EDN europe

JULY - AUGUST 2016

www.edn-europe.com

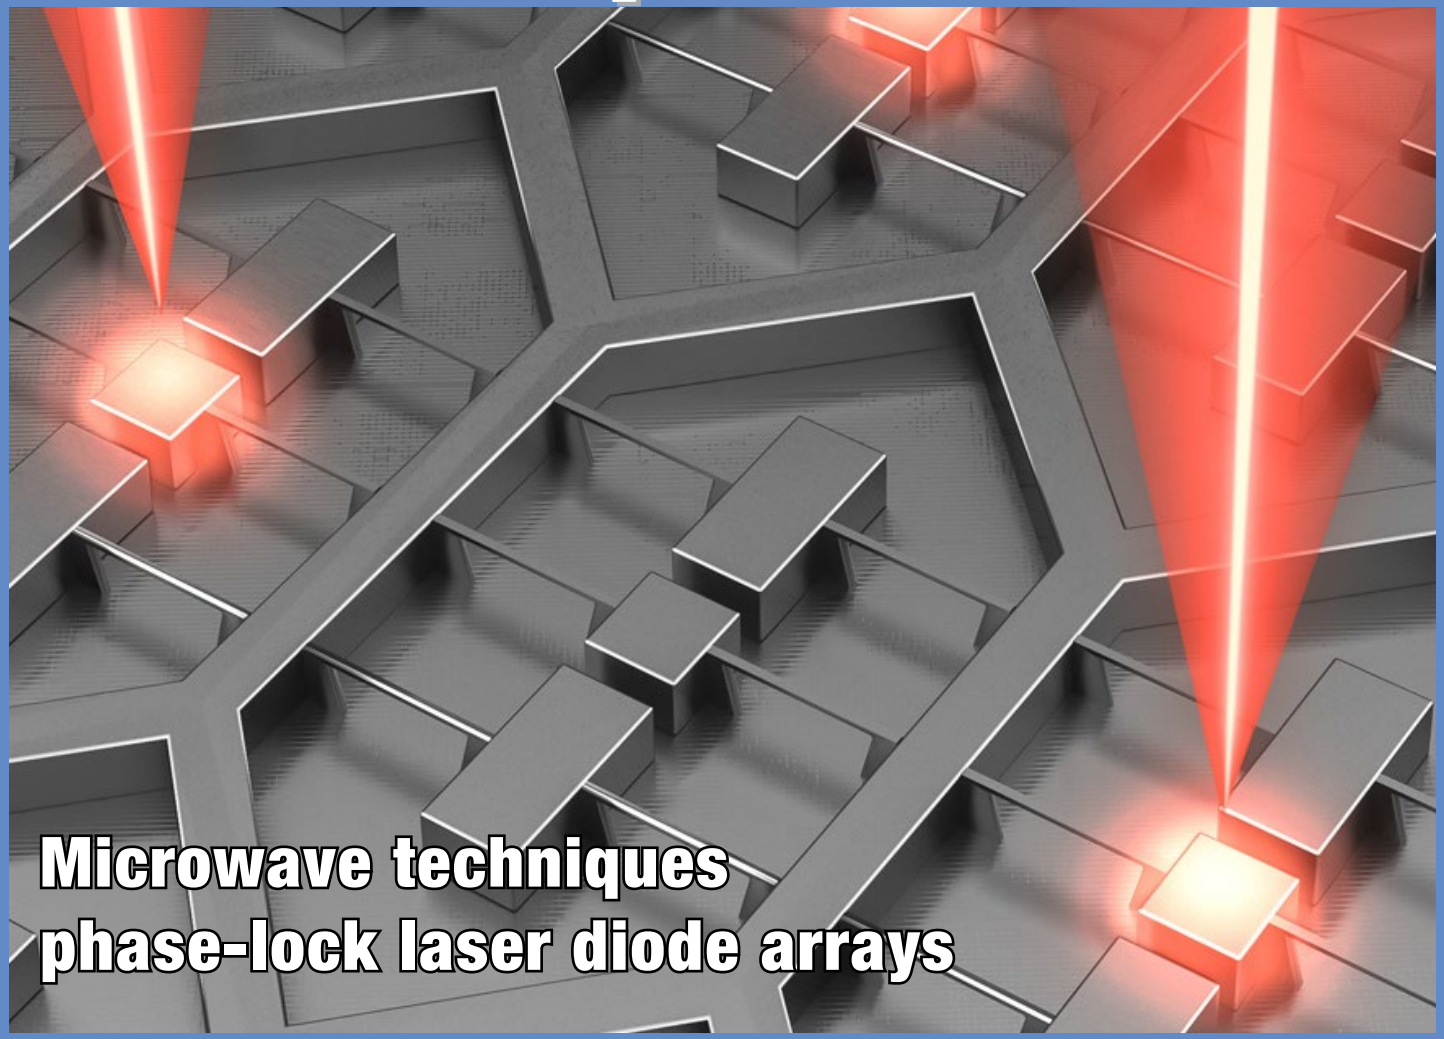

Microwave techniques phase-lock laser diode arrays

pulse

designideas

COVER

Microwave beam-forming techniques shape micro-laser arrays

Researchers at MIT and Sandia National Laboratories have designed a device that is an array of 37 microfabricated lasers on a single chip. Its power requirements are relatively low because the radiation emitted by all of the lasers is phase locked. The development, say the researchers, represents a new approach to microlasers, with the prospect of using such arrays in, among other applications, terahertz security scanners. Diagnostic or security applications in that area would require compact, low-power, high-quality terahertz lasers. The researchers describe a new way to build terahertz lasers that could significantly reduce their power consumption and size, while also enabling them to emit tighter beams. The work also represents a fundamentally new approach to laser design, which could have ramifications for visible-light lasers as well. "This whole work is inspired by antenna engineering technology," says Qing Hu, a distinguished professor of electrical engineering and computer science at MIT, whose group led the new work. "We're working on lasers, and usually people compartmentalize that as photonics. And microwave engineering is really a different community, and they have a very different mindset. We really were inspired by microwave-engineer technology in a very thoughtful way and achieved something that is totally conceptually new." Full story

FEATUREARTICLES

ONLINE THIS MONTH

Near-field scanners let you see EMI

White Rabbit networks for PTP users

Sensing current on the high side

EDN's columns

NOT VR BUT MR

At the Technology Forum presented by the imec research centre in late May ("ITF"), an invited presentation was given by Microsoft on the topic of its Hololens headset product. Hololens is conceived as a means of providing its wearer (operator? host? user? - there is, perhaps, a whole new niche of terminology to evolve here) with what Microsoft terms "mixed reality". Rather than transport the wearer (I'll stick with "wearer" for now) into a different space, it aims to insert convincing holographic images into the space that he or she is actually occupying. So; not virtual reality (there is lots of that around at present), but mixed reality. The Hololens is worth outlining, if for no other reason, to convey some measure of the scale of engineering effort that is required for a project of this ambition.

The headset inserts an image, with full depth information and focussed at the correct distance (hence, "holographic") into the wearer's visual field. The image therefore needs to be solidly located in that space. As well as capturing that space and placing its image data into it, the device accepts input from the wearer in gestures, movements and voice; it is a head-mounted, full-featured computer; an autonomous unit that doesn't rely on external markers or beacons, cameras, wired connection or host PC.

The Microsoft presentation included an exploded view which I can't bring you here, but those who know anything of preparing a consumer product for production would recognise that the plastics moulding tooling costs alone must have been spectacular – before even considering the electronics NRE. The headset

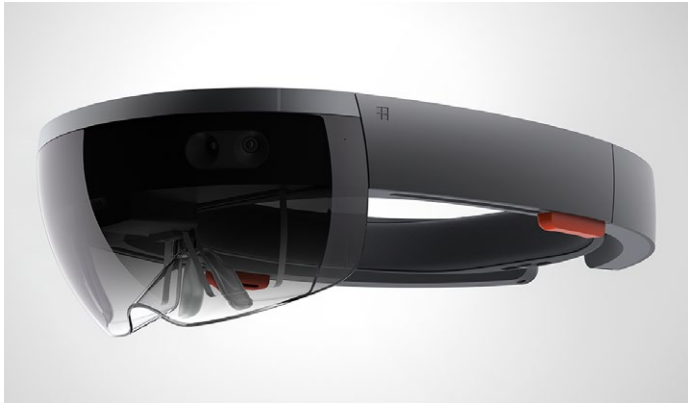

contains four environmental cameras (imaging the surroundings) plus a depth camera aided by time-of-flight sensing – and a video camera and ambient light sensor. Depth information is gathered at short range to capture the wearer's gestures.

You might suppose that the image would be some form of miniature projection head-up-display; it is not. The lenses carry fine gratings, with R/G/B feeds via optical waveguides, comprising 2.3 million light points. Resolution has to be expressed with a new metric, as 2,500 pixels/radian. That's enough, Microsoft says, to render small fonts on virtual documents. Almost incidentally, there's an audio channel with speakers at each ear.

We come, inevitably, to the silicon. The basic architecture is X86, with a CPU plus HPU – that is, a holographic processing unit which is an accelerator with an instruction set and functionality focussed on the holographic image generation task. There is 64 GB of flash memory, and 2 GB of RAM, with a custom memory architecture that can handle the flood of data incoming from the array of sensors. It's always-on, has no sleep mode and has no heatsinking or fan. Naturally, this simply would not be possible were it not for the latest generations of silicon process technology – which is somewhat the point of imec inviting the presentation as part of its ITF. The main CPU is in 14 nm FinFET, while the HPU is in (mature by these standards) 28 nm – this will be migrated to a newer process, Microsoft assures us.

It is perhaps a statement of the obvious that only the largest corporations can contemplate the research and NRE investment needed to realise such a product – especially in the context of so much activity around VR and with no certainty which of the many VR use cases will "fly". For those of us not in that league, we can reflect that a whole new space is opening up in HMI terms (human/machine interfaces) that could create a major opportunity in add-on products and accessories.

pulse


blox
Model: SARA-N200
900 MHz NB-IoT

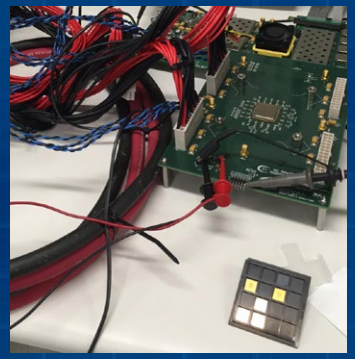

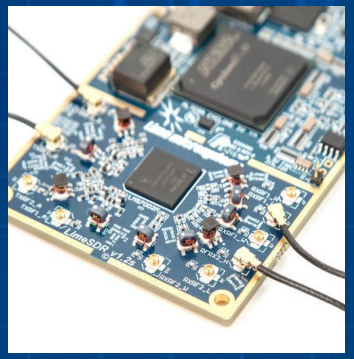

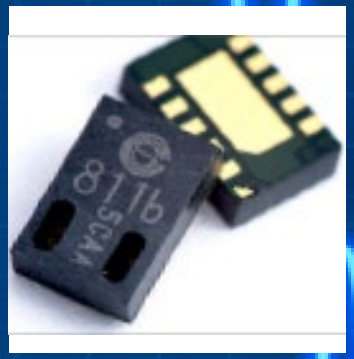

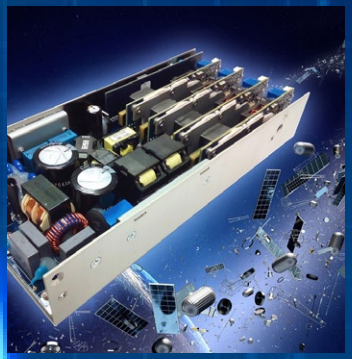

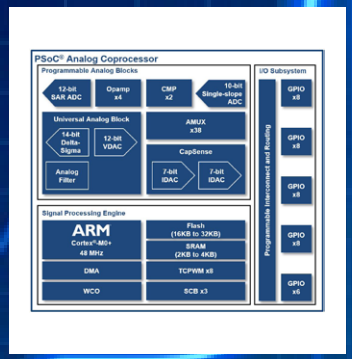
PSoC® Analog Coprocessor
Programmable Analog Blocks
12-bit SAR ADC
Opamp x4
CMP x2
10-bit Single-slope ADC
Universal Analog Block
14-bit Delta-Sigma
12-bit VDAC
Analog Filter
AMUX x38
CapSense
7-bit IDAC
7-bit IDAC
I/O Subsystem
GPIO x8
GPIO x8
GPIO x8
GPIO x8
GPIO x6
Programmable Interconnect and Routing
Signal Processing Engine
ARM
Cortex®-M0+
48 MHz
Flash (16KB to 32KB)
SRAM (2KB to 4KB)
DMA
TCPWM x8
WCO
SCB x3


AD7768

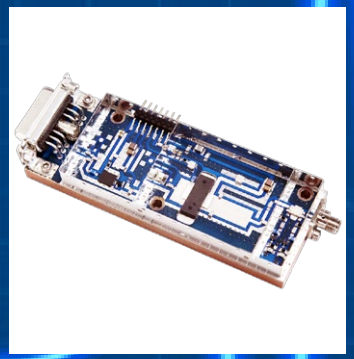


World's smallest ToF ranging sensor
VL53L0X

u-blox' first module for narrow-band cellular IoT, low data rate comms

Swiss company u-blox has set out the specifications of its forthcoming SARA-N2 Narrow-band IoT (NB-IoT) module, calling it the world's first cellular radio module compliant to the 3GPP Release 13, Narrowband IoT (LTE Cat. NB1) standard.

NB-IoT will provide a low-data-rate service as an extension of existing cellular (LTE) services. It promises very high levels of coverage and of penetration – reach of signal into difficult, 'deep-indoors' areas; and low power, with (depending on duty cycle) 10-year-plus battery life. It will also be structured to be 'cost-optimised' according to u-blox, and will have the benefits of being a broadly-based industry standard, in licensed bands, with quality-of-service guarantees.

The u-blox module has been designed for applications such as smart buildings and cities, utilities metering, white goods, asset tracking, and agricultural and environmental monitoring, it will operate for between 10 and 20 years from a single-cell primary battery. Its 16 x 26 mm LGA form factor, using u-blox nested architecture, facilitates simple upgrades from u-blox GSM, HSPA or CDMA modules and ensures future-proof, seamless mechanical scalability across technologies.

The SARA-N2 module provides secure, private communications over licensed spectrum with guaranteed quality of service. It supports peak downlink rates of up to 227 kbps and uplink rates of up to 21 kbps. Simultaneous support for three RF bands means that the same module may be used in most geographic regions.

The benefits of NB-IoT over other cellular radio technologies include lower device complexity, ultra-low power operation and support for up to 150,000 devices per single cellular cell. The technology offers a 20 dB link budget improvement over GPRS to give excellent performance under poor coverage conditions such as underground or inside buildings. This margin, u-blox says, can potentially be traded for additional battery life.

The company offers its comparison with with Low Power Wide Area (LPWA) services operating in unlicensed spectrum; "NB-IoT offers greater security and freedom from interference because it uses a licensed spectrum based network.

Other advantages include lower latency than mesh networks, thanks to its point-to-point topology, the ability to run it adjacent to existing 2G and LTE networks - it needs just 200 kHz of bandwidth - and a higher transmit power limit, which improves reliability and range. It also allows for robust 2-way communication which means that features such as firmware upgrade over the air are achievable. Furthermore, global roaming is possible with NB-IoT, which is not the case with localized unlicensed spectrum based technologies."

Complete article, here

Sub-GHz transceiver from ADI boosts link budget, cuts power

Analog Devices presents its latest transceiver IC as providing reliable radio connections and extended battery life for IoT and other wireless applications; it claims more reliable wireless radio connections with fewer retries and packet losses as well as longer battery lifetime.

The ADF7030-1 features a sensitive receiver and superior immunity to signal interferences as well as low power modes, including the lowest-available current in sleep mode. The transceiver is aimed at Internet of Things (IoT) devices, smart metering, security and building automation, industrial control, and wireless sensor networks.

The chip supports narrowband and wideband operation across sub-GHz ISM bands and data rates from 0.1 kbps to 300 kbps using 2GFSK modulation. The radio transceiver also supports Wireless M-Bus (WMBUS) and IEEE 802.15.4g-based protocols. The ADF7030-1 transceiver features a smart wake mode that provides low overall power consumption enabling long lifetime in battery operated systems while maintaining RF performance. The device can enter a low power sleep mode consuming only 10 nA with memory retained.

It features an on-chip ARM Cortex-M0 radio processor that performs radio control and calibration, including time sequencing that reduces engineering development time. A highly programmable packet handler saves more valuable time by simplifying programming interfaces and code development on the host microprocessor. In a 6 × 6 mm, 40-lead, LFCSP, the device is priced at $1.99 (1000). The ADF70301-868EZKIT is $599.00.

Complete article, here

LoRa low-power WAN evaluation kits from Microchip, in distribution

Distributor RS Components (RS) has Microchip's range of LoRa development kits; LoRa implements low-power wide-area networks, and the kits integrate gateway, sensors and local server application, and come with certification for Europe and North America, enabling fast and easy development possibilities for the IoT.

LoRa technology, overseen by the LoRa Alliance, targets tracking and monitoring IoT applications with low data rates and a low-duty cycle for a variety of markets such as energy, location, utility infrastructure, environment, agriculture and public safety. Predominantly used for the uplink of sensor data, the bi-directional nature of the communications allows real-time acknowledgement of mission-critical data and downlink control of remote actuator nodes.

The technology is capable of securely delivering two-way communication at data rates from 0.3 to 50 kbps, and over distances of up to 2 to 5 km in an urban environment and up to 15 km in a suburban environment.

Using chirped spread-spectrum modulation, in licence-free bands, enables LoRa links to maintain reliable communications over km distances for very low supply currents, with signals "in the noise floor". With extended range, much of the necessity for mesh networking that is typical of 802.11 derivative schemes is avoided, and point-point or star networks can

be used. Three modes of operation are available at increasing levels of supply power, and correspondingly increasing levels of access by node to gateway. The first enables a node to transmit at any time, and its transmission is also a request for a return data 'slot' so it receives only after transmitting. The second uses a network beacon to synchronise transmission slots from gateway to node(s), with nodes able to transmit between those times; and the third uses more power as receivers are active all the time the node is not transmitting.

The Microchip kits provide all the required components necessary for a developer to create a low-power LoRaWAN network, including: a LoRaWAN gateway; two Motes, which are LoRaWAN sensors based on Microchip's RN2483 or RN2903 LoRa modules; and a local LoRaWAN server application.

Complete article, here

1000 processor cores on a single chip; Californian researchers build 'KiloCore'

A team from University of California, Davis, Department of Electrical and Computer Engineering, has designed a 1000-core processor, with an ultimate throughput rate of 1.78 trillion instructions per second and containing 621 million transistors.

"To the best of our knowledge, it is the world's first 1,000-processor chip and it is the highest clock-rate processor ever designed in a university," said Bevan Baas, professor of electrical and computer engineering, who led the team that designed the chip architecture.

While other multiple-processor chips have been created, none exceed about 300 processors, according to an analysis by Baas' team. Most were created for research purposes and few are sold commercially. The KiloCore chip has been fabricated and run; it was built by IBM using its 32 nm PD-SOI CMOS technology.

The basic architecture is MIMD (multiple instruction/multiple data) and each of the 7-stage-pipelined cores is a general purpose unit with a 72-instruction set, single instruction/cycle. The team says that none of the instructions is 'algorithm-specific' - so distinguishing it from a GPU-class device. The 1.78 trillion instructions/sec figures comes with a clock speed of 1.78 GHz, at 1.1V: running at 0.84V and 1 GHz consumes 13.1W, while peak power efficiency of 5.8 pJ/Op is quoted at 0.56V and 115 MHz.

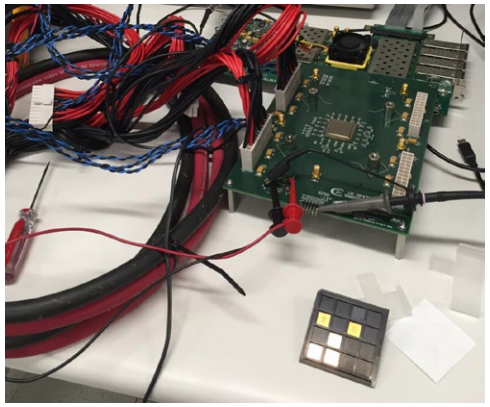

Each core is independently powered and can shut down to leakage-only power if it has no task to perform. Rather than a cache architecture, every processor can store instructions and data in a hierarchy of locations; local memory, one or more nearby processors, on-chip independent memory modules, or off-chip memory. Each processor communicates via a high-throughput circuit-switched network plus a packet-switched network (both on-chip). The team says there is little energy overhead to source operands from companion processors some way across the chip, as 'wormhole' routing is employed. That is, messages from an adjacent or nearby core will be routed via the 'circuit' network; those from further away in the processor matrix will travel via the packet network. Each core has north-south-east-west comms buffers plus a fifth channel for host-processor traffic; maximum throughput is 45.5 Gbps per router and 9.1 Gbps per port at 1.1V. At 0.9V, maximum throughput is 27.1 Gbps at 3.36 mW and at 0.67V, it is 8.1

Gbps at 429 µW.
A major challenge of working with high-number core arrays is scheduling tasks and keeping all the cores busy. The team has created a programming model and compiler; they say that programming is by a multi-step process that allocates programs to processors. However, to make use of available packaging, only the central 160 cores have been powered in tests; figures for full-chip performance are presumed to be extrapolations.

Complete article, here

Software radio apps are open-source on Ubuntu App Store

Lime Micro (London, UK) has announced that Ubuntu is putting together an App Store for LimeSDR that can be accessed once the LimeSDR crowd funding campaign successfully reaches its $500,000 pledge goal. The Snappy Ubuntu App Store will ensure the software defined radio (SDR) apps developed with the LimeSDR board are downloadable and those developed by Lime remain completely open-sourced.
Developers are, Lime says, already using the LimeSDR to develop apps with LTE, Bluetooth and LoRa, and future applications including IoT Gateway, 2G to 5G cellular Network in a box, drone command and control, utility meters, home automation and media streaming are anticipated by the team behind the LimeSDR.

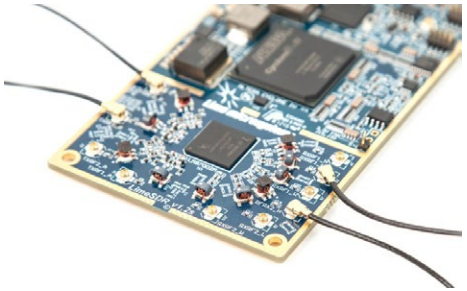

The LimeSDR platform is a low cost application-enabled software defined radio (SDR) platform that can be programmed to support virtually any type of wireless standard. From Wi-Fi, ZigBee and Bluetooth through to cellular standards such as UMTS, LTE and GSM and to the emerging IoT communication protocols such as LoRa, the platform offers multiple wireless connectivity opportunities.
The LimeSDR campaign, due to run until June 21st but still open at the time of writing, has generated the support of the individual backers as well as large corporations lead by EE, the UK's largest [cellular] operator. Individual boards can be pledged for $299 and corporate sponsorship bundles are available for $99,000. For more information on the LimeSDR campaign visit https://www.crowdsupply.com/lime-micro/limesdr

Complete article, here

Bluetooth SIG outlines Bluetooth 5 for 2017

The Bluetooth SIG (Special Interest Group) has announced that its next release, coming late 2016 to early 2017, will be called Bluetooth 5. It will quadruple range, double speed, and increase data broadcasting capacity by 800%.
The SIG says that the forthcoming version of Bluetooth technology will deliver "connectionless" IoT, advancing beacon and location-based capabilities in home, enterprise and industrial situations. Extending range will deliver robust, reliable Internet of Things (IoT) connections that make full-home and building and outdoor use cases a reality. Higher speeds will send data faster and optimise responsiveness. Increasing broadcast capacity will propel the next generation of "connectionless" services such as beacons and location-relevant information and navigation.

Complete article, here

ams builds gas/IR sensing portfolio with CCMOSS acquisition

Analogue, mixed-signal and sensor semiconductor maker ams (Premstaetten, Austria) is to buy Cambridge CMOS Sensors Ltd (CCMOSS) (Cambridge UK), specialist in micro hotplate structures for gas sensing and infrared applications; ams says that the deal positions it as [a] 'world leader in gas and infrared sensing for automotive, industrial, medical, and consumer applications.'
CCMOSS' micro hotplates are MEMS structures that are used in gas sensors for volume applications in the automotive, industrial, medical, and consumer markets. This is now added to ams' prior technology in MOX gas sensing materials to detect gases such as CO, NOx, and VOCs. CCMOSS manufactures these MEMS structures on CMOS wafers allowing the creation of complete monolithically integrated CMOS sensor ICs.

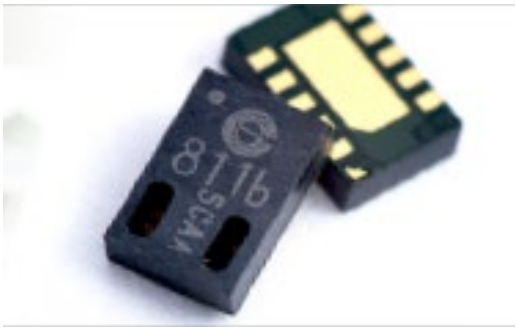

CCMOSS also brings a portfolio of IR technology comprising high performance IR radiation sources and detectors for sensor applications. ams sees this as complementary to its spectral sensing strategy for next generation optical sensor technologies,. CMOSS' IR sensing is based on the same monolithic CMOS structures as for gas sensing, enabling miniaturized implementations and efficient integration with other on-chip functions. Applications include CO_2 gas sensing and human presence detection and will extend into spectroscopic identification of organic materials.

Complete article, here

Fanless modular PSU outputs 600W from small 1U profile

Reliability is increasingly specified as a key parameter in power supply selection; and the single biggest detractor from reliability is a cooling fan; according to Excelsys, introducing its CoolX600 Series modular supplies that are rated for full output of 600W from a 8.5 x 4.5 in., 1U package.
The company positions the product as the only such modular, fanless supply available; and quotes an MTBF figure of 400,000 hours, with a 5-year warranty. It is a multiple-output, open-frame design that is rated at up to 94% efficiency: efficiency is preserved over a wide range of delivered power, only dropping to 90% below 25% of full load rating.

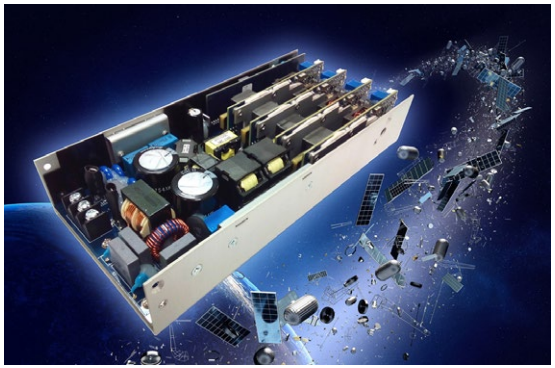

600W output is available with natural convection cooling, no airflow required, and also does not depend on heatsinking or attachment to a cold wall. The PSU provides higher input surge protection of 4 kV Line to PE for operation in harsh environments, reverse energy protection without the use of external blocking diodes as well as safety certified operation at altitudes of up to 5000m. A 24W, medically isolated, auxiliary supply is available as a standard feature, offering effectively another output for system intelligence, control, displays etc. The company notes the very many equipment designs that, even in standby, are 'always-on' to the extent of requiring a monitor to be powered. The auxiliary output

satisfies this need. Control of the supply can optionally be analogue or digital, via PMBus.
Excelsys (Cork, Ireland) says it has been able to achieve the product's specification by careful design without using "exotic" devices or topologies – the supply employs silicon power switches and the company has not yet found it necessary to adopt, for example, GaN devices. It categorises the architecture as 'modular resonant'; the design uses a high (380V) intermediate voltage, and advanced, multi-layer, planar transformer designs.

Complete article, here

Distributor Premier Farnell to be sold to Swiss company Datwyler

Premier Farnell plc (Leeds, England), distributor of electronic components (also operating as element14), has agreed a sale of the company to Datwyler Holding AG (Altdorf, Switzerland) for a cash offer of 165 pence (UK£1.65) per share.
This values the share capital of Premier Farnell at approximately £615 million (about $870 million) and implies an enterprise value of £792 million (about $1.1 billion). Datwyler is also an electronic component distributor and claims that the combination with Premier Farnell will be a strategic fit.
Premier Farnell had sales in the 2015/2016 financial year of £903.9 million (about $1.3 billion) up 2 percent from the previous year's sales of £886.6 million.
The board of directors of Premier Farnell have recommended the offer, which represented a 51 percent premium on the Premier Farnell closing price of 109.3p on June 13, 2016. The combined group will be able to achieve economies of scale compared to Datwyler's stand-alone strategy the companies said.
The combined group, with approximately 4,900 employees, is expected to generate revenue of approximately CHF1.8 billion (about $1.8 billion).

Complete article, here

Programmable-analogue SoCs for multi-sensor connected devices

Cypress' PSoC analogue coprocessor can be configured to provide custom interfaces for a variety of sensors in industrial, home appliance and consumer systems. The PSoC Analog Coprocessor, that comes in a 3.7 x 2.0-mm chip-scale package, is based on a 32-bit ARM Cortex-M0+ signal processing engine. With flash-based programming, it delivers a low-cost, fully programmable analogue front end with opamps, programmable gain amplifiers, analogue multiplexers, analogue-to-digital converters, analogue filters and digital-to-analogue converters.
Many IoT applications require multiple sensors and can benefit, Cypress proposes, from dedicated coprocessors that offload sensor processing from the host and reduce overall system power consumption. The PSoC Analog Coprocessor integrates programmable analogue blocks, includ-

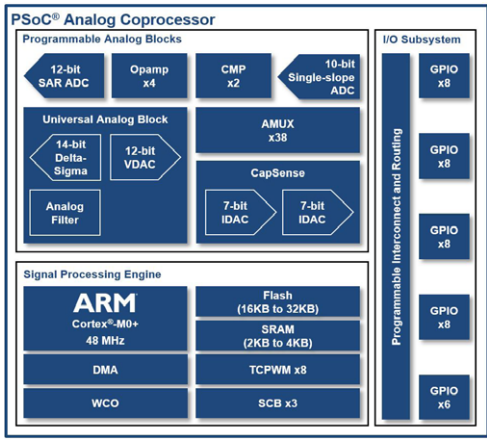

ing a new Universal Analog Block (UAB), which can be configured with GUI-based software components. This combination simplifies the design of custom analogue front ends for sensor interfaces by allowing engineers to update sensor features quickly with no hardware or host processor software changes, while also reducing BOM costs.
For example, in home automation applications, engineers can configure the PSoC Analog Coprocessor to continuously monitor multiple sensors, such as temperature, humidity, ambient light, motion and sound, allowing the host to stay in a standby low-power mode. Future design changes to support new sensor types can also be implemented by reconfiguring the programmable analogue blocks.

Complete article, here

Design win; Infineon ToF sensors add to 3D sensing in Lenovo smartphone

Lenovo has implemented "Tango" technology in a consumer product; Google's technology that allows devices to understand spatial information is an exclusive feature in the PHAB2 Pro smartphone. Based on the Time-of-Flight principle, the REAL3 image sensor chip from Infineon equips the smartphone with a 3-dimensional perception of the surroundings in real-time.
Infineon says it has the the only image sensor chip available that can meet Google's specifications. "Through the Tango-enabled PHAB2 Pro, we are opening a completely new era in mobile applications," said Hua Zhang, Vice President, Lenovo Android/ Chrome Computing Business Group. "The REAL3 image sensor chip from Infineon is a key component for Lenovo's PHAB2 Pro." Its users have instant access to Augmented Reality, benefiting from the phone's motion tracking feature which responds to physical movements, depth perception that measures the distance between objects and area learning capability which enables the device to recall the data recorded from previously visited locations.
The sensor chip integrates the pixel array, the control circuitry, ADCs and the digital high-speed interface on a single chip. The development was made in cooperation with pmdtechnologies, provider of the Time-of-Flight (ToF) technology.
This measuring principle involves infrared light. For each of its pixels, the 3D image sensor chip measures the time the light takes to travel from the camera to the object and back again. Each pixel detects the brightness value of the objects. Compared to other technologies, the ToF technology from pmdtechnologies offers the best spatial resolution and highest robustness both mechanically and concerning background light.

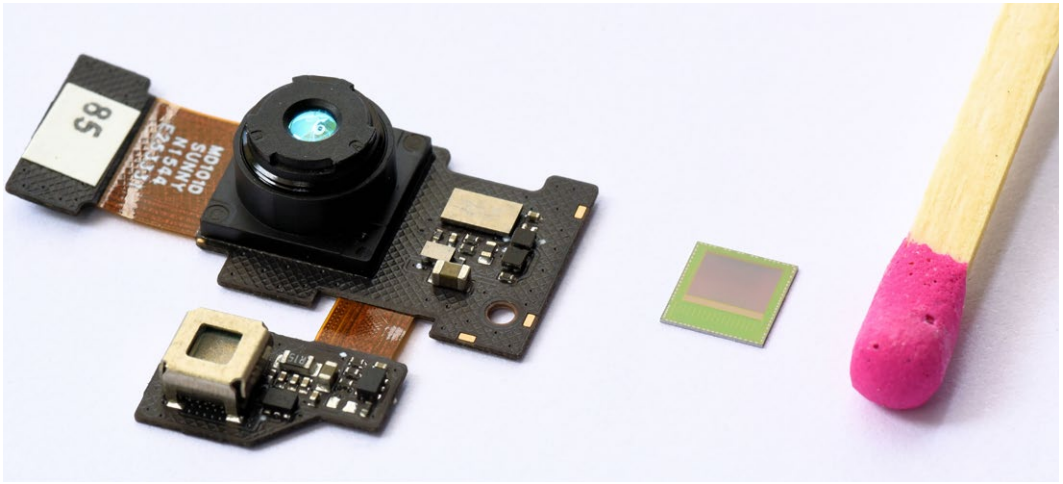

Complete article, here

USB-C & power from STM32-based design

STMicroelectronics has a certified embedded software solution based on its STM32 MCUs. The free STM32 USB-C and PD Middleware Stack is compliant with USB Type-C 1.2 and USB Power Delivery 2.0 specifications, allowing rapidly deployment into end-products.

The firmware stack, X-CUBE-USB-PD, initially based on the STM32F0 entry level Cortex-M0 processor, allows designers to upgrade their USB legacy devices to provide significant benefits to their end users. A benefit in terms of cost and PCB footprint is that solutions using the STM32 with the stack require only a very simple Analogue Front End comprising a few passive components, because it fully exploits STM32 embedded features such as comparators, ADCs, timers, and Direct Memory Access. Other features include; support for up to two USB-C ports (provider, consumer, or dual role); cable-insertion detection, plug orientation; identification of the role of the port partner attached and its current capability; Vbus Power negotiation via Power Delivery communication protocol; Vendor-Defined Messages are handled to identify device or cable ID or to manage Alternate Mode commands; maximum flexibility and adaptability versus evolving specification changes, as firmware upgrades are possible during the application lifetime.

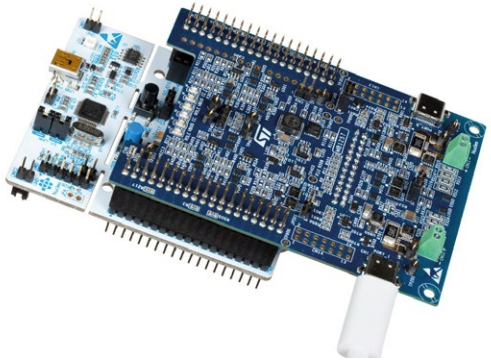

The processing bandwidth and available resources allow the MCU to perform other application-specific tasks, such as power-management control, USB2.0 communication, and/or voltage and current monitoring, on top of its USB-C functionality.

Complete article, here

Sigma-Delta ADCs step up acquired signal quality

Aimed at signal quality monitoring in instrumentation, energy and healthcare applications, Analog Devices' latest 24-bit simultaneous sampling sigma-delta A/D converters are configured for use in wide-bandwidth, high-density instrumentation, energy and healthcare equipment.

The AD7768 series includes a power scalable modulator and digital filter on each channel to enable the synchronized, precise measurement of both AC and DC signals in instrumentation applications, including modular data acquisition, audio test, and asset condition monitoring. The high throughput, fast settling response, and simultaneous sampling of the AD7768 series enables faster test times, which reduces testing costs and allows more efficient instrumentation design. The AD7768 series' high channel count provides healthcare devices, such as clinical vital signs monitoring equipment, with the means to significantly expand channel density while maintaining low power and high input bandwidth. The converters also deliver improved power quality monitoring through

the ability to detect harmonic distortion over a wider bandwidth for detection and diagnosis of grid imbalance.
The 24-bit, 8-channel AD7768 and the 24-bit, 4-channel AD7768-4 claim a 6-dB dynamic range advantage over the nearest competing products and deliver the best-available integral non-linearity (INL) performance across the widest available bandwidth, in addition to achieving 10 times better offset, a 30 times reduction in gain error, and a 2 times improvement in gain drift.

Complete article, here

RF energy/heating applications boosted by 300W GaN transistor

MACOM Technology Solutions recently announced its MAGe-102425-300, a 300W gallium nitride-on-silicon RF power transistor in plastic packaging. MACOM sees this as a significant step forward in enabling rapid growth in heating and energy-delivery applications; both for electrical efficiency reasons, and because of the economics of the technology/package combination. The applications that MACOM cites include; a new generation of domestic and commercial microwave ovens in which a GaN-based PA replaces the magnetron; light-emitting-plasma-based lighting; automotive ignition; a range of drying and curing processes in industry and even domestically, with the ubiquitous laundry drying a possible market; medical applications such as tumour ablation (targeted delivery of heat), tissue warming, and sterilisation by heat.

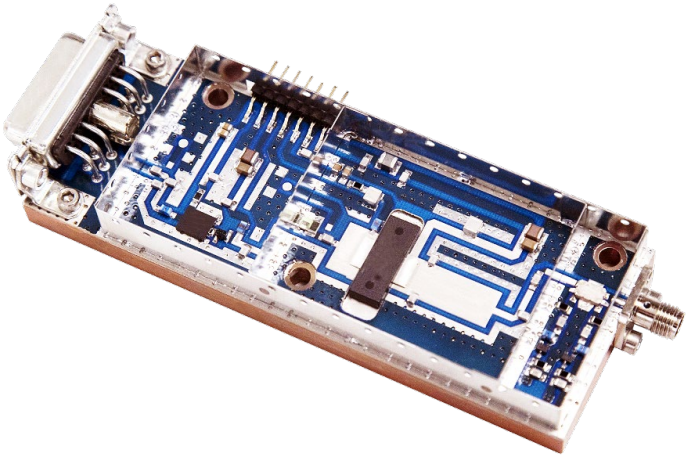

Where RF energy has been generated by solid-state amplifiers, the prior technology has been LDMOS silicon: GaN has been seen, MACOM explains, as too expensive because it has been based on GaN on silicon carbide substrates. (GaN devices are invariably fabricated in an epitaxial layer of GaN grown on a carrier substrate: in earlier generation, SiC was technically easier. Silicon is much cheaper in that role, not least because larger wafer sizes are possible, allowing the normal learning curves of semiconductor production, moving to 150 mm then 200 mm wafers, to be applied.)
MACOM has produced its current generation of parts not only using GaN-on-silicon, but moving to a moulded (plastic) package. There are four devices in this release, at 50, 100, 200 and 300W output; each offers gain of 17dB and efficiency of over 70%. The efficiency figure is a 'true' measure of DC-in to RF power out.
Using the 300W part, MACOM has also produced a PA module, a fully-assembled substrate of about 50 x 100 mm, for incorporation into a final product. In many cases – for example, manufacture of domestic microwave ovens – the appliance designers will have little experience of semiconductor amplifiers, or of routing RF power on PCBs: so a drop-in solution will be required. 300W out of a gain block with 17-dB added still requires significant drive, and MACOM has also produced an integrated module (above) with the complete RF path, and a digital control interface. The company expects to take these products, also, along

a familiar electronics integration/shrink/cost reduction path. Providing 300W output power and 70% efficiency at 2.45GHz, the MAGe-102425-300 meets the core technical requirements for next generation power amplifiers proposed by the RF Energy Alliance, a non-profit technical association dedicated to unlocking the full potential of RF energy. Meanwhile the cost structure and volume supply chain benefits achieved with MACOM's Gen4 GaN technology position the MAGe-102425-300 to meet aggressive cost targets on a par with LDMOS.

Complete article, here

Tiny time-of-flight ranging sensor has 2m range & improved speed

STMicroelectronics has added a faster and longer-range time-of-flight (ToF) ranging sensor, the second generation of its FlightSense technology, in the form of the VL53L0X laser-ranging module.

The VL53L0X extends the ToF measurement range to two metres, and is accurate to within ±3%. It is also faster, measuring the distance in under 30 msec, and energy-efficient, consuming 20 mW in active ranging mode and drawing 5 µA in standby. At 2.4 x 4.4 x 1 mm, it is also the smallest such device available, ST believes.

Unlike conventional infrared proximity sensors, the VL53L0X gives accurate distance measurement in millimetres, and is unaffected by the colour or reflectivity of the sensed target. Its FlightSense technology, used for laser-assisted camera auto-focus featured in several smartphones from top brands, can also distinguish movement towards or away from the sensor or from side to side. Distance is calculated internally and communicated over I²C, minimizing demands on the system host controller.

The sensing capabilities of the VL53L0X can support a wide range of functions including gesture sensing or proximity detection for new innovative user interfaces, wall detection, cliff detection and collision avoidance for robotic appliances such as vacuum cleaners, as well as hands-free actuation of home appliances or washroom devices such as taps, soap dispensers, hand dryers, and flushers. Other applications include user-presence detection for laptops or monitors for power-on/off control, and drones.

Complete article, here

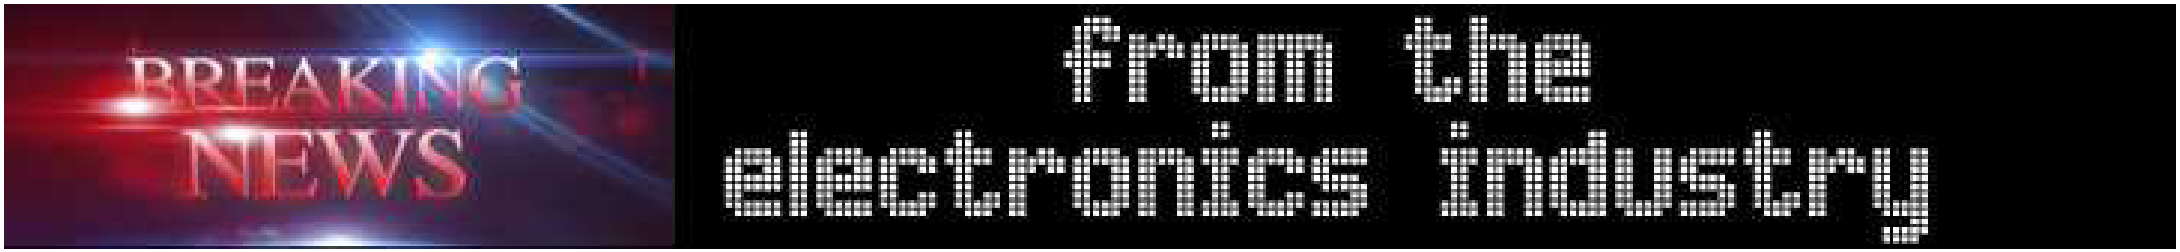

A CHECKLIST FOR DESIGNING RF-SAMPLING RECEIVERS

By Thomas Neu, Texas Instruments

The modern, advanced CMOS direct radio frequency (RF)-sampling data converter has been eagerly awaited by system design engineers for several major end-equipment manufacturers. This includes manufacturers of communications infrastructure, software-defined radios (SDRs), radar systems, or test and measurement products.

Recently introduced data converters are delivering the high dynamic range comparable to high-performance intermediate frequency (IF)-sampling data converters. Additionally, these converters integrate on-chip digital filtering (DDC), which reduces the output data rate from 3-4 Gsamples/sec sampling rate to something more manageable, and similar to traditional IF-sampling data converters.

Two major factors are driving the quick adoption of these ultra-high-speed data converters. The ever increasing demand for wider bandwidth naturally requires faster sampling rates, while higher density and integration is accomplished by removing one down conversion stage from the receiver, for example. Modern SDRs or cellular base stations need to be able to cover multiple frequency bands simultaneously, for example, to support carrier aggregation across multiple licensed Long-Term Evolution (LTE) bands to enable faster data traffic. Rather than expending one radio-per-band system, designers want to shrink the product form factor and build a multiband-capable radio. The RF sampling data converter removes the intermediate frequency (IF) stage saving printed circuit board (PCB) area and power consumption, while its wide Nyquist zone enables sampling multiple bands simultaneously.

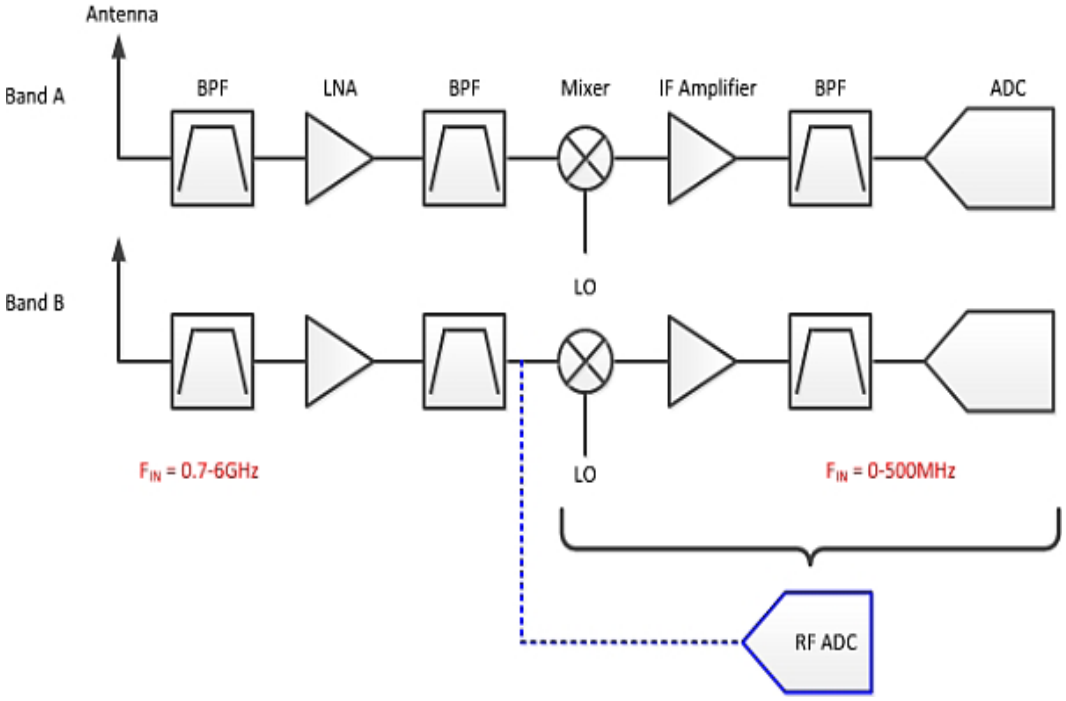

Figure 1. *One RF-sampling analogue-to-digital converter (ADC) can replace multiple IF-sampling signal chains.*

System designers who are considering switching from IF- to RF-sampling need to solve four primary challenges on their checklist:

- Receiver sensitivity
- Radio performance in presence of in-band interferer
- Filter requirements for out-of-band blocker
- Performance of the sampling clock source

Depending on their application, some may be more critical than others.

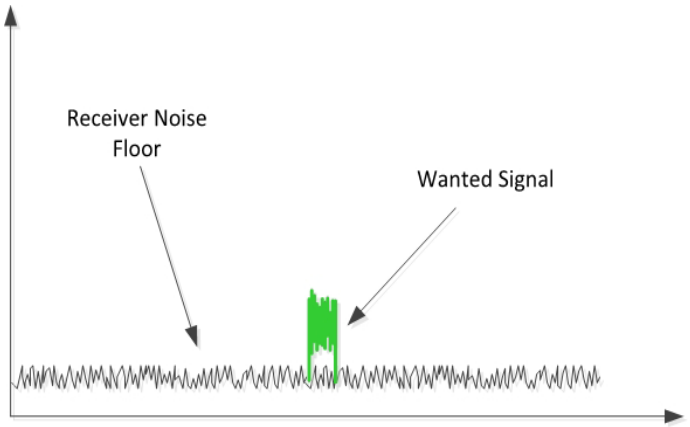

Figure 2. *Small wanted signal in sensitivity case.*

Let us examine these challenges using two different types of analogue-to-digital converters (ADCs) and compare the results. The first data converter is the ADS4249, a 14-bit, 250 Msample/sec ADC used for an IF-sampling system. The second is the ADC32RF45, a 14-bit, 3 Gsample/sec ADC for a RF-sampling system.

Receiver sensitivity

One basic performance metric of the receiver is its sensitivity, effectively the weakest signal power that it can successfully recover and process. Weak input signals cannot be demodulated if the noise of the receiver within the demodulated bandwidth is larger than the received signal itself. The noise floor of the receiver typically is expressed as a noise figure (NF) in decibel (dB), or the difference to the absolute thermal noise normalized to 1 Hz bandwidth. The most common way to improve an ADC's noise figure is to add an amplifier before the ADC.

Parameter	IF sampling ADC	RF sampling ADC
Sampling rate (FS)	250 MSPS	3 GSPS
Input full-scale (V_{PP})	2 V_{PP}	1.35 V_{PP}
Thermal noise	72.8 dB	62 dB
Input impedance (Zin)	200 Ω (external)	50 Ω (internal)
Calculated noise figure	**24.2 dB**	**26.8 dB**

Table 1. *Noise figure comparison between IF- and RF-sampling data converters.*

The noise figures of the IF- and RF-sampling converters can be calculated to yield the figures shown in Table 1.

While the noise figures of both converters are close, the IF-sampling data converter has significant external gain from the mixer and IF digital variable gain amplifier (DVGA), which substantially reduces the impact of the ADC noise figure to the receiver sensitivity. Hence, the RF-sampling ADC requires additional front-end gain (an additional low-noise amplifier or LNA) to minimize its impact on receiver sensitivity as well.

The article continues by considering in-band blocking performance, and the other topics in the above list of four. Click for pdf.

NOISY UNDER-SAMPLING IN SMART SENSORS

BY MARK LOONEY, ANALOG DEVICES

Many system developers are motivated to use fully-integrated, "smart sensors" as they typically offer convenient connectivity to accurate digital sensor data, which saves them from the perils of analogue circuit design, which can be a tedious challenge for the occasional user. While the motive to bypass analogue circuit design problems is understandable, it is also important for system integrators to understand important analogue behaviours in their smart sensors, such as bandwidth and noise, as they can influence important system-level decisions, such as data sampling and processing rates.

Consider an autonomous vehicle (AV) platform, which will use the gyroscopes from the ADIS16460 as feedback sensing elements in its Guidance Navigation Control (GNC) system. If the developer fails to consider the 330 Hz of bandwidth in these gyroscopes, they might be tempted to set the sample rate in their angular rate feedback loop, based solely on the motion profile that they expect form their AV platform. For example, if the GNC engineer believes that the AV platform's motion profile can be confined to less than 4 Hz of spectral content, collecting data at a rate of 40 SPS might seem like a conservative approach for the sample rate in the GNC's angular rate feedback loop. Unfortunately, without any pre-filtering, that "conservative" approach will actually under-sample the 330 Hz of bandwidth, which has a number of disadvantages. Figure 1 illustrates one of these disadvantages, which is in the reshaping distribution of noise energy over the resulting Nyquist frequency band (20 Hz).

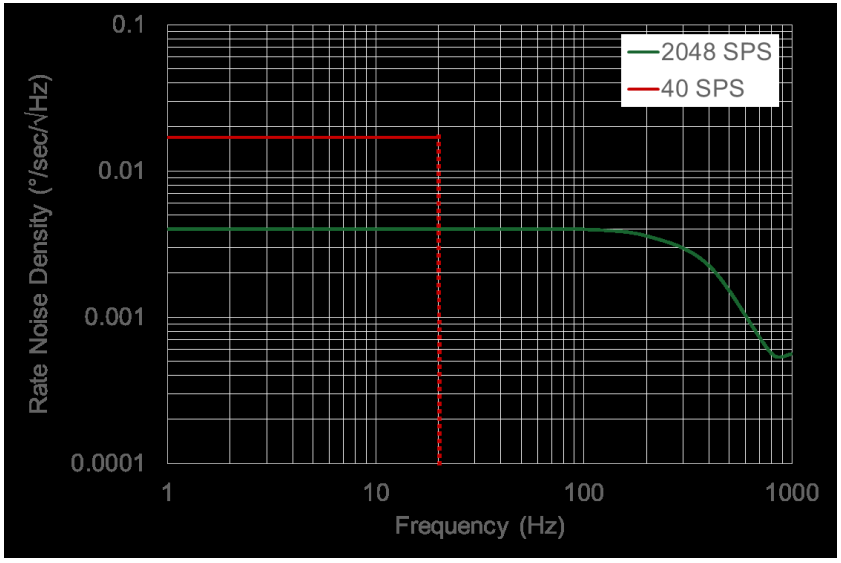

Figure 1. *ADIS16460 angular rate noise density*

[In Figure 1, the green curve illustrates the natural Rate Noise Density (RND), while the red curve illustrates the result of spreading the same total noise energy across the narrower, 20 Hz bandwidth. Assuming that total noise will be evenly distributed across the Nyquist band of the reduced sample rate (40 SPS), the following relationship predicts that the resulting Rate Noise Density will be ~0.017 °/sec/√Hz:

$$RND = \frac{TN}{\sqrt{f_{noise}}} = \frac{0.075\,^{\circ}/_{sec}}{\sqrt{20Hz}} = \sim 0.017\,^{\circ}/_{sec}/\sqrt{Hz}$$

This means that any digital filtering of the 40 SPS data will result in ~4x more noise than a comparable filter would, when using the full sample rate of 2048 SPS. The bottom line is that it is wise for system integrators to consider key analogue attributes in their smart sensors, as they may present an opportunity for performance optimization through appropriate sample rate selection and digital filter design.

CONQUERING CONCURRENCY PROBLEMS IN MULTICORE SYSTEMS

By Paul Anderson, GrammaTech

Concurrency takes on a new dimension in multicore platforms, since true parallelism comes into play, and communication between threads is often achieved using shared memory. Writing a correct concurrent program is notoriously difficult and the advent of multicore architectures makes it significantly harder again due to their added complexity.

Concurrency bugs, which include race conditions, deadlocks, livelocks, and resource starvation, are difficult to find when they manifest and even more difficult to diagnose. They require a new approach to verification that specifically addresses concurrency errors. The most effective way to reduce the risk of these bugs is to take a multifaceted approach that includes peer code reviews, testing, and most of all advanced static analysis that incorporates sophisticated models for concurrency.

Programming language support

Since most embedded developers are relatively new to multicore programming, the risk of introducing concurrency bugs is very significant. Today, C and C++ are still the most popular programming languages for embedded systems. However, one of the fundamental weaknesses of these languages is that they were not designed for concurrency. The most recent versions, C11 and C++11 introduced standardized support for multi-threading. Three features were added to address concurrency: a memory model that defines the behaviour of multi-threaded programs; atomic data types that can be safely accessed by concurrent threads; and several synchronization primitives such as locks and condition variables. Notwithstanding these improvements, the languages retain many of the core features of their ancestors that make writing multi-threaded programs very hazardous.

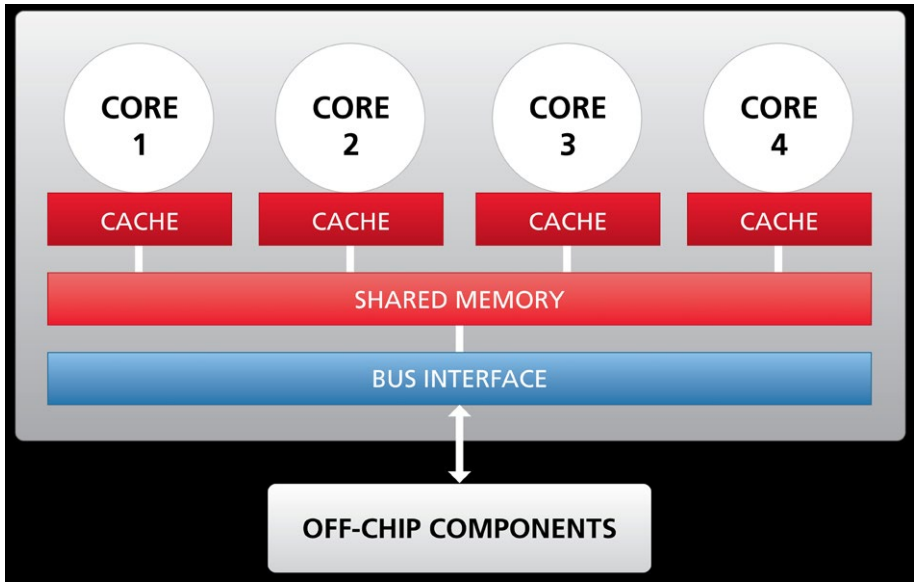

Figure 1. *Multicore processor platforms are vulnerable to concurrency bugs resulting from unsynchronized access to shared memory.*

At the same time, Java is increasingly popular with embedded developers. With 28% using it today it is now the third most popular language for embedded systems. In contrast to C and C++, Java has always had built-in support for multithreading within the programming language syntax, source compilers, and standard libraries. Additionally, Java 5 added the **java.util.concurrent** library, which was extended in Java 6 and Java 7 to provide extensive support for concurrent and parallel programming.

Many embedded designs use a combination of C or C++ and Java. For example, Java is very popular for automotive applications because it offers an easy way to program a user interface for a touch screen display or an entertainment system. Such applications may have

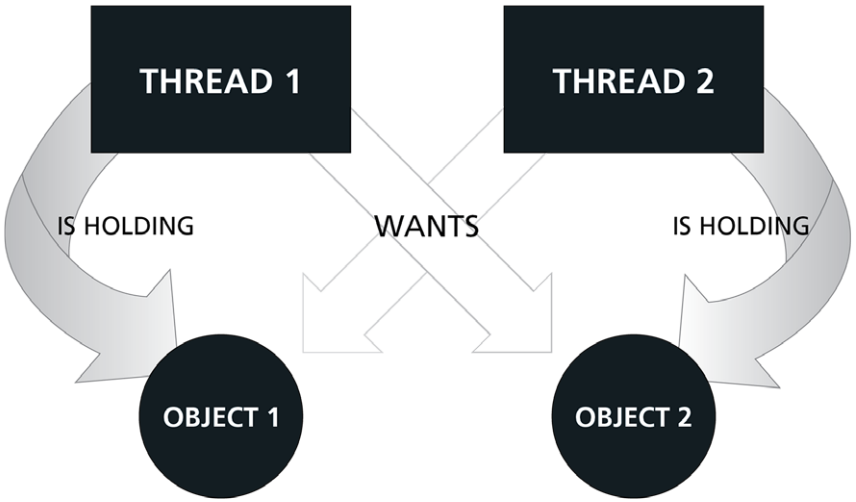

Figure 2. *In a deadlock, both threads are completely stuck, both unable to carry out its operations or get to the point where it can release its lock.*

many layers with safety critical code written in C, communicating with non-safety critical Java code running on a user interface.

Static analysis tools

Perhaps the biggest challenge with concurrent programs for multicore platforms is that no amount of testing can guarantee to find all concurrency bugs. It is the relative order in which instructions are executed in real time that is the main source of defects in multi-threaded programs. As multiple threads run, the relative order in which their instructions are executed varies depending on what other threads are active at the same time. If bugs are introduced through programming errors, non-deterministic interleaving can lead to unpredictable results. The number of possible interleavings increases enormously as the number of instructions grows, a phenomenon known as combinatorial explosion. Even the smallest threads have many possible interleavings. Real world concurrent programs have astronomical numbers of legal interleavings, so testing every interleaving is infeasible. Likewise, it is impossible to explore every potential execution path using peer code reviews or walkthroughs. This is where advanced static analysis tools excel.

These analysis tools use symbolic execution engines to identify potential problems in a program without actually having to run the program. They work much like compilers, taking source code as input, then parsing it and converting it to an Intermediate Representation (IR). Whereas a compiler would use the IR to generate object code, static analysis tools retain the IR, also called the model. Checkers perform analysis on the code to find common defects, violation of policies, etc., by traversing or querying the model, looking for particular properties or patterns that indicate defects. Sophisticated symbolic execution techniques explore paths through a control-flow graph, which is a data structure representing the order in which statements are executed during a program's execution. Algorithms keep track of the abstract state of the program and know how to use that state to exclude consideration of infeasible paths. The depth of the model determines the effectiveness of the tool. That depth is based on how much knowledge of program behaviour is built in, how much of the program it can take into account at once, and how accurately it reflects actual program behaviour.

This article continues with commentary on the respective roles of open-source and commercial analysis tools in identifying concurrency defects as they manifest themselves – click for pdf.

MAKE THE MOST OUT OF PRINTF

BY JACOB BENINGO

The oldest, tried-and-true debugging technique in embedded development is to sprinkle *printf* statements throughout the software in the hope of gleaning insight into system behaviour. Using *printf* is not always advisable, however, as it can have unforeseen real-time implications. Let's examine the fundamental issues with *printf* and then a few techniques that can be used to get the most performance from it.

The issues with *printf*

The use of *printf* comes with a few problems that developers often overlook. The first such problem is that the technique requires a developer to bring a standard C library into the software. This will undoubtedly increase ROM and RAM usage.

A second problem is that every time a *printf* statement is used, the system becomes blocked until all characters have been transmitted. This blocking can result in significant real-time performance degradation. Take, for example, the output of a simple string such as "Hello World!" to be printed out through a UART at 9600 baud (still a very common occurrence). I performed a simple timing measurement on an STM32 and, as shown in Figure 1, it took 12.5 milliseconds for the string to be formatted and printed to the terminal. During this time the system can do nothing else.

Adding any string formatting makes the situation even worse! Printing the system state to the terminal using *printf("The system state is %d", State)* results in a 21 millisecond application delay as the string is formatted and transmitted. One might argue that running at 9600 baud is ridiculous but even increasing to 115200 would still result in 1.05 and 1.75 milliseconds respectively to transmit these two messages. That's a lot of processor bandwidth and potential real-time performance hits for minimally useful information.

Now, on to how to address these problems.

Performance Technique #1 – Create a non-blocking *printf*

Every *printf* version that I have ever encountered in the wild has been the blocking type. Once a call to *printf* is made the application stops execution until every character has been successfully transmitted. Amazingly inefficient! An alternative, then, is to create a non-blocking version. A non-blocking *printf* version will:

- format the string
- stuff the formatted string into a transmit buffer
- initiate the transmission for the first character
- let an interrupt service routine handle the remaining characters in the transmit buffer
- continue executing code.

The big hit for a non-blocking *printf* is the setup time, which on the STM32 at 9600 baud I found

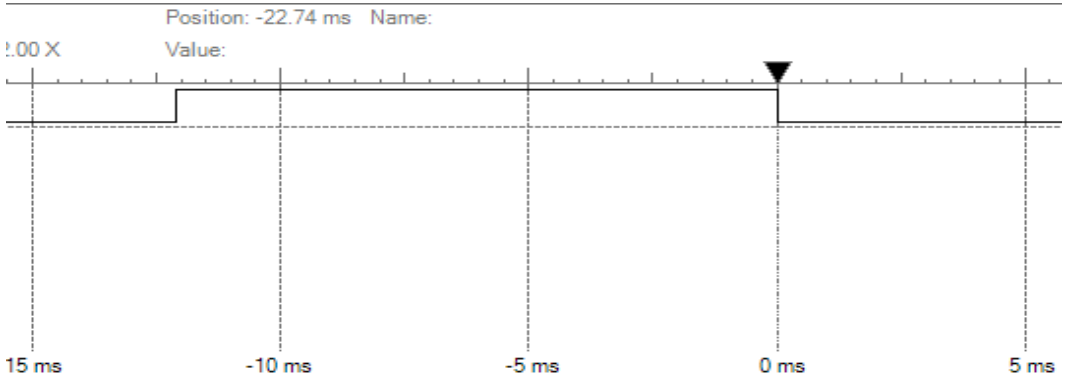

Figure 1. *Printing "Hello World!"*

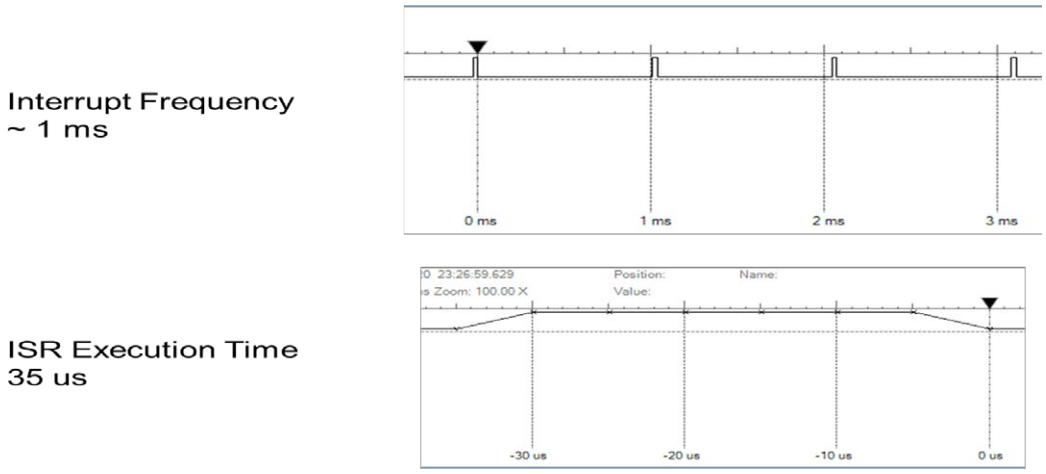

Figure 2. *Non-blocking printf performance*

to vary between 0.8 and 1.8 milliseconds. After the initial setup time, a transmit interrupt occurs approximately every one millisecond. The routine then requires only 35 microseconds to stuff the next character into the UART transmit register before getting back to doing useful work. Figure 2 shows the periodic interrupt and also the interrupt execution time. Keep in mind that the execution time does not include the interrupt overhead, which is less than 25 clock cycles in this case.

Performance Technique #2 – Increase the baud rate

It blows my mind that so many developers still will default their UART to 9600 even though serial hardware today can handle baud rates of 1 Mbps or better! Occasionally, I'll encounter someone bold enough to set the baud rate at 115200. But unless there is a potential electrical or hardware related issue with running up the clock, there is nothing wrong with setting the baud to 1 Mbps and getting debug messages out as fast as possible in order to minimize real-time performance issues. At 1 Mbps the original blocking *printf* for "Hello World!" would only block for 120 microseconds. That's far more acceptable than 12.5 milliseconds.

Performance Technique #3 – Use SWD

Modern day microcontrollers had the *printf* performance issues in mind when their creators developed the silicon. For example, developers that take advantage of the ARM Cortex-M's debugging capabilities can skip the UART altogether and use the internal debug module to transmit *printf* messages back through the debugger to the IDE. Skipping the UART in this manner not only saves setup, the internal hardware mechanism minimizes software overhead. An internal buffer gets filled with the message and the debug hardware automatically handles transmission to the debug probe, which results in minimal impact on the application's real-time performance.

Conclusion

Few developers are going to toss out their favourite, tried-and-true *printf* debugging techniques. In today's modern microcontroller hardware, though, there exist multiple options for improving the performance and efficiency of *printf* that minimize impact on real-time performance. For developers looking to try these improvements themselves, I've put together a Keil project for the STM32 that demonstrate how to use these techniques. The project can be found here. (Examine **uart_app.c** for the most interesting code).

Jacob Beningo is an embedded software consultant who currently works with clients in more than a dozen countries to transform their businesses by improving product quality, cost and time to market. He has published more than 200 articles on embedded software development techniques, is a speaker and technical trainer and holds three degrees which include a Masters of Engineering from the University of Michigan. Feel free to contact him at jacob@beningo.com, at his website www.beningo.com, and sign-up for his monthly Embedded Bytes Newsletter here.

MYO ARMBAND: WEARABLES DESIGN FOCUSES ON PACKAGING

By Patrick Mannion

Designing wearable devices can be as much about dealing with packaging issues as it is about using cutting-edge silicon, sensors, and developing "secret sauce" software. Such was the case with Thalmic Labs' Myo gesture-control armband, which went through at least seven iterations before settling on the current version. Let's go inside and find out why.

The Myo is one of many new devices for gesture control, a market that Global Industry Analysts expects to top $12.7 billion by 2020. It could be argued that movies like Minority Report may well have inspired many with its portrayal of mid-air computer control, but the attraction of gesture control is magnetic, as we look for different ways to communicate and interface with our machines.

The main market drivers are gaming, healthcare, automation, consumer devices, and automotive. However, the key enablers are the availability of low-cost, low-power silicon for processing and wireless communications; the falling cost of advanced sensors and sensor-fusion capabilities; the wide availability of open-source hardware and software; and rapidly advancing embedded vision.

Some inhibitors continue to be the lack of standards for gesture movements, a lack of awareness of the possibilities, and the high costs of development. However, those inhibitors crumble in the face of the fertile imaginations of the many innovators and up-and-coming designers who are dreaming up the next man-machine interface methodology.

Two of the most recent intriguing examples are VocalZoom, an Israeli startup that has invented a means of optically converting human voice to digital signals that paradoxically get more accurate in the presence of loud ambient noise. The second is Ultrahaptics, which has developed a means of adding mid-air touch-less haptics to gesture recognition and control systems.

In the specific realm of gesture control, there are many startups, such as ChiTronic with its smart

Figure 1. *Sci-fi movies are always fun sources of inspiration, and for gesture control, Minority Report's creative mid-air antics certainly stirred the imagination.*

Figure 2. *There are many companies using IMUs to detect motion and translate that to a control signal. ChiTronic has its smart ring and Apotact's Gest puts a sensor on all four fingers. (Images courtesy of Chitronic and Apotact.)*

Figure 3. *The current form of the Myo (right) is substantially lighter and also more sensitive to bioelectric signals, having moved to medical-grade stainless steel sensors instead of copper-based capacitive PCB sensors. (Image courtesy of Bashny.Net)*

ring and Apotact with its Gest, which puts an inertial measurement unit (IMU) sensor on every finger (Figure 2).

However, for their gesture-control design, the team at Thalmic Labs took another approach, choosing to dive into the world of electromyography (EMG) to directly sense the bio-electric pulses that run down the arm as the hand is moved, and then correlate those signals with specific hand movements, such as moving the hand side to side, up and down, or making a fist.

The idea was to combine the EMG signals with those from an IMU that tracks that motion of the arm, and then "fuse" the signals to control devices. After many (many) iterations, they ended up with the Myo (Figure 3).

The current form of the armband weighs 93g, is 11.5 mm thick, uses medical-grade stainless steel sensors and has an LED lighting up the logo, as well as one indicating status at the bottom. It connects to the controlled device via Bluetooth.

Beside the novelty of the device, the interesting aspect of Myo is the SDK that the developers provided, which allows anyone to develop a Myo application. Originally intended for controlling presentations on laptops or gaming, the community has taken the SDK and run with it, with over 100 Myo-controlled devices now on the market, including the Parrot 2.0 drone, GoPro Hero camera, and the Oculus Rift.

The Myo has clearly gained traction, so let's take a look inside... *Patrick Mannion continues his investigation in the full version of this teardown, click for pdf*

Download PDF of Article

Find Wearables Design on EETsearch

TAKE SCAN TEST OUT OF THE CRITICAL PATH

BY RON PRESS, MENTOR GRAPHICS

Integrated circuit complexity and integration continuously advances, posing challenges to the development process. Market profitability, however, demands that products be designed and produced as fast as possible. DFT (design for test) tools are used to make the designs more easily tested and to produce production test patterns. Traditionally, much of the DFT work that changes the design occurred late in the design cycle and pattern generation was in the critical path of design completion. In recent years, DFT and pattern generation is undergoing a shift to occur earlier in the design development flow.

Logic is tested by configuring sequential elements in the design into many shift registers called scan chains that a tester then loads and unloads. This lets ATPG (automatic test pattern generation) efficiently and automatically test any type of design. But, as designs kept growing in size, the test time and data increases as well. As a result, about fifteen years ago embedded compression logic was added to the scan chain interface, which provided 100× reduction in test time and data [F. Poehl et al, ITC, 2003. 10.1109/TEST.2003.1271110]. The general setup of scan compression is shown in Figure 1.

Traditionally, compression logic was often added after the gate-level design was complete, so designers would know exactly how many scan chains exist. So the first "shift left" of DFT is to create embedded compression logic at the RTL (register transfer level) design or earlier, as illustrated in Figure 2.

With recent features in embedded compression, you can estimate the maximum range of internal scan chains used for embedded compression and then slightly over-specify them. Then, the RTL for embedded compression can be completed early in the design flow, even before design RTL is ready. If some of the scan chains or scan channels aren't used, the compression and pattern-generation tool can still work effectively. This provides flexibility to

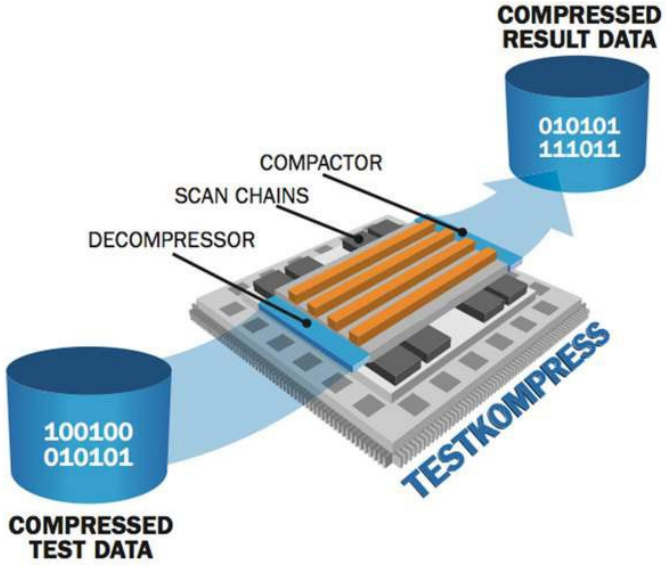

Figure 1. *Embedded test compression is now a standard methodology for testing ICs.*

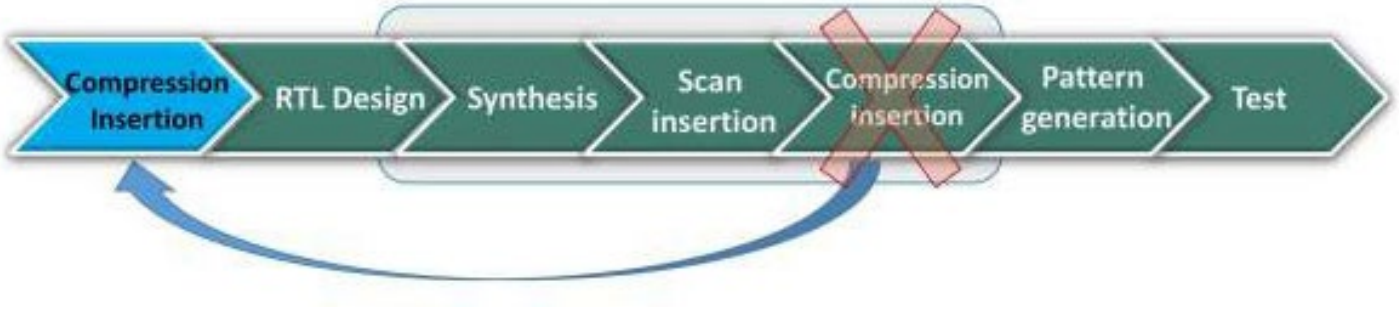

Figure 2. *Compression insertion can be moved before RTL, making it independent of synthesis.*

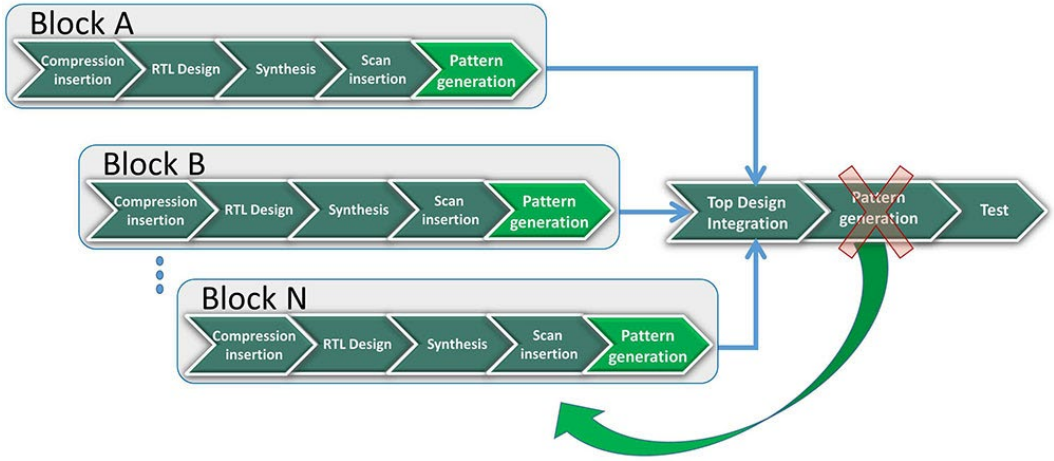

Figure 3. *Move pattern generation to the block level for large designs.*

handle additional scan chains that occur late in the design process due to ECOs or if chains are added for test points such as EDT Test Points. The test compression logic is also flexible enough to use different numbers of input channels, which could change due to packaging or due to tester limitations.

Another important industry problem is how to deal with growing design sizes. With RTL embedded compression insertion, the DFT work is done earlier but the pattern generation still needs the gate-level design. As designs keep growing, it isn't uncommon for designs of 40 million gates to require a week for pattern generation and we are now seeing designs in excess of 500 million gates.

What most of the industry has done for these designs is another shift left; instead of waiting for the complete top-level IC design to be implemented before pattern generation—which puts it in the critical path—pattern generation is moved to individual blocks as they are ready, as illustrated in Figure 3. This way pattern generation is on a much smaller portion of the design so it is faster and requires a smaller workstation. But more important for many is that it can be completed much earlier in the design flow. Completion of DFT and pattern generation at the block level is referred to as hierarchical DFT. It is a plug and play approach which shifts the pattern generation out of the critical path.

Moving the DFT insertion and pattern generation earlier in the flow takes DFT out of the critical development paths. Thus, working out DFT and test coverage issues can be done at a more practical time. In addition, when compression logic (and BIST) is inserted at the RTL then it is independent of the synthesis tool providing even more flexibility.

Also see:

- Design for Test Boot Camp, Part 1: Scan Test
- Design for Test Boot Camp, Part 2: Test Compression
- Design for test boot camp, Part 3: Advanced fault models and cell-aware test
- Design for test boot camp, part 4: Built-in self test
- Hierarchial test improves pattern application efficiency
- Test and Diagnoses Strategy Metrics: A New Perspective, Part I

Ron Press is the technical marketing manager of the Silicon Test Solutions products at Mentor Graphics. With 25 years' experience in test and DFT (design-for-test) he has presented seminars on DFT and test throughout the world. He has published dozens of papers in the field of test, is a member of the International Test Conference (ITC) Steering Committee, and is a Golden Core member of the IEEE Computer Society, and a Senior Member of IEEE. Press has patents on reduced-pin-count testing and glitch-free clock switching.

POWER SYSTEM MANAGEMENT ADDRESSING

By Michael Jones, Linear Technology Corporation

The foundation of all PMBus applications, including Linear Technology's power system management (PSM), is the ability for the PMBus master (system host) to communicate with all PMBus slaves (PSM controllers, PSM managers, PSM µModules, and PMBus monolithic devices) on the bus. Every slave on the bus must have a unique address that does not conflict with other devices.

The bus master must also be able to communicate with PSM slaves in a few less obvious situations:
- Address discovery
- Global actions
- Multiphase rails
- Invalid NVM
- Bus MUXes

Device addressing is achieved with a combination of base registers plus external address select (ASEL) pins, as well as special global, rail, ARA, and other special addresses.

This article presents the fundamental design principles underlying the Linear Technology PSM family, details on product family differences, as well practical examples and advice. Special cases, such as invalid nonvolatile memory (NVM), will also be discussed.

The benefit is a design that works on day one, and works even when things go wrong. For example, if you are writing to the NVM with LTpowerPlay software, and power is lost during the NVM write, your design will be recoverable. Furthermore, it will be recoverable in the field should you choose to implement "In Flight Update" found in Linear's Linduino reference code. [Note 1] Finally, you will be able to recognize symptoms of degenerate systems and fix them.

Once you have an understanding of how to implement PSM addressing, you will be able to design reliable systems quickly.

Basic PMBus operation

PMBus is a serial communication standard that is an extension of SMBus, which is similar to I²C. Two open-drain wires, SCL and SDA, support a bidirectional communication bus with masters and slaves. Masters are devices that control communication and are typically a microcontroller or FPGA. Slaves are devices that the master controls, and they are typically a small integrated circuit, in our case a power supply manager such as the LTC2977 or a power supply controller such as the LTC3880.

A system can have more than one master, but it is rare in practice. Usually there are multiple slaves. A master directs communication to a single slave at a time by using an address, even in a system with one slave. This means that every slave must have a unique address for proper system function.

Linear's power system management devices use an EEPROM along with resistors on pins to set the unique address of each device (slave). Therefore, part of addressing is ensuring that if any EEPROMs do not have valid data, a master can repair the system to the state where each device (slave) has its unique address.

The following sections explain these address configuration mechanisms in detail, including how to select, design, and repair addresses.

Basic PMBus addressing

PMBus addressing is defined by the SMBus standard that the PMBus standard refers to. SMBus addressing is the same as the I²C standard. For clarity, this discussion is limited to the SMBus standard.

Consider Read Byte protocol, defined in Figure 29 of the SMBus 3.0 standard (see Figure 1). The address of any transaction proceeds after the start bit (S), and concludes before the

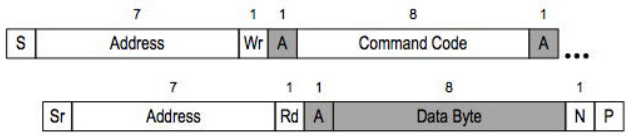

Figure 1. *Read Byte protocol (SMBus 3.0 standard, Figure 29)*

ACK bit (A). In between the (S) and (A) are 8 bits; the first 7 bits are the address, and one bit is used to indicate write (Wr) or read (Rd).

Seven bits means there are 128 potential addresses. In this article, addresses are written without the (Wr/Rd) bit as follows:

0x00 to 0x7F (7-bit addressing) [Note 2]

Sometimes programmers like to write the addresses with the extra (Wr/Rd) bit holding a zero, such as:

0x00 to 0xFE (8-bit addressing)

These are all the even numbers. Oscilloscopes and spy tools, like the Total Phase Beagle, use the same form as this application note, and it is natural to use 0x00 to 0x7F, so we shall do the same. However, be careful when engineers give you an address and don't indicate the format. The given address may be left-shifted by one, thus appearing twice the size it really is.

Note that Read Byte protocol uses the Address twice, but with a repeated start (Sr) preceding the second address. A repeated start is part of all read transactions. In this case, the same address must be used for (S) and (Sr).

Address map

Not all addresses are available for PSM slaves, because the SMBus standard reserves some of them. Appendix C of the SMBus 3.0 Specification has a table of pre-assigned addresses. It is not necessary to know the intended use of them, so a simpler table will do.

The SMBus Address Map shown in Table 1 uses a simple coding scheme. Conservative designs only use addresses that have a white background in the Description column. This avoids all the reserved and special addresses. PSM designs can use all the addresses with a white background in the Address column, with the possible exception of addresses 0x28 and 0x37 that were added for PMBus 3.1 zone operations.

The continuation of this article describes how Linear Technology implements some of the detail of PMBus addressing, and how a system can be made resilient against errors. Click for pdf.

Table 1. SMBus Address Map

Address	Description
0x00	General Call or START
0x01 to 0x0B	Battery etc
0x1C	Alert Response Address (ARA)
0x1D to 0x27	Open for use
0x28	Zone Read
0x29 to 0x2B	Open for use
0x2C to 0x2D	Reserved
0x2E to 0x36	Open for use
0x37	Zone Write
0x37 to 0x3F	Open for use
0x40 to 0x43	Reserved
0x44	Reserved
0x45 to 0x47	Open for use
0x48 to 0x4B	Prototyping
0x61	Default SMBus address
0x62 to 0x77	Open for use
0x78 to 0x7B	10 bit address
0x7C to 0x7F	Reserved

FIVE TECHNIQUES FOR IMPROVING ANALYZER MEASUREMENT SENSITIVITY AND ACCURACY IN NOISE

By Cherisa Kmetovicz, Keysight Technologies

Noise is a fact of life for today's engineer. Defined as undesirable electrical signals that distort or interfere with a desired signal, it can be caused by a variety of sources, either internal or external to the system. It results in random disturbances of useful information between the system's transmitter and receiver and is particularly detrimental in digital wireless or communication systems, where its presence can affect digital signals, cause errors in data sampling and limit the system's overall performance. It is so common in electronic circuits that it is considered a fundamental parameter; one that must be properly tested in all transmitter and receiver components.

Typically, the signal or spectrum analyzer is the engineer's tool of choice when it comes to observing signals. Such instruments are high-performance broadband receivers - and that means the presence of any internal noise can severely limit their ability to measure very low amplitude signals while maintaining high accuracy and fast sweep speed. Often these small signals, along with any spurs or anomalies, simply go unseen, hidden in the noise generated by the analyzer itself. To measure these signals quickly and accurately, especially when they are near noise, the analyzer's measurement sensitivity has to be improved. Fortunately, this can be accomplished using five key techniques.

To aid in this discussion and better visualize how these techniques impact the analyzer's measurement sensitivity, consider the measurements in Figure 1, which were made by a signal analyzer with a 100-kHz resolution bandwidth (RBW) filter and 10 dB front-end attenuation. The baseline amplitude of the signal being measured, depicted by the yellow trace, is very close to the analyzer's displayed noise floor. A marker placed at the peak of the signal reports an amplitude of -85.1 dBm. Comparing the peak to the displayed noise floor, there is a roughly 5-dB Signal-to-Noise Ratio (SNR); a low value that will surely impact the accuracy of the amplitude measurement. All other traces in the Figure show measurements made after various sensitivity improving techniques have been applied.

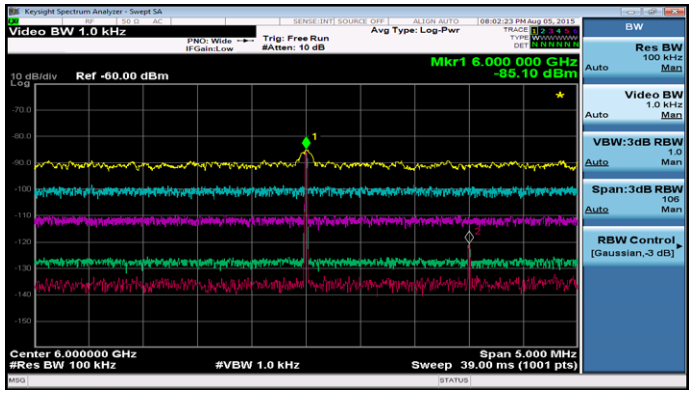

Figure 1. *The traces shown here depict a baseline measurement and four other measurements made once various techniques for improving analyzer sensitivity were implemented. All measurements were made using Keysight Technologies' N9010A EXA X-Series signal analyzer. Utilizing these techniques, the ENA provides the highest sensitivity and fastest sweep rates of any economy mid-range analyzer*

Read the five hints and suggestions for instrument setup that will minimise noise effects; click for full-article pdf.

Download PDF of Article

Find Minimizing Noise effect in Analysers on EETsearch

A TEST-BENCH EVALUATION; COEXISTENCE OF LTE AND S-BAND RADAR

By Dr. Steffen Heuel & Darren McCarthy, Rohde & Schwarz

A test method is available for measuring the effects of S-band radar systems on LTE wireless networks operating within the same frequency range, and how wireless signals affect the radars.

Bandwidth is precious and limited, so much so that the evolution of microwave/wireless applications will involve careful planning and accurate measurements to avoid and overlapping of frequencies. For example, potential issues may arise with frequencies occupied by existing S-band (2 – 4 GHz) radar systems and allocations for Long Term Evolution (LTE) cellular/wireless commercial-communications systems. Ideally, the different systems will remain in their proper frequency bands and will not interfere with each other. But performance degradation and system malfunctions can occur, such as excessive spurious levels, that can cause problems. Proper measurement methodologies can help avoid such problems and ensure the coexistence of S-band radar systems and LTE networks.

Multiple applications within the same frequency band can face coexistence problems. As an example, consider the performance of LTE base stations and mobile devices and how they might have coexistence issues with recently allocated frequency bands defined by the Third Generation Partnership Project (3GPP) standard, 3GPP TS 36.101 (Release 13, December 2015). Proper measurements can determine if interference exists, notably through in-field measurements at critical locations for both applications, such as airports.

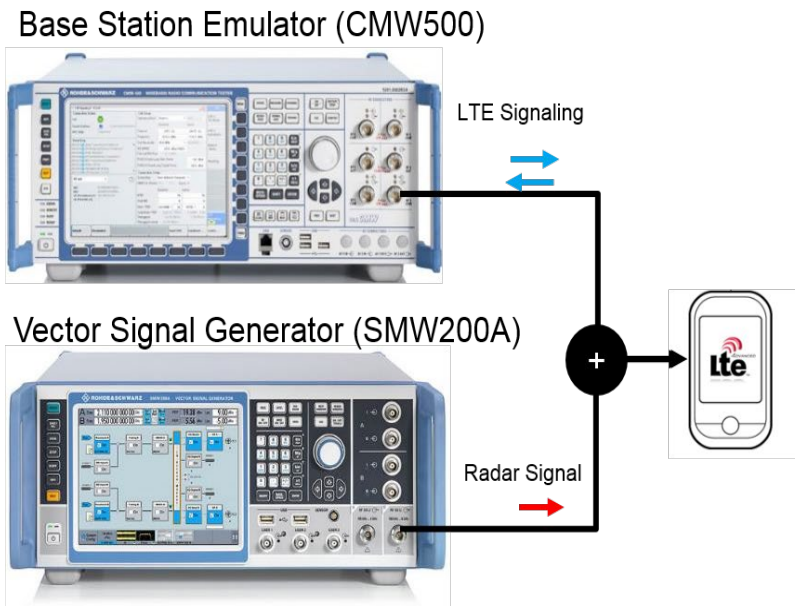

The S-band frequency range has been defined by the IEEE as all frequencies between 2 to 4 GHz. Along with aviation and weather forecasting systems, a number of different maritime radar systems worldwide also operate at S-band frequencies. S-band radar systems include air-traffic-control (ATC) radars (typically between 2700 and 3100 MHz) and AN/SPY-1 Naval Air Surveillance Radar (ASR) systems operating between 3100 and 3500 MHz. Coexistence is a concern for different S-band radar systems, especially in the US, because of LTE networks operating in Band 42, from 3400 to 3600 MHz, Band 43, from 3600 to 3800 MHz [with time-division-duplex (TDD) single-frequency operation for transmit and receive functions], Band 7, from 2620 to 2690 MHz in the downlink and 2500 to 2570 MHz in the uplink, and Band 22, from 3510 to 3590 MHz in the downlink and 3410 to 3490 MHz in the uplink. Frequency-

division-duplex (FDD) systems falling within this frequency range are also a concern for coexistence with radar systems at S-band frequencies.

While the number of proposed LTE bands has increased from 11 to more than 50 operational bands in the last four years, 150 MHz of spectrum in Bands 42 and 43 is anticipated for auction in the US. This spectrum, from 3500 to 3650 MHz, has been proposed for subdivision as 50 MHz for Tier 1 operators, 50 MHz for Public Safety use, and 50 MHz for Citizens Broadband Radio Service (CBRS). Any use of the spectrum by mobile radio networks must not disrupt primary users within the allocated spectrum, but no agreement of national and international authorities has been reached on guidelines to assess potential interference issues. At the recent World Radio Congress 2015 (WRC 2015) in Geneva in November, 2015, Recommendation 75 (REV.WRC-15), "Study of the boundary between the out-of-band and spurious domains of primary radars using magnetrons," proposes assessing the improvement of measurement methodologies and performance levels of radars. (Reference 5). Also, recommendation 207 (REV.WRC-15), "Future International Mobile Telecommunications (IMT) Systems", proposes a continuous study of the necessary technical, operational, and spectrum-related issues to meet objectives for developing future mobile telecommunication systems, taking into consideration requirements for other services. (Reference 2).

The FCC requires use of cognitive radio technologies and consultation of the national spectrum database, the Spectrum Access System (SAS) with this spectrum. Cognitive radio technologies are certainly not new, and detect and avoid (DAA), transmit power control (TPC), and dynamic frequency selection (DFS) techniques have been employed for many unlicensed radio technologies, including the IEEE 802.11 wireless-local-area-network (WLAN) standards, to avoid interference with weather radar and military applications in the 5-GHz band. However, licensed technologies, such as LTE, have not previously had to assess performance for coexistence with scanning radars.

Most of these types of radars apply pulse and pulse compression waveforms. After transmitting a pulse, the radar switches to receive mode to receive radar echo pulses from illuminated targets. The high sensitivity needed to acquire low-level returning pulse echoes also makes radar receivers susceptible to interference signals. LTE networks using nearby frequencies can cause this interference and may significantly degrade radar performance. Developing capability for an LTE network or mobile device receiver to detect and avoid a scanning radar may also represent a significant challenge.

Spectrum sharing

Disturbances to an LTE network can occur due to performance degradation of an S-band radar system resulting in an increase in the LTE network's block error rate (BLER). This loss of LTE network performance and poor spectral efficiency may not be a major drawback to a mobile communications network operator, but the reduction of power can result in an increase in operating costs to maintain the performance expected by mobile customers. The 3GPP specifications may define solutions for the problem, such as the use of dynamic frequency selection or transmit power control, that do not disturb other signals, but the challenge in developing a suitable receiver capable of detecting a scanning S-band radar should not be underestimated. *The illustration shows the outline of a suitable test setup, and the methodology is explained in the continuation of the article, click for pdf.*

Download PDF of Article

Find Coexistence of lte and s-band radar on EETsearch

designideas

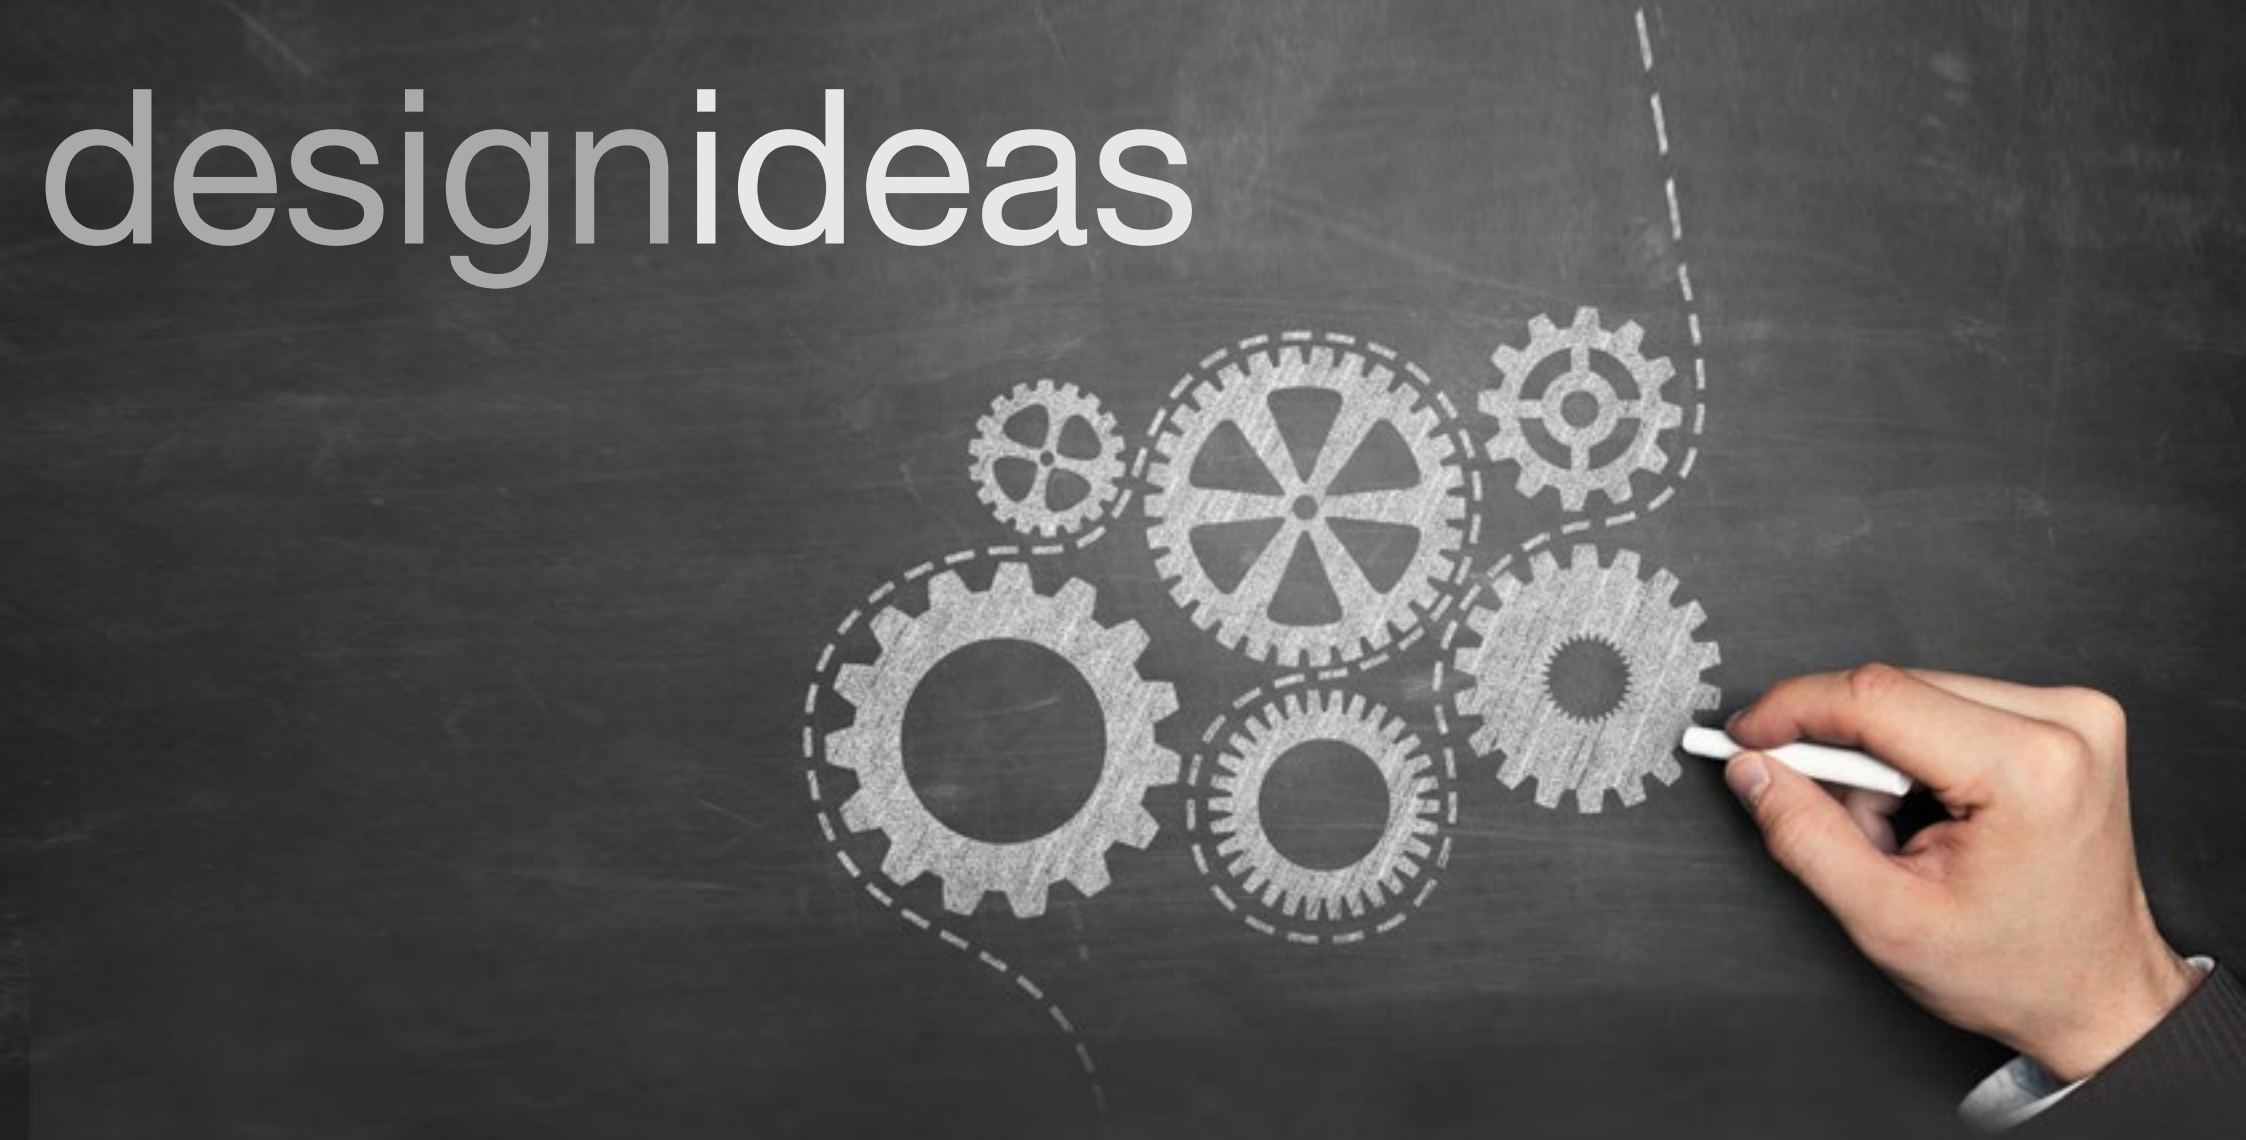

- Moving averager rejects noisy outlier values
- LED current regulator has low dropout

Moving averager rejects noisy outlier values

By David Vincenzoni

Measurements made in a noisy environment can exhibit sporadic disturbances. This Design Idea describes a digital circuit that removes outlier spikes without compromising bandwidth.

Smart moving average

The moving average is a process that continuously computes the average over N samples of data flowing through a FIFO (First In First Out) buffer. Every new sample added to the buffer will remove the oldest sample used to compute the previous mean value.

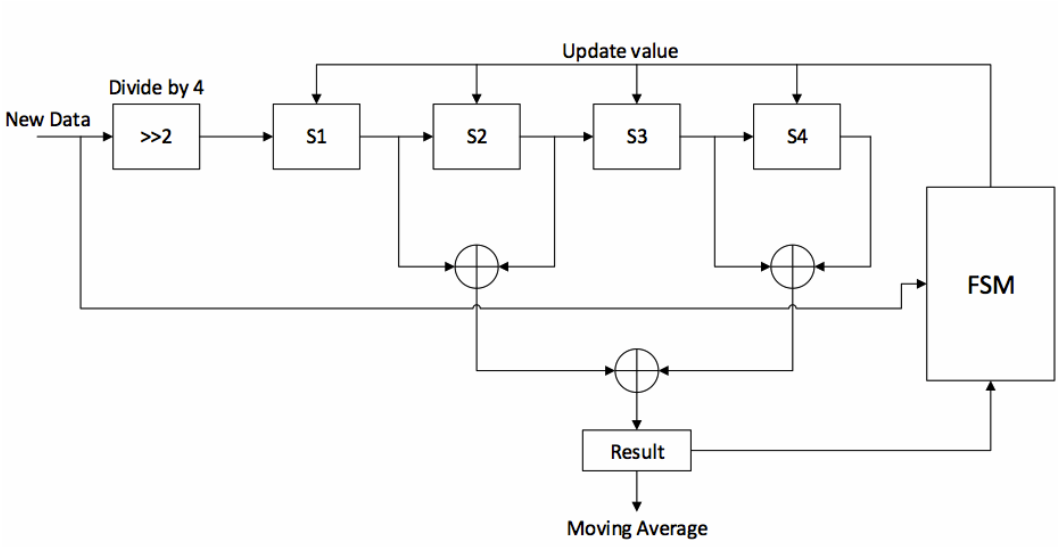

Figure 1. *The smart moving-average circuit determines when to reject new data.*

The smart moving-average is a variation on the theme; a digital circuit (Figure 1) where the idea is to collect the latest N measurements like a classical moving average, but new data will be added into the buffer (Sx) only if its value is within set limits of the actual average of the previous N samples.

The Finite State Machine (FSM) of Figure 1 manages this task. Every new data sample is compared against a maximum and minimum which depend on the current mean value. The new value is discarded when it is beyond the set limits.

Keep the number of stored samples (Sx) to a power of two to minimize of the size of the adders and avoid a general-purpose divider. If we use 2^p samples, we can shift-right the data input by p bits performing a zero-cost division, and the adders' size is reduced by the same p bits.

In this example, we have four (2^2) samples and we right-shift the input data by 2. The samples will be of size M-2, where M is the bus size of input data.

The computation circuit is formed by three full adders with carry features. The result is used by the FSM for checking the new input data.

At startup, since there is no mean value to be used as a comparison point, the first data are stored in the Sx buffers; then the mean value is computed. The FSM steady state is AVERAGE, shown in Figure 2. Here the FSM is waiting for new data input that will be checked against the mean value on the next state: CHECK DATA. This state's implementation depends on what we are going to measure. In the case of temperature

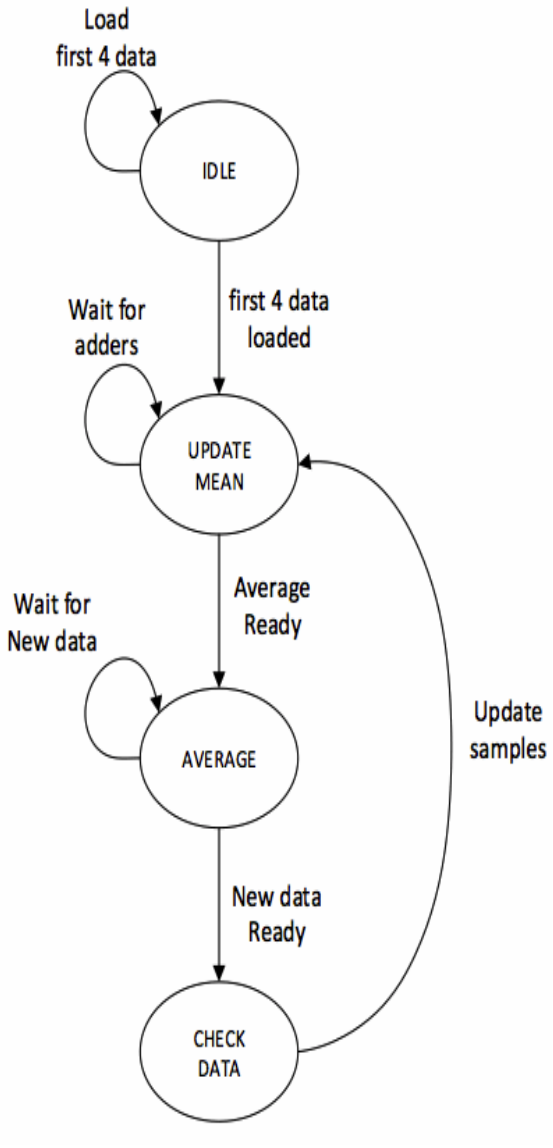

Figure 2. *FSM state diagram*

data, there is relatively slow variation; we do not expect new data to be very dissimilar from previous, nor from the mean. For this specific case, the data comparison can be done in the following way:

1. Only the most-significant bits of the mean and input (MSb) will be compared.
2. If the values are within a defined difference range, the new data will be stored and the mean value updated.
3. If the values have a greater difference, the new data will be discarded and the mean value kept.

Figure 3 shows an example comparison table. The range of possible values for the new data input is divided into four rows; only two MSb are used in the comparison (for some conditions the third bit is also checked). For example, when the new data MSb are “00”, it is accepted if the mean value of the previous four samples has MSb “00” or “01”. Otherwise, the data is discarded.

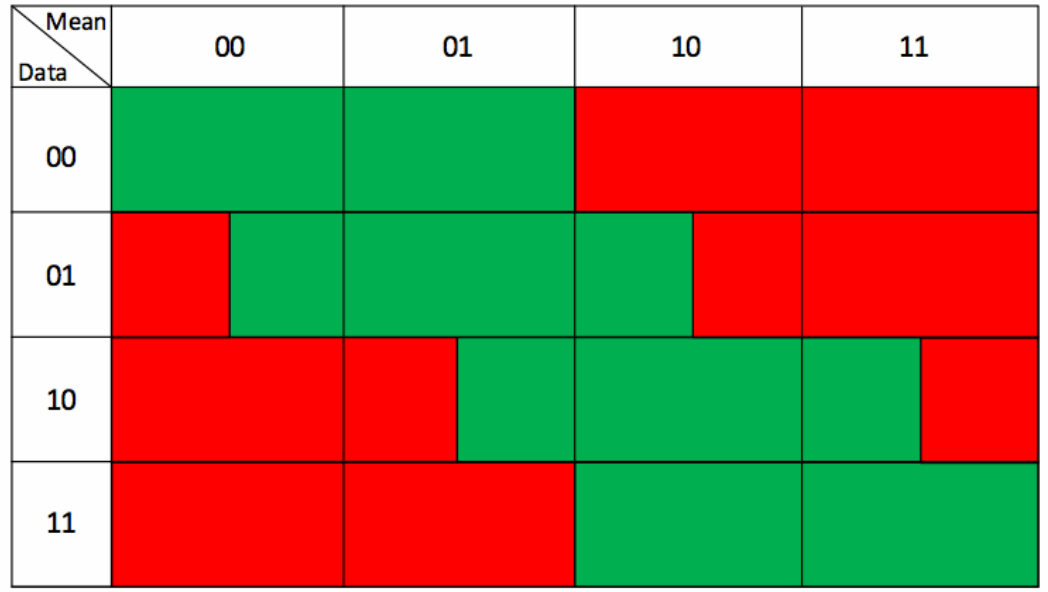

Figure 3. *Comparison scheme*

The number of elements used to compute the moving average can depend on the occurrence of incorrect data. In the case of sporadic events, this can be low (e.g., four); if the occurrence is high, then it will be necessary to increase the buffer size to eight or 16 elements.

David Vincenzoni *is R&D Design Manager at STMicroelectronics, responsible for the design and verification of new chips for Broadband Power Line Modems and for new families of devices for industrial applications.*

LED current regulator has low dropout

By Peter Demchenko

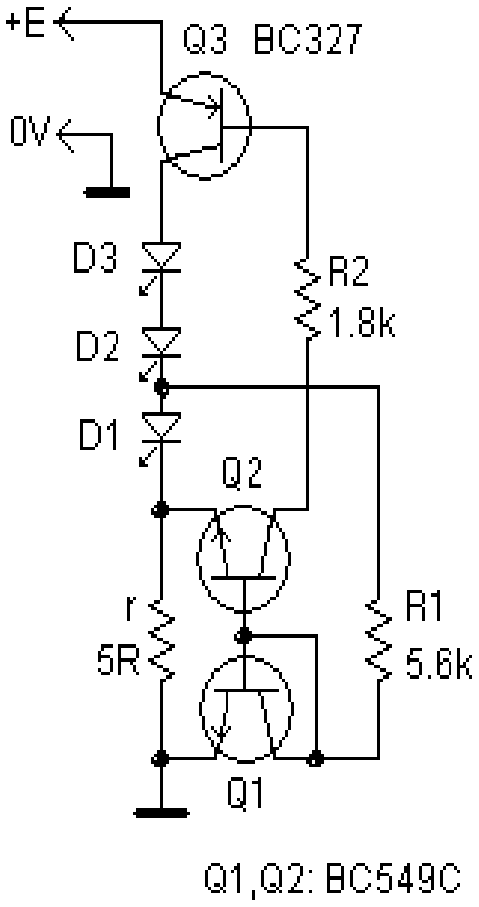

Figure 1. *LDO constant-current LED driver*

This Design Idea circuit regulates the current through one or more LEDs, making it almost independent of supply voltage. Its main advantage is a very small dropout voltage, which can be less than 100 mV. The design could find use on LED strips, where the voltage can vary along the length due to resistive drop, and small voltage changes result in large current changes.

The voltage drop on the current sensing resistor r is less than 40 mV. The rest of the drop depends on Q3's parameters.

The nominal LED current here is 7.2 mA at 9V. Increasing to 20V causes a current change of +15%, giving a dynamic resistance of about 10 kΩ.

The value of R1 is suitable for a blue/white LED with a voltage drop in the range of 2.9V - 3.4V. To maintain this current level at other forward voltage drops, change the value of R1 proportionally to the voltage drop change.

The current through the LEDs is inversely proportional to the value of r. Current can be roughly changed with this resistor, and finely tuned by varying R1.

To obtain good current stability over temperature, Q1 and Q2 should be in good thermal contact. Ideally, they should be on the same die, but the discretes here show good results when placed close to each other.

The circuit performs well with only one LED. The maximum number of LEDs in the string is only limited by the parameters of the circuit components.

productroundup

productroundup

Reinforced isolated amplifier for HV motor drives, inverters

The AMC1301 is presented by Texas Instruments as the first isolated amplifier with a working breakdown voltage of 1,000 Vrms for a minimum insulation barrier lifetime of 64 years, which exceeds VDE0884-10 requirements. The IC has no relevant lifetime degradation, unlike optical components. At 55% less high-side supply current and 45% less low-side supply current than competitive devices, the isolated amplifier improves power efficiency, simplifies power supply design and reduces thermal drift issues.

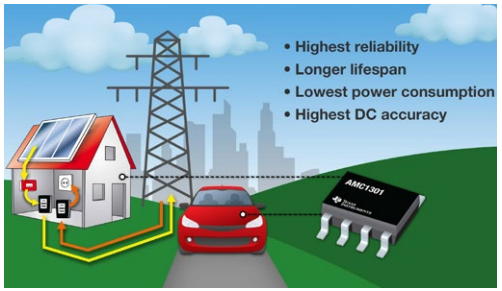

Complete article, here

Lowest power PIC32 family features core-independent peripherals

Microchip's latest PIC32MM family of MCUs for consumer, industrial control and motor control applications has a number of core-independent peripherals, Microchip's term for functional blocks that can operate and interoperate autonomously without waking the processor core, to save power. The PIC32MM family bridges the gap between the company's PIC24F XLP and PIC32MX families. The devices are supported by the Microchip MPLAB Code Configurator (MCC): they support applications with power and size constraints, with sleep modes down to 500 nA and 4 x 4 mm package options.

Complete article, here

Quad-channel, 28 Gbaud linear driver for PAM-4 comms links

MACOM Technology Solutions, maker of RF, microwave, millimeter wave and photonic products, recently introduced the MAOM-003419, a quad channel linear externally modulated laser (EML) driver for 28 Gbaud PAM-4 solutions for 200G and 400G. The MAOM-003419 is a low power electroabsorptive modulated laser (EML) driver with high gain, high bandwidth, and output voltage capability up to 2 Vpp while it consumes less than 500 mW of power. The device has differential inputs to provide common-mode rejection and single-ended output to drive industry standard EMLs. The MAOM-003419 is available in a small form factor surface mount package while also integrating the high frequency bias T for the EA modulator.

Complete article, here

Seeed's BeagleBone Green Wireless;Wi-Fi & Bluetooth 4.1

Distributor Mouser has the BeagleBone Green Wireless, a Linux-based development board from BeagleBoard.org and Seeed Studio. It is the first board in the BeagleBone community to feature a wireless connection, with both 2.4GHz Wi-Fi 802.11b/g/n and Bluetooth 4.1 low energy connectivity in place of the BeagleBone Green's Ethernet connection. Like the BeagleBone Green it is based on the AM335x 1GHz ARM Cortex-A8 processor from Texas Instruments. The board has 512 MBytes of DDR3 RAM, 4 GBytes of eMMC flash storage, microSD slot, and two PRU 32-bit microcontrollers, plus a 3D graphics accelerator and NEON floating-point for enhanced multimedia.

Complete article, here

productroundup

Mixed-signal programmable IC integrates logic, timing functions

Silego Technology has added a device to its GreenPAK programmable mixed-signal line, to aid designers in "mopping up" functions that might otherwise require a variety of IC packages and passive components. In its 1.0 x 1.2 x 0.55 mm, 6-GPIO STQFN package, and operating from 1.8 V to 5 V VDD, the SLG46108V contains: 4 look up tables (LUTs); 7 combination function macrocells; 2 selectable D-Flip Flop (DFF)/Latch or 2-bit LUTs; 2 selectable DFF/Latch or 3-bit LUTs; 1 selectable pipe delay or 3-bit lut; 1 selectable counter/delay or 4-bit lut; 1 programmable Delay/Deglitch Filter; 3, 8-bit counter/delays with external clock/reset; RC oscillator (25 kHz/2 MHz); and power-on-reset.

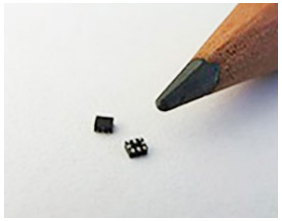

Complete article, here

Qt release 5.7 adds Qt 3D Module, applies C++11 features

The Qt Company (Helsinki, Finland) has issued Qt version 5.7, an update to its cross-platform application and user interface (UI) development framework. Qt 5.7 revises UI creation, including a new Qt 3D module, as well as new ways to create modern touch-based UIs with Qt Quick Controls 2 and Qt Quick Designer. Qt 5.7 follows the evolution of modern C++, making use of C++11 in Qt APIs. With Qt 5.7 and the new Qt 3D module it is now easy to create 3D UIs and interact with 3D objects using high-level Qt C++ and QML APIs. Qt 5.7 also introduces Qt Quick Controls 2, a new library of UI controls (buttons, sliders, dials, etc).

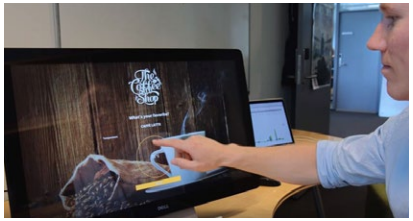

Complete article, here

Micro Bluetooth SMART module with built-in antenna

ALPS' low-power UGMZ2AA Bluetooth SMART communication module with built-in antenna is positioned as among the smallest and with one of the lowest power consumption levels in the field. The module has dimensions of 4.7 × 4.7 × 2.0 mm, with a custom-designed antenna circuit pattern on the top of the unit, in a single package. This removes the need for end product manufacturers to separately prepare an antenna. Power consumption is 5.1 mA (0.6 µA in sleep mode), making the module suitable for devices running off a button cell or other battery.

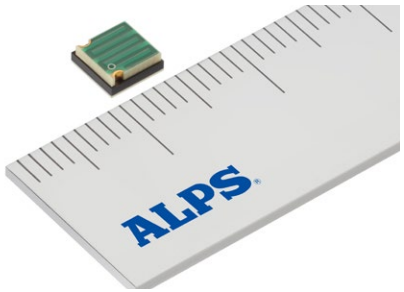

Complete article, here

Dual-ARM-core variants lower entry-point for Xilinx Zynq UltraScale+ MPSoCs

Xilinx has added streamlined dual-core members to its Zynq UltraScale+ MPSoC family of devices. The dual-core "CG" family members expand the Zynq MPSoC portfolio scalability, to include dual application and real-time processor combinations. These dual-core devices add processing scalability at a lower cost entry point to the current Zynq UltraScale+ family, which offers quad ARM Cortex-A53s, dual Cortex-R5s, a graphics processing unit, and a video codec unit.

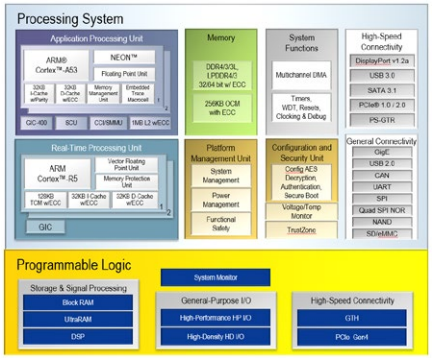

Complete article, here

productroundup

Smallest rotary switches for PCB configuration

C&K Components' smallest diameter rotary switch series provides design engineers with a switch option that reduces board space by 50%. Available with two to eight positions, the RM Series single-pole rotary switches have 9.2 mm body diameter. The switches are equipped with either an extended shaft with one of two standard knob options, or a flush actuator with a screwdriver slot for adjustment. Gold-plated contacts provide operating lifespans of 2,000 cycles, with operating force of 270±100 gf cm. Contacts are rated for 0.5A at 24 VDC and 0.2A at 48 VDC.

Complete article, here

Use coax or STP for in-car data, with high-speed deserialisers

Maxim Integrated has posted details of its MAX9276A/MAX9280A, which are 3.12 Gbps GMSL deserializers for coax or STP (shielded twisted pair) input, and parallel output. They have been designed to allow automotive data, especially for infotainment, to be routed on lighter-weight, lower cost cabling. Use these parts, Maxim suggests, in high-resolution automotive navigation systems; megapixel camera systems; and rear-seat infotainment. The gigabit multimedia serial link (GMSL) deserializers receive data from a GMSL serializer over 50Ω coax or 100Ω shielded twisted-pair (STP) cable and output deserialized data on LVCMOS outputs.

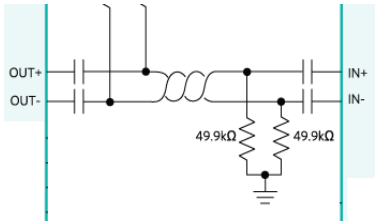

Complete article, here

Stackable PCIe/104 boards gain Xilinx Kintex FPGA-based module

Sundance has added a Xilinx UltraScale FPGA to its EMC2 family of PCIe/104 "OneBank" compatible I/O boards; the EMC2-KU35 is a stackable FPGA module with Gen2 PCI Express interfaces that are "OneBank" compatible and has a VITA57.1 FMC I/O slot controlled by a Kintex UltraScale KU35 FPGA, and is fully supported by the latest Xilinx Vivado tools. The board hosts either a Xilinx Zynq SoC or Xilinx Artix/Kintex FPGAs. The EMC2-KU35 has two banks of 16-bit DDR4 with close to 2 Gbytes/sec bandwidth each. It also has four-lanes of PCI-Express and a VITA57.1 FMC-LPC expansion module.

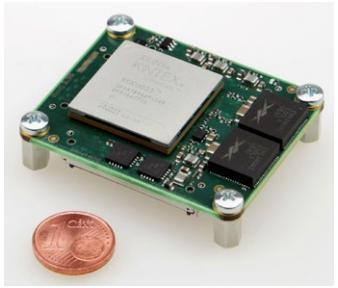

Complete article, here

16x16 LED display reference design is Arduino and ARM mbed compatible

MAXREFDES99# integrates four MAX7219 LED drivers to provide designers a 16x16 display with 256 LEDs, allowing them to create signs in multiple languages. The MAX7219 driver simplifies the process of controlling multiple LEDs in matrix or numeric display form. For fast prototyping, software code developed for the MAXREFDES99# reference design works with both ARM mbed and Arduino platforms. Developers are encouraged to build upon and enhance the code base for more signage possibilities.

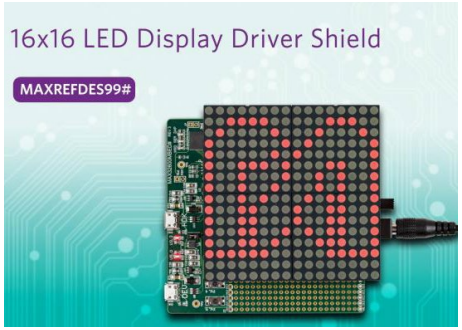

Complete article, here

productroundup

Software-defined programming of Xilinx' Zynq SoC from C/C++

Xilinx has extended the scope of its SDSoC development environment, enabling software defined programming of the 16nm Zynq Ultrascale+ MPSoC. The release also accelerates C/C++ based programming with system level profiling tools and a 50% reduction in end-to-end compile time. The 2016.1 release of the SDSoC development environment enables software defined programming for the Zynq family of SoCs and multi-processing (MP) SoCs using C and C++ languages. The new release includes support for the recently introduced 16nm Zynq UltraScale+ MPSoC.

Complete article, here

ARM enables ASICs-for-all: eases route to SoCs for embedded & IoT markets

ARM has expanded its ARM DesignStart initiative to offer simpler and faster access to EDA tooling and design environments from Cadence and Mentor Graphics. The new partnership builds on the benefits of free access to the ARM Cortex-M0 processor IP offered through the DesignStart portal. The new ARM Approved Design Partner programme also provides DesignStart users with a global list of audited design houses for expert support during development. Via the DesignStart portal, SoC designers can gain free access to ARM Cortex-M0 processor IP for design, simulation and prototyping with the option to buy a simplified and standardized $40,000 fast track licence. ARM cites costs to fabricate test chips starting at $16,000, using figures from the Europractice multi project wafer pathway, for a batch of 45 samples using 25 mm² silicon per chip in a 'mature' process at 180nm.

Complete article, here

Memory compiler targets 40nm ultra-low-power process for IoT

SureCore (Sheffield, UK), provider of IP for low power SRAM, has released its TSMC 40nmULP process technology memory compiler, that provides a vehicle for its recently announced 40nm Ultra Low Voltage SRAM IP that effectively operates at a record-setting 0.6V across process voltage and temperature. The 40nm ULP compiler supports synchronous single port SRAM with operating voltages ranging from 0.6 to 1.21V and capacities from 8 kBytes to 576 kBytes with maximum word lengths of 72 bits, with up to 80% savings in dynamic power consumption and up to 75% reduction in static power.

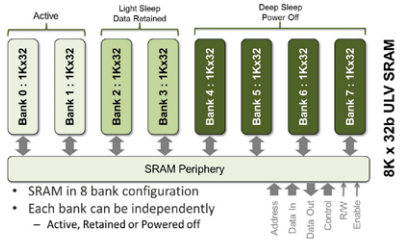

Complete article, here

Small-motor drive IC provides speed control versatility

Offering PWM or DC-voltage speed control, and featuring a bridge-tied-load (BTL) drive architecture to minimize audible switching noise and electromagnetic interference (EMI), the AM9468 and AM9469 brushless DC motor drivers from Diodes Inc. drive fans in notebook and desktop computers, instrumentation and similar equipment, as well as supporting a wide range of other medium-voltage, low-power, motor-driving applications. Flexible motor speed control is achieved using either an external PWM signal or DC voltage, or a voltage derived from a thermistor network.

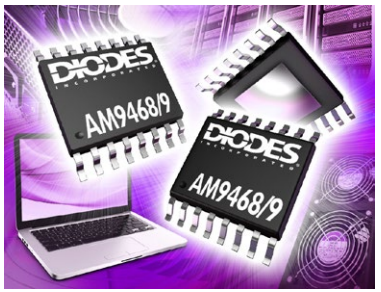

Complete article, here

EDN europe

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

CONTACTS

PUBLISHER

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

EDITOR-IN-CHIEF

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

Patrick Mannion
Brand Director EDN Worldwide

european business press

CIRCULATION & FINANCE

Luc Desimpel
luc.desimpel@eetimes.be

ADVERTISING PRODUCTION & REPRINTS

Lydia Gijsegom
lydia.gijsegom@eetimes.be

ART MANAGER

Jean-Paul Speliers

ACCOUNTING

Ricardo Pinto Ferreira

SALES CONTACTS

Europe

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

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

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

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

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

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

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

USA & Canada

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

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

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

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

Asia

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

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

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