EDN europe

APRIL 2015

www.edn-europe.com

Simulation predicts LED lighting effects

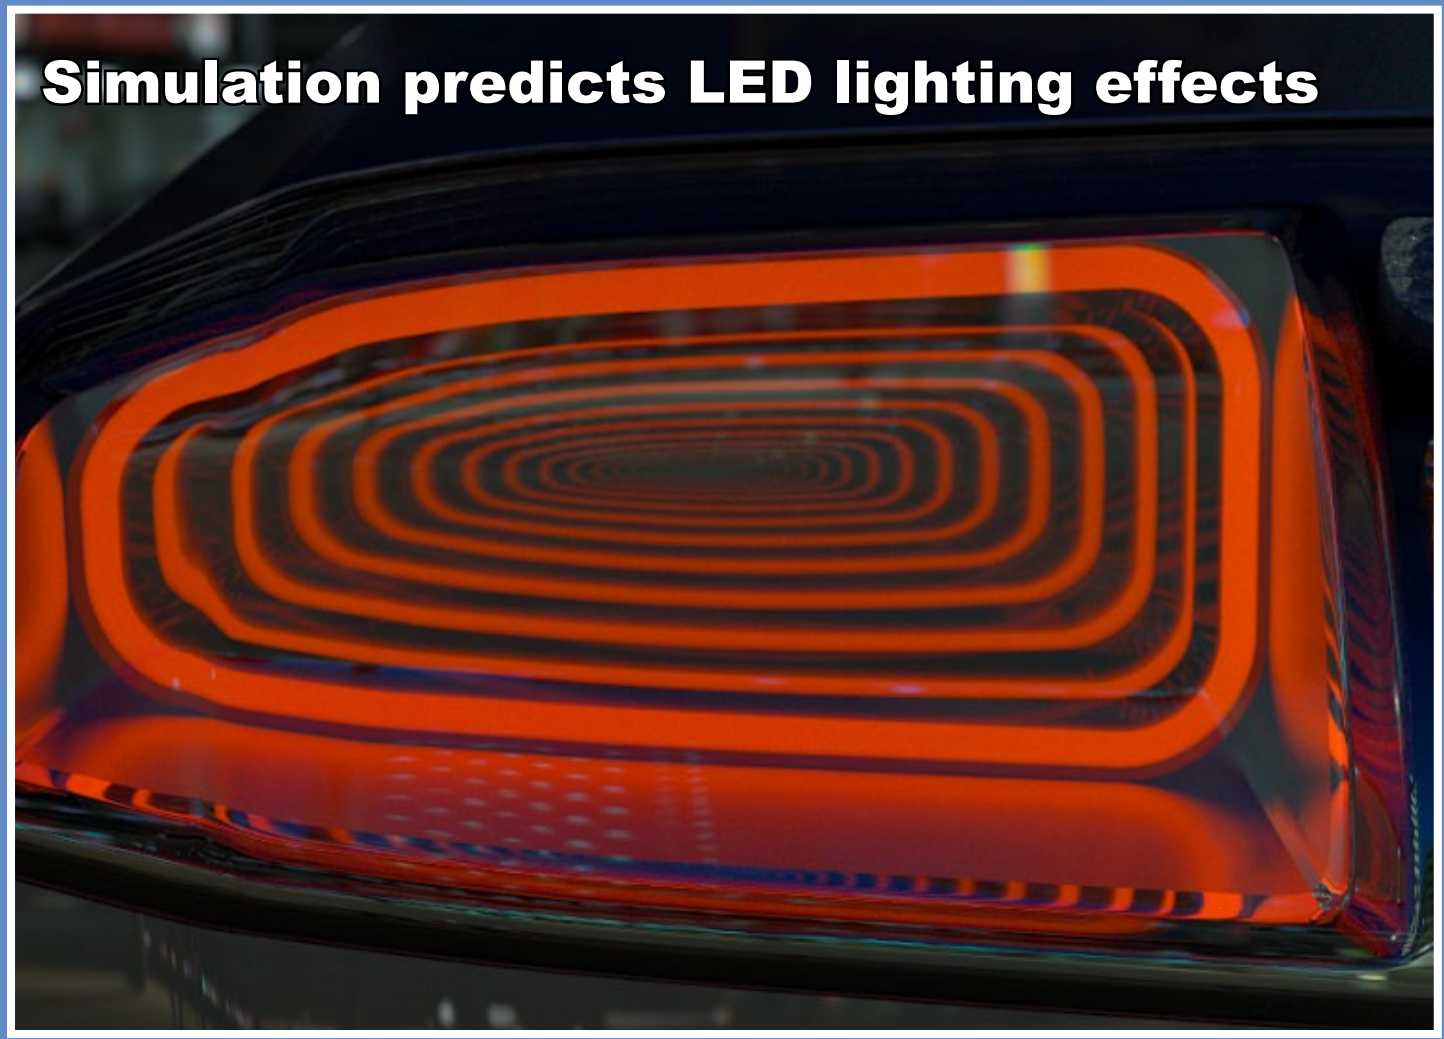

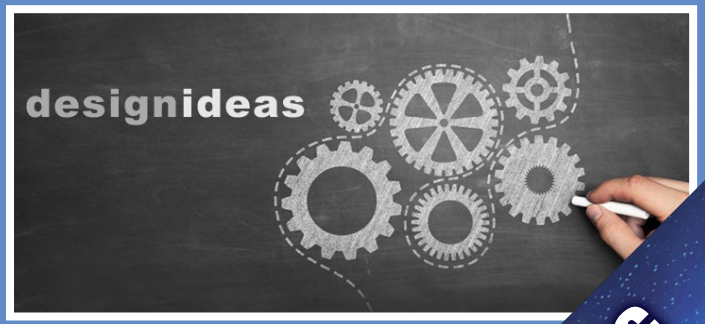

GLOBAL COVERAGE!

FREE SHIPPING ON ORDERS OVER €65!*

THE WORLD'S LARGEST SELECTION OF ELECTRONIC COMPONENTS

Available for Immediate Shipment!™

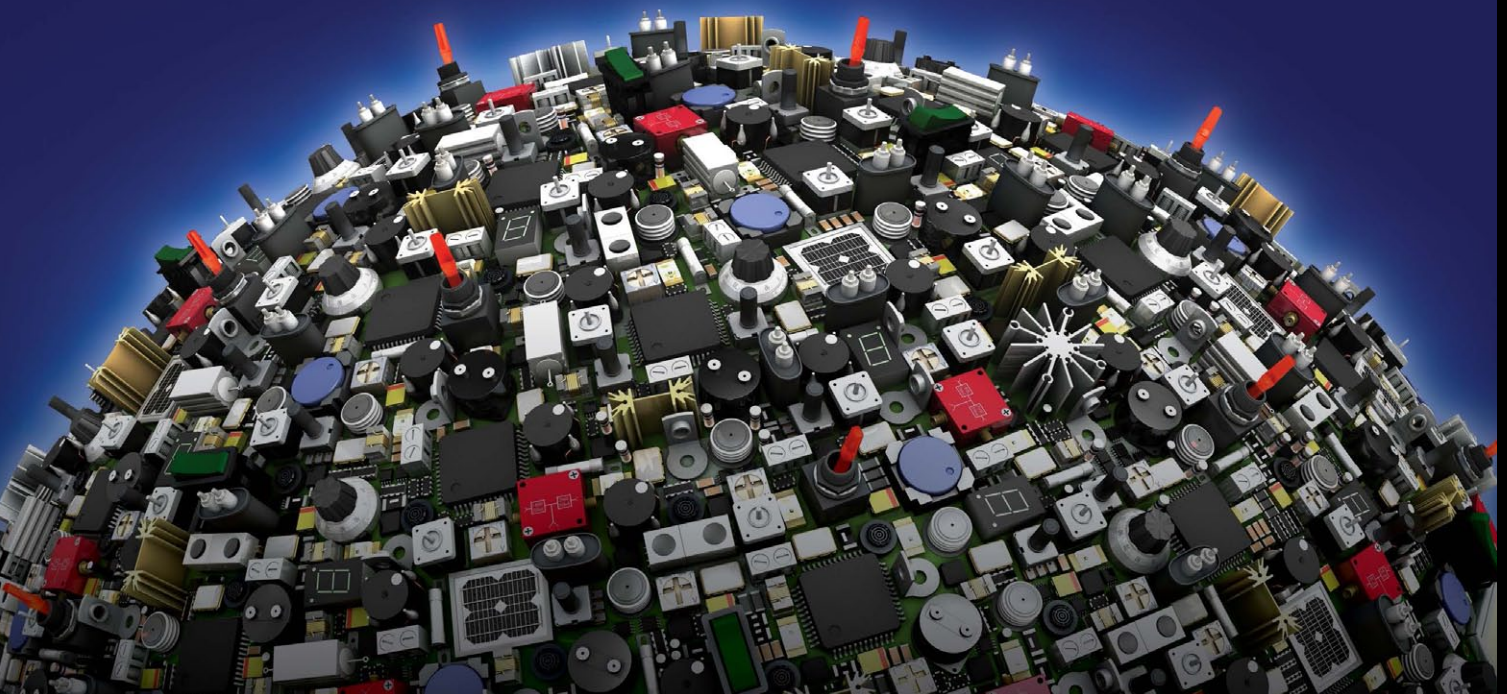

1,000,000+ PRODUCTS IN STOCK

3.9 MILLION PARTS ONLINE

650+ INDUSTRY-LEADING SUPPLIERS

100% AUTHORIZED DISTRIBUTOR

Digi-Key® ELECTRONICS

FIND CONTACT AND ORDERING INFORMATION FOR YOUR REGION AT

DIGIKEY.COM/EUROPE

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

ecia MEMBER ecsn member CEDA

COVER

Automotive lighting design software models perception

This image is an application of a software package from EDA vendor Synopsys; version 2.0 of Synopsys' LucidShape lighting design package enables more accurate model creation and contains a diagnostic tool for evaluating how a lighting system will be perceived by the human eye. This is particularly important, Synopsys says, for evaluating automotive components such as turn signals and brake lights.
The LucidShape Visualize Module provides photorealistic visualisations that help designers troubleshoot and improve automotive lighting designs at any stage in the development process. Its Ray Data Viewer allows designers to preview vendor light sources before building model geometry, and its Light Pipe Tool enables fast creation of complete light pipe geometry. An expanded material library with measured BSDF data delivers accuracy for modelling automotive lighting surfaces.
LucidShape's Visualize Module delivers high speed photorealistic images of an automotive lighting system's lit appearance, which demonstrate luminance effects when light sources in a model are illuminated.
(Full story here)

FEATUREARTICLES

ONLINE THIS MONTH

EDN's columns

HOW GOOD ARE YOUR MODELS?

The information that reaches you on these pages comes via a variety of channels; often, it reaches me in the form of a presentation from the designers or makers of the product, component, technology, software or concept. Most often, these days, such presentations are conducted remotely; Webex, Skype, Go-to-Meeting – all of them play their part, they all get the job done. In passing, I often remark that we (as an industry) have expended a vast amount of effort in expanding the bandwidth for such communication – video calls are readily available but in practice don't add much. The simple audio channel still carries the greatest part of the information. To put that another way, an old-fashioned phone call is often entirely adequate to the task.

Conversely, a proportion of the information is imparted to me in face-to-face presentations, and occasionally, a meeting organiser will select a high-profile (pun intended, read on) venue. So it was, recently, that I found myself on the 34th floor of London's Shard building, viewing the cityscape on a sunlit spring morning. (Sure, it's tough, but someone has to do it.) Engineering is never far away, though; on entering the room there was a considerable amount of wind noise, and a maintenance operative was summoned; with a hex key he was able to release a complete floor-to-ceiling glazed panel, swing it on its hinges, re-seat it on its seals, and restore tranquillity. It's not hard to imagine how this comes about. With glazing panels of over 3m height, dissimilar materials (glass, steel, alloy, polymer seals) and temperature cycling, it's unsurprising that joints might "creep" to a small extent. But (we say) did the architects and structural engineers not model this? Don't they have databases of material properties to predict such effects and design them out? Of course they do, but it's one more (literal) reminder of the gaps that exist everywhere in engineering between the simulated and the real.

My colleague Bill Schweber has been reflecting, in the North American edition of EDN, on another aspect of the way we use modelling tools; to reduce costs by reducing design margins. Once you have a model of the system you propose to build, you can pare it down to the bare minimum, reducing every part and component to the absolutely lowest-cost specification that still gets the job done. This is fine, Bill observes, if your model is entirely accurate. If it is not, then you have a part or component that might be outside margins from the start; then add manufacturing tolerances, and wear, and ageing effects, and abuse by the end-user... we know how it all too often ends up. Not to mention design revisions; over the years, EDN's Tales from the Cube column has carried many accounts of working designs rendered problematic by a component substitution that looked innocuous but took a design beyond its tolerances.

"In the end," Bill's take on the topic continues, "it's all about the input data and underlying assumptions you must make. In some ways, it's scary how much we rely on these models not only to validate designs, but to allow us to modify them with the goals of reducing cost, weight, power, and other factors.

"This thinking represents a substantial shift from design in the not-so-distant 'old days,' when models and tools such as we have now didn't exist or were not relied upon to such a degree. Instead, humility was combined with the mistakes learned via experience to force designers to add a little (or lot) of design "insurance" in the form of extra mass, current drive, or "whatever" in case the calculations used in the design did not fully reflect reality. We know there's a fuzzy border between a smart, minimalist design and one which is inadequate, either long term or even soon after deployment."

The consequences of going that small step beyond "finely tuned" or "cost optimised" may take some time to appear, and may range from uncomfortable meetings called to discuss a sudden upswing in warranty returns, to serious damage to brand values.

Faced with the slick user interface and powerful features of today's modelling tools, it's always possible to slip into the assumption that what we are looking at is 100% equivalent to reality; do you "add a bit for peace of mind"?

pulse

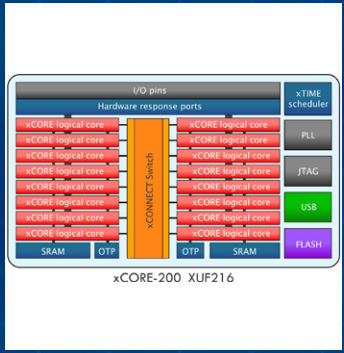

xCORE-200 XUF216

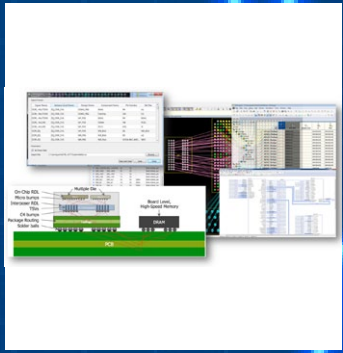

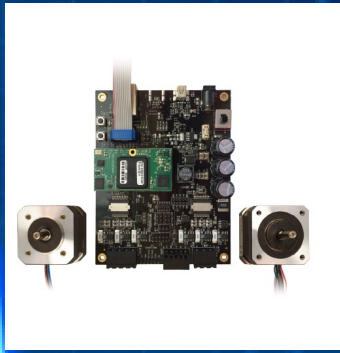

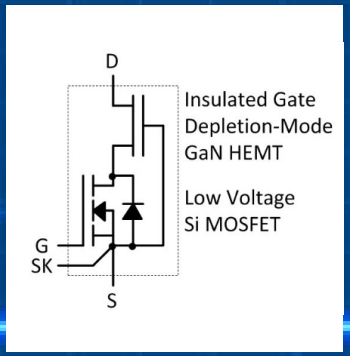

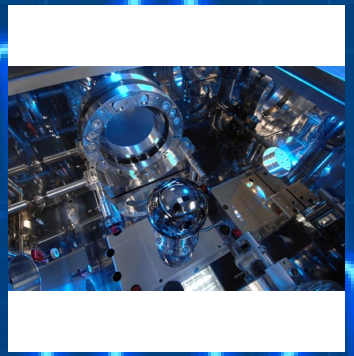

70 GHz real-time oscilloscope from Tek claims lowest noise, highest ENOB

Tektronix' claims for its DPO70000SX 70 GHz ATI Performance Oscilloscope include the lowest-noise and highest effective bits of any ultra-high bandwidth real-time oscilloscope available on the market. Tek is aiming it at high-speed coherent optical systems or any leading-edge research workers who need capture very high-speed – especially single-shot – events.

Other prospective purchasers of the unit will include those for whom the ability to capture the widest possible segment of bandwidth in one shot is a necessity. One of the first things you notice about the DPO70000SX is its compact format; in a reversal of recent trends to larger scope screens, this unit has only a small – if detailed – LCD screen. With this format, Tek recognises that most users will view the measurements on an attached monitor; and that at multi-GHz frequencies, short connections to the unit under test are essential and the benefits of the "compact" format will dominate.

Tek cites key features of the scopes as including;

- first 70 GHz real-time oscilloscope with Tektronix' Asynchronous Time Interleaving (ATI) technology that preserves signal-to-noise ratio for higher fidelity. This means, the company says, that users can more accurately capture and measure their signals at higher speeds than is possible with any other oscilloscope available today.
- 200 Gsample/sec sample rate with 5 psec/sample resolution for improved resolution and timing.
- compact form factor that enables the instrument to be positioned very close to the device under test (DUT) while providing flexible display and control for analysis. Reducing the distance to the DUT enables shorter cable lengths which in turn help to preserve signal fidelity at very high frequencies.

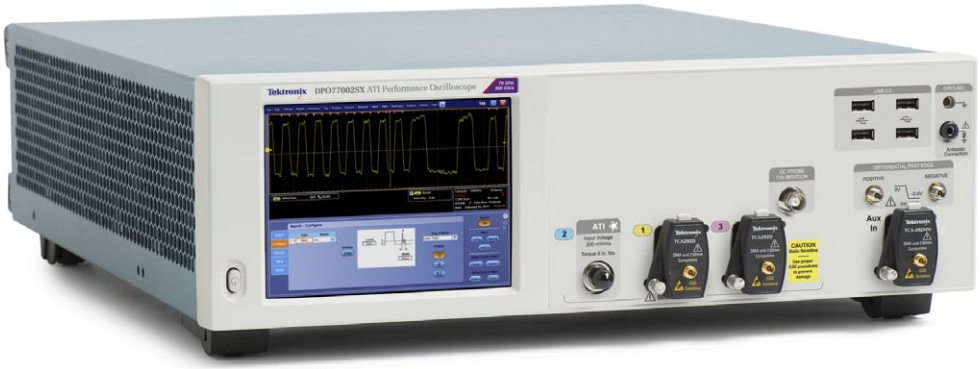

- scalable oscilloscope system with Tektronix' UltraSync architecture that provides precise data synchronisation and convenient operation of multi-unit systems. Precise multi-channel timing synchronisation is key to meeting acquisition requirements in applications such as 100G and faster coherent optical modulation analysis.

Tek's ATI approach to high-bandwidth digitising sends the incoming signal down two identical sampling channels, where they are digitised separately. Sample interleaving is in the time, rather than frequency domain; the full signal is digitised at a lower rate in both paths but the information to represent the full bandwidth of the incoming signal is still present. In an allusion to alternative approaches that split the incoming signal into two (or more) bands, then down-convert the higher band before digitising them separately, Tek asserts that there are issues of phase, amplitude, and other mismatches which are very difficult to overcome when the two bands are reassembled.

Precise multi-instrument timing synchronisation is required for high-speed coherent optical applications. UltraSync meets this requirement with a high-performance synchronisation and control bus that allows multiple DPO70000SX oscilloscopes to work together to deliver the precise high-speed multi-channel acquisitions needed for these ap-

plications.
DPO70000SX oscilloscopes come in a 5.25-in. tall form factor while retaining full high-performance functionality. This compact packaging means that engineers can stack two oscilloscopes in the same space as a standard bench oscilloscope.
The DPO70000SX also provides the flexibility to use one channel at 70 GHz, 200 Gsample/sec or two channels at 33 GHz, 100 Gsample/sec. Available dual unit systems synchronise two DPO70000SX units providing two channels at 70 GHz, 200 Gsample/sec each, or four channels at 33 GHz, 100 Gsample/sec operating as a single, coordinated instrument. DPO70000SX oscilloscopes are available worldwide with export restrictions for certain countries.
70 GHz ATI units start at €250,000 with a two-unit system starting at €381,000. A 33 GHz non-ATI unit is available starting at €226,000 with a two-unit system starting at €295,000.

Complete article, here

"Ideal chip for wearables" claims Dialog, for Bluetooth Smart IC

Dialog Semiconductor has designed what it terms the first Bluetooth Smart "Wearable-on-Chip"; wearables is the fastest-growing sector of the Bluetooth Smart market, the company says, and the configuration of the DA14680 SmartBond chip focusses on power saving and power management.
The IC integrates all the functions needed to create high performance wearable products with longest battery life. DA14680 includes the functionality to create a fully hosted wearable computing product, with an ARM Cortex-M0 core. It features flexible processing power, flash memory for virtually unlimited execution space, dedicated circuitry for sensor control, analogue and digital peripherals optimised for wearable products, and an advanced power management unit. DA14680 eliminates several external chips from wearable product design, Dialog says, allowing smaller form factors, lower system cost and lowest power consumption.
DA14680's ultra-low power 30 µA/MHz ARM Cortex-M0 application processor will run at a clock frequency of up to 96 MHz. Security features include a dedicated hardware crypto engine with elliptic curve cryptology (ECC), delivering

end-to-end banking-level encryption, ensuring personal data security. The device integrates 8 Mbit flash memory, audio support with PDM and I²S/PCM interfaces, two separate I²C and SPI buses, three white LED drivers, a temperature sensor, multi-channel DMA, and an 8-channel, 10-bit ADC. Intelligent power management, including system power rails and a Li-ion/LiPo battery charger and fuel gauge are also on-chip.
Although not having sensor fabrication in-house, Dialog says it addresses all of the key aspects of wearable design; RF, application processing with the resources to "pre-digest" data, and sensor fusion. Running the M0 core at up to 96 MHz gives the power to handle arrays of sensors; the DA14680 is also a suitable choice for smart-home designs. Dialog sees the M0 core, in a low-power low-leakage technology, but run at the higher clock rate, as the optimum choice of processor.
Dialog cites its RF experience in the product's link budget specification, and the inclusion of a balun on-chip; you can, "put a paperclip on the output [as an antenna]" a spokesman added, noting, "the challenges are in the software and sensors – the RF [with this degree of integration] is straightforward. The DA14680 will be sampling in the second quarter of 2015.

Complete article, here

XMOS boosts performance & connectivity with 2nd-generation multicore MCUs

Multicore microcontroller maker Xmos has refreshed its product line with two series; xCORE-200 is a general purpose series that steps up performance, memory and I/O count over its predecessors, and xCORE-Audio is a family of devices specifically optimised for audio applications.

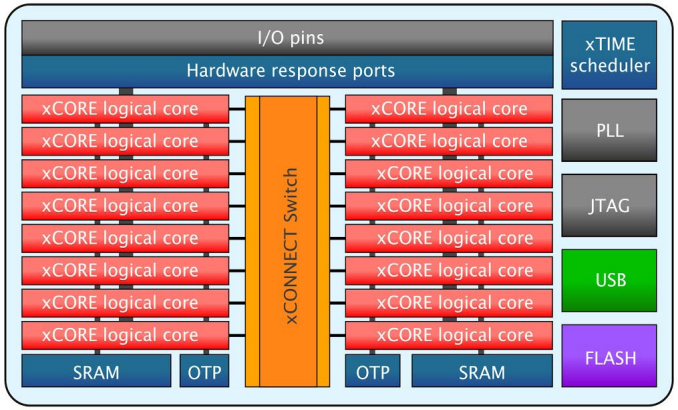

xCORE-200 XUF216

The latter family are "Hi-resolution Audio connectivity processors", which Xmos says will enable "high resolution [sampling rate] audio to be deployed in a wider range of consumer and professional product designs. Asked if – in the light of the xCORE Audio introduction, and the company's recent visibility in audio-centric designs – audio has become the main focus area for its activities, and Xmos spokesman says that to the contrary, its intent is to be a supplier of MCU products to the broad market; "Audio happens to be one of the sectors where designs have progressed faster [making that activity higher profile]."

He comments that the Xmos processors – with their multi-threaded and deterministic behaviour – have been gaining acceptance in many areas, and are, "finding a space between [conventional] microcontrollers and FPGAs; the flexibiity of FPGAs, but at the price point of MCUs."

The xCORE-200 multicore microcontroller family is "Gigabit Ethernet enabled" - integrating 16 32bit-RISC-processor-cores in a

single device, the first xCORE-200 device delivers up to 2000 MIPs of real-time computing power and is 10/100/1000 Gigabit Ethernet-equipped, with programmable MAC layer and Internet webserver support. This, Xmos, says, opens a path to Gigabit-speed Internet-of-Things applications. It has a dual-issue pipeline, that contributes to the "2x performance" claim; four times the RAM of its predecessors (up to 512 kB SRAM) and up to 2Mbytes of on-chip flash. There is a programmable USB 2.0 interface (the interface is emulated by a process running in the xCORE multicore fabric; the Ethernet MAC is also a software implementation).

The XE216-512-TQ128 will cost less than $4.75 in high volume; the family map will be filled in with parts configured to provide 8, 10, 12 and 16 cores. Advanced xCORE-200 products with 24 and 32 high-performance 32bit-processor cores will be available in the second half of 2015. A new release of the xTIME-Composer Studio development tools provides full support for xCORE-200 family, and is available through a free download.

xCORE-AUDIO is an optimised variant of the XMOS xCORE-200 multicore microcontroller, with the stereo Hi-Res-2 sampling now, and will cost less than $2.00 in high volume. The first xCORE-AUDIO Live products will follow in April 2015. The devices offer the combination of a high-grade USB 2.0 implementation, to audio standards, with the higher resolution needed to handle high-sampling-rate audio streams. Xmos comments that the demand for high-resolution audio content for consumption on the move, and the complexities of creating the same content in the studio and concert hall, are well served by its USB 2.0 and networked audio solutions. Selecting output formats from Xmos' audio I/O libraries, and using multicore power to implement complex DSP functions, can produce professional audio solutions at consumer prices.

Complete article, here

TI takes low-power MCU line to 32-bit with an ARM core

Texas Instruments has introduced its 32-bit MSP432 microcontrollers. The numbering scheme links to the company's low-power MSP430 16-bit series: however, the new parts employ ARM cores and are based around Cortex-M4F.

The original MSP designation, TI recalls, stood for "mixed-signal microcontroller" and the 32-bit parts follow that path with features such as a 14-bit ADC. TI has developed a customised process at 90 nm geometry, that it runs in its own fabs, to provide the performance/power characteristics it needs for these MCUs.

The devices run at clock speeds of up to 48 MHz, achieving 95 µA/MHz in active power and 850nA in standby power, and returning 3.41 on the EEMBC CoreMark benchmark. TI says this is the lowest power Cortex-M4F MCU implementation available anywhere; when deciding on a 32-bit core, the company says it considered

a Cortex-M0+ but quickly determined that its process allowed the use of the higher-performance M4F with no power penalty. The MSP432 MCUs also deliver a best-in-class ULPBench (also by EEMBC) score of 167.4 – outper-

forming all other Cortex-M3 and -M4F MCUs, TI says.

TI uses is own flash memory on the chips, organised in banks that allow operate-from-one/erase-another. The process does not support FRAM, and TI's FRAM MCUs continue on their own, separate path. As part of its support, TI provides a driver library – customer profiles for its prior devices showed this to be sufficiently popular that the company opted to place it in ROM on the 432 series, providing faster run-time operation.

An integrated DC/DC optimises power efficiency at high-speed operation, while an integrated LDO reduces overall system cost and design complexity. The 14-bit ADC operates at 1 Msample/sec and achieves and ENOB (effective number of bits) of 13.2: it consumes 375 µA at 1Msps. MSP432 MCUs include a selectable RAM retention feature that provides dedicated power to each of the eight RAM banks needed for an operation, so overall system power can be trimmed by 30 nA per bank. MSP432 MCUs can also operate as low as 1.62V and as high as 3.7V with full-speed operation to lower overall system power (but operation direct from a single lithium cell is not intended). Declared product plans include additional analogue functions and up to 2MB flash memory; the first part in the series includes 256 kB flash and 64 kB RAM. An advanced encryption standard (AES) 256 hardware encryption accelerator enables developers to secure their device and data, while IP protection features on MSP432 MCUs ensure data and code security.

The introduction does not in any sense, TI adds, replace the MSP430 line; "We will continue to invest heavily in the 16-bit line, and the market for the 16-bit devices is still growing." If you wish to move from the 430 to the 432 line, as long as your code is in standard C, the tools will make it a simple operation; standard ARM tool sets from Keil and IAR Systems will support the parts.

Complete article, here

Smoothing the design process across chip/package/PCB handovers

Mentor Graphics' Xpedition Package Integrator Flow for IC-Package-Board Design is intended as a solution for integrated circuit (IC), package, and printed circuit board (PCB) co-design and optimisation.

Mentor has addressed what it sees an issue of information exchange between design teams working at the interfaces between the physical aspects of integrated circuit/ASIC pinout; and the design and optimisation of custom IC packages; and the process of entering that package and pinout into the PCB place-and-route process.

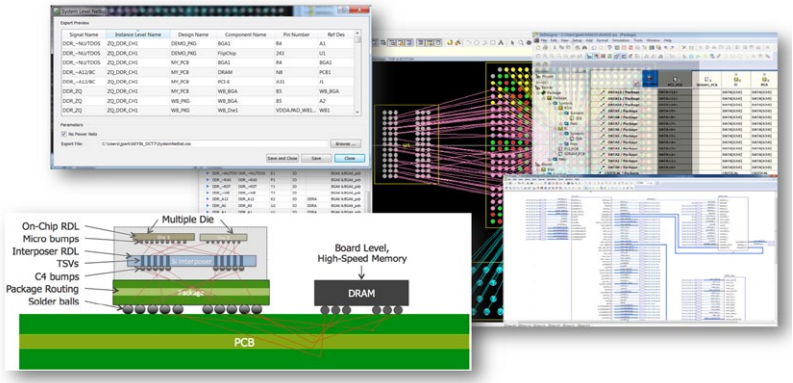

New packaging techniques – stacked dice, through-silicon vias and very high bump/pin counts – are forcing change, Mentor says. There is necessarily a great deal of data that can be fed both forwards and backwards along that chain. An initial IC die-bump layout may call for a very complex pin assignment and interposer design within the IC's package; looking at the constraints of the package design can generate feedback to the the IC pin placement that can subsequently simplify and cut costs in the package design. Similarly, an initial pin assignment on the package might demand a very complex routing escape pattern; a well-informed revision could simplify the board place-and-route, and save layers and costs. Where a

change is necessary, it feeds back the required information to invoke a fresh run of an analysis tool, for example in areas such as signal integrity or thermal design Mentor's introduction addresses these "domain boundaries". It does not link, in an active, "live windows" sense, the EDA packages used in each domain; but it facilitates extracting from each, just the data that is needed to pass key design information and constraints up and down the chain – automating a process that already takes place but with (sometimes cumbersome) manual and simple (e.g. spreadsheet) tools. It employs standardised data formats already established for exchange of data, and Mentor says it is vendor-agnostic in respect of any of the design tools used.
This solution ensures that ICs, packages and PCBs are optimised with each other to reduce package substrate and PCB costs by efficient layer reduction, optimised interconnect paths, and streamlined/automated control of the design process.

Complete article, here

Dual-axis motor control kit uses mixed-signal programmable IC

Microsemi has used its SmartFusion2 SoC FPGA as the hardware basis of a dual-axis motor control development package, that includes an IP suite and reference design to simplify motor control designs, and provide 30,000 RPM performance.
The kit, which simplifies motor control designs using a single SoC FPGA, is scalable across a range of designs. Typical applications include factory and process automation, robotics, transportation, avionics and defence motor control platforms. The SmartFusion2 SoC FPGA dual-axis motor control kit comes ready to use with a compact evaluation board (hardware), access to Microsemi's encrypted motor control IP, and a Libero Gold Edition (software) licence to jumpstart motor control designs. The kit uses a modular

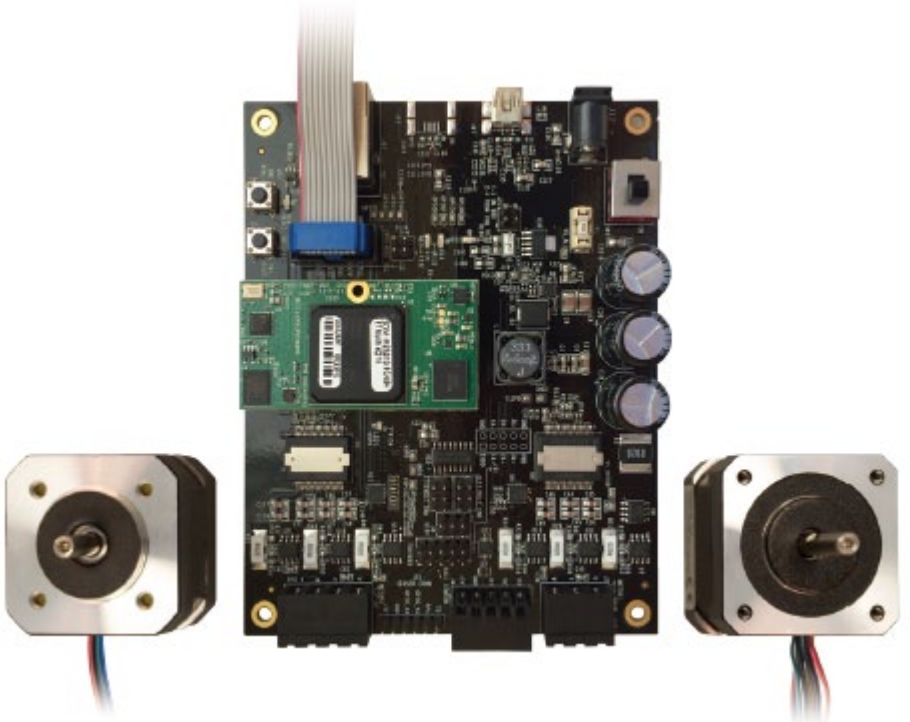

SmartFusion2 SoC FPGA daughter card that enables customers to partition their motor control solutions for hardware only or hardware/software implementation using the modular motor control IP suite and the ARM Cortex M3 MCU in SmartFusion2.
SmartFusion2 with Microsemi's motor control IP enables designers to offload CPU/DSP processing to the FPGA for faster parallel processing. The modular IP suite running on the FPGA is able to drive two BLDC/stepper motor channels with the motor controller kit (while the IP can scale to six-axis) or drive motor performance beyond 30,000 RPM. SmartFusion2 SoC FPGAs can lower total (silicon) device power consumption by more than 50% and lower static power by 10 times over competing SRAM-based FPGAs. The security features available in SmartFusion2 also safeguard connected motors against hackers and intellectual property theft.

Complete article, here

64-bit capable credit-sized dev kit for the makers

by Julien happich

Qualcomm's DragonBoard 410c has been designed as a low-cost development board

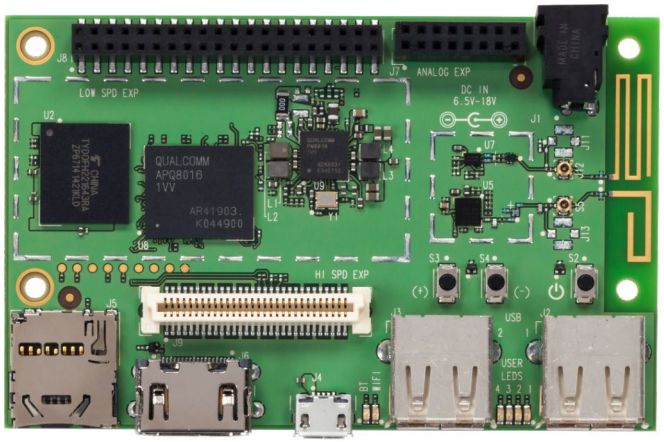

based on the company's 64-bit Snapdragon 410 processor. The credit card-sized development kit is designed to support rapid software development and prototyping for commercialising new inventions and products. The Snapdragon 410 processor is a powerful enablement technology for embedded solutions providers to support off-the-shelf or custom system-on-modules, support and design services for commercial deployments. The DragonBoard 410c is anticipated to be made commercially available by third party distributors this summer.

It hosts quad-core ARM Cortex-A53 CPUs at up to 1.2 GHz per core and 64-bit capability, a Qualcomm Adreno 306 GPU with 400 MHz high quality graphics, 1080p HD video playback and capture with H.264 (AVC), support for 13 megapixel camera with Wavelet Noise Reduction, JPEG decoder, and other post-processing techniques done in hardware. The board also comes with LPDDR2/3 533 MHz Single-channel 32-bit (4.2 GBps) non-POP/ eMMC 4.51 SD 3.0 (UHS-I), Qualcomm's VIVE 802.11 b/g/n, Wi-Fi, Bluetooth and FM connectivity and integrated location through the Qualcomm IZat Gen8C GPS services.

Complete article, here

Magnetic position sensor replaces optical rotary encoders

Unparalleled accuracy over full temperature range at high speed is claimed for the latest version of ams' 47 series; the AS5047P with incremental ABI output is a replacement for optical encoders and resolvers in motor and motion control systems, enabling lower system costs while providing high accuracy performance.

AS5047P features ams' DAEC (dynamic angle error compensation) technology, which produces very accurate angle measurements even at very high rotation speeds. The 47 series magnetic sensors' high-speed performance and ABI incremental outputs make them ideal as robust replacements for optical encoders. The total system cost of a design based on the 47 series is also normally far lower than the cost of an equivalent optical encoder or resolver.

In the AS5047P sensor, the maximum rated measurement speed has been increased from 14,500 rpm to 28,000 rpm, thus enabling it to support many high speed rotating shaft applications that previously could only be addressed with optical encoders.

ams has also increased the number of steps per revolution in the device's incremental ABI outputs. The 47 series' ABI outputs are equivalent to the outputs of a standard optical encoder, so motor control system designers can replace an optical encoder without any change to their control software or interface. In the AS5047P, the maximum resolution of the ABI output is raised to 4,000 steps/1,000 pulses per revolution in decimal mode, and 4,096 steps/1,024 pulses per revolution in binary mode.

AS5047P is highly immune to interference from stray magnetic fields, and is also unaffected by the dirt, dust, grease, humidity and other contaminants. DAEC technology in the AS5047P provides compensation for the dy-

namic engle error caused by the propagation delay as the sensor processes raw measurements of magnetic field strength. DAEC eliminates the need for system designers to implement discrete error-correction circuitry in an external DSP or MCU. With DAEC, the AS5047P achieves very high accuracy, rated at a maximum 0.34° at a constant rotation speed of 28,000rpm (excluding integral non-linearity). AS5047P is priced at $4.63 (1,000).

Complete article, here

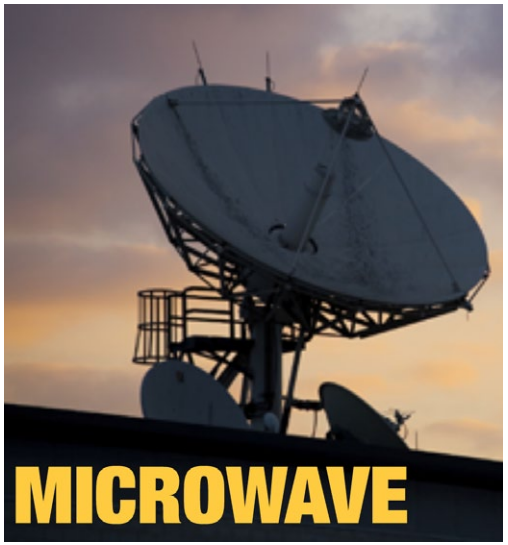

GaN Systems uprates its custom GaN transistor package with topside cooling

Enabling simpler PCB design, GaNPX packaging enables use of conventional heat dissipation techniques with GaN Systems' (Ottawa, Canada) high-power enhancement-mode gallium nitride power switching semiconductors. Topside cooling enables engineers to use conventional, well-understood PCB cooling techniques when incorporating GaN Systems' semiconductors.

GaN Systems' gallium nitride power transistors are based on its proprietary Island Technology - the die consist of islands rather than traditional fingers, which brings significant advantages in terms of better current handling, lower inductance, scaling, isolation and thermal management, as well as enabling smaller die and lowering cost. GaN Systems' enhancement-mode devices with current ratings ranging from 8A to 250A are delivered in its proprietary GaNPX packaging: the die is embedded within a laminate construction and a series of galvanic processes replace conventional techniques such as clips, wire bonds and moulding compounds. These near-chipscale high power switching transistors are now packaged to be cooled via the topside of the chip using a heat sink or fan – conventional techniques that are well-understood and familiar to design engineers who may be unfamiliar with using GaN devices or using them for the first time.

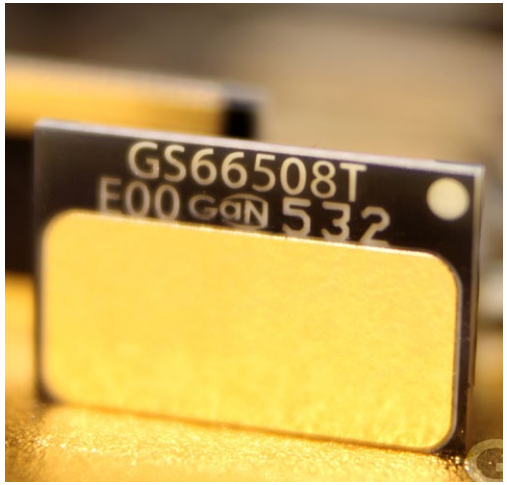

Complete article, here

Infineon confirms dual-strand approach to GaN product portfolio

Infineon which, by internal development and acquisitions, now has a range of power device

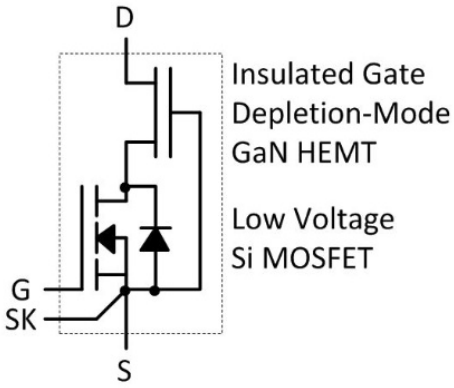

architectures available to it, has confirmed it intends to continue to offer multiple options and has disclosed plans for a portfolio of energy-efficient enhancement mode and cascode configuration GaN-on-Silicon transistors.

The company has announced the expansion of its Gallium Nitride (GaN)-on-Silicon technology and product portfolio; Infineon now offers both enhancement mode (normally-off transistor) and cascode configuration (depletion mode, normally on with companion silicon driver transistor - illustration) GaN-based platforms.

"Infineon's GaN-on-Silicon portfolio combined with the acquisition of International Rectifier's GaN platform together with our partnership with Panasonic clearly positions Infineon as the technology leader in this promising GaN market," claims Andreas Urschitz, President of the Power Management & Multimarket Division of Infineon Technologies AG. "In line with our 'Product to System' approach, our customers can now choose enhancement mode or cascode configuration technologies according to their application/ system requirements. At the same time, Infineon is committed to developing Surface Mount Device (SMD) packages and ICs that will further leverage the superior performance of GaN in a compact footprint. As a real world example, using our GaN technology, a laptop charger found on the market today could be replaced by one that is up to four times smaller and lighter," he added.

Infineon says its expanded offering will include dedicated driver and controller ICs which enable the topologies and higher frequencies that fully leverage the benefits of GaN. Its patent portfolio includes, the company notes, the GaN-on-Silicon epitaxy process and 100V-600V technologies resulting from the acquisition of International Rectifier. Additionally, through a strategic partnership with Panasonic Corporation, Infineon and Panasonic will jointly introduce devices utilising Panasonic's normally-off (enhancement mode) GaN-on-Silicon transistor structure integrated into Infineon's Surface Mount Device (SMD) packages, providing a highly efficient, easy-to-use 600V GaN power device with the added benefit of dual sourcing (reported here).

Complete article, here

4th-gen mixed-signal programmables add analogue features, boost accuracy

Silego Technology's SLG46620V GreenPAK4 is the the fourth generation of its GreenPAK (GPAK) mixed-signal matrix ICs, a family of configurable mixed-signal devices that can be programmed to yield system functions such as comparators, ADCs, logic, delays, counters, resets, power sequencing, voltage sensing, and interface circuits while minimising component count, board space, and power consumption.

The SLG46620V has nearly double the on-chip resources when compared to previous generation parts. This new device also includes new functions, such as a hardware reset pin and a DAC. The chip is available in a 2.0 x 3.0 x 0.55 mm 20-pin STQFN.

The SLG46620 is a user program-

mable device with One-Time-Programmable (OTP) memory elements that are able to construct combinatorial logic elements. Three of the I/O Pins provide a connection for the bit patterns into the OTP on board memory. A programming development kit allows the user the ability to create initial devices. Once the design is finalised, the programming code is forwarded to Silego to integrate into a production process.

Complete article, here

Will a silicon sphere become the new kilogram standard?

By Christoph Hammerschmidt

For many years, scientists around the globe have been considering what a new definition for the kilogram might look like. A definition less prone to ageing and impermanence than the current standard kilogram (a small cylinder made of an alloy of platinum and iridium) is thought to be needed. The German institute for metrology PTB now claims to have the formula. In the pursuit of creating standards that really are constant by definition, the General Conference on Weights and Measures has agreed to define four basic physical quantities by creating a reference to natural constants. These four constants are Ampere, Mol, Kelvin and the kilogram. The new definition for the kilogram will be based the Avogadro constant and the Planck constant. In order to determine these two constants in a complex manufacturing and analysis chain, the scientists have to count the number of atoms in a silicon sphere. This shape of this sphere has to be extremely exact - the deviation from the "ideal" geometric sphere needs to be less than 100 nanometres (nm). The PTB scientists can measure the diameter of such a sphere at an accuracy of three atom diameters; an UHV reflectrometre at the Braunschweig, Germany based institute can determine the thickness of oxide layers on the surface with an exactness of 1 nm.

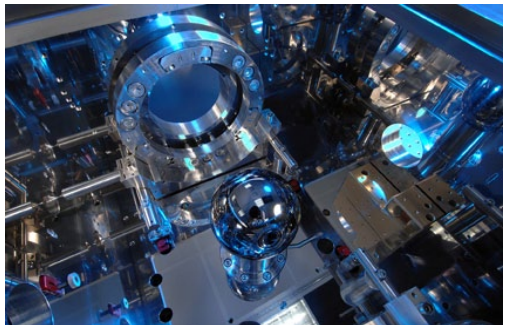

Very high requirements also apply as to the purity of the material, provided by the Electrochemical Plant in Zelenogorsk (Russia). This silicon will exhibit a purity of better than 99.998 percent with respect to its isotopic composition; in use is polycrystalline silicon-28. Out of a large silicon-28 crystal, the PTB will produce two spheres. During the course of the next couple of months, the Leibniz Institute in Berlin will receive another six-kilogram crystal of the same material. Out of this crystal, the Berlin scientists will create a monocrystal with extremely "clean" internal structure and deliver it to the PTB. Ultimately, the PTB will have material for four spheres at its disposal. These four crystal balls with their extremely spherical shape will enable the scientists to deduce the combination of the mass of the spheres and the mass of an atom: They will measure mass and volume of their Si sphere as well as the order of the atoms in the crystal and the incidence of the three existing silicon isotopes - which will enable them to deduce the Avogadro constant and the Planck constant. They can count the number of atoms at high exactness: The deviation is 2 atoms on 100 million atoms counted.

Complete article, here

Automatically-compensated digital PoL controller matches analogue loop performance

Maxim Integrated has posted details of the MAX15301AA02, an efficient digital point-of-load (PoL) controller with advanced power management and telemetry (PMBus) features. It is, Maxim states, the only digital power IC to meet transient and light-load efficiency performance metrics set by analogue controllers.

Unlike PID-based digital power regulators, Maxim continues, the MAX15301AA02 uses the company's InTune automatically compensated, state-space control algorithm. The InTune control law is valid for both the small- and large-signal response and accounts for duty-cycle saturation effects. These capabilities result in fast loop transient response and reduce the number of output capacitors compared to digital controllers.

The MAX15301AA02 includes multiple features to optimise efficiency. An internal switch Buck regulator generates the gate drive and the internal bias supplies for the controller with low power loss. An advanced, high-efficiency MOSFET gate driver has adjustable non-overlap timing and load-variable gate-drive voltage to minimise switching losses over the full range of voltage, current, and temperature.

The IC has an on-board PMBus-compliant serial bus interface for communication with a supervisory controller for monitoring and fault management. A full suite of power management features eliminates the need for complicated and expensive sequencing and monitoring ICs. Basic DC-DC conversion operation can be set up via pin strapping and does not require user configuration firmware. This allows for rapid development of the power-supply subsystem before board-level systems engineering is completed.

Complete article, here

Fastest 8051 IP core runs at 29x original 80C51 performance

Polish IP company Digital Core Design's DQ8051 IP, for FPGA or ASIC implementation, is an extremely-fast implementation of the 8051 MCU core that has a Dhrystone 2.1 performance rating of 0.27292 DMIPS/MHz, which therefore enables a 29.01 times speed-up over the original 80C51 chip operating at the same frequency.

Speed, says DCD, is not everything; the DQ8051's dynamic power consumption can be as low as 1.2 µW/MHz, which rivals not only all other 8051-compatible cores but also low-power 32-bit processors. The 8051, DCD comments, although first introduced more than 30 years ago, is believed to be one of the most popular, or even the most popular, microcontroller core in history. Many engineers have a lot of expertise with this CPU and there's a great amount of legacy code and 8051-based tool chains around. But the 8051 from 1981 and the DQ8051 from 2015 are very different entities. "Using our fifteen years of experience on the market, we've mastered a great portfolio of 8051 IP cores," says Piotr Kandora, Vice President at Digital Core Design, "if our DT8051 is the World's smallest 8051, then the DQ8051 is the World's fastest 8051 for sure. The nearest competition stopped at 26x, with the power consumption almost two times higher than DCD's."

DCD's DQ8051 Dhrystone score rates at 29.01x the original at the same frequency, with the size of 7.5 kgates. The nearest solution consumes almost 12 kgates and achieves no more than 26x speed improvement. The DQ8051 is available with USB, Ethernet, I²C, SPI, UART, CAN, LIN, HDLC, and Smart Card interfaces. The core is also equipped with the company's DoCD hardware debugger with Instruction Smart Trace technology.

Complete article, here

TESTING AUDIO ADCS AND DACS

BY DAVID MATHEW, AUDIO PRECISION

In 2015, there's not much question about audio storage, transmission or streaming: it's digital. Apart from rare sightings of vinyl or open-reel tape in boutique sales or creative enclaves, audio is digital. Done right, digital audio is flexible, robust, and of very high quality. PCM recording, lossless surround formats, and even lossy compression (at least at high data rates) provide the soundtrack for our lives.

But, of course, sound in air is not digital. The pressure waves created by a human voice or a musical instrument are recorded after exciting a transducer of some sort, and the transducer responds with an electrical voltage that is an analogue of the pressure wave. Likewise, at the end of the chain the digitised audio signal must eventually move air, using a voltage that is the analogue of the original sound wave to drive a transducer that creates a pressure wave.

Near the beginning of a digital chain, then, we must use an ADC (analogue to digital converter) to transform the analogue electrical signal to a digital representation of that signal. Near the end of the chain, we must use a DAC (digital to analogue converter) to transform the digital audio signal back into an analogue electrical signal. Along with the transducers, these two links in the chain (the ADC and the DAC) are key in determining the overall quality of the sound presented to the listener.

Testing audio converters

The conventional measurements used in audio test can also be used to evaluate ADCs and DACs. These measurements include frequency response, signal-to-noise ratio, inter-channel phase, crosstalk, distortion, group delay, polarity and others. But conversion between the continuous and sampled domains brings a number of new mechanisms for non-linearity, particularly for low-level signals. This article looks at problems seen in audio A-to-D and D-to-A conversion, and some methods that have evolved to address these issues.

Of course, ADCs and DACs are used in a great number of non-audio applications, often operating at much higher sampling rates than audio converters. Very good oscilloscopes might have bandwidths of 33 GHz and sampling rates up to 100 Gsample/sec, with prices comparable to a Lamborghini. Although audio converters don't sample at anywhere near that rate, they are required to cover a much larger dynamic range, with high-performance ADCs digitiing at 24 bits and having SNRs over 120 dB. Even a high-end oscilloscope typically uses only an 8-bit digitiser. 24-bit conversion pushes the measurement of noise and other small-signal performance characteristics to the extreme edge of what is possible; consequently, measurements of such converters require an analyser of extraordinary analogue performance.

Test setups

The typical test setups are straightforward. For ADC testing, the analyser must provide extremely pure stimulus signals at the drive levels appropriate for the converter input. For converter ICs, the analyser must have a digital input in a format and protocol to match the IC output, such as I²S, DSP, or a custom format. For a commercial converter device, the digital format is typically an AES3-S/PDIF-compatible stream. For devices that can sync to an external clock, the analyser should provide a clock sync output.

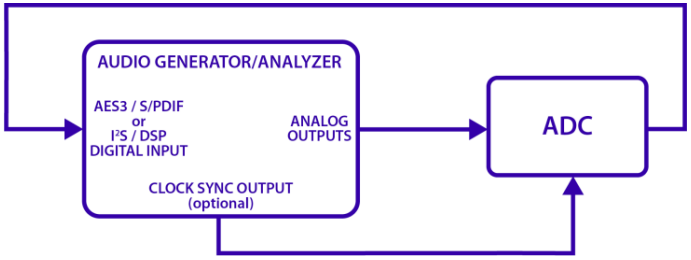

Figure 1. *ADC test block diagram*

For DAC testing, the analyser must have a digital output in the appropriate format, and analogue inputs of very high performance.

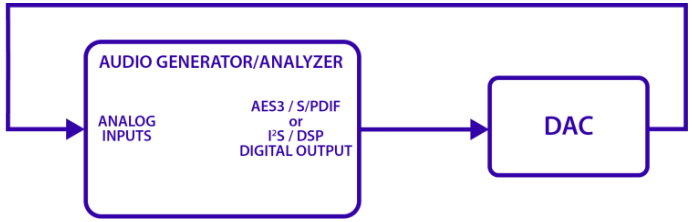

Figure 2. *DAC test block diagram*

The graphs in this article were created by testing commercial converters, using the AES3 digital transport. The analyser is the Audio Precision APx555.

As mentioned previously, ADCs and DACs exhibit behaviours unique to converters. The Audio Engineering Society has recommended methods to measure many converter behaviours in the AES17 standard. The following examples examine and compare a number of characteristic converter issues.

Idle tones

Common audio converter architectures, such as delta-sigma devices, are prone to have an idling behaviour that produces low-level tones. These "idle tones" can be modulated in frequency by the applied signal and by DC offset, which means they are difficult to identify if a signal is present. An FFT of the idle channel test output can be used to identify these tones.

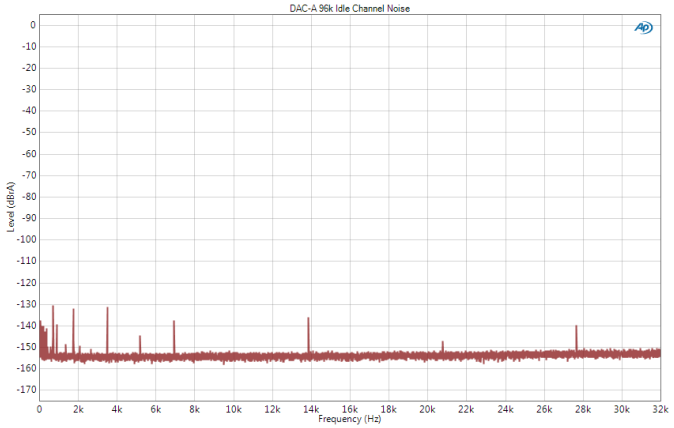

Figure 3. *FFT idle channel noise, DAC "A"*

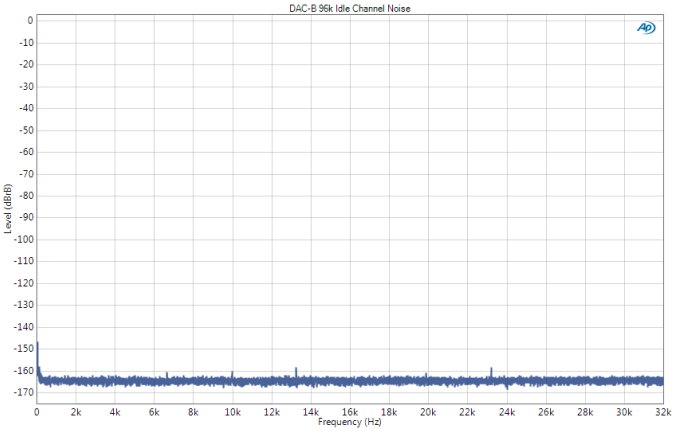

Figure 4. *FFT idle channel noise, DAC "B"*

The DAC in Figure 3 shows a number of idle tones, some with levels as high as –130 dB. The idle tones (and the noise floor) in Figure 4 are much lower.

Signal-to-noise ratio (dynamic range)

For analogue audio devices, a signal-to-noise ratio measurement involves finding the device maximum output and the bandwidth-limited RMS noise floor, and reporting the difference between the two in decibels.

With audio converters, the maximum level is usually defined as that level where the peaks of a sine wave just touch the maximum and minimum sample values. This is called "full scale" (1 FS), which can be expressed logarithmically as 0 dBFS. The RMS noise floor is a little tricky to measure because of low-level idle tones and, in some converters, muting that is applied when the signal input is zero. AES17 recommends that a –60 dB tone be applied to defeat any muting and to allow the converter to operate linearly. The distortion products of this tone are so low they fall below the noise floor, and the tone itself is notched (filtered) out during the noise measurement. IEC61606 recommends a similar method, but calls the measurement "dynamic range".

Below is a comparison of the signal-to-noise measurements of two 24-bit DACs operating at 96 ksample/sec, using this method. As can be seen, some converter designs are much more

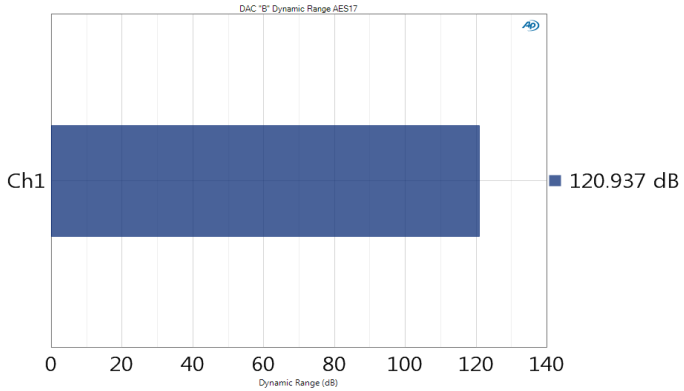

Figure 5. *SNR for dac "b"*

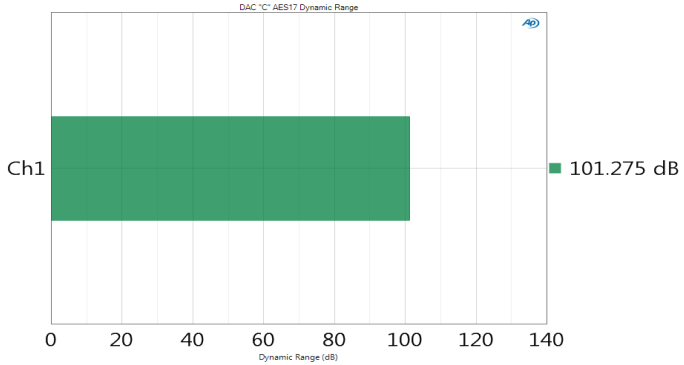

Figure 6. *SNR for dac "c"*

effective than others.

The author continues his discussion with examples of further measurements specific to converters targeting high-performance audio, including jitter considerations; and with an explanatory section on interpretation of noise in FFT power spectra. Click to download the pdf.

Download PDF of Article

Find ADC on EETsearch

USB 3.1 TESTING: START TO FINISH, PART 1

By Randy White, Tektronix

SuperSpeed USB 3.1 is poised to slowly take over as the newest version of this popular interface. It boosts data-transfer rates to 10 Gbits/sec, compared with 5 Gbits/sec for USB 3.0. Along with the higher speeds, additional changes such as the use of 128b/132b encoding and more complex equalisation increase overall design complexity and introduce new PHY validation and debug challenges. Other recent updates include the release of the updated USB Power Delivery 2.0 specification and the new reversible USB Type C connector.

In this first of three articles, I'll outline the challenges associated with bumping the data rate to 10 Gbits/sec and then provide an overview of transmitter and receiver precompliance testing. A second article will take a closer look at the Type-C connector and potential test implications. The third article will explore the new power-delivery specification and delve into test considerations.

In contrast to more enterprise-oriented specs such as PCI Express Gen3 and the upcoming Gen4, the top consideration for USB 3.1 is cost. USB is a consumer-grade interface and it must be inexpensive yet provide reliable operation and keep up with application demands—faster solid state drives with 600 MBits/sec and faster transfer speeds. These considerations resulted in a number of tradeoffs to reduce costs, such as the move to a 1-m maximum cable length for 10 Gbits/sec data rates (compared to 3m for previous versions of USB).

Charting the changes

Table 1 shows the differences between USB 3.1 and the previous generation from a PHY validation perspective. With the release of the 3.1 specification, USB 3.0 technically no longer exists. Instead, the 5 Gbit/sec data rate is now referred to as USB 3.1 Gen1. Regardless of what it's called, USB must maintain backward compatibility with previous versions.

While the most important difference between Gen1 and Gen2 is the higher date rate, those of you familiar with PCI Express will recognise that the encoding for Gen2 has changed to something closer to PCIe (128b/130b). All previous USB PHYs use 8b/10b encoding, which has been around a long time and is used by many serial-bus standards. Now, we have a new state machine that speaks another language,

	Gen1	Gen2
Data Rate	5 Gb/s	10 Gb/s
Encoding	8b/10b	128b/132b
Target Channel	3m + Host/Device channels (-17dB, 2.5 GHz)	1m + board ref channels (-23dB, 5 GHz)
LTSSM	LFPS, TSEQ, TS1, TS2	LFPSPlus, SCD, TSEQ, TS1, TS2,
Reference Tx EQ	De-emphasis	3-tap (Preshoot/De-emphasis)
Reference Rx EQ	CTLE	CTLE + 1-tap DFE
JTF Bandwidth	4.9 MHz	7.5 MHz
Eye Height (TP1)	100 mV	70 mV
TJ@BER	132 ps (0.66 UI)	71 ps (0.714 UI)
Backwards Compatibility	Y	Y
Connector	Std A	Improved Std A with insertion detect

Table 1. *Higher data rates and a new 128b/132b encoding scheme headline the changes in the USB 3.1 Gen 2 PHY.*

128b/132b. Why the difference? The biggest driver was faster data rates under real-world conditions. The use of four bits in the header results in better reliability and better error handling. For example, if there is a single bit error, it can actually be corrected without a need to retransmit the data. Not only does USB 3.1 Gen2 transmit twice as many bits compared to USB Gen1, it makes better use of those bits with better error handling and less overhead (20% for 8b/10b vs. 3% for 128b/132b).

In addition to encoding, USB 3.1 Gen2 has new compliance patterns, although they're not

as difficult to implement as the new encoding. Data is scrambled using a higher-order scrambler (the same polynomial that is used for PCIe). This same pattern is used for measuring jitter and eye height as well as for receiver testing. The jitter measurements also require a separate clock pattern in addition to the data pattern. When a Gen 2 device goes into transmitter compliance mode it will initially transmit the Gen1 CP0 clock pattern. To get the CP9 scrambled data pattern and CP10 clock pattern you can use the standard pattern toggle method (ping.LFPS) to step through each pattern.

With the higher data rate, more loss is to be expected, which is one of the reasons for the change in channel length. While the target reference channel loss for Gen1 was under -20 dB at 2.5 GHz, the loss is now targeted at about -23 dB at 5 GHz. While much of the signal loss is recoverable, there is some relief from requiring a more complex receiver equaliser model and limiting expected cable length to 1m, which is equivalent to -6 dB at 5 GHz. The target channel length certainly could have stayed at 3m, but this would likely have added design complexity and higher power requirements.

One other change worth discussing is Tx EQ (transmitter equalisation). Tx EQ for Gen1 was required using de-emphasis. Now, a 3-tap model is used as the reference for Tx EQ and includes a normative Tx EQ setting. This will ensure margin to support the long 23-dB channel where much of the (now higher) deterministic jitter can be compensated for.

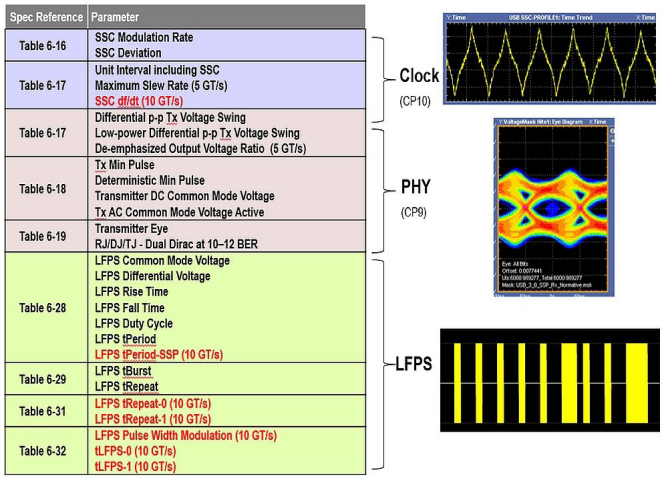

Figure 1. *The new transmitter tests for USB 3.1 Gen 2 are highlighted in red.*

We all know that long cables significantly degrade high-frequency signals, which is why you always check the eye diagram to make sure that the signal can be read. But with USB 3.1 Gen2, you should look at short-channel scenarios, too. So what is it about the short channel that would cause a Gen2 system not to work?

As data rates increase, equalisation is becoming the "secret sauce" to make things work. Equalisation is very sensitive to a signal that overdrives the receiver input and is also very sensitive to the signal-to-noise ratio. These factors make the short channel a must-test scenario for 10 Gbits/sec. It's also required for 5 Gbits/sec, but I have never seen anyone fail the short channel case for 5 Gbits/sec. At 10 Gbits/sec... we'll see.

Transmitter testing

Figure 1 shows the required transmitter tests for USB 3.1 Gen2. There are three main groups: clocking with spread-spectrum clocking, traditional PHY measurements such as jitter, voltage and eye mask, and LFPS (low-frequency periodic signal) timing and voltage swing measurements. The tests that are unique to 10 Gbits/sec are highlighted in red. While there aren't many new measurements, the Gen2 measurements have predictably tighter limits.

In the continuation of this article, Randy White goes into more detail on measurement set-ups for USB transmitter testing, covers receiver testing, and calibration requirements. Click for the PDF download.

NOISE FIGURE IN ANALOG-TO-DIGITAL CONVERTERS

BY ROB REEDER, ANALOG DEVICES

With ADCs, noise figure (NF) and signal-to-noise ratio (SNR) are interchangeable. NF is great for understanding noise density, while SNR measures the total amount of noise in the band of interest. Let's take a closer look at NF, though. Some tradeoffs can be misleading, and low NF does not always translate into lower front-end noise seen by the converter.

NF is easy to use when trying to understand the dynamic implications of cascaded signal chains. As the source resistance is quadrupled, the NF will improve by 6 dB, but the increased resistance also increases the Johnson noise seen by the converter. With higher source resistance, or more than one half full-scale input signal across the analogue front-end (transformer or amplifier), noise becomes more difficult to manage over the band of interest, ultimately degrading the converter's performance.

Why is this? If the full-scale input is lowered, the gain must be increased. This looks fine on paper, but transformers are more gain-bandwidth dependent than amplifiers, so optimising the NF using a high-gain transformer makes it difficult to realise common high-IF applications above 100 MHz.

The problem with amplifiers is similar: as the gain is increased, both the signal and the amplifier noise are increased, which degrades the converter's performance. A high-order anti-aliasing filter using lossy resistive components is required to preserve the performance.

When designing a converter front-end, keep noise spectral density (NSD) in mind. Usually specified in nV/√Hz, NSD is what's really important to the converter, as it will be processed in the digital domain to differentiate and ultimately pick out the signals of interest within band.

In summary, make sure the input and output full-scale signals are maximised throughout the signal chain by positioning gain where appropriate. Attenuation, padding, or resistance is not a good NF tradeoff in any signal chain, as it wastes power and increases noise due to resistors.

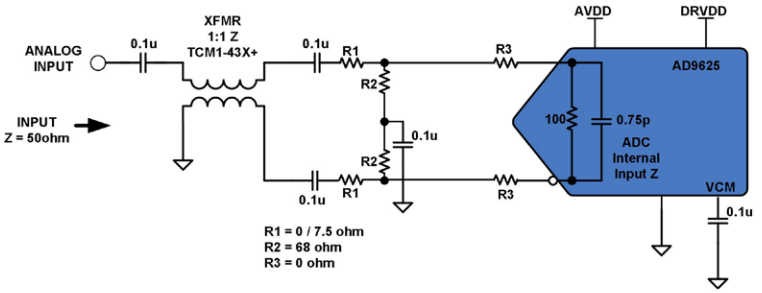

Figure 1. *AD9625 front-end design example.*

The noise figure, NF, is equal to Pfs – SNR – 10 log BW + 174 dBm, where Pfs is the full-scale power of the input network, SNR is measured for the input network, BW is the –3 dB bandwidth of the input network, and 174 dBm is the thermal noise floor = kTBW, where k is 1.38 $\times 10_{-23}$ and T is 300K for room temperature.

References

Kester, Walt. Analog-Digital Conversion. Analog Devices. 2004. ISBN 0-916550-27-3.

MT-006: ADC Noise Figure – An Often Misunderstood and Misinterpreted Specification

MT-052: Op Amp Noise Figure: Don't Be Misled

Brannon, Brad. Analyzing ADC Noise Impacts on Wireless System Performance. EE Times. 2003

Rob Reeder *[rob.reeder@analog.com] is a senior system application engineer with Analog Devices in Greensboro, N.C. He has published numerous papers on converter interfaces, converter testing, and analog signal chain design for a variety of applications. Rob received his MSEE and BSEE from Northern Illinois University in DeKalb, Illinois.*

DEFEND ENCRYPTION SYSTEMS AGAINST SIDE-CHANNEL ATTACKS

By Pankaj Rohatgi, Rambus

From its ancient origin as a tool for protecting sensitive wartime or espionage-related messages, cryptography has become a foundational building-block for securing the systems, protocols, and infrastructure that underpin our modern interconnected world. But the physical mechanisms used in performing encryption and decryption can leak information, making it possible to bypass this security. Protecting designs against such side-channel attacks starts with understanding how such attacks operate.

At its very essence, cryptography is a branch of mathematics dealing with efficiently computable transforms that convert inputs to outputs using additional data known as a cryptographic key. These transforms have the property that, despite observing many input/output pairs, it remains infeasible to compute or invert the transform without the knowledge of the key.

An example of a cryptographic transformation is the symmetric-key based Advanced Encryption Standard (AES-256). An AES-256 encryption device that has access to a 256-bit secret cryptographic key, can transform any sensitive message - known as plaintext - into an unintelligible form known as the ciphertext. Anyone observing the ciphertext, without knowing the plaintext or the key, cannot recover the plaintext or the key. Further, even an observer who knows or can choose the plaintext and can observe the corresponding ciphertext can still not recover the secret key being using within the encryption device. However, any AES decryption device that has access to the same 256-bit secret key as the encrypting device, can readily recover the plaintext from the ciphertext.

Another example of a cryptographic transformation is a public-key based RSA (Rivest-Shamir-Adelman) digital signature algorithm. This algorithm uses pairs of cryptographic keys consisting of a non-secret public key and a secret private key. A signing device that has access to a secret private key can attach a "tag" or digital signature to any message. This RSA signature has the property that without knowledge of the private key, it is infeasible to calculate the digital signature to a message. Anyone who receives a message with a digital signature on that message can use the corresponding public key to establish the authenticity of the message by verifying that the digital signature corresponds to that message.

Strong mathematical guarantees make cryptographic primitives (established, low-level cryptographic algorithms) highly popular as building blocks for securing systems and infrastructure. Encryption is widely deployed to protect confidential data during storage or transmission over insecure networks. Digital signatures are widely used for validating the authenticity and integrity of software, software updates and the data that systems rely upon. Other cryptographic primitives such as message authentication codes, key agreement protocols, and hash functions are also widely deployed for protecting information and systems from attacks.

However, successful attacks on fielded cryptographic systems have also highlighted the pitfalls of relying on purely mathematical guarantees for securing physical systems. It may be infeasible to extract keys mathematically from message traffic, but monitoring message traffic is only one of many possible approaches to breaking encryption.

One common attack vector is exploiting deficiencies in protecting secret cryptographic keying material. Real world systems need to be carefully designed so that secret keys cannot be easily recovered by malicious software or via a simple hardware attack. Unfortunately, incidents where systems get compromised due to poorly protected secret keys are still common.

Another source of problems has been poor communication between the cryptographers, who are mostly mathematicians, and the engineering community that actually develops these systems. If cryptographers do not properly convey all the requirements needed for the mathematical proofs of security - such as the non-reuse of certain parameters or the quality of certain random inputs – to the system designers, the resulting implementations may be vulnerable to a mathematical attack. For example, hackers were able to recover the digital signature key used for signing code for the Sony PlayStation 3 because designers reused a once-per-signature parameter across multiple signatures.

Side-channel attacks

Even if a system protects keying material and meets all the mathematical requirements of the security proofs, there is a class of attacks on all cryptographic implementations that can easily and non-invasively recover secret keys from a system. These attacks, known as side-channel attacks, rely on the fact that any physical realisation of cryptography in hardware or software cannot be an atomic black-box transform as assumed by the mathematical proofs of security. A physical system must necessarily leak information about the process of computing the transform into the environment.

The article continues with examples of this "side-channel" information, and outlines how hackers can extract critical details from them. It goes on to describe the history of side-channel attacks as directed at Smart-cards, the "myth" that such vulnerabilities are limited to Smart-cards, and the growing list of devices that might come under attack as the IoT develops.

DEMONSTRATING ASIC IP PERFORMANCE AND QUALITY DEMANDS AN FPGA-NEUTRAL DESIGN FLOW

By Ali Osman Örs and Daniel Reader, Cognivue, Quebec, Canada

"Companies designing complex SoCs [are moving] increasingly toward licensing IP cores for a majority of the building blocks of their designs instead of building their own in-house custom versions. Selecting the right IP cores is the fundamental challenge of this developing paradigm."

Companies designing new system-on-chip (SoC) products are subject to ongoing market pressure to do more with less and achieve higher returns. The result is shrinking engineering teams, reduced design tool budgets and shortened time lines to get new products to market. This has led companies designing complex SoCs to move increasingly toward licensing IP cores for a majority of the building blocks of their designs instead of building their own in-house custom versions. Selecting the right IP cores is the fundamental challenge of this developing paradigm and the means of evaluating and presenting it is as important to the purchaser as it is to the developer.

The reality is that IP cores are offered with a huge variety of features and options. And, even once you've sorted through the catalogue of potential vendors and products, there is still a vast range in IP quality. The trick is to separate the truly robust and capable from IP that is buggy, insufficiently tested, and lacking in real world performance and a wide and active set of successful users.

CogniVue is innovating with embedded vision enabling small smart cameras that see and react to the world around them, cars that see and avoid accidents, cameras on our TVs that recognise our faces and gestures, and smart phones that see and give us an augmented view of the world around us. CogniVue with its Image Cognition Processing is enabling dramatically new levels of embedded vision, making previously impossible small smart cameras possible. When it comes to vision processing, CogniVue aims not only to have the highest quality IP to offer, but also ensure that it meets the needs of the widest range of application, both for today and tomorrow. This is a field where use cases are still developing and where many won't know their real needs until the design project is well underway.

CogniVue's APEX image cognition processing core, shown in Figure 1, is designed for efficient pipelining of embedded image and vision processing algorithms. The Image Cognition Processor (ICP) is in production and used in many applications including automotive cameras such as Zorg Industries for AudioVox as well as some new wearable type consumer products such as the Neo.1 smart pen from NeoLAB Convergence Inc., as shown in Figure 2. Its deployment in these kinds of consumer applications comes thanks to the ability to deliver 100x better performance per area per power for vision processing compared with conventional processor architectures. For the NEO.1 product the APEX core was able to provide processing at rates of 120 frames per second while still maintaining very low power dissipation,

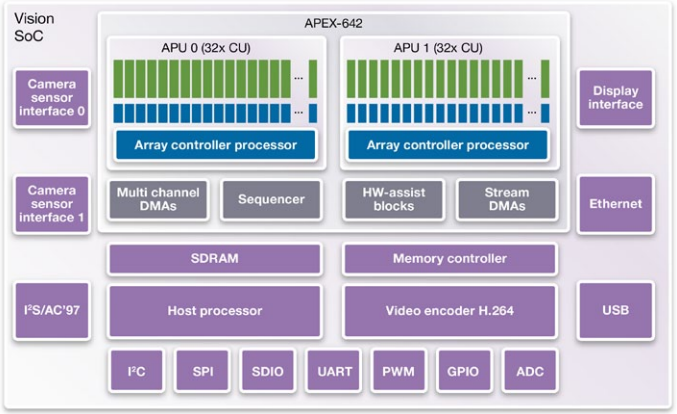

Figure 1. *Example of vision-enabled SoC architecture with a CogniVue APEX2-642 Core*

Figure 2. *The CogniVue APEX core powers the NeoLAB Convergence Inc. NEO.1 Smart Pen*

allowing this battery-powered device to last for many days on a single charge.

This kind of success is achieved both through a fundamental knowledge of image processing requirements and through an exhaustive testing and demonstration approach that targets customer needs within their industrial landscape. Before any core is delivered, extensive validation is needed, especially in markets such as automotive where compliance to industry standards for safety (e.g. ISO 26262 "Road vehicles – Functional safety") is required.

Although testing is necessitated by such requirements, there is also ancillary motivation for IP companies to provide validation and evaluation platforms that not only show functionality and compliance, but that also perform at levels capable of highlighting their true value to prospective customers.

As an example of this motivation, consider the fact that it is less difficult to create vision IP that performs well for narrow, targeted applications that are currently known. Building vision usefulness and flexibility into the technology from the ground up is, however, what will ensure that the IP can perform at the highest levels across multiple applications. And we know that talk is cheap; the IP quality and fit for the application may not be apparent without a real-world "eyes-on" demonstration to prove that the IP's quality and capabilities exist.

The challenge for the fabless IP provider looking to enable their partners and customers is to demonstrate a real IP application running in the real world. Thankfully FPGA platforms continue to leap forward alongside the rest of the technology world, providing a vehicle for this demonstration. In other words, FPGAs can provide the necessary capacity and performance to demonstrate what is possible if the IP is selected for use in the next generation custom ASICs. In spite of this, it seems that we always operate on the edge, pushing the limits of FPGA capacity and performance, and always wanting a little bit more.

FPGA vendors are getting very good at software tool development. But such tools tie use of IP to an individual FPGA company. A demo running on one FPGA vendor today should stand ready to work and move to an entirely different FPGA vendor tomorrow. This can be driven by internal teams or the end customer, and can be due to a combination of factors such as preference/familiarity, legacy infrastructure (hardware and software components), and sometimes the availability of new faster, less costly, better-sized platforms. Moreover, a common RTL code base must work in both the eventual ASIC design flow and in the FPGA "IP demonstration" design flow, as shown in Figure 3.

The authors continue this article with a further description of their methodologies for migrating from prototyping to final ASIC implementation – click for PDF download.

TRADING SIGNAL COMPLEXITY FOR BANDWIDTH

BY RANSOM STEPHENS

In days of yore, PLLs (phase locked loops) recovered clocks from data signals. Since the clock is embedded in logic transitions, PLL clock-recovery circuits couldn't handle waveforms with long strings of identical bits. To assure that the clock recovery circuit would never have to face more than five identical bits in a row, we introduced 8B/10B encoding, which dedicated an extra pair of bits to every eight bits of data. The 20% overhead seemed necessary. But technology moves forward.

PLL-based clock-recovery circuits take up too much chip space, use too much power, and engineering innovation waits for nothing. Phase interpolators came along and reproduced PLL performance—more or less—with lower power while taking up less area. Then the phase interpolator evolved into a DLL (delay-locked loop). DLLs can maintain a synchronised data-rate clock signal even with dozens of bits without a transition. Of course, even DLLs need some transitions, so we implemented 64B/66B (100 GbE and FibreChannel) or 128B/130B (PCIe), and chanted, "Hallelujah!" at our gain of nearly 20% bandwidth.

In going from 8B/10B to 64B/66B encoding, FibreChannel kept the overhead in their rate-doubling naming scheme. While 1GFC runs at 1 Gbit/sec, 2GFC at 2 Gbits/sec, and so on, up to 8GFC at 8 Gbits/sec, in switching to 64B/66B they doubled the payload data rate but 16GFC runs at 14 Gbits/sec and 32GFC at 28 Gbits/sec. I guess it makes more sense to confuse engineers than analysts and users.

The nearly 20% bandwidth reclaimed by switching from 8B/10B to 64B/66B or 128B/130B wasn't free.

The waveform of an 8B/10B signal uses bandwidth from 1/10th the data rate to about 3/2 the data rate; for a 10 Gb/sec signal, that's 1 GHz to 15 GHz. The waveform of a 128B/130B signal covers a bandwidth from not quite 1/64 to 3/2 the data rate; for 10 Gbits/sec, that's about 160 MHz to 15 GHz. To see where those numbers come from, let's start with a square wave.

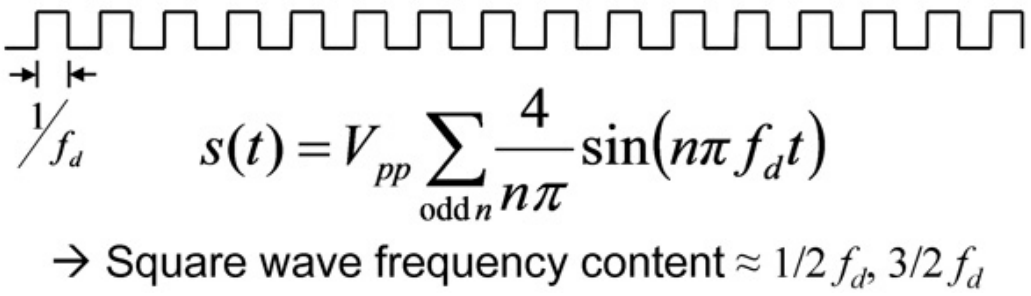

Figure 1. *A square wave with its Fourier series representation—all odd harmonics*

Figure 1 shows a square wave with its Fourier series representation—all odd harmonics.

At the high end, [the figure of] 3/2 the data rate, or $1.5f_d$, comes from the fact that high speed serial standards typically set minimum required rise and fall times and that limits the square wave Fourier components to just the first and third harmonics, though elements of the fifth can show up, too. At the low end, you can see how CIDs (consecutive identical bits) introduce low frequency content by picturing the digital waveform and drawing in the fundamentals for each sub-sequence, as I did in this graphic (figure 2).

Limiting the number of CIDs, like the 5 CIDs for 8B/10B, puts a floor on the low frequency content. Remove that limit and the floor falls into the basement.

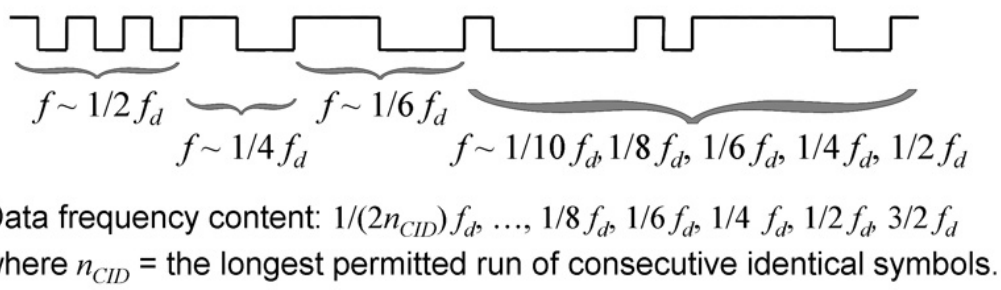

Figure 2. *A digital signal with frequency content indicated for different subsequences (Copyright Ransom Stephens).*

Who cares about extra low frequency content?
Your equaliser cares and, if it cares, than so does your BER (bit error ratio). The equaliser's job — whether at the transmitter in the form of pre- or de-emphasis; or at the receiver in a combination of CTLE (continuous time linear equaliser) and DFE (decision feedback equaliser)—is to correct ISI (inter-symbol interference). ISI is caused by the frequency response of the channel. The wider the signal bandwidth, the greater the ISI.

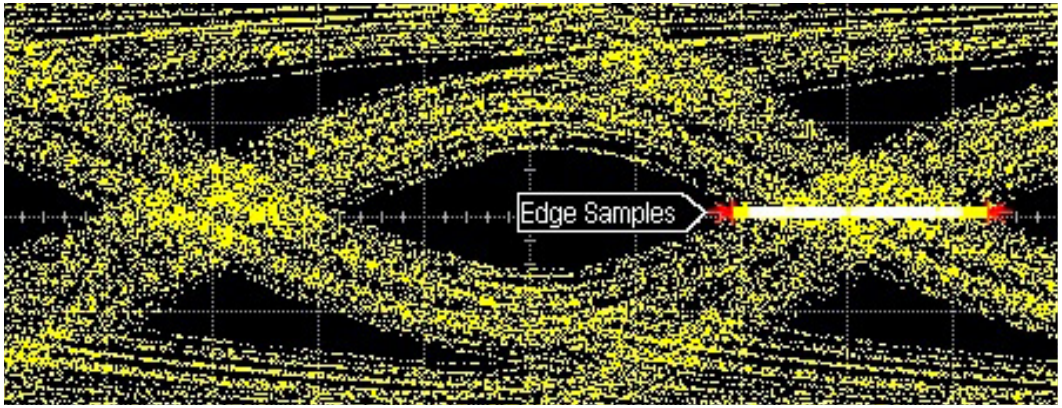

Figure 3. *Time averaged eye diagram so you can see all the ISI.*

The tradeoff in going from 8B/10B to 64B/66B or 128B/130B to reclaim that 20% payload bandwidth is that more sophisticated equalisation schemes are required to maintain a low BER.
Once the equalizer is implemented in the serdes, it's cheap.
Well, not so fast. Don't forget your looming crosstalk problems; these high-speed buses employ multiple fast serial channels in parallel.
When so many high-speed serial lanes reside in a single system, they're bound to interfere with each other. Increased pre-/de-emphasis blasts high-frequency noise through the circuit board making aggressors ever more aggressive.
CTLEs selectively amplify high frequency content which makes victims ever more vulnerable.
DFE, on the other hand, with its magical, nonlinear logic feedback, doesn't affect crosstalk, but will DFE be enough? And we're back at making the receiver ever more complicated.
Once you have four lanes at 28 Gbits/sec or eight at 50 Gbits/sec, that extra bandwidth might not look so cheap.

A SOFTWARE INFRASTRUCTURE FOR THE INTERNET OF THINGS

By Tuukka Ahoniemi, The Qt Company

The rise of the Internet of Things (IoT) marks an important and radical shift in the nature of embedded systems development. By allowing disparate devices to work together across the Internet, it is possible to build more intelligent and responsive systems than ever, create fascinating innovations for users' lives and revolutionise numerous industries.

For example, a logistics company has real-time access to all its vehicles and packages; different parts of an off-shore factory line can be remotely monitored, optimised and controlled from anywhere in the world; or city traffic systems can intelligently share information on load and interruptions, providing real-time guidance for public transit or digital road signs.

IoT can also be applied to create individual smaller scale systems of connected sensors, for instance in applications for fitness and health. Multiple sensors distributed around the body can detect motion and other body functions, such as heart rate and blood pressure. The sensors use their own built-in intelligence to only relay important changes to a host controller worn on a belt or carried in a bag. The controller further processes the data to work out whether changes in the user's condition need attention. But the controller does not need to perform all of the processing itself. Some of it might be devolved to servers running in the cloud, allowing the use of highly complex signal-processing techniques.

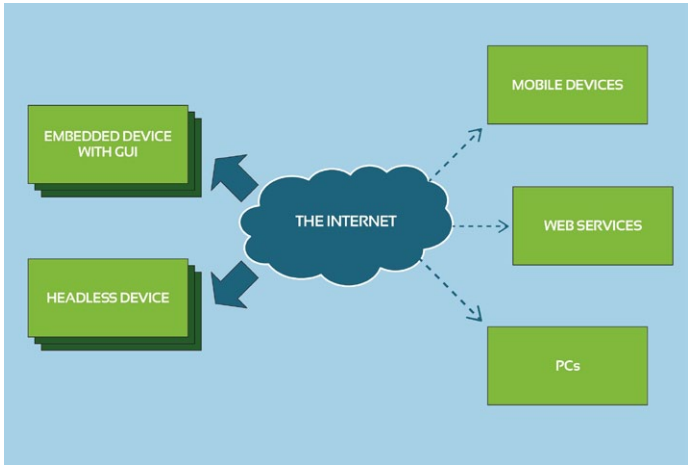

The use cases and innovations that Internet of Things can enable are fascinating and the technical pieces for creating such systems and their cloud back-ends are already in place. However, looking at this from the perspective of embedded software design, creating these systems has increased the complexity and technical requirements of an embedded system dramatically. In order to avoid future pitfalls in integration, one would need to manage a whole portfolio of different front and back end technologies and perceive the whole system when starting an architecture design.

A whole new world through cloud connectivity

By allowing devices to cooperate seamlessly, the IoT architecture makes it possible to deploy complex services – some of which may not even have been considered during the initial rollout – using comparatively inexpensive and power-efficient devices. Not only that; services can be upgraded over time without needing to disturb software running on every device, especially if the technologies used allow dynamic changes to software or have flexible, well-designed external APIs that use, for instance, plugin-based design patterns to allow easy extension of services.

The cloud servers can be used as permanent, device-independent storage and archive content and hold important information about the state of each device. In a consumer-entertainment context, this can be realised as the ability for a user to pause music they listen to on a mobile device and then play it from that moment on when they sit down in front of their PC. Through the cloud, systems can push notifications down to devices as well as receive sensor data, process and analyse usage statistics.

As the data storage and complex data analysis are centralised into the cloud backend, the capacity and computational power can be dynamically scaled up without affecting the

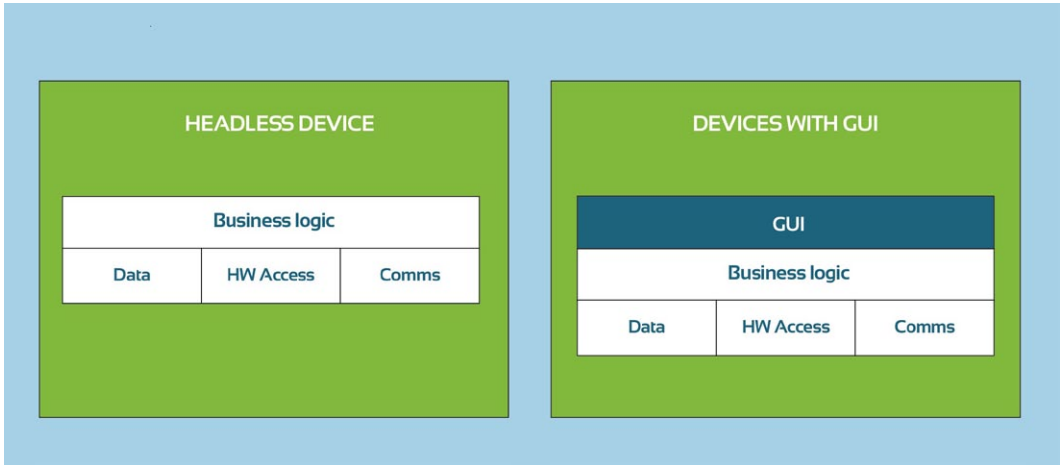

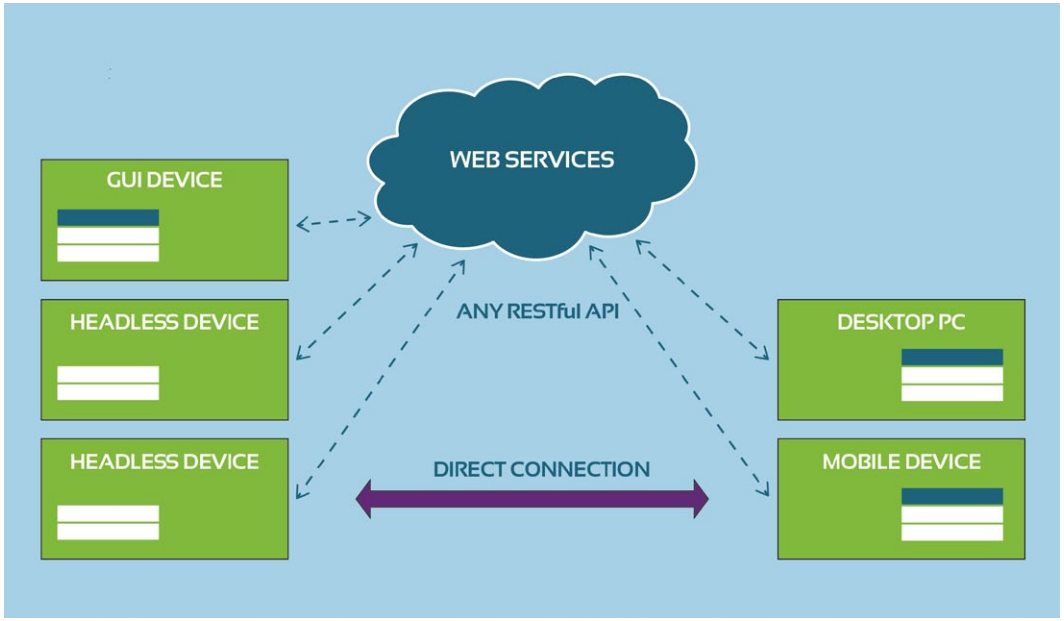

individual connected devices. Also, random loss of connection to a device doesn't make a critical difference for other parts of the system. The requirement for availability is of course still there, just for the cloud backend. So, to ensure the functionality of an individual device without connection to the backend, an offline cache and logical offline behaviour needs to be well designed.

REST assured: higher level functionality

Although the IoT concept brings with it the potential to build flexible, future-proof networked systems, there are associated development challenges that need to be addressed. The architecture of the Internet and World-Wide Web provide facilities to ease some aspects of IoT creation. For example, the representational state transfer (REST), which is an architecture based on the simple HTTP protocol operations, offers higher degrees of flexibility than traditional client-server architectures, such as SOAP.

A RESTful API allows clients and servers to interact in complex ways without any requirement for the client to know details about possible services before it connects to the server. This is very important when we consider the extensibility of a system with arbitrary devices or services. A simple universal resource identifier (URI) is all that is required to allow the client to find the server. From this point on, server and client perform a series of negotiation and data transfer steps using the HTTP operations that provide information about the services on offer and how to access them.

Although the core REST mechanisms available within HTTP are flexible, they operate at a low level. An applications framework such as Qt can abstract these low-level details through convenience classes, while also providing high-level object-oriented classes for other parts of the application.

Read the continuation of this article, which includes a flexible-end-device use-case, in the PDF download, click right.

CRITICAL REQUIREMENTS IN HIGH SPEED SIGNAL GENERATION APPLICATIONS

By Clarence Mayott, Linear Technology

In high speed signal generation applications bandwidth and resolution are the critical requirements. Modern signal generation applications use high speed digital-to-analogue converters (DACs) to produce various types of waveforms – from single tones to complex multichannel waveforms with several hundred megahertz of bandwidth. These applications demand high speed DACs that are fast enough to produce these waveforms without sacrificing analogue performance. In many signal generation applications phase noise will limit the number of channels and the spacing of the channels that is possible. Phase noise is traditionally set by the clock that is driving the DAC clock inputs, but any phase noise added by the DAC will show up in the output spectrum and can limit the signals that it is possible to generate. The ideal DAC for any general purpose signal generation application should be as fast as possible, with low noise, high linearity, and very low additive phase noise. If any one of these specifications is lacking, then the generated waveform will not be adequate to meet the application's needs.

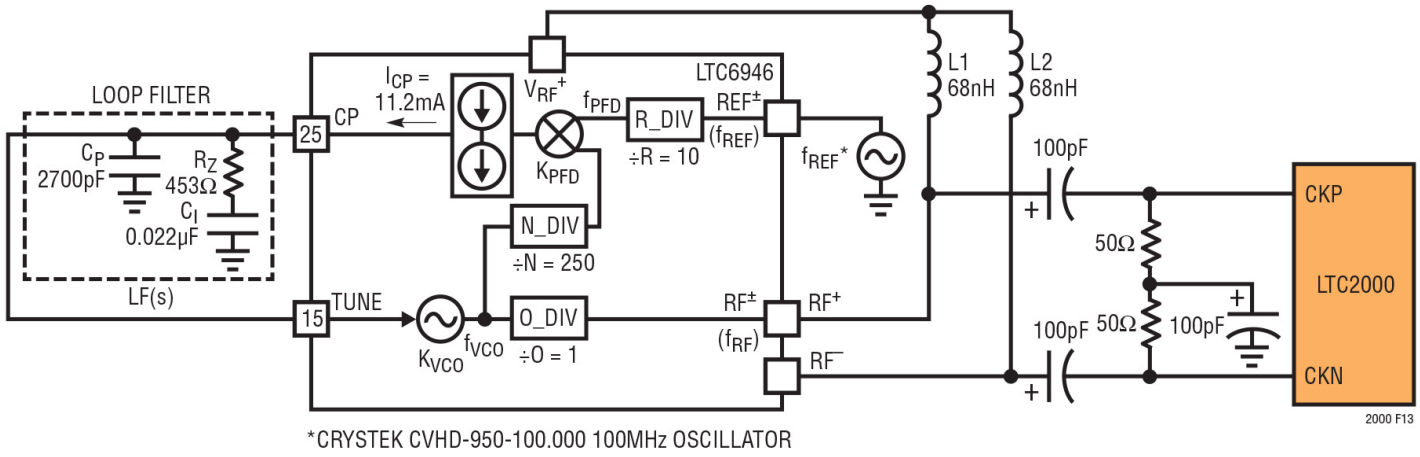

Figure 1. *LTC2000 and LTC6946 block diagram*

Bandwidth

In any signal generation application the most important design criteria is bandwidth. The first question any designer will ask is: How much bandwidth do I need to generate the signal of interest? The designer may need a certain amount of bandwidth for a particular signalling protocol or for a particular application. The bandwidth that the designer is trying to achieve can only be generated with a DAC that is at least twice as fast as the bandwidth of interest. This relationship between bandwidth and sample rate (fs) was defined by Harry Nyquist and describes how signals behave in sampled systems.

While it is possible to generate a signal that spans from DC to the $f_s/2$ limit, it is often not practical to do so due to the images of the generated signal that appear in the output spectrum. Images will occur at $N*f_s \pm f_{out}$, where f_{out} is the frequency of the generated signal. In practice, reconstruction filters are required to attenuate any images of generated signal that may appear in the output spectrum. Even if the generated signal does not extend up to $f_s/2$ but closely approaches it, the images will be difficult to filter out due to constraints on the filter. The reconstruction filter is implemented

in the analogue domain with real components. Unlike digital filters, these components are not ideal, which leads to non-ideal passband with ripple and insertion loss. Generally, the higher order these filters are, the more ripple and insertion loss they will have, making an ideal filter more challenging to design. The closer the bandwidth of the signal is to $f_s/2$, the higher order the filter must be to attenuate the images produced by the sampling process. This higher order filter requires more components, and will have more insertion loss and passband ripple.

By using a DAC with a faster sample rate, the usable bandwidth will increase, which will ease the requirements of the filter. This allows for a filter with fewer components and less complexity, which simplifies the design and produces better results. The LTC2000, which is a high performance, 16-bit, 2.5 Gsps high speed DAC, for example, has an $f_s/2$ frequency of 1.25GHz. Thus, for a signal bandwidth of 800 MHz there will be an image of the signal beginning at 1.7 GHz. There will be 900 MHz between the frequency band of interest and the image frequency. With 900 MHz of guard band the image can easily be filtered out with a simple low pass filter. A DAC with a slower sample rate would produce images that are closer to the frequency of interest requiring a more aggressive and complex filter.

Another issue with generating a signal that extends out to $f_s/2$ is that with any DAC there will be a SINC (sin(x)/x) roll-off that will attenuate the generated signal as the frequency increases. This roll-off has a zero at the sample frequency (f_s), making it impossible to produce a signal that appears at exactly the sample rate; the resulting signal would just be a DC voltage. For practical applications, about 60% of the Nyquist zone (DC to $f_s/2$) can be used without much SINC attenuation. If 0 dB is the signal level at DC, then at 60% of Nyquist the signal level would be down by 6 dB. The inverse of this roll-off is frequently implemented in the digital domain, which corrects for the natural roll-off of the generated signal. This allows DACs to produce waveforms with constant amplitude over frequency. Using a higher speed DAC will reduce the roll-off of the SINC function as the output frequency of the DAC increases.

Phase noise

Another important consideration in signal generation applications is the phase noise of the output. The phase noise that is present on the output signals limits how closely the signals can be spaced and can limit the order of modulation that is possible to produce. The more phase noise that is added in the signal generation process, the lower the SNR and the higher the bit error rate of the generated signal will be.

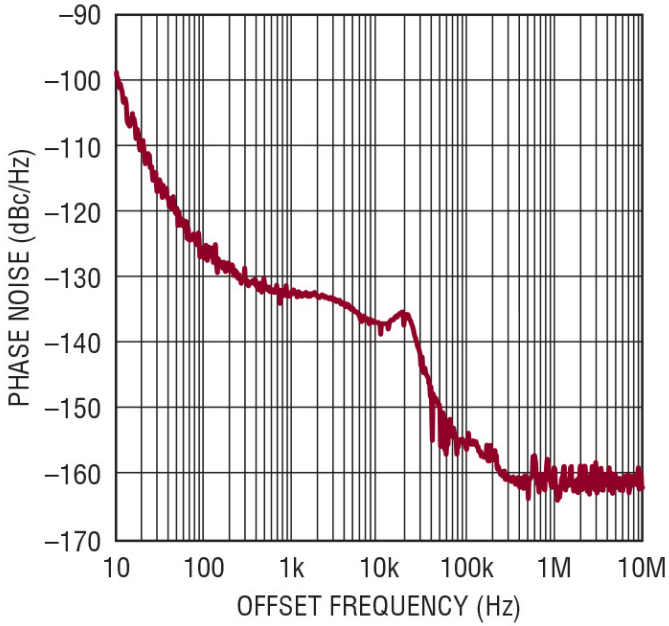

Figure 2. *LTC2000 and LTC6946 phase noise Fout = 80 MHz*

Jitter is a measurement of the accuracy of the zero crossings of a signal in the time domain. A perfect signal will have a zero crossing at the same point in time each period. In reality there will be a certain distribution of times that these zero crossing occurs. If this distribution is translated to the frequency domain, then phase noise can be seen as spectral leakage around the fundamental tone. If there are several tones closely spaced, the SNR of a tone can be degraded by the spectral leakage of its neighbours, which will reduce the bit error rate of the

signal and the accuracy of the signal produced. This loss of signal integrity can be avoided by reducing the phase noise introduced into the generated signal.

The easiest way to avoid introducing phase noise into a signal generation system is by starting with an extremely low phase noise clock. A low phase noise clock will transpose less phase noise onto the generated signal. It is also important to note that there will be an attenuation of the phase noise of the clock on the generated signal that is proportional to the ratio of the generated signal to the clock sample rate. This relationship means that generating low frequency signals with a high sample frequency clock produces less phase noise on the output signal than generating a high frequency signal with the same clock. If the generated spectrum is wideband, the generated signals at the high end of the spectrum will have more phase noise than the signals at the lower frequencies.

The LTC6946 is a frequency synthesiser that can produce signals from 370 MHz up to 5.7 GHz without an external VCO. It has excellent phase noise and very low spurious content, making it ideal for use as a clock source for signal generation applications. When the LTC6946 is used to drive the LTC2000 high speed DAC, the resulting phase noise is low enough for most demanding signal generation applications. A typical schematic of the LTC6946 driving the LTC2000 is shown in Figure 1. Figure 2 shows a plot of the phase noise of the LTC6946 and the LTC2000. The LTC6946 uses an internal VCO which trades convenience for phase noise. Even lower phase noise is achievable by using the LTC6945 with an external VCO. For both the LTC6945 and LTC6946 frequency synthesisers, the dominant source of phase noise is the VCO. The LTC2000 has an additive of -165 dBc/√Hz at a 1 MHz offset when producing a 65 MHz output tone. This ensures the phase noise of the clock will dominate the additive phase noise of the LTC2000 itself. To avoid other noise degrading the output signals, care should be taken to use proper layout techniques in the analogue output section – *which is where the article continues in the PDF download, click below..*

designideas

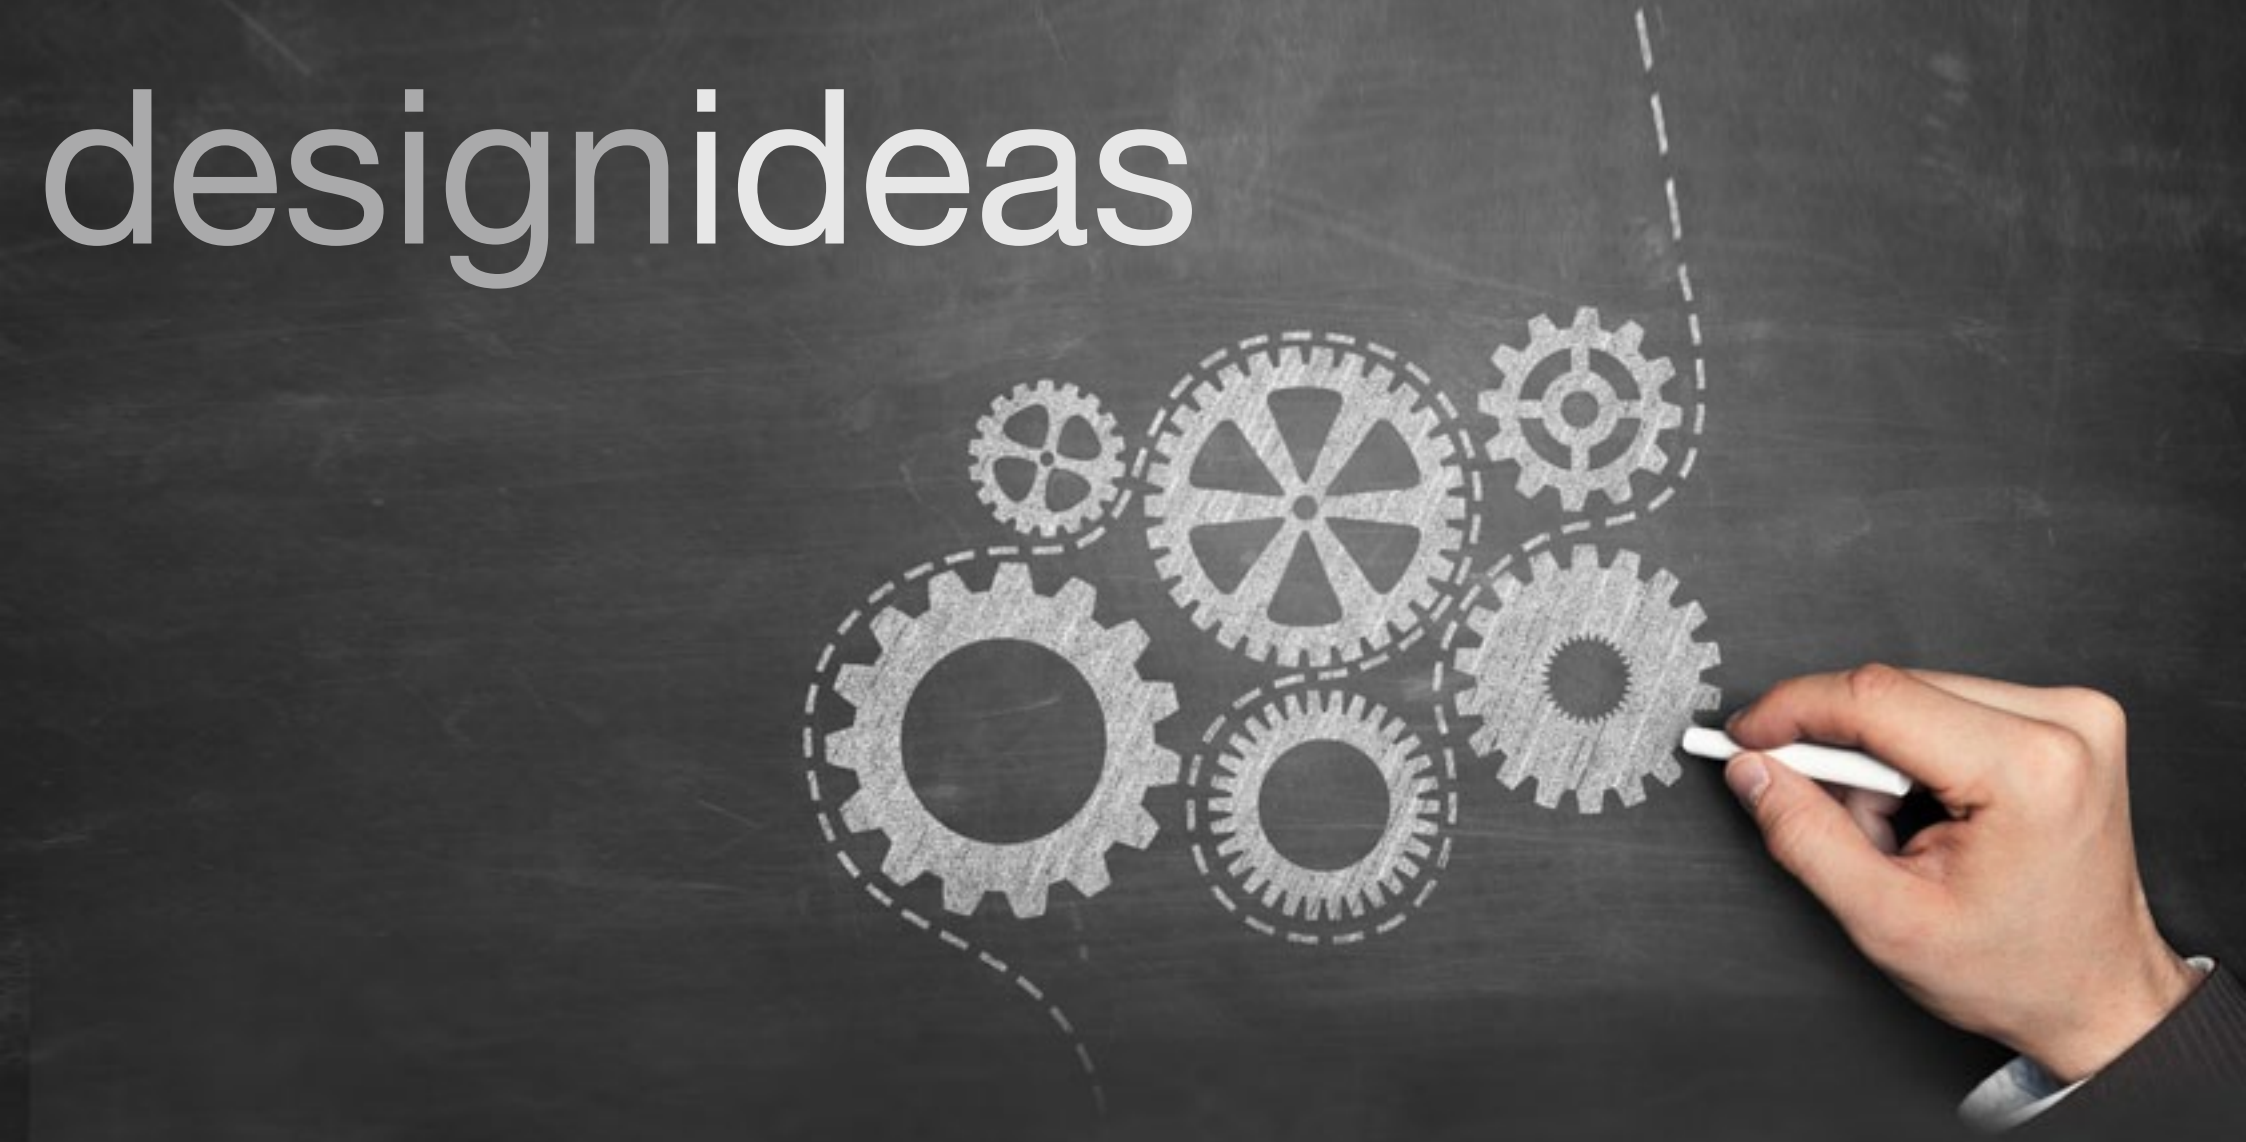

- Schmitt trigger adapts its own thresholds
- Instrumentation amp makes an accurate transimpedance amp too

Schmitt trigger adapts its own thresholds

By Anthony Smith

In 1938, the Journal of Scientific Instruments published details of a comparator circuit that converted a slowly varying input signal into an abrupt change in output voltage. Based on cross-coupled thermionic valves, the circuit was developed by American scientist O. H. Schmitt. Since then, the Schmitt trigger has become a staple building block in many signal processing circuits. Hysteresis – the difference between upper and lower voltage thresholds – is inherent in the Schmitt trigger's operation. Provided the input signal crosses both thresholds, the circuit will reject noise contained in the input signal and will produce a rectangular output signal at the same frequency as the input.

Whether you implement a Schmitt trigger using transistors, op amps, or comparators, you need to determine how much hysteresis is required and what the threshold voltages should be. This is usually a simple process if you know the amplitude of the input signal and how much noise it is likely to contain. However, if these parameters are variable or largely unknown, setting the thresholds to produce reliable triggering can be tricky: too much hysteresis can prevent the input signal from crossing one or both thresholds; too little hysteresis can lead to false triggering if the input signal has large noise content.

The Design Idea shown in Figure 1 solves these problems by implementing a circuit which automatically adjusts the trigger thresholds to suit the amplitude of the input signal. Comparator IC1A together with analogue switch IC2B and capacitor C1 form a positive peak detector. When the input signal rises above the voltage stored on C1 at the comparator's inverting input, the comparator output goes high causing IC2B to switch into the position shown in the schematic. The detector now samples the input signal and tops up the charge stored on C1. When the input signal drops below the voltage on C1, the switch changes state such that the voltage VU stored on C1 is a DC level corresponding to the upper peak value of the input signal.

Comparator IC1B, analogue switch IC2C and capacitor C2 form a negative peak detector. This operates in the same way as the positive peak detector described above, but instead sampling on the negative peaks of the signal, such that the voltage V_L stored on C2 is a DC level corresponding to the lower trough value of the input signal.

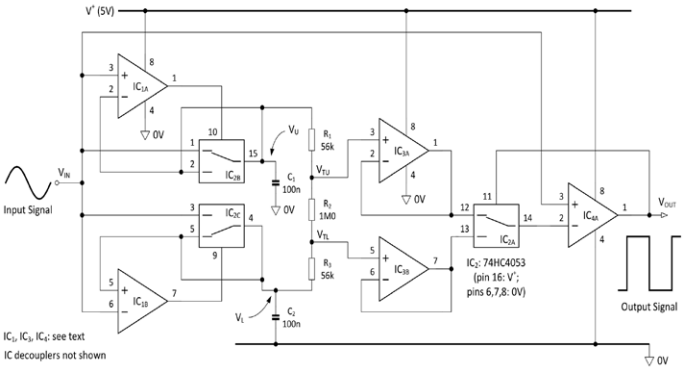

Figure 1. *Self-adaptive Schmitt trigger*

The resistor network formed by R1, R2, and R3 provides a discharge path for the charge stored on the sampling capacitors and also sets the upper and lower threshold voltages, V_{TU} and V_{TL} respectively, for the final comparator IC4A. The resistor values are chosen such that V_{TU} is just slightly less than V_U, and V_{TL} is just slightly greater than V_L. If we set R1 equal to R3, the voltage difference in percentage terms is given by:

$$\text{Voltage difference} = [R1/(2R1 + R2)] \times 100\%$$

With the values shown in the figure, V_{TU} is 5% less than V_U, and V_{TL} is 5% greater than V_L. In

this way, the thresholds are constantly adjusted to track the input signal amplitude and DC level. For example, a 1V peak-peak signal riding on a DC level of 2V (i.e., V_U = 2.5V and V_L = 1.5V) would produce thresholds of V_{TU} = 2.45V and V_{TL} = 1.55V. It can be seen that the hysteresis voltage V_H given by $V_H = V_{TU} - V_{TL}$ (in this example 0.9V) is always just slightly less than the peak-peak amplitude of the input signal.

The threshold voltages are buffered by IC3A and IC3B before being fed to analogue switch IC2A. To understand how the final part of the circuit works, assume that IC2A is in the state shown in the figure such that threshold voltage V_{TU} is fed to the comparator's inverting input, and that the input signal at the comparator's non-inverting input is rising up from its negative peak. The digital output signal, V_{OUT}, is currently at its low level. At the moment when the input signal just crosses V_{TU}, the comparator output immediately goes high causing IC2A to change state and feed V_{TL} to the comparator's inverting input. This positive feedback – typical of Schmitt trigger behaviour – ensures rapid, clean switching of the digital output signal. Buffers IC3A and IC3B are necessary (particularly at high frequencies) to prevent stray capacitance around IC4A's inverting input from introducing aberrations into V_{TU} and V_{TL} when IC2A changes state.

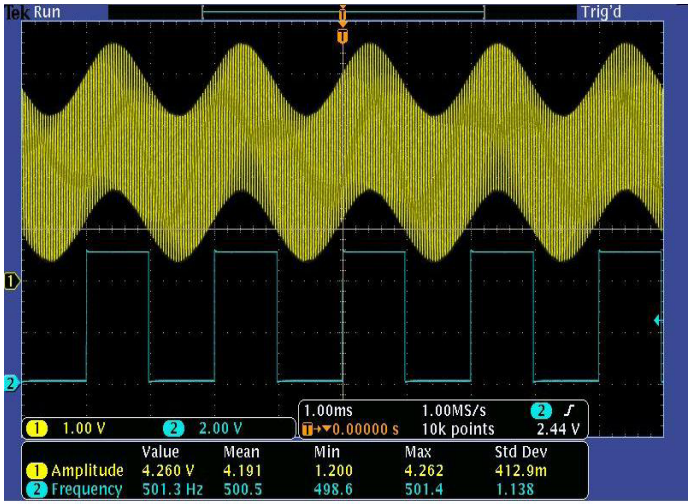

Figure 2. *500Hz signal with modulation "noise"*

The oscillographs in Figures 2 and 3 show the performance of a test circuit built using comparators IC1 and IC4 = TLC3702 and op amp IC3 = TLC2272. These fairly extreme examples illustrate the circuit's ability to handle widely differing input signals.

In Figure 2, the source signal was a 500Hz sinewave at 1.56V P-P modulated by a 100 kHz sinewave at 2.88 VP-P, resulting in a composite signal of around 4.4 VP-P riding on a DC level of 2.5V. Despite the "noise" amplitude being almost double the source amplitude, the circuit output switches cleanly at the source frequency and is completely unaffected by the HF modulation.

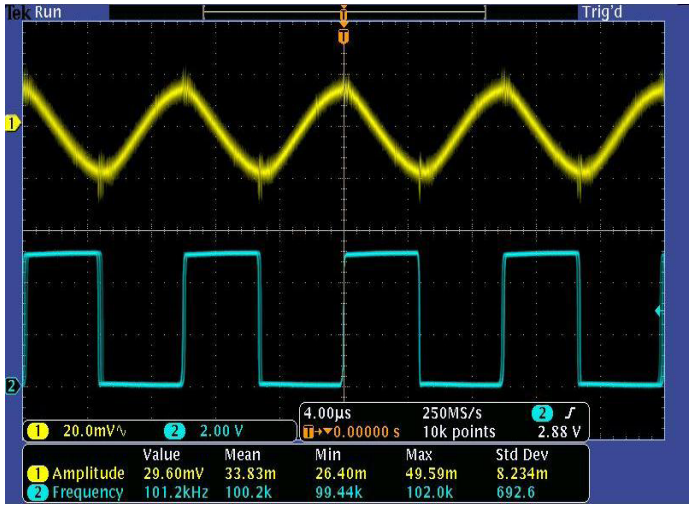

Figure 3. *Low amplitude input*

Figure 3 shows the circuit's response to very small input signals. Here, the source signal is a 100 kHz sinewave at around 30 mVP-P riding on a DC level of 400 mV. The presence of switching spikes in the input signal (due to less-than-perfect breadboard layout) results in some jitter in the output signal. Note that the input signal in Figure 2 is over a hundred times larger than that in Figure 3. In fact, provided the input signal remains within the common mode range of the comparators and buffers (in this case, around 0 to 4V), the circuit can handle different signal amplitudes that vary by as much as two orders of magnitude. AC coupling is only required if the signal's DC level is outside the input common mode range.

You should choose C1 and C2 to suit the anticipated frequency range. Values of around 100 nF are suitable for frequencies above 300 Hz or so. Below this level, the sampling capacitance should be increased to prevent excessive decay ripple appearing on V_U and V_L. The TLC3702 comparators work well up to around 100 kHz, but beyond this level you may need faster devices.

Be aware that when the positive peak detector samples the input signal, capacitor C1 takes a gulp of charge from the input; the associated current that flows from the input is limited only by the on-resistance of IC2B. The same process occurs on the negative peaks with C2 and IC2C. If the input source impedance is significant, these current pulses can produce spikes in the input signal which can lead to erratic triggering. Therefore, it may be necessary to buffer the input signal to avoid these problems.

The circuit is not a panacea for all triggering applications, but it should prove useful for cases where the fixed thresholds of the conventional Schmitt trigger are not suitable. If it were not for my lack of temerity, and the confusion it might cause, I would suggest calling it the Smith Trigger.

Instrumentation amp makes an accurate transimpedance amp too

By Stefano Salvatori

Transimpedance amplifier circuits are often used to convert photodiode current to a voltage signal. In the conventional transimpedance amplifier schematic of Figure 1, the current-to-voltage conversion factor is given by the value of the op-amp feedback resistor R_F. Obviously, neglecting any non-ideal behaviour of the op-amp, the accuracy of the I-V converter is tied to that of R_F (on both its absolute value and thermal coefficient).

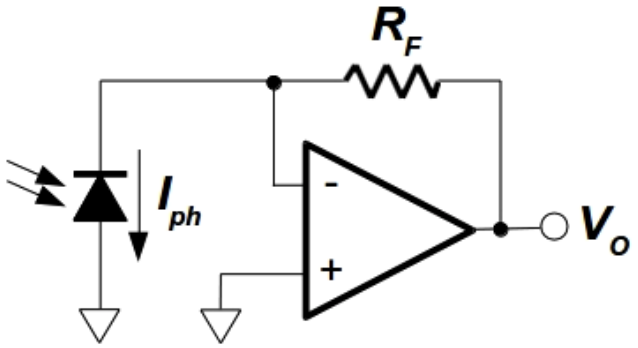

Figure 1. *A classical transimpedance amplifier for photodiode signal conditioning. The current-to-voltage conversion factor is given by R_F*

High precision thin film resistors are available, either as a single components or as an array of matched-value resistors. In both cases, typical tolerance of commercially available precision resistors is ±0.1%, and ratio tolerances are commonly found in the 0.02% - 0.05% range. The cost of such components may dominate, being easily greater than that of a low-noise, low-offset, precision op-amp. This Design Idea presents an alternate approach.

Monolithic three-op-amp instrumentation amplifiers (IAs) usually integrate laser-trimmed precision resistors, with accurate absolute values for the first-stage feedback parts, and high matching-accuracy for those around the output difference amplifier (see, for example, the AD620 or INA129). The overall gain error and gain non-linearity can be as low as ±0.01% and ±0.001%, respectively. Also thermal coefficient is very low – on the order of tens of ppm/K.

Programmable IA integrated circuits usually have a design as in Figure 2. The connections shown allow one to realise a highly accurate current-to-voltage converter having a transimpedance dependent only on the R_F value of

the IA input stage. Indeed, neglecting OA1 and OA2 inverting-input bias currents, the same photodiode I_{ph} current flows in the two R_F feedback resistors. Since the IN(+) and IN(-) input pins are connected to ground, the OA3 difference amplifier will see $R_F \times I_{ph}$ and $-R_F \times I_{ph}$ voltages at its inputs. The IA output voltage V_O is:

$$V_O = 2 \times R_F \times I_{ph} \quad (1)$$

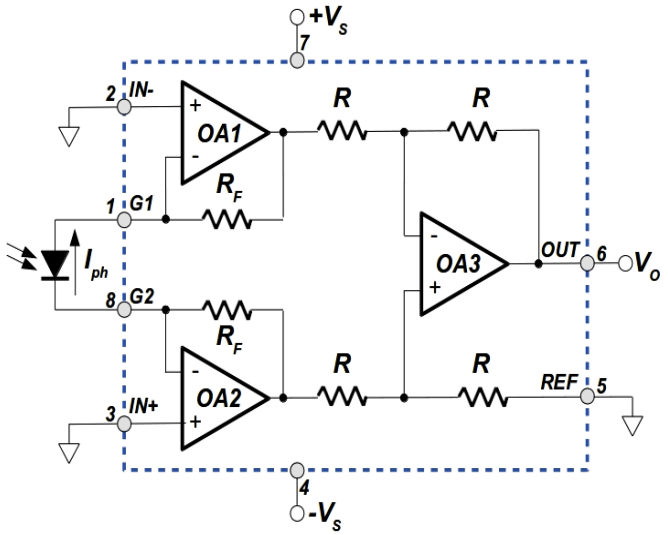

Figure 2. *An instrumentation amplifier integrates precision matched resistors. Feedback resistors RF typically have absolute tolerances as low as ±0.01%. This configuration turns the IA into a transimpedance amplifier for photodiode signal conditioning with transimpedance equal to 2 × RF.*

where I_{ph} is the photodiode photocurrent. The output, V_O, will be positive, as the I_{ph} current flows in the direction shown in Figure 2.

If a bias voltage is necessary for photodetector biasing (a reverse bias for a photodiode, or a polarisation for a photoresistor), the component can be connected with opposite polarity at the G1/G2 pins compared to that in Figure 2. Then, V_{bias} can be applied to the IN(-) input of OA1 such that the same potential is virtually maintained by OA1 at the detector cathode. If the same voltage is applied to the IA's REF input, Equation (1) still remains valid, though watch the allowed common-mode voltage amplitude on the IA inputs.

A typical circuit's V_O will range from a few millivolts to a few volts; fixed gain, but adequate for many applications. In order to overcome this limitation, a precision programmable-gain amplifier (e.g., the PGA204) can follow the circuit of Figure 2.

productroundup

productroundup

30A high integration DC/DC for FPGA PoL regulation

A DC/DC conversion module from Altera's Enpirion acquisition features adaptive digital control and is optimised to power FPGA cores: specifically, Altera's Generation 10 devices of which the Arria 10 variants are now shipping. Altera claims the EM1130 as having the densest-available footprint, with close output regulation, and fast transient response. The digitally-controlled modules are equipped with PMBus which provides the interface for Altera SmartVID technology. PMBus also allows the EM1130 to communicate with the system to provide telemetry—measurement of key parameters such as current, voltage and temperature.

Complete article, here

80-V half-bridge FET module offers drop-in GaN evaluation

From Texas Instruments, the first 80-V, 10-A fully-integrated gallium nitride (GaN) field-effect transistor (FET) power-stage prototype consists of a high-frequency driver and two GaN FETs in a half-bridge configuration – all in an easy-to-design quad flat no-leads (QFN) package. The module will help first design-ins of next-generation GaN power-conversion solutions; TI says, "One of the biggest barriers to GaN-based power design has been the uncertainties around driving GaN FETs and the resulting parasitics due to packaging and design layout....with power-conversion topologies with frequencies up to 5 MHz."

Complete article, here

Complete embedded UI has A8 processor, OS and display

In andersDX' entry-level DX series Embedded Display platforms, the DX1 is priced at under £100 (UK£ or equivalent) and based on the latest 600MHz ARM Cortex A8 processor from Texas Instruments. DX1a has a pre-integrated TFT touch screen, and a pre-loaded Linux, Android or Windows CE operating system, and becomes a fully working prototype system with the addition of the application software. Booting the DX is no harder than booting a new desktop PC, and developers just need to load their application to create a prototype platform for their design that can be used to prove a concept and deliver a working demonstration.

Complete article, here

LED light-engine-maker endorses thermal substrate

Lumichip, developer of LED packaging solutions and light engines, has designed-in Cambridge Nanotherm's thermal management technology for LEDs. Using the material has contributed to the performance of Lumichip's newly-released range of high performance UVA LED modules. Thorough tests by Lumichip revealed Cambridge Nanotherm's technology to have "extraordinarily" high levels of thermal performance, well beyond that of comparable solutions and approaching the thermal performance of exotic ceramic materials such as AlN (aluminium nitride). Lumichip has outlined its test results in a free-to-download whitepaper.

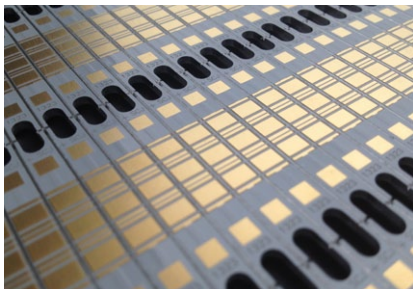

Complete article, here

productroundup

Atmel adds ARM Cortex-M based MCU at 35µA/MHz

Staking a claim for the lowest power microcontroller based on the ARM Cortex-M core series, the Atmel | SMART SAM L Family uses one-third the power of alternative offerings. With power consumption down to 35 µA/MHz in active mode and 200 nA in sleep mode, the SAM L21 combines ultra-low power with Flash and SRAM that are large enough to run both the application and wireless stacks. The family integrates the company's proprietary ultra-low power picoPower technology. While running the EEMBC ULPBench benchmark, the SAM L21 achieves a score of 185, the highest publicly-recorded score for any Cortex-M based processor or MCU. Atmel's SAM L21 family consumes less than 940 nA with full 40 kB SRAM retention, real-time clock and calendar and 200 nA in the deepest sleep mode.

Complete article, here

LED driver delivers up to 40A pulsed

LT3744 is a synchronous step-down DC/DC converter designed to deliver constant current to drive high current LEDs. Its 3.3V to 36V input voltage range suits it for applications including industrial, DLP projection and architectural lighting. It uses two external switching MOSFETs, delivering up to 20A (80W) of continuous LED current from a nominal 12V input. In pulsed LED applications, it can deliver up to 40A of LED current or 160W from a 12V input. Delivering efficiencies up to 95%, it can eliminate the need for external heat sinking.

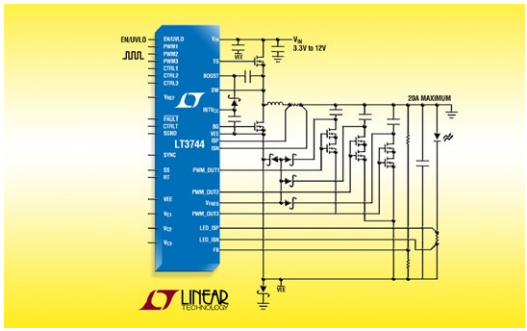

Complete article, here

10-Ω, solid-state switches build multiplexers, with fault detection/protection

A solid-state switch for a wide variety of signal-routing applications, including relay replacement, in areas such as analogue I/O, process control, data acquisition, instrumentation, avionics, ATE, and communications systems, the ADG5436F contains two independently controlled single-pole/double-throw (SPDT) switches. They conduct equally well in both directions when on, and have an input signal range that extends to the supplies. They can safe handle tens of mA in the switched channel, up to 113 mA depending on drive and temperature conditions, and switch in around 0.5 µsec.

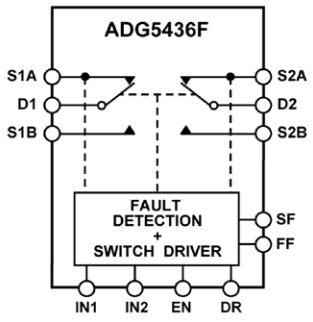

Complete article, here

3D stacked structure flash memory at 16 GB

Toshiba has announced the development of the world's first 48-layer three dimensional stacked cell structure flash memory called BiCS, a 2-bit-per-cell 128-Gigabit (16 GigaBytes) device. The BiCS is based on a 48-layer stacking process, which Toshiba says enhances the reliability of write/erase endurance and boosts write speed, and is suited for use in diverse applications, primarily solid state drives (SSD). The structure stacks flash memory cells in a vertical direction, from a silicon plane, which allows significant density improvement over conventional NAND flash memory, where cells are arrayed in a planar direction on a silicon plane.

Complete article, here

productroundup

IO-Link temperature sensor reference design

Maxim Integrated's IO-Link smart temperature sensor reference design lowers cost and increases uptime for industrial control and automation; the company says it is more flexible, robust, and lower cost than any discrete alternative, while consuming minimal power and providing better than ±0.5°K accuracy. The temperature appears immediately on an LED display for a quick and convenient snapshot of coarse temperature values. It supports all three IO-Link speeds. The MAXREFDES42# is available for use with a 2-wire, 3-wire, or 4-wire PT100 RTD over a wide temperature range.

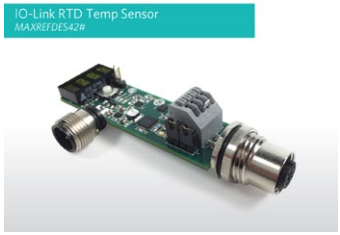

Complete article, here

EDA tool automates exploration & design of application-specific instruction sets

Synopsys' ASIP Designer tool speeds the design of application specific instruction set processors (ASIPs) and programmable accelerators, with a language-based approach that allows the automatic generation of synthesisable RTL and software development kits (SDKs) from a single input specification, accelerating the processor design and verification effort by up to five times compared to traditional manual approaches. ASIP Designer enables users to explore multiple processor architecture alternatives in minutes. Using a single input specification in the nML language, the tool automatically generates both the synthesisable RTL of the processor as well as an SDK that includes an optimising C/C++ compiler, instruction set simulator, linker, assembler, software debugger and profiler. The compiler generation technology includes an LLVM compiler front end and support for the OpenCL kernel language.

Complete article, here

Software simulation of ST power devices

A free software tool from STMicroelectronics includes accurate dynamic simulation for VIPower devices. TwisterSIM is a free software suite that includes a Smart Interactive Selector and a Dynamic Electro-Thermal Simulator for Automotive Low/High-side driver/switches and H-bridges for motor control built using ST's VIPower technology. The tool allows complex evaluations with accurate dynamic simulations of load-compatibility, wiring-harness optimisation, fault condition impact analysis, diagnostic behaviour analysis and dynamic thermal performance.

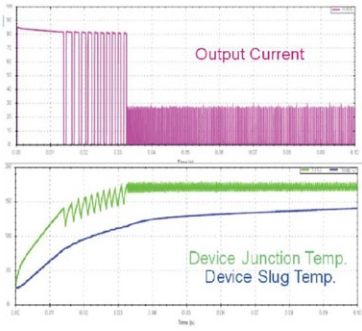

Complete article, here

Formal-methods checking of FPGAs in hi-rel designs

Microsemi and OneSpin Solutions (Munich and San Jose, California) are targeting high-reliability design verification with a formal-based FPGA equivalency checking solution; OneSpin's 360 EC-FPGA is used to verify Microsemi FPGA designs, including SmartFusion2 SoC FPGAs and IGLOO2 FPGAs, for safety-critical applications. OneSpin 360 Equivalence Checking (EC)-FPGA verification software now fully supports Microsemi's Libero System-on-Chip (SoC) design flow. Equivalency checking, Microsemi notes, has become a critical component in the verification of high-reliability designs such as safety critical components, ensuring that no issues are introduced during the design refinement process. OneSpin's EC-FPGA product augments Microsemi's Libero design flow to ensure the functional consistency of high-reliability designs throughout design refinement, reducing the risk of an end-product fault.

Complete article, here

productroundup

600V GaN transistor in a TO-247 package

This low on-resistance "Quiet Tab" TPH3205WS enables up to 3 kW high-efficiency inverter designs and titanium class power supplies without the need for parallel-connected transistors.

Transphorm is now offering engineering samples of its TPH3205WS transistor, the first 600V GaN transistor in a TO-247 package. Offering 63 mOhm R(on) and 34A ratings, the device uses the company's Quiet Tab source-tab connection design, which reduces EMI at high dv/dt to enable low switching loss and high-speed operation in power supply and inverter applications.

Complete article, here

5 GHz power amplifier module extends WLAN range

A PA module from Microchip cuts the cost and complexity of implementing wireless LAN installations based on the IEEE 802.11ac Wi-Fi standard. SST11CP22 is a 5 GHz power amplifier module (PAM) for the IEEE 802.11ac ultra high data rate Wi-Fi standard. This PAM delivers 19 dBm linear output power at 1.8% dynamic EVM with MCS9 80 MHz bandwidth modulation. It delivers 20 dBm linear power at 3% EVM for 802.11a/n applications, is spectrum mask compliant up to 24 dBm for 802.11a communication, and has less than - 45 dBm/MHz RF harmonic output at this output power.

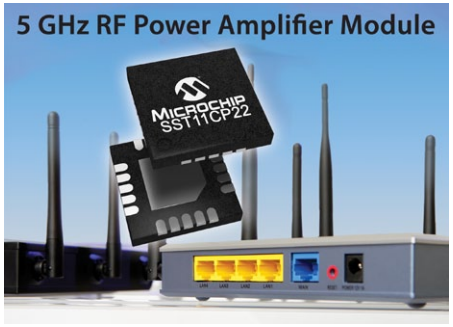

Complete article, here

Low-cost arbitrary/function generator

Tektronix' AFG1022 offers a high level of price/performance, and integrates with the TekSmartLab Wireless Lab Management configuration. Compared to scope-based signal generators or stand alone AFGs in its price class, the AFG1022 offers better performance and greater flexibility, Tek asserts. Key performance specifications include dual-channel, 25 MHz bandwidth with 1 mVpp to 10 Vpp output, 14-bit vertical resolution and 1 µHz frequency resolution. It provides a 125 Msample/sec sample rate along with 64 MB of built-in non-volatile memory and USB memory expansion for user-defined waveforms.

Complete article, here

Rogowski current probes measure power semiconductor switching currents

Power Electronic Measurements' CWT MiniHF is an AC probe with improved common mode immunity to local high voltage transients, and more precise measurement delay which can be compensated for to give improved power loss measurement. Combining a novel shielding technique, low sensitivity coil and low-noise signal-conditioning circuitry, the CWT MiniHF wide-band screened probe boosts immunity to local dV/dt transients while maintaining small size, flexibility, and 3 dB bandwidth of up to 30 MHz for a 100 mm coil. The probes feature a coil 4.5 mm thick with 5 kV insulation voltage, and can handle maximum current slope of 100 kA/µsec.

Complete article, here

BY RICHARD TOMKINS

Phone a friend

I worked a number of years ago, in the '70s, for a large telephony manufacturer. When I joined the firm we were shipping the last Strowger Switch CO (Central Office) out the door and just starting to build the first digital CO.

I worked in the final test area at a test set, which had a –48V DC supply, ring generator, card interface, plug bank, switch bank, light bank, patch bank, rotary dial, and some other things I cannot remember. The test process involved interconnecting the various hardware for the card under test and then working your way through a manual switch procedure, observing that specific lights came on and others went off, and at times measuring relay closure make and break times with an oscilloscope. Sometimes, there was no need to measure the timing; simply listening for the relay closure was sufficient.

An intermediate transition point on the way to the full digital CO switch was the use of reed relay matrix cards. We made our own reed relays on-site. We had an automated machine that cut glass tubes and cut raw material wire into proper lengths, magnetized the wire, positioned the two contacts, and surrounded them with the glass envelope, sealing them inside with an inert gas and annealing the ends of the glass. It was a fast machine kept in a clean room and it worked flawlessly. On another line, we made injection moulded plastic bobbins. Ultimately, two of those reed relays would find themselves inserted by another machine into these plastic bobbins and this assembly went on to another production line to wind coils around the plastic bobbins connecting the wire ends to pins inserted into the plastic bobbins by another machine.

All these in the thousands went to assembly and found their way into a module by the hundreds. All those coils on all the bobbins were driven by a matrix by silicon controlled rectifiers (SCRs).

We made PCBs that had approximately 100 relays on them. These PCBs were four layer and had no silk screen, top or bottom. Initially, all of our production was hand assembly, but to reduce costs and time to build, we introduced pick and place machines into the operation.

We started having intermittent problems with these relay boards shortly after our first batch had had their through-hole components, resistors, diodes, and some ICs, inserted by the pick and place machines. The remaining components, transistors and relays, were all hand inserted, some hand soldered.

I decided to work on this problem personally before calling in the engineering support staff. They had shut the production line down earlier and that was costly. I found that certain boards failed at certain test stations and after a few hours of testing (remember this was on telephony test set and the test took close to an hour per module), I had also found that some boards failed when slightly twisted, even though they had passed 100% earlier in the evening.

In those days, we didn't know anything about BPST (basic problem solving techniques), but we did have a break in the cafeteria over coffee and cigarettes, and that was pretty close to the same thing. We brainstormed on this amongst ourselves and decided that we would look at how the components were being inserted by the pick and place machines. Right away, we noticed that hand placed components were mounted slightly above the board, not completely flush on the board, with one end or both ends of a two wire through-hole component being off the board surface. The pick and place machines punched the components into the board, at least that's what the rapidity of its operation looked like to us, and components were flush with the PCB board surface.

It looked brutal, so we started to surmise that the machines were physically damaging components. We got the component laboratory involved to remove components and test them and there were no failures. The diodes on these boards were wired across each relay's primary coil to clamp the CEMF (counter electromotive force) and prevent the drive transistors from blowing.

While one of the guys was probing transistors and diodes with a scope probe (by now I had the whole test group interested in solving this problem before the engineers did), he noticed he could make the fault occur when he placed more force on the oscilloscope probe against the diode's PCB trace. I pulled out my pliers and gently lifted the lead on the diode and the fault permanently went away.

I told the engineers that the diodes were shorting out the PCB traces, and instead of scrapping everything we re-worked the boards by lifting the diodes off the PCB board surface on one end.

Richard Tomkins *is a consultant, involved in technical writing, web development, computer support and a founder of PiXCL Automation Technologies, developing a language for ARM MCUs. https://www.linkedin.com/in/tomkinsrichard*

www.electronics-eetimes.com

April 2015

EE Times europe

For the creators of technology

Wearables for therapy adherence

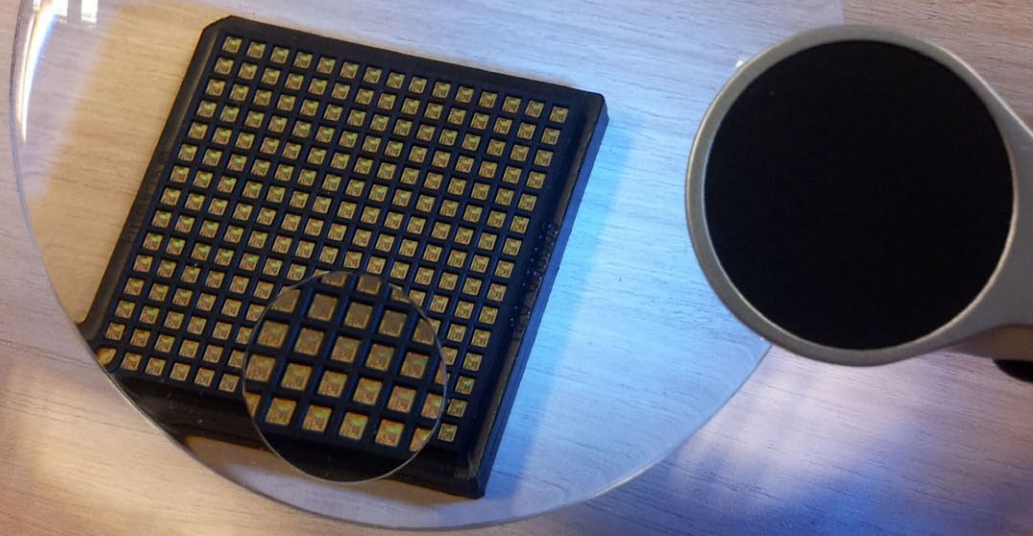

Special focus:
Optoelectronics

Iconic Insights:
Jaguar Land Rover's Wolfgang Ziebart

EDN europe

EDN-EUROPE is published 11 times in 2015 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

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

SALES CONTACTS

Europe

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

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

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

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

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

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

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

USA & Canada

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

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

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

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

Asia

Keita Sato
Japan
+81-3-6824-9386
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