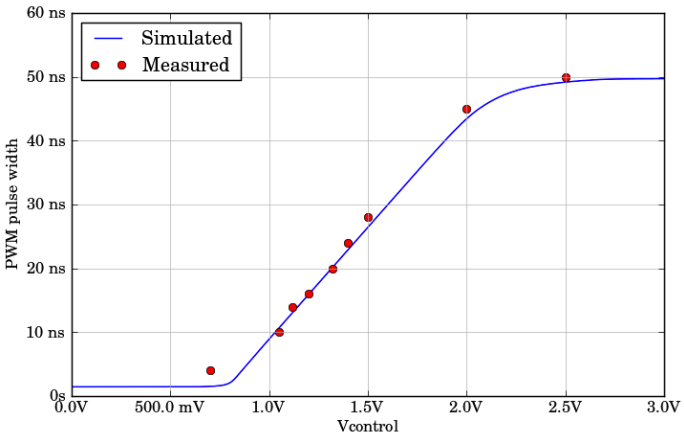

Figure 4. *Simulated and measured pulse width for the integrated PWM generator*

Flyback switcher works down to 1.1V, flashes HBLEDs

R O Ocaya

This Design Idea shows a way to drive low-power electronic circuits using a single 1.5V cell. The design is based on a free-running oscillator that drives a flyback transformer to generate a controllable higher voltage. It can be used to power analogue circuitry, microcontrollers, and any other light loads.

The power circuit was designed, simulated, and constructed. It was shown to operate reliably with a power draw of less than 50 mW, even when the supply drops to 1.1V. The regulated outputs tested for the given circuit values were between 6V and 24V by the adjustment of a single resistor value. The output power is sufficient to drive a PIC microcontroller in low power mode (15 µA @ 32 kHz). With no modification, the circuit will work as a strobe, and flash a string of LEDs – or a power LED – at rates from 0.1 Hz to 20 Hz. It might seem a lot of circuitry for otherwise "simple" tasks, but it is worthwhile given the low supply voltage, and the fact that the output can be regulated or the flashing controlled.

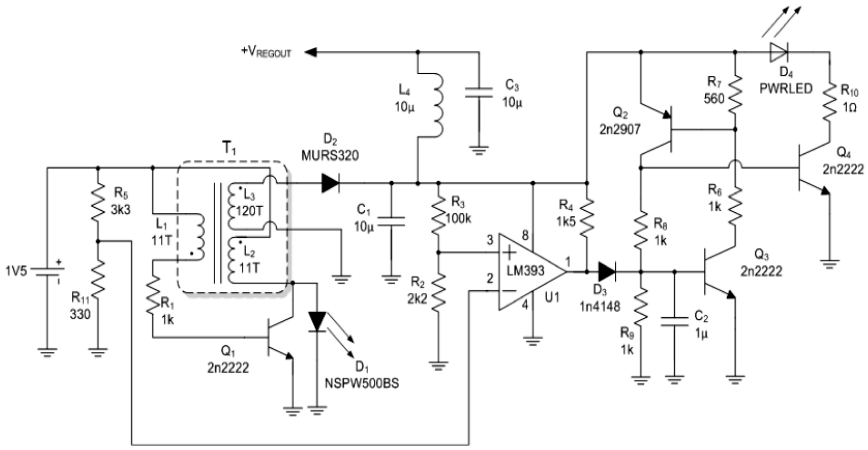

Figure 1 *Circuit diagram of the 1.5V inverter. Inductors L1 and L2 (10 turns, 22 AWG), and L3 (130 turns 32 AWG) are wound on a Fair-Rite Products Corporation core part number 5961001801.*

To use the circuit for a regulated output at $+V_{REGOUT}$ in Figure 1, for instance 20.7V in the following discussion, resistor R2 was set to 680Ω and R3 to 100 kΩ. The voltage across R11 is set by the R5/R11 potential divider to about 140 mV. The oscillator made up of Q1 runs continuously on application of power, and couples energy magnetically into L3 from L2. Once the voltage across C1 rises above 2V, comparator U1 is effectively powered. Initially, the voltage across R2 rises gradually as the capacitor charges up, until it just exceeds the voltage across R11. The connection is such that it will be compared constantly with the 140 mV by the open-collector comparator U1. The thyristor comprised of Q2/Q3 are similarly unpowered and also untriggered by the output of U1. The capacitor continues to charge unhindered. Therefore, depending on the value of the R2-R3 combination, the output of U1 will go high when:

$$\frac{R_2}{R_2+R_3}V_{C1} \approx \frac{R_{11}}{R_{11}+R_5}V_1$$

The core of T1 can be of different shapes and sizes. To duplicate the results above, a core having an inductance factor of around 80 nH and relative permeability (µ) around 80 is recommended. The ETD core from EPCOS, order number B66361G0100X1 and the toroidal core TN33/20/11-2P80 from Ferroxcube or similar are suitable.

For the component values in Figure 1 (with the exception of R2=680Ω and R3=100kΩ), the thyristor triggers when V_{C1}=20.7V. The output filter is made up of L4 & C3. C1 has a double role in the circuit: First, it is the charge storage "bucket"; second, it provides some loop stability together with R10. Resistor R4 is necessary to pull-up the open collector of U1, while D3 prevents upsetting the output bias of U1.

To use the circuit as a three-LED flasher/strobe, R10 should be 100Ω. For the component values shown, the thyristor triggers when V_{C1} = 6.33V. As soon as the thyristor has triggered, a pulse of about 2V is applied to the base of Q4, turning it fully on. This causes a large Q4 collector current, which causes a quick discharge of C1. If this collector current occurs through a series of LEDs or a single power LED, it gradually brightens, until a brilliant flash is seen at the moment Q4's base is driven hard. The discharge also turns off U1 (i.e., its supply voltage falls below 2V). Therefore, VR2 reaches a minimum of about 43 mV before charging of C1 resumes and the process repeats. When flashing LEDs, a limiting resistance is required (R10) as LEDs will be destroyed without it. The simulation shows peak current pulses of 3.3A with 50 µsec duration measured at 50%. The rate of flashing can be adjusted by changing R2 or R11.

Download the simulation file (saved as a plain text file; remove the .txt extension).

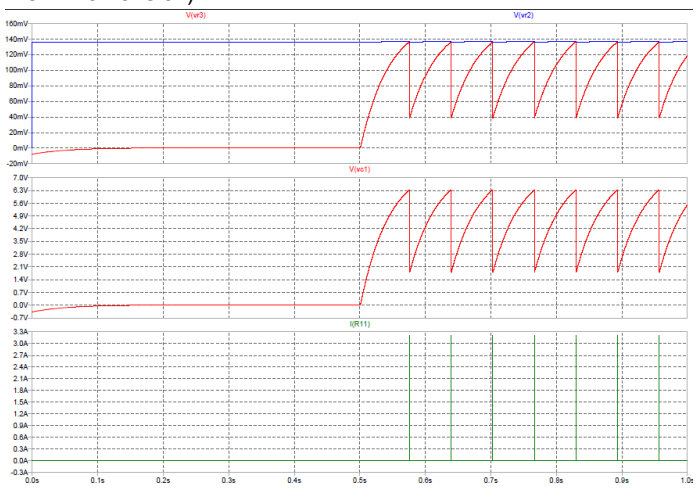

Figure 2 *LTspice-simulated timing waveforms in the LED flasher/strobe. The current in R10 is measured under the assumption that power LED D4 is not connected.*

Monitor 15 contacts with one PIC input

Benabadji Noureddine

Several previously published Design Ideas and appnotes show how to use many pushbuttons with a minimum number of inputs. They require an RC circuit where the timing can be measured to identify which pushbutton has been pressed, or an ADC input, with resistors forming a divider for each pushbutton pressed.

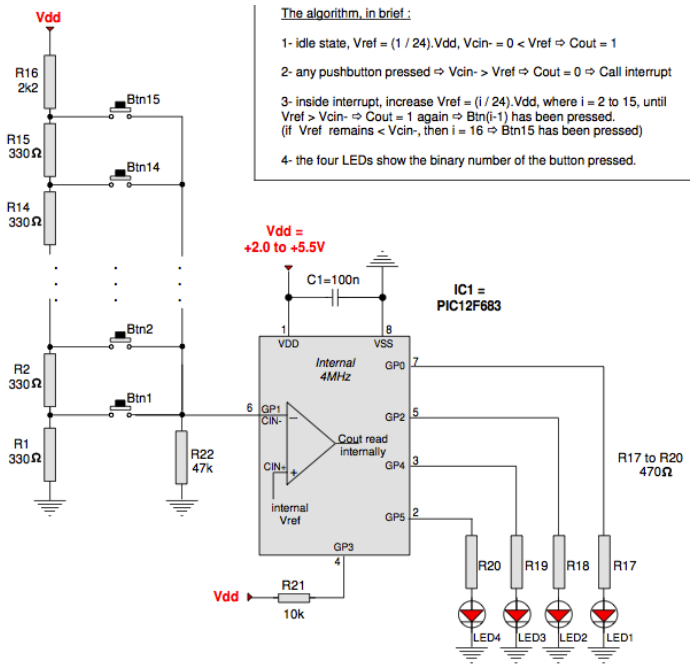

Figure 1. *Sense 15 buttons with one input*

The following Design Idea shows another simple way to use up to 15 pushbuttons with only one I/O. The microcontroller chosen must contain an internal comparator with selectable values for the internal voltage reference VREF. The majority of 8-bit PIC microcontrollers in the mid-range family feature one or two internal comparators with just this.

As a simple and practical example, the circuit in Figure 1 uses an 8-bit PIC controlling a 4-bit LED display to show which button has been pressed. A comparator interrupt routine reacts immediately to any pushbutton pressed. The routine cycles through V_{REF} values until the comparator output C_{OUT} returns high, which will indicate the pushbutton pressed.

The fully commented assembler source code takes less than 100 words. It is not overly optimised, so as to simplify understanding and make translation to other controllers easy.

References:

[1] Bernhard Linke: Sense multiple pushbuttons using only two wires
[2] Petre Petrov: Measure two resistive sensors or multiple switches with a single Schmitt
[3] AN234: Hardware techniques for PICmicro microcontrollers (MICROCHIP | 2003)
[4] Tips'n tricks for 8-pin PICmicro microcontrollers (MICROCHIP | 2003

productroundup

Integrated gate drivers with on-chip galvanic isolation

The STGAP1S advanced single-channel gate driver from STMicroelectronics integrates galvanic isolation with analogue and logic circuitry in the same chip to help simplify driver design while ensuring high noise immunity for safe and reliable power control.

STGAP1S is the first in ST's new generation of gapDRIVE gate drivers, which combine proprietary bipolar-CMOS-DMOS (BCD) process technology with an isolation layer grown on-chip to allow greater system integration. Up to 1500V can be present on the high-voltage rail without interfering with other circuitry, ensuring a level of robustness that makes this device suitable for use in industrial drives, high-power 600V or 1200V inverters, solar inverters, and uninterruptible power supplies. With signal-propagation delay of 100 nsec across the isolation layer, the STGAP1S is capable of transmitting high-accuracy PWM signals. The integrated driver stage can sink or source up to 5A.

Complete article, here

FPGA-based functional safety boards, IP and reference designs

Altera FPGAs and IP, with a functional safety development board from NewTec, speed time to market for IEC61508 functional safety applications: reference designs reduce SIL 3 development and certification costs for industrial designs. Altera has announced a functional safety development board and FPGA reference designs through NewTec, provider of safety-related systems. The board and reference designs, coupled with Altera's TÜV Rheinland-qualified FPGAs, tools, and IP, enable users to accelerate development and certification of FPGA-based system designs. Altera's functional safety data package includes device support, enhanced software design flows and IP for functional safety system design, allowing designers using Cyclone FPGAs to reduce risks and cost-effectively meet safety-critical system requirements across product lines. The functional safety board and accompanying reference designs are designed by Altera and NewTec to reduce customer design effort for safety designs requiring IEC 61508 certification up to SIL 3 (safety integrity level 3).

Complete article, here

J-Link on-board adds Drag & Drop applications programming

Segger has introduced a Drag & Drop "intuitive programming" feature to its J-Link OB single-chip on-board debug probe. A J-Link OB with Drag & Drop capability appears both as a debug probe and as a flash drive on the user's computer. The J-Link's performance and functionality for programming and debugging remain unchanged; the additional flash drive capability simplifies programming the target to dragging a Motorola S-record, Intel Hex, or plain binary file onto the J-Link Drive using the desktop. J-Link Drive enables manufacturers, and third-party software developers, to deploy ready-to-go demonstration applications to customers who can then quickly evaluate the capability of the board, software, and device—all without installing complex development software.

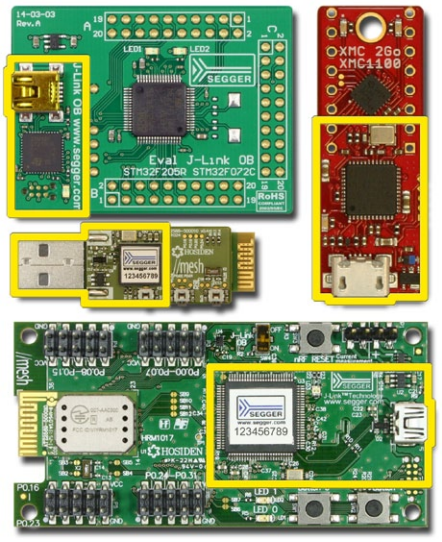

Complete article, here

120A from a TO-247 package; TO-247PLUS for IGBTs

Infineon's TO-247PLUS package carries up to a 120A IGBT co-packed with a full rated diode in the same footprint and pin-out as JEDEC standard TO-247-3. The TO-247PLUS can be used in industrial applications such as UPS, welding, solar, industrial drives as well as automotive applications such as powertrain inverter to upgrade existing designs for higher power output or to improve the thermal conditions in the application. The higher current capability of the TO-247PLUS allows for reducing the number of devices in parallel thus enabling more compact product designs. Due to the absence of the mounting hole the TO-247PLUS package may accommodate a 70% larger silicon die area compared to standard TO-247. Additionally, the 26% larger thermal pad area contributes up to a 20% lower thermal resistance Rth (jh) compared to standard TO-247.

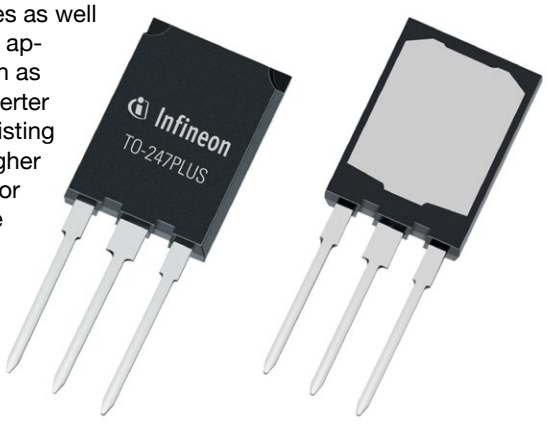

Complete article, here

Magnetically shielded power inductors

Coilcraft's XAL15xx rugged, magnetically-shielded power inductors are offered in two versions – one providing saturation current ratings up to 111A and the other offering inductance values as high as 33 µH. The XAL1510 Series offers a combination of high inductance values (up to 33 µH) and current rating as high as 39.0A, in a package that measures 15.4 x 16.4 mm and has a maximum height of 11.0 mm. For higher current applications, the XAL1580 Series offers Isat ratings up to 111A, with a 15.2 x 16.2 mm footprint and maximum height of 8.2 mm. Both series offer very low DCR - down to 0.70 mΩ for the XAL1580 - and their soft saturation characteristics make them suitable for VRM/VRD applications. They feature RoHS compliant tin-silver (96.5/3.5) over copper terminations and offer a maximum reflow temperature of 260°C. XAL1510 and XAL1580 Series feature a proprietary composite core that is not subject to thermal ageing issues.

Complete article, here

Added integration in ST's latest ARM Cortex-M0 MCU

STMicroelectronics' STM32F091 ARM Cortex-M0 MCU overcomes the resource constraints imposed by similar devices positioned for cost-sensitive applications by providing large on-chip memories and up to eight USARTs that save multiplexing communication ports in applications such as in-car audio or three-phase power metering. With up to 256 kBytes of flash on-chip and a 32 kByte SRAM - enough to implement a Java stack - the STM32F091 can host applications that deliver user experiences normally associated with larger or more expensive devices. All variants, including those in economical 48-pin packages, have at least six USARTs. Three of the USARTs are able to support smartcard, LIN, IrDA, and Modbus modes. The device has a flexible power architecture that allows analogue circuitry, including the 12-bit 1.0 µsec ADC and two-channel 12-bit DAC, to be operated at up to 3.6V for maximum dynamic range.

Complete article, here

Warm to cool white colour-temperature tunable LEDs

LED manufacturer Everlight has introduced what the company claims to be the world's first colour-temperature tunable LEDs in a simple COB (chip on board) package. After brightness dimming, tunable colour temperature is a feature that allows end users to tune the warmth of the light they receive. Typically, this feature was implemented through the use of multiple LEDs binned from cool white to warm white, behind a diffuser. With its CHI3030 27V/29W series, Everlight claims to have a very compact solution, with LEDs packaged behind concentric layers of phosphors offering different colour temperatures of white. Depending on how much of warm white or cool white you choose to light up, you can get a precise colour-temperature mix. A multichip solution, the CHI3030 is the largest such colour-temperature tunable COB LED, measuring 30 x 30 mm and drawing 29W at 27V.

Complete article, here

PC software for EMI pre-compliance testing on a €4000 budget

Rigol has introduced PC-based software that automates standard EMI precompliance tests. Used with a Rigol DSA800 or 1000 spectrum analyser, the software forms an economical test system for EMC parameter measurements. The simplest form of pre-compliance measuring of radiated emissions can be performed with a Spectrum Analyser such as Rigol's DSA815 with its bandwidth of 1.5 GHz, and near field probes for E-fields (electrical) or H-fields (magnetic). All Rigol Spectrum Analysers can be equipped with EMI filters and Quasi Peak Detectors. Option DSA800-EMI is available for the models DSA815 (1.5 GHz), DSA832 (3.2 GHz) and DSA875 (7.5 GHz) and enables the EMI Filter and the Quasi Peak Detector. By combining the right measurement instruments with the software and standard accessories, it is possible to build a flexible and reasonably priced EMI Pre-Compliance test system. The basic investment remains in an affordable price range, Rigol says, below €4,000.

Complete article, here

7-GHz ADC drivers bring AC performance to DC-coupled applications

Texas Instruments has added two fully differential amplifiers (FDAs) that provide DC-coupled applications with best-in-class AC performance; the LMH3401 and LMH5401 FDAs provide higher bandwidth and slew rate, and lower distortion than existing ADC drivers. The LMH3401 provides 16 dB of gain with higher performance than similar FDAs, while providing a 30% reduction in both size and power consumption. It delivers 7 GHz of -3-dB bandwidth at 16-dB gain, a slew rate of 18,000 V/µsec, and harmonic distortion of -77 dBc at 500 MHz. The LMH5401 can be configured for 6 dB of gain or more, delivering 6.2 GHz of -3 dB bandwidth at 12-dB gain. It has a slew rate of 17,500 V/µsec, harmonic distortion of -90 dBc at 200 MHz, and IMD3 of -90 dBc at 200 MHz.

LMH3401 and LMH5401 deliver highest bandwidth and slew rate at low power; they also eliminate the need for baluns, converting a signal from single-ended to differential.

Complete article, here

Debug for Cortex-M: full instruction trace in real time

IAR Systems' I-jet Trace is a probe providing extensive debugging and trace functionality, delivering large trace memory capacities and high-speed communication via USB 3. The probe supports all ARM Cortex-M cores, including the ARM Cortex-M7 core, that have Embedded Trace Macrocell (ETM) capabilities. Tracing every single executed instruction, ETM provides developers with insight into the microcontroller's activities and enables them to find critical bugs that are difficult or even impossible to find any other way. I-jet Trace has a memory capacity of 32 Msamples. It allows real-time trace clocking at up to 150 MHz and Serial Wire Output (SWO) using Manchester and UART, clocking at up to 200 Mbps. The trace data is collected by the C-SPY Debugger in IAR Embedded Workbench and can be visualised and analysed in various windows.

Complete article, here

Integrated MCU/Wi-Fi module for IoT edge nodes

In the Atmel | SMART portfolio of SmartConnect low-power, secure Wi-Fi solutions, this FCC-certified module offers a complete standalone edge node solution that promises design flexibility and security for IoT developers. Positioned as the first fully integrated FCC-certified Wi-Fi module with a standalone MCU and hardware security from a single source, the SmartConnect SAM W25 module includes Atmel's 2.4 GHz IEEE 802.11b/g/n Wi-Fi WINC1500, along with an Atmel | SMART SAM D21 ARM Cortex M0+-based MCU and Atmel's ATECC108A optimised CryptoAuthentication engine with ultra-secure hardware-based key storage for secure connectivity.

The fully integrated SAM W25 delivers a secure 'plug and play' solution integrating wireless technologies with the design flexibility required for Internet of Things (IoT) designers.

Complete article, here

Software-controlled magnetic sensing gives HMI-design flexibility

Calling it a software-defined sensor, Melexis has introduced a magnetic sensing product that it describes as ultra-flexible, and that gives engineers great freedom in defining HMI implementations through its breadth of programmable parameters. Melexis is aiming this device beyond the automotive market to general sensing applications; it gives, the company says, three-axis sensing without a great deal of signal processing. You can access "raw" measurements from its SPI /I²C interface. This fully programmable, extremely compact sensor IC is capable of accurately measuring changes in magnetic flux density along its X, Y and Z axes. Based on the company's Triaxis technology, the MLX90393 provides "almost unlimited" scope with which all manner of human machine interface implementations can be accomplished - from joystick, slide switches, push/pull switches, levers, linear swipe switches and rotary knobs, right through to complex 3D position sensing systems.

Complete article, here

Wideband IF receiver chip provides highest dynamic range

The AD6676 "IF subsystem-on-a-chip" uses bandpass sigma delta A/D converter technology to digitise IF signals between 70 and 450 MHz with bandwidth up to 160 MHz. The enabling technology behind the performance of the AD6676 is a highly programmable continuous-time bandpass sigma delta A/D converter (CTBPSD) topology which combines highly programmable analogue filtering functionality with high dynamic range conversion. The implementation in deep sub-micron CMOS allows the advantages of oversampling by providing digitised bandwidth of up to 160 MHz. This provides a combination of features that enables "reconfigurable radio" architectures for high-performance heterodyne systems. Dynamic range can be configured to achieve a noise figure of 13 dB, IIP3=36 dBM, and noise spectral density (NSD) as low as -160 dBFS/Hz.

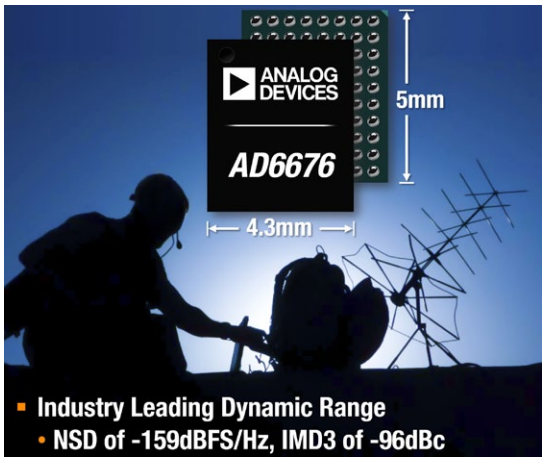

Complete article, here

Reinforced isolation meets highest industrial performance needs

Texas Instruments has introduced data converter and digital isolation device families that enable high levels of reinforced isolation ratings along with precision performance for high-voltage applications. Digital isolation and delta-sigma modulator families offer reinforced isolation to help protect electronic equipment from high line voltages. The ISO7842 family provides the highest immunity reinforced isolators and are claimed as first to offer an isolation barrier with a withstand-working breakdown voltage of 1,500 Vrms for a minimum lifetime of 40 years. The AMC1304 delta-sigma modulator family offers precision performance and the lowest power consumption with reinforced isolation for current shunt monitoring. The performance of the Δ-Σ modulator permits the use of smaller-value current measurement shunts in data-acquisition and motor-control boards, reducing losses and resistive heating. The ISO7842 family provides superior overall reinforced isolation breakdown ratings: 8,000 V maximum transient overvoltage, according to standard DIN V VDE 0884-10; or an isolation rating of 5,700 Vrms, according to standard UL 1577, surpassing typical ratings of 5,000 Vrms.The AMC1304 family provides precision with offset error maximum of 150 µV and gain error of ±0.3%, claimed as the industry's best DC performance.

Complete article, here

High stability MEMS accelerometer/ gyroscope for automotive, industrial uses

Murata's SCC2000 series of combined accelerometer and gyroscope sensor devices is aimed at automotive and industrial applications, with optimum temperature dependency, shock sensitivity and bias stability characteristics. It consists of a low-g 3-axis accelerometer with two angular rate sensor options of either X or Z-axis detection, together with a 32-bit digital SPI interface. The sensor has a software selectable 10 or 60 Hz low pass filter that can be configured via SPI. Gyro range is ±125 degrees per second with sensitivity of 50 LSB per degree per second. Typical accelerometer offset temperature drift is ±6 mg for the 2g sensor and ±12mg for the 6g version. Gyroscope offset temperature drift is typically in the range ±0.5°/s for the 125°/sec X & Z -axis product versions. The gyroscope has a typical offset short-term bias stability of 1°/hour for the 125°/sec X-axis device and 2°/hour for the 125°/sec Z-axis product.

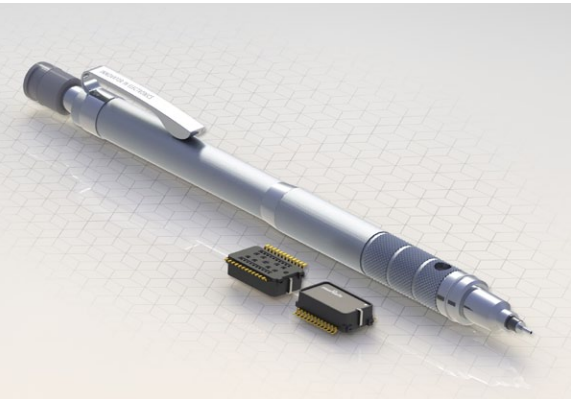

Complete article, here

High-integration Bluetooth Low Energy IC

Cypress has introduced two highly integrated, single-chip Bluetooth Low Energy products that it says will simplify design of low-power, connected, sensor-based systems. The PSoC 4 BLE Programmable System-on-Chip delivers is a customisable device for IoT applications, home automation, healthcare equipment, sports and fitness monitors, and other wearable smart products; the PRoC BLE Programmable Radio-on-Chip provides a turnkey solution for wireless Human Interface Devices (HIDs), remote controls and applications requiring pure play wireless connectivity. These Bluetooth Low Energy products enable complete systems by integrating a Bluetooth Smart radio, a 32-bit ARM Cortex-M0 core with ultra-low-power modes, programmable analogue blocks, and Cypress's CapSense capacitive touch-sensing. This combination of technology claims prolonged battery life, customisable sensing capabilities, and intuitive user interfaces.

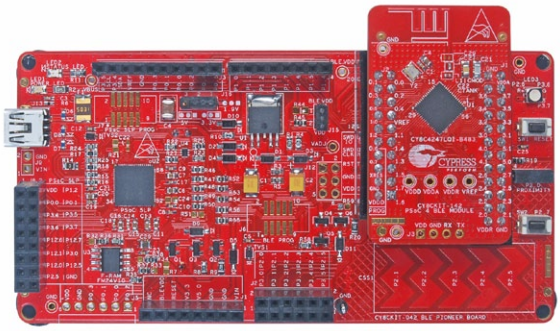

Complete article, here

Why HTTP won't work for IoT

Are you trying to fit devices into the Internet? If so, HTTP isn't going to work. There are much better options. In this short article I will introduce how protocols are used in IoT and analyse the current options.

HTTP is for the Old World Internet. The New World Internet (the IoT) is made up of unseen devices that require very little interaction. They consume very little power and frequently have poor network connectivity. HTTP is too heavy to be a good fit for these devices. An HTTP request requires a minimum of nine TCP packets, even more when you consider packet loss from poor connectivity, and plain text headers can get very verbose. And even with all this overhead HTTP doesn't guarantee message delivery!

The HTTP overhead also adds to IoT operating expenses. Current costs of wireless connectivity are exorbitant. My cellphone provider charges $10/GB, and I've seen as much as $1/MB! This is to say nothing of the wireless spectrum shortage. Bandwidth conservation is especially important with enterprise customers that often have hundreds of thousands or millions of devices deployed.

Fortunately there are suitable alternatives to HTTP for communicating with IoT devices.

IoT network traffic falls into two categories: telemetry and telecommand. Telemetry is the act of gathering telemetrics, or sending data over long distances. Usually telemetry involves sending data from many dumb sensors to a smart hub of some sort. The dual of telemetry is telecommand, or the act of sending commands across a network.

Most telemetry protocols are modelled as publish/subscribe architecture. Sensors connect to a broker and periodically publish their readings to a topic. A central cluster of servers (the cloud) will then subscribe to the topic and process sensor readings in real-time. A typical enterprise arrangement will have thousands or millions of sensors sending telemetrics to a handful of servers that split up the task of processing the data.

MQTT (Message Queue Telemetry Transport) is the leading telemetry protocol. It typically runs on top of TCP, adding only a 2-byte header. Usually clients are expected to have poor network connectivity -- either a fault of the wireless technology or because the device sleeps as a means of saving power. MQTT solves this problem by having the recipient acknowledge that it processed a message. This is called Quality of Service (QoS) and allows clients to opt out of this behaviour by using a lower QoS level if the normal guarantees of TCP are suitable.

The protocol allows the subscriber to encode information into the topic. Sensors will often publish to a topic such as "device123/temp" or "device456/temp"; consumers can use a wildcard and subscribe to "+/temp" to receive messages from all devices, yet still retain the device ID information separate from the message payload.

MQTT has a couple variants. First, you can run MQTT over SSL/TLS to create a secure connection that's safe against man-in-the-middle attacks. Second, MQTT-SN (MQTT for Sensor Networks) is made to run over a non-TCP/IP stack. This could mean using UDP, SMS, or a variety of packet-based wireless protocols. I believe MQTT-SN is promising, as recent results may indicate problems in certain TCP-based protocols over cellular networks.

MQTT provides a lot of functionality for such a low overhead. It's oriented toward devices with low bandwidth requirements. However, I think that the debugging experience is sacrificed. My main complaint is that there isn't an easy way to indicate to the client that it did something wrong. When topic security is enforced, the broker can acknowledge a SUBSCRIBE request, but it can't actively deny one (i.e., lack of permission). The only way some MQTT requests can be denied is by the broker not responding with an acknowlegement.

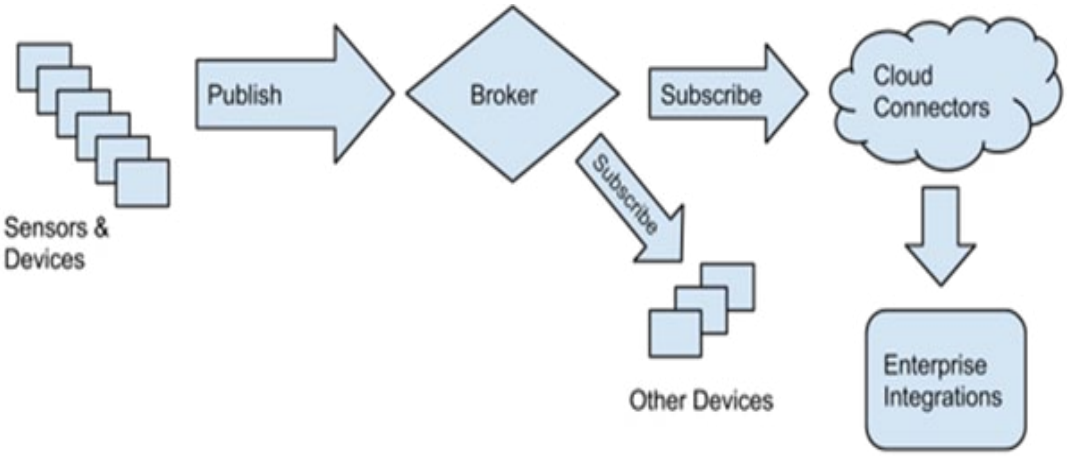

There are alternatives to MQTT: STOMP, XMPP, and ActiveMQ all are sometimes used for telemetry. They all meet M2M requirements to varying degrees, but I think MQTT and variants are superior for low-power devices and are most likely to become ubiquitous.

So, don't use HTTP for the IoT. HTTP was created for the "command ->" response required by web pages and applications. IoT has different requirements. Wireless bandwidth is placed at a premium. Connectivity can't be expected to be available all the time. Most importantly, architecture is much cleaner and easier to understand when publishers and subscribers don't have to be simultaneously available.

Tim Kellogg is a software engineer working on the platform, protocols, and cloud services at 2lemetry, a provider of IoT services.