EDN europe

NOVEMBER 2016

www.edn-europe.com

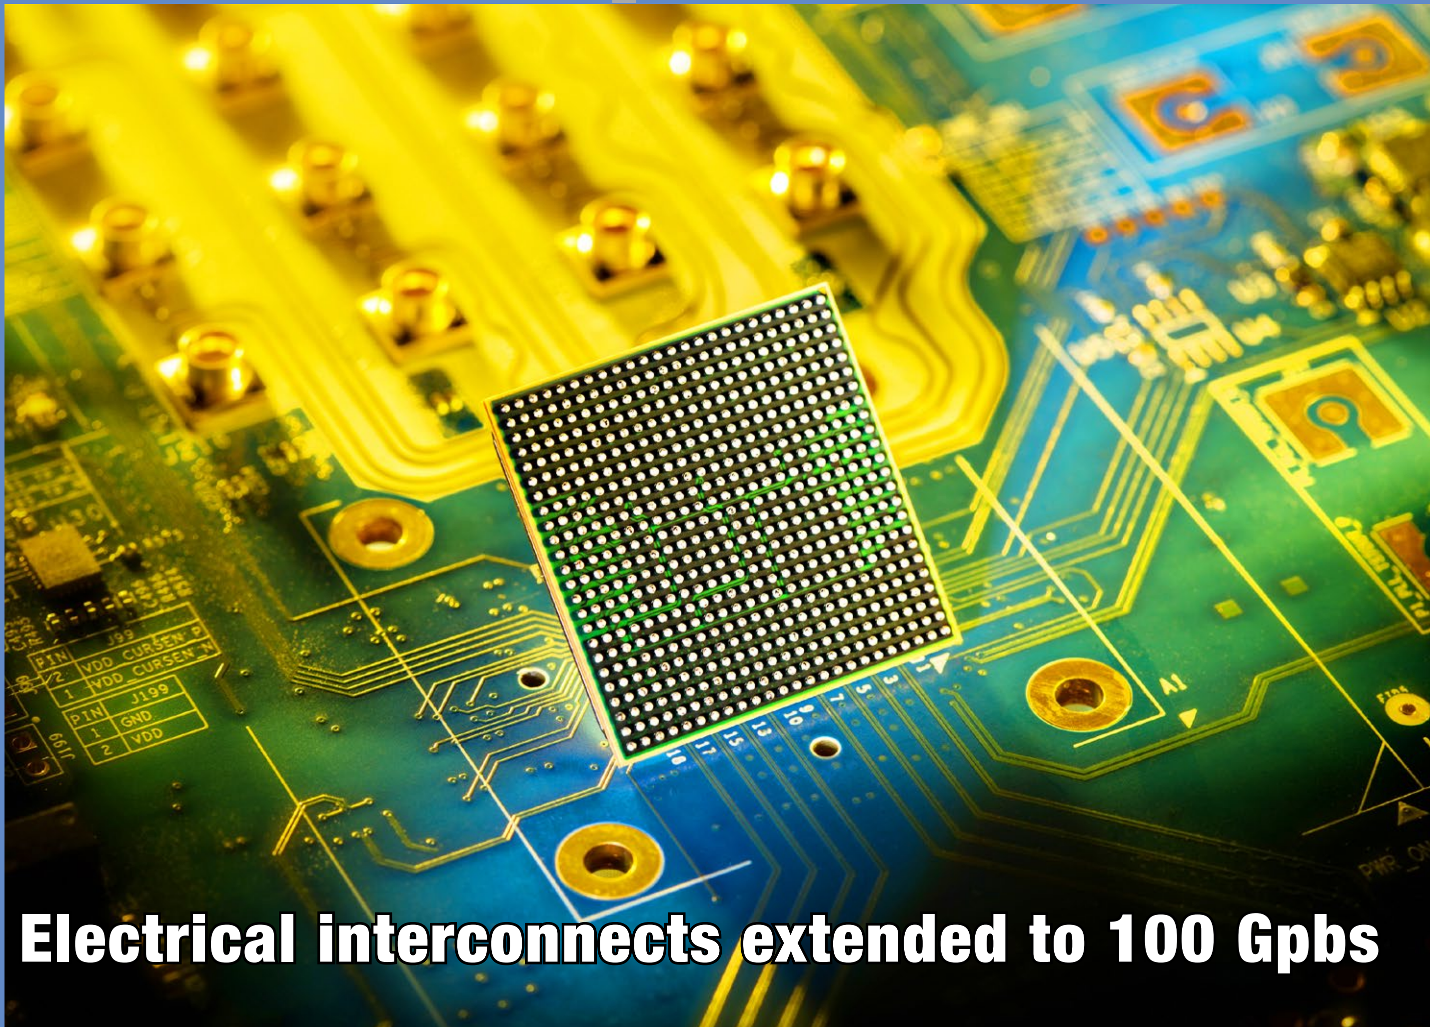

COVER

Aquantia & GlobalFoundries team to take copper interconnect to 100G

For many years, the bandwidth limits of copper interconnect have been repeatedly forecast, and those forecasts have repeatedly been confounded. Now, high-speed-signalling fabless semiconductor company Aquantia (San Jose, California), in collaboration with chipmaker Global Foundries, has announced it has pushed that limit to 100 Gbit/sec, and anticipates applications in "hyperscale" data centres, for storage and for provision of cloud computing resources.

The increased throughput will, of itself, help to reduce costs; and enabling it without conversion to all-optical interconnect will extend that benefit. The announcement is, specifically, that technology Aquantia calls QuantumStream can delivers the first single-lane, 100 Gbit/sec solution by building on GlobalFoundries' 56 Gbit/s SerDes. QuantumStream, Aquantia says, applies "architectural innovations" in interconnect technology and delivers low latency performance for next-generation hyperscale architectures. QuantumStream technology is aimed at inter- and intra-rack connectivity up to a few metres; beynd that, longer-reach optical connectivity solutions will continue to be used in hyperscale data centres. System vendors and data centre operators will be assisted in their, "push towards higher performance and newer topologies in hyperscale architectures". The target of the collaboration is to enable 400 Gbit/sec Ethernet Direct-Attach Cable connections; read the full item here.

FEATUREARTICLES

ONLINE THIS MONTH

EDN's columns

BIG DATA, MY DATA, YOUR DATA...OUR DATA?

The notion of Big Data has become one of the defining characteristics of the interconnected world that many of us are creating. All those billions of connected, sentient (in the literal sense of, being able to sense something) devices, constantly monitoring their environment and making that data available for our use. "Big data is about connected things leaving a trace as they interact."* The language, the terms in which the topic is discussed, constantly evolves. We realise that in passing data to the cloud, that the cloud doesn't want, and can't cope with, the unfiltered torrent. So we look to pre-process it closer to source: and thus we have fog computing, and other such ideas. But underlying everything is the belief that the data has inherent value, and that there is revenue to be extracted. The IoT is either not [only] about big data: or it is all about big data.

However, anyone looking to create and then generate revenue from a data stream based on information harvested from some form of connected sensor node or nodes, would have heard a cautionary theme expressed at the recent 2016 World Congress on Intelligent Transport Systems, Melbourne, Australia. Over three days, the Congress' Plenary sessions considered; autonomous vehicles; smart cities; and big data. And, a repeated point made by speakers in each of these sessions was; data sharing is vital to the success of each of those domains. The quote* above is from speaker Adam Game, Director of Strategy, Intelematics; he noted that there are already consumer expectations around connectivity, but they are bound up with issues such as privacy. And, he used a term that was repeated by several other speakers; silo. "The key challenge is accessing data [that is] often in silos." As speaker Leon van de Pas, Head of IoT Business Group, HERE, expressed it, "Data silos are holding us back, [when] everything is connected... data should benefit the community." He went to remark that connectivity and big data are not enough in themselves, that there must be an environment of open systems in which the data can be translated into value; there being a cost to "ingest" data, to normalise it and to enrich it, and to make it available.

The notion that data must be shared is intuitively obvious in the case of autonomous/assisted-driving vehicles and their interaction with infrastructure and with pedestrians and other road users. But the ITS concept takes that much further and sees those data sets merged with others about the functioning of the city as a whole. This is difficult enough at the practical level; accessing data sets from multiple hosts, in multiple formats, extracting the salient and useful information, without compromising user privacy – especially as, to be useful, much of this will have to be done in near-real-time. Securing the agreements to share may represent a different order of difficulty. Describing some aspects of the EU's ITS Directive, the European Commission's Claire Depré (DG, MOVE) noted that the task is, "All about a public/private partnership... not all about data, but governance." Some of the data is in private hands (captured by in-vehicle systems, for example: some, public, such as infrastructure and road signage).

From the US arm of HERE (location data specialists), ITS Director Monali Shah reiterated that collaboration is essential, citing an agreement between Audi, Mercedes and BMW to share captured data for more effective ITS outcomes. (Those three companies being joint owners of HERE.) She listed four problems that the extended data from the vehicle can contribute to; safety; congestion (and by extension parking and emissions/air quality); road signage (real-time update of central databases by in-car systems); and traffic data now supplied by GPS probes. She concludes, "[What is needed] ...cannot be a supplier/vendor relationship – it needs to be a partnership."

This highlights what seems likely to be an ongoing area of tension; between a commodity in which there is an expectation of a developing market – not least, because the value traded in such a market will underpin the investment to build the networks: and which at the same time, works best if it is freely shared. The challenges of creating the sensor-to-cloud path might be minor in comparison.

Excellent Efficiency, 36V, 3A Synchronous Step-Down Converter

MPQ4423A

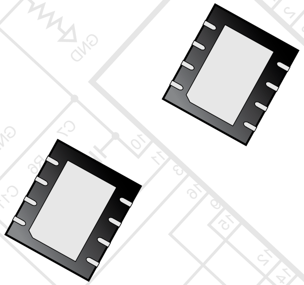

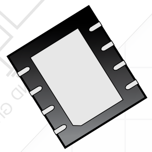

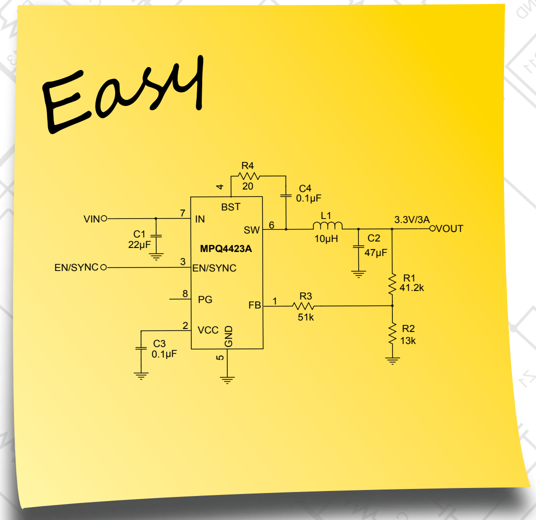

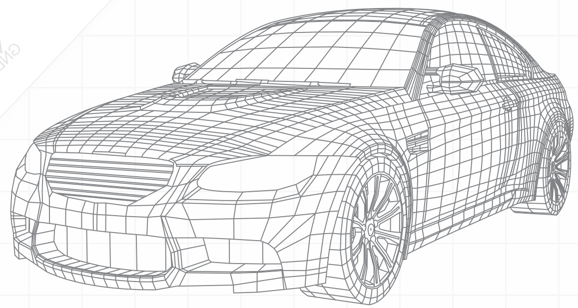

Key Highlights

- AEC-Q100 Qualified
- Wide 4V to 36V Continuous Operating Input Range
- 95% Efficiency Synchronous Mode Operation
- Good EMI, Compliant with CISPR25 Class 5
- Low Dropout Voltage for Automotive Cold Crank

MonolithicPower.com

Follow us:

MPS®

Simple, Easy Solutions™

pulse

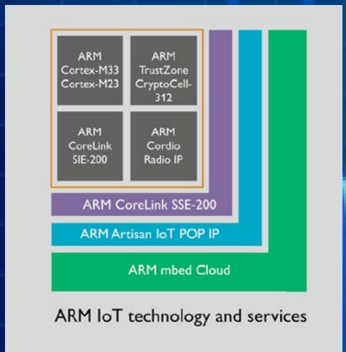
ARM Cortex-M33 Cortex-M23
ARM TrustZone CryptoCell-312
ARM CoreLink SIE-200
ARM Cordio Radio IP
ARM CoreLink SSE-200
ARM Artisan IoT POP IP
ARM mbed Cloud
ARM IoT technology and services


ework
s and their interactions can be described quickly and simply


SoC

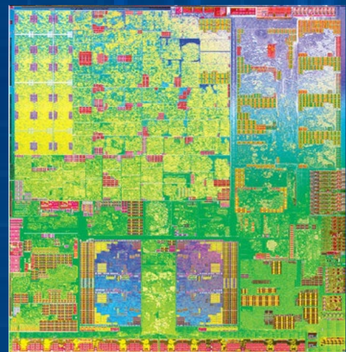


DIGITAL INPUT
DIGITAL OUTPUT
MOTION
ENCODER
HS TERM
IO-LINK
PROTOCOL
USER
Pocket IO

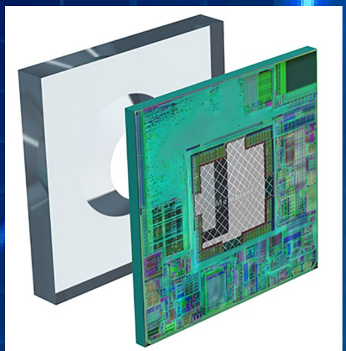

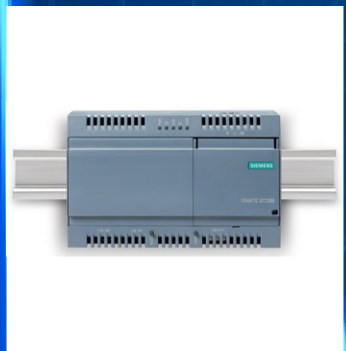

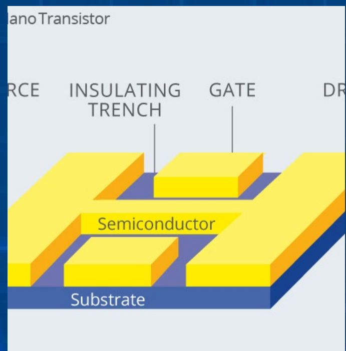
ano Transistor
RCE
INSULATING TRENCH
GATE
DR
Semiconductor
Substrate


V2X Solutions for V2V and V2I Communication
RENESAS
R-Car W1R

ARM bids to capture IoT business from silicon IP to data services

In one move, ARM has both created a new focus on secure IoT applications and systems; and has broadened the scope of its operations by offering a package that starts with its 'home ground' processor IP and extends through to provision of cloud data services.

ARM summarises its campaign as "accelerating secure IoT from chip to cloud". Enabling the IoT revolution, ARM says, needs efficiency – devices will need to be battery powered or run from energy harvesting – security, to protect data from sensor to cloud, and scale. The package of product and services includes;

- two specific Cortex-M processors based on the ARMv8-M architecture, to be known as M23 and M33, which come with ARM's TrustZone technology;

- IP to build an IoT subsystem on custom SoC designs: the CoreLink SSE200, with ARM CoreLink system IP for, the company asserts, the fastest, lowest-risk path to silicon;

- secure SoC designs that embody ARM's TrustZone CryptoCell technology;

- a wireless solution with ARM Cordio radio IP for IEEE802.15.4 (ZigBee and Thread) and Bluetooth 5. time of introduction, ARM has collaborated with foundry TSMC to offer an optimized implementation on ARM Artisan IoT POP IP, for TSMC's 40ULP process. ARM says that 40 nm, while some way behind today's leading edge, is likely to be the "sweet spot" process in which IoT chips are built.

ARM says that this release is its most comprehensive product suite ever, intended to deliver new levels of security, efficiency, low-power connectivity and device life cycle management, adding, "..the IoT already runs on ARM but the goal now is scale, which we are enabling today through a uniquely comprehensive set of technologies and services built to work together seamlessly."

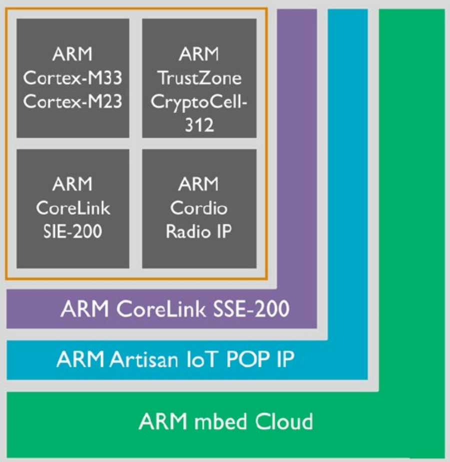

ARM IoT technology and services

Bringing TrustZone to Cortex-M processors, and to resource-limited IoT nodes, the Cortex-M23 and Cortex-M33 are the first embedded processors based on the ARMv8-M architecture. The majority of the top ten global MCU suppliers have already licensed one or both processors. Lead partners include Analog Devices, Microchip, Nuvoton, NXP, Renesas, Silicon Labs and STMicroelectronics.

Cortex-M33 has configuration options including a coprocessor interface, DSP and floating point computation, with increased performance and efficiency relative to Cortex-M3 and Cortex-M4. In 40 nm an M33 core is "very small" ARM says, at under 0.1 square mm. The M33 is a general-purpose processing unit, with its range of optional features.

The Cortex-M23 is even smaller, by a further 75%, and is 50% more power efficient, but includes security for the most constrained devices. Both cores are backwards compatible with ARMv6-M and ARMv7-M architectures for direct and fast porting.

Complete article, here

Fast UI prototyping for Raspberry Pi

ByteSnap Design's user interface development framework SnapUI has been extended to support Raspberry Pi and i.MX6-based embedded devices. SnapUI already allows Windows embedded and Linux developers to build and test user interfaces before hardware development of the target device is complete.

ByteSnap Design's latest iteration of SnapUI brings high performance to UI prototyping, with the emphasis on CPU embedded chipsets. The small, focused SnapUI codebase and easy integration of new features makes adoption of the framework painless for software engineers, enabling them to extend built-in custom graphical objects easily.
When designing new hardware devices, keeping the user interface design in step with hardware development can, ByteSnap asserts, be difficult; worse, when the user interface is overlooked or is the last consideration in the design process, projects can suffer significant delays in getting product to market. SnapUI mitigates this risk by enabling development of the user interface independently of the hardware, on a reference hardware platform or readily available Raspberry Pi, ready for deployment to the target device when the hardware is ready. The compiled SnapUI library contains a powerful OpenGL engine, which renders content optimised for Raspberry Pi hardware and achieves high frame rates with low CPU usage.

Complete article, here

Avnet completes acquisition of Premier Farnell

Avnet, Inc. (based in Phoenix, Arizona) has announced that has completed its acquisition of Premier Farnell plc (based Leeds, UK) in an all cash merger for £1.85 per share, which equates to an equity value of approximately £691 million.
This completes a process commenced in June this year; on 15th June 2016 we reported Distributor Premier Farnell to be sold to Swiss company Datwyler; at that initial stage Premier Farnell (also operating as element14), had agreed a sale of the company to Datwyler Holding AG (Altdorf, Switzerland) for a cash offer of 165 pence (UK£1.65) per share.
That offer represented a premium to the prior share price of Premier Farnell of 51%. At the start of August, Avnet launched its counterbid, as we reported on 2nd August 2016; Avnet counterbids for Premier Farnell and it is this bid that has ultimately prevailed.
"The combination of Premier Farnell with Avnet's components business will create a truly unique distribution model that supports customers at every stage of the product lifecycle. By pairing our deep expertise in large volume broadline distribution with Premier Farnell's specialization in proof of concept and design, we can offer true end-to-end solutions that accelerate a customer's time-to-market and moves their products seamlessly from prototype through to volume production," said William Amelio, Chief Executive Officer of Avnet.

Complete article, here

FPGA IP to embed in SoCs for function acceleration

Over time, a number of efforts have been made to integrate a block of programmable logic into an otherwise fully-diffused system-on-chip integrated circuit, with limited success. Now, Achronix Semiconductor says it has resolved the issues preventing the concept being applied.

Where the intention has been (for example) to provide for late design changes, a barrier that has frequently been cited as holding back the concept is the disparity in density between FPGA structures and fully-diffused cells and blocks. Put simplistically, if you lay down enough FPGA to be useful, you use a disproportionate amount of silicon area. (The converse, embedding specific blocks of diffused IP, such as processors cores – e.g. Xilinx Zynq – has been much more successful.)

Now, Achronix (Santa Clara, California) has announced its Speedcore embedded FPGA (eFPGA) IP for integration into designer's SoCs. Speedcore is designed for compute and network acceleration applications and is based on the same architecture that is in Achronix's Speedster22i FPGAs that have been shipping in production since 2013.

With Speedcore, designers specify the optimal die size, power consumption and resource configuration required for their end application. They define the quantity of look-up-tables (LUTs), embedded memory blocks and DSP blocks. Additionally, designers define the Speedcore aspect ratio, IO port connections and can make tradeoffs between power and performance. Achronix delivers a GDS II of the Speedcore IP that engineeers integrate directly into their SoC, and a custom, full-featured version of the ACE design tools that customers use to design, verify and program the functionality of the Speedcore eFPGA. Speedcore eFPGA products are fully supported by Achronix's ACE design tools.

Complete article, here

Intel Atom E3900 processors: 'computing nearer the sensor'

Intel has announced the latest generation of its Atom processors for IoT applications. The Intel Atom processor E3900 series is designed from the ground up to support the rapid development and the growing complexity of IoT businesses. The result, Intel aserts, is a processor exceptionally capable of delivering on performance, processing and scalability. Intel claims new levels of security, determinism, and image and video processing power for the MPU; there is also a variant specifically intended for automotive applications. The A3900 automotive series will enable a complete software defined cockpit solution that includes in-vehicle infotainment (IVI), digital instrument clusters and advanced driver assistance systems (ADAS) - all in a single, compact and cost-effective SoC. The A3900 series will allow car makers to offer new levels of determinism for real-time decision-making required in next-generation cars. It is currently sampling with customers and will be available in Q1 2017.

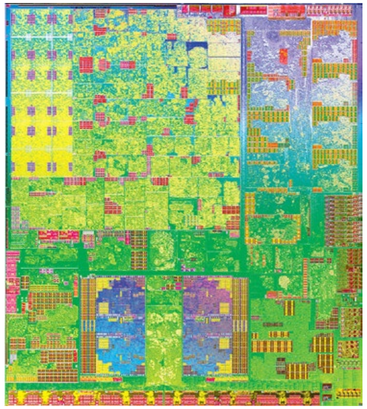

The Intel Atom processor E3900 series will make the edge and fog more intelligent - enabling many of the processing needs to take place at or near the data sensor and alleviating the need to push all processing to the data centre. Fog computing, also known as fog networking, is a decentralized computing infrastructure in which computing resources and application services are distributed in the most logical place at any point from the data source to the cloud.

With more than 1.7 times more computing power over the previous generation the E3900 Series is designed to enable faster memory speeds and memory bandwidth to provide the efficient processing capability needed for edge to cloud network computing. Built into a compact flip chip ball grid array (FCBGA) and featuring 14 nm silicon technology, the Intel Atom processor E3900 series is suitable for a wide range of IoT applications, where scalable performance, space and power are at a premium.

Complete article, here

EnOcean to bring Bluetooth LE into energy-harvested operation

EnOcean, developer of energy-harvesting wireless technology, ha sannounced its first energy-harvesting wireless modules for 2.4 GHz BLE (Bluetooth Low Energy) systems. The new modules complement the existing EnOcean wireless technology in the sub-1 GHz band and ZigBee wireless products in the 2.4 GHz frequency band.

The first product in the EnOcean 2.4 GHz BLE portfolio is the batteryless PTM 215B switch module, which also comes with an NFC function for pairing and set-up, making it easier to train and configure the switch. The PTM 215B is based on the established form factor of the PTM 21x module, invented by EnOcean, and can be integrated into a large number of existing switch designs. The switch is also available as a white label end product. EnOcean aims to enable product manufacturers to develop reliable and self-powered solutions on the 2.4 GHz frequency band for use in smart homes and modern light control all over the world.

Manufacturers of BLE-based systems for the 2.4 GHz band can now incorporate energy harvesting technology from EnOcean into their portfolios and develop bat-

terylless, room-based wireless controllers. If they choose to use the standardized PTM 21x form factor, switch manufacturers can integrate the new 2.4 GHz module into their existing product ranges and use maintenance-free BLE systems that produce energy from motion.

EnOcean has also disclosed a change in its branding for energy harvesting wireless modules; the company will market its energy harvesting wireless modules and white label products under the Dolphin brand, with the tag line, "Dolphin – Self-powered IoT by EnOcean".

Complete article, here

Industrial PLC reference design aimed at 'Industry 4.0'

Maxim Integrated has configured a programmable logic controller (PLC) development platform and reference design that demonstrates all of the functional blocks the company has to offer for the erea in which every node in the factory will be connected, with access to the 'cloud'.

As well as incorporating an increased function set for the 'Industry 4.0' context, the design demonstrates the much-reduced footprint in both space and power that is possible relative to previous-generation products. The Pocket IO PLC Development Platform can, Maxim says, help transform traditional manufacturing processes with real-time intelligence, adaptive manufacturing, and distributed control, giving designers the ability to achieve the smallest form factor and highest power efficiency for next-generation PLC designs. As well as being called on to deliver always-on operation with real-time data gathering, PLCs require fan-less operation due to harsh industrial environments. As a result, highly efficient power solutions are required to minimize heat dissipation.

Maxim's Pocket IO PLC supports real-time intelligence to quickly and effectively make decisions, adaptive manufacturing to avoid potential downtime, and distributed control to provide redundancy; features include; The Pocket IO provides the following key advantages to increase productivity:

- Real-time intelligence: Fast data processing provides the necessary data to make intelligent decisions quickly and effectively to optimize yield.
- Adaptive manufacturing: Manufacturing flexibility allows for real-time changes and adjustments to avoid potential downtime.
- Distributed control: Ultra-small footprint of less than 10 cubic inches and smart energy

Complete article, here

Linaro organisation, with ARM, aims for end-end open source

With the objective of producing reference software for more secure connected products, ranging from sensors and connected controllers to smart devices and gateways, for the industrial and consumer markets, Linaro has announced LITE: Collaborative Software Engineering for the Internet of Things (IoT).

Linaro and the LITE members will work to reduce fragmentation in operating systems, middleware and cloud connectivity solutions, and will deliver open source device reference platforms to enable faster time to market, improved security and lower maintenance costs for connected products. Industry interoperability of diverse, connected and secure IoT devices is a critical need to deliver on the promise of the IoT market, the organisation says. "Today, product vendors are faced with a proliferation of choices for IoT device operating systems, security infrastructure, identification, communication, device management and cloud interfaces."

Initial technical work will be focused on delivering an end to end, cross-vendor solution for secure IoT devices using the ARM Cortex-M architecture. This will include a bootloader, RTOS platform, security, communications, middleware and a choice of application programming tools. LITE will also work on Cortex-A based smart device and gateway solutions for IoT using Linux.

Complete article, here

Innovative use of MEMS for mid-range pressure sensing

Melexis' (Tessenderlo, Belgium) MLX90819 sensor IC is targeted at measuring mid pressure (10-50 bar) levels; it employs a compact, robust micromachined solution with accurate sensor functionality, integrated signal processing and analogue or SENT outputs.

Melexis aims to "revolutionize the sector" with this design, applying its MEMS technology to an area that had not been previously adequately served by a fully integrated sensor solution, with system design benefits.

Available as a bare die, the MLX90819 runs off a standard 5V supply. It can be employed to accurately determine fluid pressure levels in wide variety of automotive and industrial focused applications. These include the monitoring of engine oil, transmission oil, engine coolant, vehicle fuel lines, air conditioning refrigerant and truck air brake pressures. The IC delivers a relative pressure value. It has a 1 msec response time, plus ±2.5% over life accuracy and ±0.2% linearity across its entire operational pressure range.

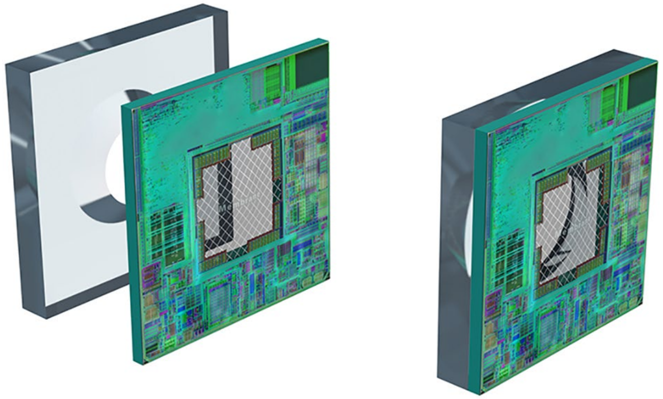

By using a more advanced process technology the MLX90819 is able to use a microcontroller based architecture. This allows a much higher degree of sophistication to be incorporated onto the semiconductor die - with all digital processing, data conversion, temperature offset, gain compensation and filtering tasks being taken care of directly. A 16-bit analogue-to-digital converter, 12-bit digital-to-analogue and 16-bit processor core are embedded directly into each sensor IC.

This robust IC is suited to both heavy duty industrial and automotive deployment, with full AECQ100 compliance and a specified temperature range of -40°C to 150°C.

Complete article, here

Siemens gateway; industrial IoT launch-point

Distributor RS Components is to be sole supplier of Siemens' SIMATIC IOT2020 IoT gateway targeting industrial engineers, education and the maker community; RS says it is extending its portfolio of industrial IoT (Internet of Things) devices with the exclusive supply arrangement.

RS comments that a key global trend in IoT technology open source software, including easy-to-use IDEs, and ever-improving hardware. Siemens and RS joined forces to offer an IoT platform for the engineers of tomorrow, and the IOT2020 encapsulates these trends, providing the simplest way for engineers to get started with the industrial IoT and meet the challenges of an increasingly connected world. The IOT2020 is an open and flexible IoT gateway that is designed for continuous industrial operation and comes with the appropriate certificates. It can be used to retrieve, process, analyse and send data to almost any kind of device or network due to its various interfaces including Ethernet, USB and micro SD. The gateway is compatible with open source software such as the Arduino IDE and Yocto Linux, to benefit from programming in high-level languages such as Java, C++ and JSON. This openness enables various communication possibilities to further automation hardware or sensors via Modbus, PROFINET or other protocols, or even the direct connection to cloud solutions via MQTT or AMQP. In addition to the on-board interfaces the IOT2020 is expandable with Arduino shields and via an on-board PCIe port.

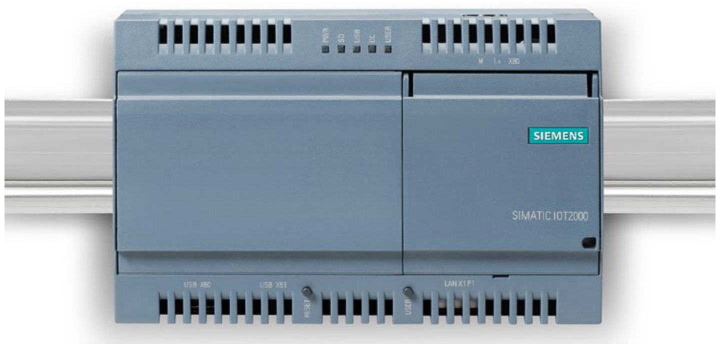

The €89.00 device, which is an industrial certified product with UL and CE approval, is suited for educational purposes.

Complete article, here

Printed electronics pioneer funded to commercialise “FlexLogIC”

PragmatIC is a Cambridge, UK-based company that has been developing a technology for printing electronics – with all circuit element, including transistors and other active devices – on flexible substrates, with the aim of producing very low cost devices for areas such as wearables, and “smart packaging”.

The company has attracted funding of £18million (UK Pounds) to take the technology to production readiness; investors include packaging and labelling giant Avery Dennison, and follow-on investment from ARM and Cambridge Innovation Capital.

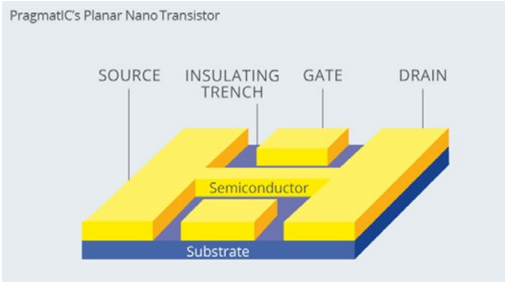

The company regards its technology as an enabler for the Internet of Things and will according to Mike Muller, CTO of ARM ‘open up a whole new world of computing'.

It is built around a “planar nano-transistor” in which the conducting channel and gate structure are co-planar. Circuitry is laid down in layers that are of the order of 100 nm thick, and have feature sizes of the same order. A type of self-alignment is inferred, and the company says this maintains critical spatial relationships even on substrates that bend and stretch. The process is capable of complexity of [at least] enough construct an RF or NFC tag. Among the company's claims for its process are that it can be 1/10th the cost of silicon, and in “many applications” can product a circuit for under 1¢.

(Images are from PragmatIC's website)

Complete article, here

SoCs and development kits for vehicle-to-everything comms

Renesas Electronics has announced vehicle-to-vehicle and vehicle-to-infrastructure communication system solutions; they implement real-time processing and and pre-warning of information covering 360 degrees around the vehicle.

Renesas' introductions include two system-on-chips (SoCs) that will ease the development process for vehicle-to-vehicle (V2V) and vehicle-to-infrastructure (V2I) communication systems. One of the solutions includes the R-Car W1R 760 MHz band wireless SoC for the Japanese market, and the new R-Car W2H SoC that features a high-performance security engine that is needed for V2X systems designed for the Japanese, U.S., and European markets. The other solution consists of the W2H SoC combined with the R-Car W2R 5.9 GHz band wireless communication SoC developed for U.S. and European markets. By combining these new SoCs, system developers can now develop systems that support the corresponding V2X standard for the Japanese, U.S., and European markets. Renesas also delivers development starter kits that will enable system developers to immediately start application software development, leading to significant reductions in the development period for V2X systems.

Renesas' wireless communications SoC kits comprise;

- R-Car W1R: Wireless communications SoC that supports the

Japanese communications standards (760 MHz band).
- R-Car W2R: Wireless communications SoC that supports the European and US communications standards (5.9 GHz band)
- R-Car W2H: V2X communications processor SoC that supports the Japanese, US, and European security standards.
- R-Car W1R + R-Car W2H kit: Implements V2V and V2I communication systems for the Japanese market.
- R-Car W2R + R-Car W2H kit: Implements V2V and V2I communication systems for the European and US markets.

Complete article, here

ST opens up Cortex-M7 designs with development system expansion

STMicroelectronics has introduced STM32F7 microcontroller lines and added accessories and options to the development ecosystem, easing access to high-performance embedded design based on the ARM Cortex-M7 core.
The latest STM32F722 and STM32F723 microcontrollers in the very high-performance STM32F7 series reduce memory footprint by integrating value-added features including code-execution protection and high-speed USB physical-layer (PHY) circuitry that streamline development of connected applications. The STM32F732 and STM32F733 variants come with extra cryptographic features on-chip, such as an efficient AES256 HW engine. There are package options from a 64-pin LQFP up to 176-pin LQFP or UFBGA for projects demanding high I/O count, and 256 kB or 512 kB of on-chip Flash memory with 256 kB RAM.

Pin, package, and software compatibility between the new devices and higher-end STM32F7 variants with 256 kB to 2 MB Flash memory and 256 kB to 512 kB RAM in packages up to 216-pin TFBGA simplifies design scaling for product differentiation and future-proofing.
ST is also adding a Discovery Kit, based on the STM32F769, with expansion connectors that include an 8-pin socket for Wi-Fi modules. The kit supports Power over Ethernet (PoE). Other features include a 512-Mbit Quad-SPI (QSPI) Flash interface and new accessories that enhance flexibility and extend application reach, including the B-LCDAD-RPI1 15-pin single row Flexible Printed Circuit DSI adapter board and B-LCDAD-HDMI1 DSI-to-HDMI adapter, which attach a variety of displays.

Complete article, here

Silicon Labs takes IoT RTOS in-house with Micrium purchase

Silicon Labs has announced that it is acquiring RTOS specialist Micrium. Silabs comments that this, "helps simplify IoT design for all developers by combining a leading, commercial-grade embedded RTOS with Silicon Labs' IoT expertise and solutions". Micrium's RTOS and software tools will continue to be available to all silicon partners worldwide, giving customers a wide range of options, even when using non-Silicon Labs hardware. Micrium will continue to fully support existing as well as new customers. Micrium's flagship µC/OS RTOS family is (Silabs says) recognized for reliability, performance, dependability, impeccable source code and extensive documentation. Micrium's RTOS software has been ported to more than 50 microcontroller architectures and has a global footprint with more than 250,000 downloads across all embedded vertical markets, with solutions certified to meet safety-critical standards for medical electronics, avionics, communications, consumer electronics and industrial control.

Complete article, here

Energy-efficient biomedical sensor hub for wearables

Belgian research centre imec, and the associated Holst Centre have designed a sensor hub integrated as a system-on-chip (SoC) intended for a broad range of wearable health devices and applications.

The SoC combines an "unprecedented" number of biomedical analogue interfaces into a single chip, with on-board digital signal processing, high fidelity operation, and multi-day monitoring capability with a single battery. Thanks to its small form factor, the SoC can be easily integrated in new innovative designs enabling maximum user comfort.

The biomedical analogue interfaces include three ECG channels, photoplethysmography (PPG), galvanic skin response (GSR), two multi-frequency bio-impedance (BIO-Z) channels to support new applications such as impedance-tomography, body fluid analysis and stroke volume measurements, and three reconfigurable channels. Imec's and Holst Centre's SoC combines advanced biomedical readouts, supported by an ARM Cortex M0+ controller and accelerators for sample-rate conversion, matrix processing, data compaction, and power management circuitry (PMIC). The PMIC operates from a battery source (2.9-4.5V) and generates the required voltages for the readout IC.

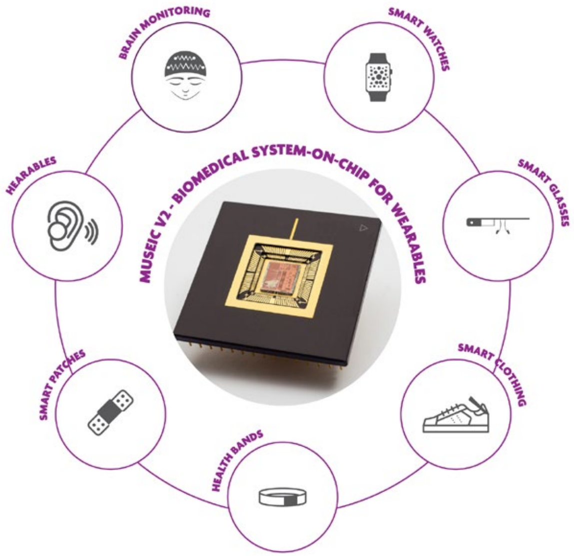

Complete article, here

WHY DO YOU NEED AN RTOS IN AN IOT WORLD?

By Stefan Ingenhaag, Renesas Electronics Europe

An RTOS is an important element in enabling rapid, responsive and reliable functionality that facilitate internet connectivity.

There are more things than ever on the web today thanks to the rise of smart devices. We have networked everything from mobile phones to home automation and security services. Even our lights and refrigerators are becoming aware of their environment and attach to the cloud with reams of data enhancing our digital ecosystem. Most recently biometric monitoring has joined this conversation with wearables and other sensing devices monitoring everyday life. This growing and evolving consumer market requires rapid turn-around times to realize product opportunities. To accelerate delivery timeframes, designers are exploiting existing, off-the-shelf software over proprietary solutions, and must overcome the hurdles associated with embedded design

RTOS, the proven solution

An ‘off-the-shelf’ RTOS provides a proven solution for critical layers of the software stack. This frees developers to focus efforts on the innovative functionality that differentiates the product. Engineering teams typically only need to configure inbuilt capabilities to suit their particular requirements and proceed with implementation. This is further accelerated by the use of well documented APIs that find direct support within integrated development environments. Using a tried and tested proven solution for multitasking support and internet connectivity reduces the debugging and testing requirements. This will go a long way towards mitigating the risk of problems with interoperability, while simultaneously easing the “time crunch” associated with bringing a product to market. As seen in the chart (Figure 1), the time that engineers and designers are given to complete a task has remained stagnant over the last five years, increasing just 1% since 2011.

RTOS vs OS

Should you use an OS or an RTOS? The line is becoming increasingly blurred. Understanding how an RTOS facilitates connecting devices to the Internet of Things requires an understanding of what an RTOS is and how is differentiates from a generic alternative. The key aspect that differentiates these is responsiveness. Responsiveness is crucial to empowering devices in meeting the performance expectations of edge applications that incorporate low level hardware interfacing. An RTOS helps devices manage software tasks in a way that improves their ability to meet timing requirements. A real time OS is a specially optimized process man-

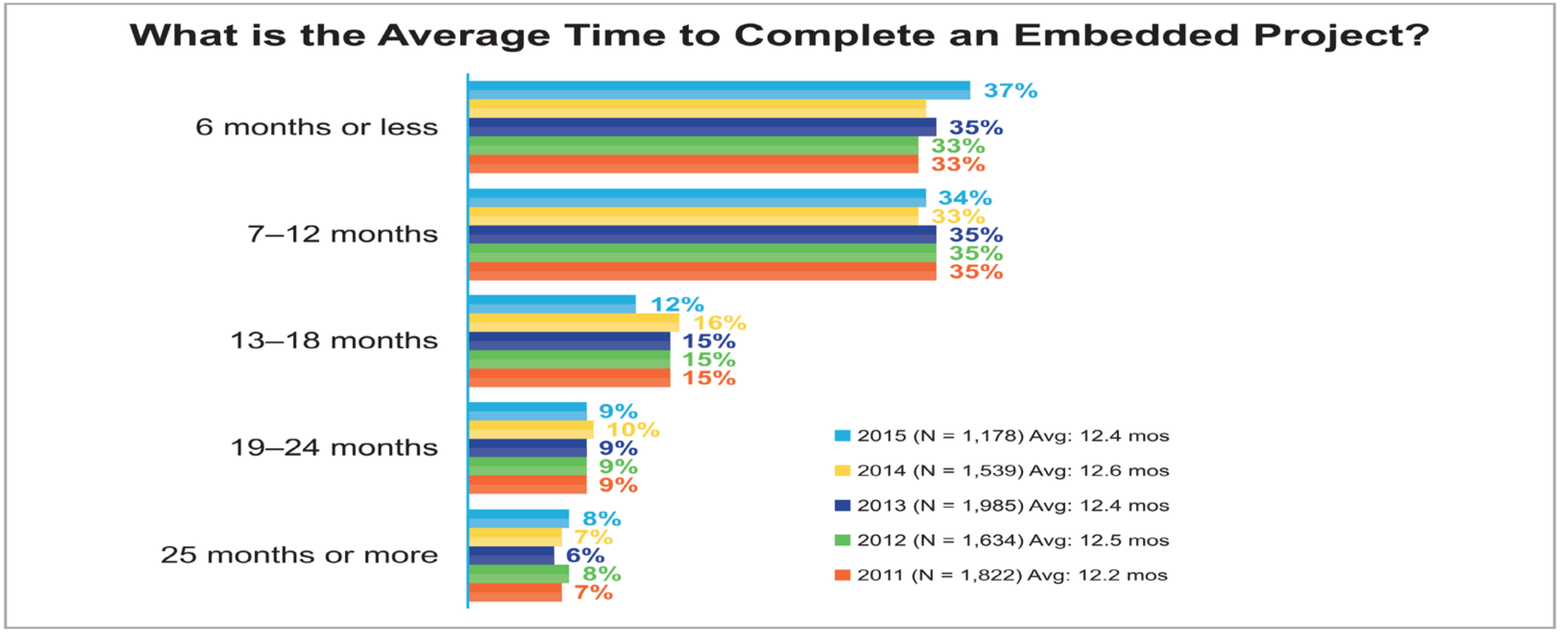

Source: UBM Electronics Embedded Markets Study 2015, slide 21)

Figure 1. *Time allocated to embedded system design tasks. Source: UBM Electronics Embedded Markets Study 2015*

agement service that enables timely task context switching. Timeliness is expressed as latency and determinism – measures of the speed of switching between tasks and the time variance or jitter in the transition. An RTOS can respond with lower latency and determinism to events than a generic OS and this is why they can be found running on devices handling time critical, hardware specific event processing and digital signal processing algorithms.

The article continues by exploring the aspects ot "real time" that are pertinent in the IoT context, and notes the advantages of using an off-the-shelf solution. Click for pdf.

Download PDF of Article

Find RTOS for the IoT on EETsearch

LOW NOISE OR LOW POWER? YES!

BY MARK LOONEY, ANALOG DEVICES

A survey of MEMS accelerometers reveals that the lowest noise and lowest power are not available in the same product at this time.

When comparing a "low-noise" accelerometer, such as the ADXL355, with a popular "low-power" accelerometer, the ADXL355 presents the following trade-offs:

- At 20 µg/√Hz, the noise density is 9x lower.
- At 338 µW, the power dissipation is ~13x higher.

For applications that will power cycle their sensors to save energy when they are not in use, the noise/power relationship can actually be quite different. This difference comes from a source that may be surprising to some: settling time. In applications that need to average a sequential array of sensor data to achieve key uncertainty criteria, the time it takes to fill that array has a direct impact on the overall settling time. For example, the Allan Variance curve in the ADXL355 datasheet suggests that an averaging time of 0.01 seconds will reduce its uncertainty to less than 100 µg. Achieving a similar level of uncertainty in the "low-power" sensor will require an averaging time that is 81 times longer than the ADXL355, since noise reduction in an averaging filter is proportional to the square root of the averaging time.

Equations 1 and 2 quantify this trade-off in terms of the energy it takes for each sensor/filter combination to support this level of precision in a single data record. The results of this estimation are quite interesting, because the much shorter averaging time (t_{LP} = 81 x $t_{ADXL355}$) causes the energy requirement for the ADXL355 ($E_{ADXL355}$) to be six times lower than the energy requirement for the "low-power" component (E_{LP}).

$$E_{ADXL355} = t_{S355} \times V_{MIN355} \times I_{355}$$
$$E_{ADXL355} = 0.01\ s \times 2.25\ \text{V} \times 0.00015\ \text{A} = 3.38\ \mu\text{J} \quad (1)$$

$$E_{LP} = t_{LP} \times V_{LP} \times I_{LP}$$
$$E_{LP} = 81 \times t_{S355} \times V_{LP} \times I_{LP} \quad (2)$$
$$E_{LP} = 81 \times 0.01\ s \times 2\ \text{V} \times 0.000013\ \text{A} = 21.1\ \mu\text{J}$$

Equation 3 captures the power dissipation that will come from these energy levels, with respect to the time between each measurement (T).

$$P_{ADXL355} = \frac{E_{ADXL355}}{T} = \frac{3.38\ \mu\text{J}}{T} \quad (3)$$
$$P_{LP} = \frac{E_{LP}}{T} = \frac{21.1\ \mu\text{J}}{T}$$

The graphical view of this relationship (Figure 1) offers a couple of interesting observations. First, at measurement cycle times (T) that are lower than 0.81 seconds, the "low-power" device will be supporting continuous operation. Second, for measurement cycle times that are greater than ~0.13 seconds, the ADXL355 solution will consume less power. The bottom line is that with an open mind, sometimes we can achieve the lowest power *solution*, by using the lowest-noise (highest performing) components.

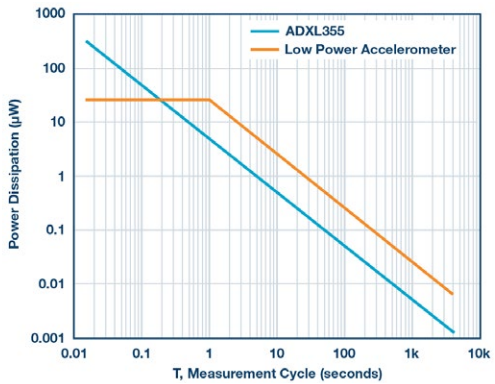

Figure 1. *Power dissipation vs measurement cycle time*

PULSE OXIMETRY USING A PROGRAMMABLE MIXED SIGNAL ARRAY IC

By Ahmed Asim Ghouri, Silego Technology

Pulse oximetry is a simple non-invasive method of monitoring the percentage of haemoglobin (Hb) which is saturated with oxygen. The pulse oximeter consists of a probe attached to the patient's finger or ear lobe which is linked to a processing unit. The unit displays the percentage of Hb saturated with oxygen together with an audible signal for each pulse beat, a calculated heart rate and in some models, a graphical display of the blood flow past the probe. Audible alarms which can be programmed by the user are provided.

The colour of blood varies depending on how much oxygen it contains. A pulse oximeter shines two beams of light through a finger (or earlobe etc.), one beam is red light (which you can see when a pulse oximeter is used), one is infrared light (which you don't see). (Netter T, 2004)

These two beams of light can let the pulse oximeter detect what colour the arterial blood is and it can then work out the oxygen saturation. However there are lots of other parts of a finger which will absorb light (such as venous blood, bone, skin, muscle etc.), so to work out the colour of the arterial blood a pulse oximeter looks for the slight change in the overall colour caused by a beat of the heart pushing arterial blood into the finger.

This change in colour is very small so pulse oximeters work best when there is a good strong pulse in the finger when the probe is on. If the peak signal value is too low the measured oxygen saturation may not be reliable and with lower signal peak value the pulse oximeter will not be able to work; to acquire a strong signal output from the IR sensor pulse oximeter increase the intensity of RED and IR LED's in successive steps.

Oxygen concentration

A Pulse-oximeter monitor displays the percentage of blood that is loaded with oxygen. More specifically, it measures what percentage of hemoglobin (Tilakaratna, n.d.), the protein in blood that carries oxygen, is loaded. Acceptable normal ranges for patients without pulmonary pathology are from 95 to 99 percent. For a patient breathing room air at or near sea level, an estimate of arterial pO_2 can be made from the blood-oxygen monitor "saturation of peripheral oxygen" (SpO_2) reading.

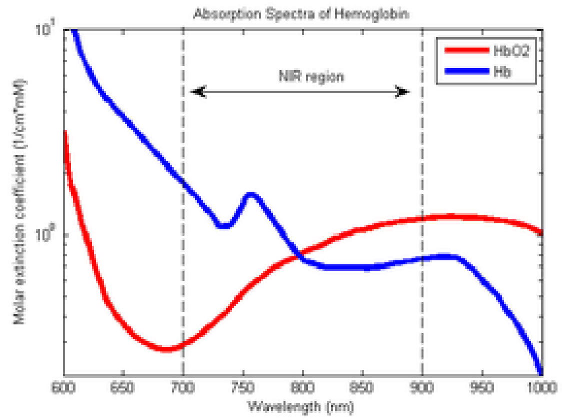

Figure 1. *Absorption of red light and infra-red light (Twonsend, 2001)*

A typical pulse oximeter utilizes an electronic processor and a pair of small light-emitting diodes (LEDs) facing a photodiode through a translucent part of the patient's body, usually a fingertip or an earlobe. One LED is red, with wavelength of 660 nm, and the other is infrared with a wavelength of 940 nm. Absorption of light at these wavelengths differs significantly between blood loaded with oxygen and blood lacking oxygen. Oxygenated haemoglobin absorbs more infrared light and allows more red light to pass through. Deoxygenated haemoglobin allows more infrared light to pass through and absorbs more red light. The LEDs sequence through their cycle of one on, then the other, then both off, about thirty times per second which allows the photodiode to respond to the red and infrared light separately and also adjust for the ambient light baseline. (Pulse oximetry, n.d.) The amount of light that is transmitted (in other words, that is not absorbed) is measured, and separate normalized signals are produced for each wavelength. These signals fluctuate in time because the amount of arterial blood that is present increases (literally pulses) with each heartbeat. By subtracting the minimum transmitted light from the peak transmitted light in each wavelength, the effects of other tissues is corrected for. The ratio of the red light measurement to the infrared light measurement is then calculated by the processor (which represents the ratio of oxygenated haemoglobin to deoxygenated haemoglobin), and this ratio is then converted to SpO_2 by

Figure 2. *Pulse-oximeter finger probe*

the processor via a lookup table based on the Beer–Lambert law. The hardware and software to acquire SpO_2 data will be included in the further extension of this application note.

Signal processing

Some initial signal processing is required when trying to extract oxygen concentration from the signal coming from a finger sensor. The calculations follow Beer-Lambert Law (Matviyenko, 2011) to assess the percentage of the oxygenated blood. Mathematically given as:

$$I = I_0 \exp(\alpha(\lambda)cd)$$

Figure 3 shows how light is absorbed in the finger.

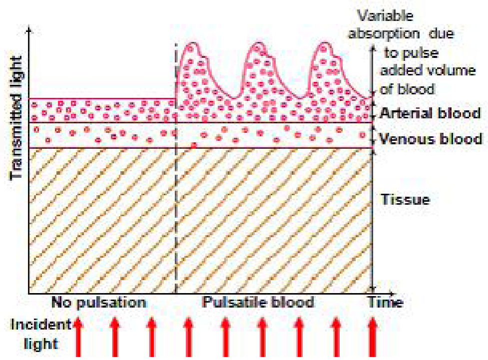

Figure 3. *Absorption of light (Matviyenko, 2011)*

There are several components which contribute to the absorption of light, listed below :

1. Oxygenated haemoglobin in the blood
2. De-oxygenated haemoglobin
3. Absorption that is not from arterial blood
4. Optical attenuation due to scattering, geometric factors etc.

Figure 4 shows the main block diagram of the Pulse Oximeter application using the GreenPak device, SLG46140.

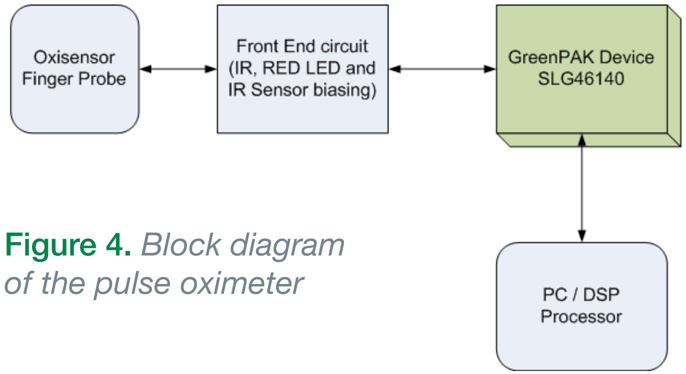

Figure 4. *Block diagram of the pulse oximeter*

This detailed article continues to show how the single mixed-signal programmable part generates drive signals for both IR and RED LEDs, and how the ADC within the device will sample the IR Sensor output and send data serially out. Click for pdf.

Download PDF of Article

Find Mixed signal array IC for oximetry on EETsearch

HIGH-PERFORMANCE GAN-BASED 48-V TO 1-V CONVERSION FOR POL APPLICATIONS

By Alberto Doronzo, Texas Instruments

48V to point-of-load (PoL) conversion is common in numerous markets, including telecommunications, industrial, aerospace, and increasingly, server environments. The prevalent system solution is a two-stage conversion from 48V to an intermediate bus (9V to 12V) and then to PoL, typically 0.8V to 3.3V.

Data centre operators are particularly interested in single-stage 48-V to low-voltage DC/DC conversion and its associated efficiency gains, because data centres may consume as much as 140 TWh of power usage in the U.S. yearly by 2020 [Ref. 1], or about 3.5%t [Ref. 2] of overall U.S. power usage. The PoL in the application discussed in this article includes general-purpose and application-specific processors, field-programmable gate arrays (FPGAs) and application-specific integrated circuits (ASICs).

While the limitations of silicon-based converters force a two-stage solution, new gallium nitride (GaN) devices enable single-stage conversion. Two key design requirements are efficient power conversion and ultra-fast load transient response of the ASIC/processor core rails.

GaN's advantages are strongly evident during single-stage conversion, with higher efficiency and smaller size than traditional solutions, while simultaneously reducing bill of materials (BOM) count and systems cost. In this article, I will discuss the design of a 48-V to 1.0-V PoL 50-A converter using a GaN half-bridge module and a high-performance controller for the single-stage topology, achieving efficiency exceeding 90%.

Introduction

GaN devices provide stronger performance and more ideal switching characteristics compared to silicon. GaN enables more efficient and smaller power systems. However, the greatest benefit of GaN does not come from a direct silicon substitution in an existing topology; the key to using GaN is to identify applications and converter topologies that specifically harness GaN's advantages. With GaN's low switching loss and zero reverse recovery, hard-switched half-bridge topologies show noticeable performance improvements. One of the strongest application areas for GaN is in the 48-V input converter space, targeting telecom and industrial customers as well as aerospace and server applications.

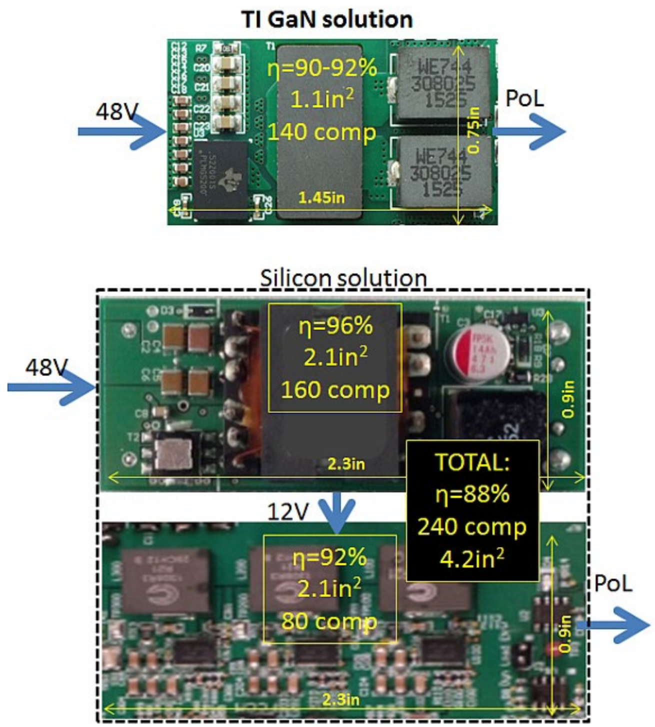

Figure 1. *GaN solution versus a silicon traditional solution.*

In 48-V telecommunications and server applications, there is strong demand for conversion down to PoL in order to power processing-intensive cores such as microprocessors, FPGAs and ASICs. Traditionally, the conversion from 48V down to PoL uses a two-stage approach. This entails a brick-type bus converter performing an isolated conversion between 48V and an intermediate voltage such as 12V or 9V using a hard-switched half-bridge or full-bridge topology, or a resonant topology such as inductor-inductor-capacitor (LLC). Next, a multiphase buck converter converts down to PoL at up to 200A, while handling the transient response requirements of a highly dynamic digital load.

The traditional solution's beginning-to-end efficiency is a product of the two individual stages. With the first stage peak efficiency of 96% and the multiphase buck stage peak efficiency close to 92%, the total efficiency is only 88% (Figure 1) at frequency ranges between 100 to 200 kHz. The two-stage solution also has a larger board-area footprint.

A single-stage approach would improve efficiency and reduce board area. While this has been the desired target for a long time, silicon-based designs have not demonstrated a compelling solution. GaN enables an elegant new topology that offers PoL-level transient response from a single, efficient 48-V to PoL conversion stage, while running at higher frequencies close to 600 kHz, further aiding size reduction.

An isolated half-bridge with a current-doubler output topology is extremely advantageous for this type of application, as shown in Figure 1. In this topology, a GaN half-bridge (using the LMG5200) connects to the transformer, which is referenced to a split capacitor network to provide voltage balancing.

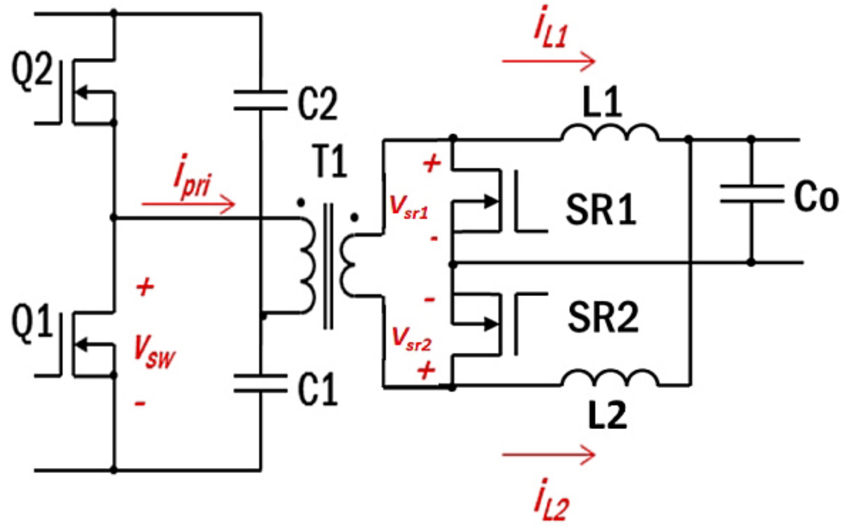

Figure 2. *This figure shows a half-bridge with current-doubler output.*

The secondary side connects in a current-doubler synchronous rectifier, which enables better heat distribution, a simpler transformer design and the advantage of foregoing high-side switches, as both are ground referenced. Since the charging current of the capacitor is the sum of the two inductor currents, which are 180° out of phase, the output voltage ripple is halved.

Figure 2 shows synchronous rectifiers instead of diodes. The average current through these rectifiers can reach 25A, which a diode would not be capable of sustaining in an application. Note that the rectifiers will be operating in continuous conduction mode and will be subject to reverse recovery; therefore, the preferred choice is GaN FETs, as they present the advantage of zero recovery loss. Alternatively, using silicon FETs requires very tight dead-time control.

This article continues by demonstrating the waveforms in the conversion path, and looking at hardware design, converter efficiency and dynamic performance – click for pdf.

Download PDF of Article

Find GaN FETs for POL applications on EETsearch

IMPLEMENTING IEEE'S 802.3BT 71W POE STANDARD - NOW

By Christopher Gobok, Linear Technology

With Power over Ethernet (PoE) being a popular and mature technology, it's no surprise that Power Sourcing Equipment (PSE) and Powered Device (PD) developers are eager to jump onto the next IEEE bandwagon and start delivering higher levels of power down Ethernet cables. However, as with many industry standards, IEEE PoE standards often finalize much later than is needed by the market.

The latest timeline from the IEEE PoE Task Force shows the 71W 802.3bt standard won't be ready until early 2018. That's a long time to wait in technology years, which can easily translate to missed business opportunities or markets. So, what are PSE and PD developers supposed to do in this case? Why, design and manufacture now of course! Developers who want to be first to the 802.3bt market have Linear Technology's LT4295 IEEE 802.3bt PD controller available to them today, compliant to draft 2.0 of the standard.

More power, more potential

Not surprisingly, the main focus for the IEEE PoE Task Force and developers has been how much power is ultimately provided to PDs. In 2003, the IEEE PoE Task Force established in the original IEEE 802.3af PoE standard that ~13W would be available at the PD input's RJ-45 jack. Since then, the market has continued to demand more power. So, in 2009, the IEEE PoE Task Force revised the standard and released IEEE 802.3at (also known as PoE+), which increased the maximum PD power level to 25.5W. Fast-forward to today and it is expected that the current revision, IEEE 802.3bt draft 2.0 (also known as PoE++ or 4PPoE), will provide PDs with up to 71W of power.

With more power, developers can easily add more features and upgrade existing products. Take the security camera as an example of an application that has evolved throughout the PoE years. With only 13W available, the first PoE-powered security cameras were simple stationary units. However, when 25.5W was allocated by 802.3at, extra power was available to drive multiple embedded motors which provided security cameras with pan, tilt and zoom capabilities. Now, with 802.3bt's 71W to tap into, pan-tilt-zoom security cameras can also integrate fans and heating elements to support operation at extreme temperatures. For some, higher PoE power levels may unlock distinctly new markets. For example, a traditional LED lighting manufacturer, who only produced ceiling units controlled by wall switches, but can now produce PoE-enabled units that will help pave the way for smart homes or buildings. Whether more power evolves or revolutionizes the end product, there is clearly more market potential with each iteration of the PoE standard.

Changes to the PSE-PD link

As you might have already guessed, 4PPoE stands for "4-pair PoE", since 802.3bt takes advantage of all four twisted pairs in the Ethernet cable to transmit power up to 100m; in older PoE standards, 4-pair power delivery was non-compliant. See Figure 1 for a typical block diagram of an IEEE 802.3bt PSE-PD link. Cabling requirements are still not completely defined, but cabling committees and manufacturers are guessing that 802.3bt will require at least Category 5E cabling in order to support the increased power levels while operating with 10GBASE-T (10 Gbit/sec Ethernet data speed). In any case, because we are now reaching the current handling limitations of the Ethernet cable, you may have to pay attention to cabling-system performance characteristics that were previously ignored.

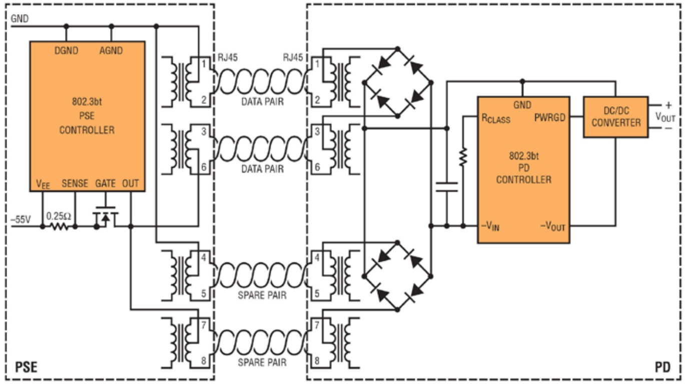

Figure 1. *Typical block diagram of an IEEE 802.3bt PSE-PD link*

802.3bt introduces two new PD topologies: single-signature and dual-signature. That is, a single-signature PD is an 802.3bt PD that shares the same detection signature, classification signature and maintain power signature (MPS) between both pairsets, whereas a dual-signature PD is an 802.3bt PD that has independent signatures between both pairsets. New 802.3bt designs will, no doubt, gravitate towards the simpler and more cost-effective single-signature topology, which only calls for a single PD interface. Dual-signature PDs require two parallel PD interfaces, one for each pairset, where the power from two PSEs are summed together after each PD interface. Essentially, the dual-signature topology uses, for example, two 25.5W PDs to make a single 51W PD – a complex solution that can cost twice as much as a single-signature 51W PD.

The 802.3bt detection process has been expanded to be able to differentiate, not only that the connected PD is an 802.3-compliant PD, but also to determine whether a single or dual signature PD is connected. As such, detection is now augmented by a Connection Check to determine the single or dual signature PD configuration.

802.3bt introduces four new high power PD classes, bringing the total number of single-signature classes to nine as shown in Table 1. Classes 5 through 8 are new to PoE and translate to PD power levels ranging from 40.0W to 71.0W. PSEs still have their choice of using the physical layer (i.e. 5-event classification for 71W) or data link layer (i.e. link layer discovery protocol, LLDP) to classify PDs, and PDs still need to be able to support both classification schemes in order to be compliant. 802.3bt PDs may also implement an optional extension of the physical layer classification, known as Autoclass, where an 802.3bt PSE measures the actual maximum power draw of a connected PD. This handy power management feature allows, for example, a PSE to allocate leftover power to additional light bulbs if it knows that a particular bulb will draw less than its class power.

SINGLE-SIGNATURE PDs		DUAL-SIGNATURE PDs	
CLASS	PD POWER AVAILABLE	CLASS	PAIRSET PD POWER AVAILABLE
0	13W	—	—
1	3.84W	1	3.84W
2	6.49W	2	6.49W
3	13W	3	13W
4	25.5W	4	25.5W
5	40W	5	35.5W
6	51W		
7	62W		
8	71W		

Table 1. *IEEE 802.3bt PD classes and power levels*

This article continues by exploring low-power configurations, how the new standard may maximise available power, and backward-compatibility issue, ansets out what is practical to market immediately. Click for pdf download.

Download PDF of Article

Find POE 802.3BT 71W Standard implementations on EETsearch

OPTIMIZED FOR ADAS APPLICATIONS: NEW COMPILER CHALLENGES

By Dr. Alexander Herz, Tasking

The compiler is a key tool for the cost-efficient design of ADAS (advanced driver assistance systems) applications. However, the tools currently available must be better adapted to this challenging task. This includes considering the code structures and specific safety requirements typical of ADAS applications.

When planning a compiler technology roadmap, you will inevitably touch the issue of advanced driver assistance systems (ADAS), which all major OEMs and software suppliers of the automotive industry are committed to. A closer look, though, will raise some questions: What requirements are placed on compilers and toolsets by ADAS applications? Are these things related at all? What are the differences between traditional automotive applications and ADAS applications?

ADAS applications as a challenge

To better support the task of driving autonomously, vehicles need to be much more aware of their surroundings. Several new sensors (Radar, Lidar, cameras, etc.) can be used to detect road markings, other vehicles, obstacles and other relevant environmental data with high resolution (Figure 1).

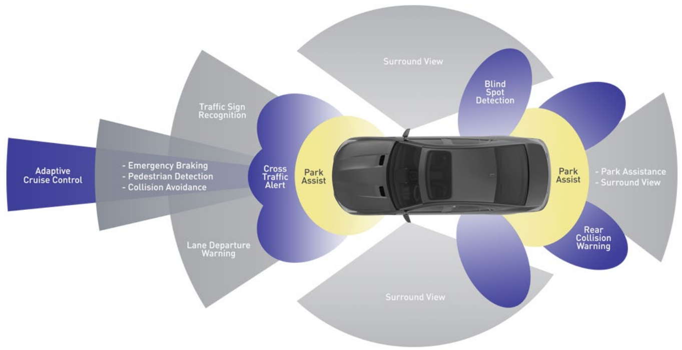

Figure 1. *The spectrum of sensor inputs that will contribute to the overall ADAS data flow.*

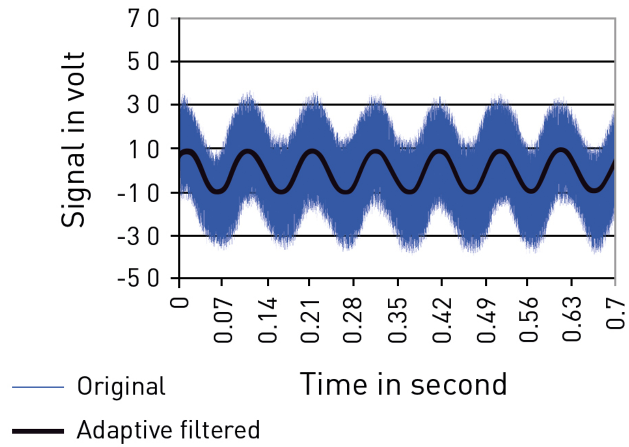

Figure 2. *Real world sensor data comprises signal plus noise, and processing is required to extract the wanted information.*

In the past, it was common to process only individual measurements from specific actuators (steering angle, pedal positions, various engine sensors, etc.) in real time. As is common with physical measurements, the environmental data acquired for ADAS applications are subject to noise (Figure 2) and measurement errors. Therefore, they require electronic post-processing by hardware and software before they can be used for their ultimate purpose, which is to automatically offload decisions from the driver.

Quite frequently, data from different sources are consolidated (sensor fusion) for reduced error susceptibility. In order to automatically make decisions on behalf of the driver, a tremendous amount of data must be processed in real time. For instance, camera images (approx. 340 kbps) or radar data (approx. 1.5 Mbps) are often provided as floating-point numbers (floats/doubles). Traditionally, only isolated sensor data (involving only some integer or fixed-point numbers with 1 - 5 kbps) needed to be processed. Therefore, it is immediately apparent that ADAS applications require a lot more processing power than traditional applications.

Currently, it is very hard to predict which high-performance hardware architectures will prevail for these kinds of applications. However, it is clear that compiler support will be required for all architectures because ADAS applications must be produced in a reusable and cost-efficient manner. This mandates the use of abstract, portable design methodologies (e.g. C++11/14), model-based design and additional technologies including parallel programming (e.g. OpenCL, Pthreads). Furthermore, highly optimized, certified libraries will be required to implement standard operations efficiently, safely and with maximum hardware independence.

As ADAS applications intervene with the driving process, these applications and the hardware used to execute them must adhere to relevant safety standards (ASIL-B or higher; ISO 26262).

Finding a suitable hardware architecture

There is a risk for companies developing ADAS applications, associated with the fact that no specific hardware architecture has prevailed until now. In general, major hardware accelerators including the NVIDIA GPU derivatives (Drive PX) provide adequate computational power in the Teraflops range for the data-parallel parts of ADAS applications. However, apart from lacking sufficient safety features, these devices are rather cost-intensive regarding their power consumption and purchasing price. Typical architectures for safety-critical applications up to ASIL-D (incl. AURIX or RH850) have not yet utilized some hardware based opportunities to achieve higher data rates because these will be hard to certify according to ASIL-D.

OEMs or large suppliers of ADAS systems are therefore in danger of selecting an architecture that may fail in the market because it is too large, too expensive or cannot meet the safety requirements. On the other hand, there is a risk to select an architecture that fully supports safety-critical applications but is too small for the more demanding computations. During the development process, it might turn out that the envisioned application cannot be implemented for efficiency reasons.

Thus, the requirements of ADAS projects are quite complex. On the one hand, it is mandatory to create very efficient, target specific code, to meet all safety goals and to minimize the risks outlined above. On the other hand, portable and high-level design methods are necessary to enable cost-effective application development. These high level design requirements mandate modifications of the embedded compilers that were originally designed for traditional embedded applications.

This brings us back to the initial question concerning which new capabilities an embedded compiler must provide in order to meet the above-mentioned requirements. A necessary, new compiler feature is the need to support the typical code structures of ADAS applications in order to create highly-efficient code for this kind of application.

This article continues with a discussion of the structure to be expected of the code in an ADAS application, and looks at compiler optimizations, especially as they relate to writing effective and dependable parallel code. Click for pdf

Download PDF of Article

Find Compilers for ADAS on EETsearch

FORMULA E SHOWCASES AUTOMOTIVE EFFICIENCY & PERFORMANCE GAINS THROUGH SIC

By Arun Kumar, Rohm Semiconductor

SiC, Silicon Carbide, which has for many years attracted attention due to its superior characteristics and performance over silicon, is being increasingly researched and adopted in a variety of applications in the industrial, automotive, railway, consumer, and other sectors in order to improve energy savings and efficiency.

Formula E, one of the world's first all-electronic motor racing series held by FIA, presented a prime opportunity for showcasing the performance advantages of SiC in EVs. With Season 3 recently underway [at time of writing] on October 9, 2016 in Hong Kong, the challenge was to find the most efficient method of transferring energy provided by the battery to the road. To accomplish this, Rohm partnered with Formula E team Venturi to incorporate SiC Schottky Barrier Diodes (SBDs) in the drive inverter (Figure 1) to enhance the performance of the powertrain system by significantly improving thermal efficiency and switching speeds. Season 4 will see the integration of SiC MOSFETs, which should result in further efficiency gains.

SiC SBDs deliver lower switching loss and forward voltage characteristics, reducing conduction losses for greater efficiency and performance. As a result they are often utilized in high-efficiency power source devices.

SiC SBDs featuring breakdown voltages from 600V and up are readily available. Compared to silicon FRDs (Fast Recovery Diodes), SiC SBDs provide much lower reverse recovery current and recovery time, which translates to significantly reduced recovery loss and noise emissions. Furthermore, unlike silicon FRDs, these characteristics do not change significantly over current or operating temperature, allowing system designers to improve efficiency, lower the size of heat sinks, and increase switching frequency to reduce the size and cost of magnetics.

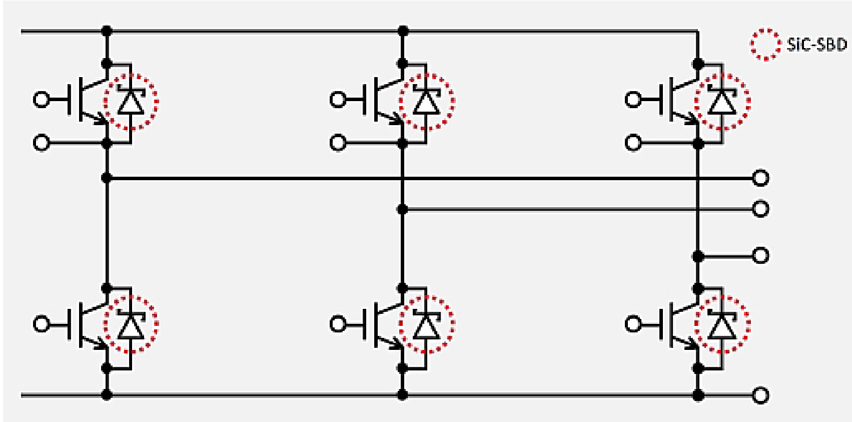

Figure 1. *Hybrid module circuit diagram utilizing Rohm SiC SBDs (for Season 3)*

SiC SBDs are increasingly applied to circuits such as PFC (Power Factor Correction) and the secondary side of bridge rectifiers in switch mode power supplies. Today's applications range from air conditioners, solar power inverters, and EV chargers to industrial equipment and beyond.

Lighter weight - SiC SBDs reduce inverter weight by 2 kg

Improved efficiency - Electric efficiency is increased by 1.7% while the volume of heat dissipation components reduces by 30%.

Beyond SiC MOSFETs, SiC modules and EDLC cell balancing ICs are ideal for future use in EVs.

The continuation of this article outlines these devices that will contribute to greater EV drive train efficiency. Click for pdf

Download PDF of Article

Find Automotive SiC components on EETsearch

designideas

- RC oscillator generates linear triangle wave

RC oscillator generates linear triangle wave

By Arturo Rivera

This Design Idea avoids the RC charging waveshape of a standard relaxation oscillator, replacing it with a linear rise/fall triangle wave. Positive feedback is used to increase the charging rate over each half-cycle and straighten out the curve.

The oscillator consists of Schmitt trigger comparator U1 and non-inverting adder U2. Oscillation is accomplished by the same principle that drives the relaxation oscillator, toggling the comparator output when the capacitor voltage reaches the threshold of the hysteresis. This hysteresis voltage, *HsV*, is defined by R1 & R2 in the comparator's positive feedback loop.

$$HsV = \pm V(U1_{Out}) * \frac{R2}{R1 + R2}$$

The hysteresis band sets the amplitude of the triangular waveform across C1.

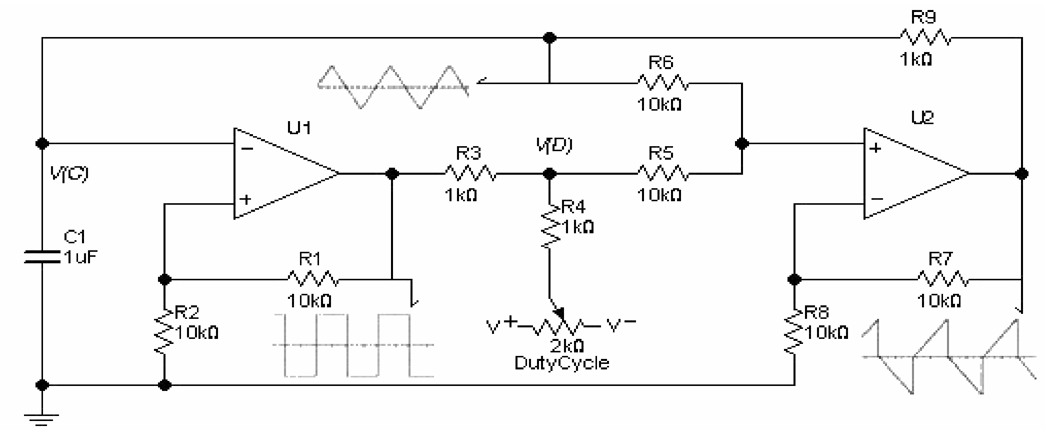

The comparator output is a rectangular wave whose amplitude is dependent on U1's output stage. This signal is reduced by R3 & R4 before the adder; otherwise, the triangular shape will revert to the RC curve.

$$V(D) = V(U1_{Out}) * \frac{R4}{R3 + R4}$$

Since *V(D)* must be stable, R3 & R4 are several times lower than the surrounding resistors; R3+R4 is the comparator's main load. Low values (here, 1 kΩ) are suggested, commensurate with acceptable U1 output loading, though all values can be scaled up.

The U2 adder sums *V(D)* and the capacitor voltage *V(C)* with a gain of two, defined by R7 & R8. Its output charges C1 through R9. With V(D) at ±½ $V(U1_{Out})$, the capacitor voltage will describe a linear-slope line, forming a triangular waveform.

The adder input resistors R5 & R6 divide both voltages by two. The Barkhausen stability criterion requires unity gain to perform oscillation, so the adder gain must recover this loss:

$$Vin = \frac{V(C)*R5}{(R5+R6)} + \frac{V(D)*R6}{R5+R6}\ ; \ \text{if R5=R6 then:}\ \ Vin = \frac{V(C)}{2} + \frac{V(D)}{2}$$

$$Vout = \left(\frac{V(C)*R5}{(R5+R6)} + \frac{V(D)*R6}{R5+R6}\right) * \left(\frac{R8+R7}{R8}\right)$$

If R5=R6 and R7=R8 then *Vout* = *V(C)* + *V(D)*

The addition process can be described as two time functions: one from T0 to T1, where *V(C)T0* = *-HsV* and *V(C)T1* = *+HsV* due to the positive value of *V(D)*, and the other from T1 to T2, which reverses the process: *V(C)T2* = *-HsV*, due to the negative value of *V(D)*. Each integral must equal zero because a DC offset is not expected.

$$V(C) = 0 = \int_{T0}^{T1} V(C) + V(D)dt \; + \int_{T1}^{T2} V(C) - V(D)dt$$

$$V(C) = \int V(D)dt \; V(C) = \pm V(D) * T$$

Solving the general form of the equation, where *V(C)* initial value is zero and *V(D)* toggles between positive and negative: The solution for *V(C)* is a linear function of *V(D)* and time. *V(D)* toggles between positive and negative when it reaches ±HsV. If *V(D)* increases, the frequency will increase too.

C1	Frequency
10µF	10-50 Hz
1µF	100-500 Hz
0.1µF	1-5 kHz
0.01µF	10-50 kHz

The op-amp's slew rate limits the frequency of this application. The comparator output must maintain a square shape, so the minimum period can be defined using the total excursion and a factor of 10:

$$\frac{SlewRate(V/\mu s)}{2*VOut*10} = Max\ Freq \quad e.g., \ \frac{13(V/\mu s)}{20V*10} = 0.065MHz\ (65kHz)$$

At this maximum frequency, the adder output current through R9 & C1 must be drivable by the op-amp. If necessary, calculate the RC impedance and C1 value – in this example, to meet 2 kΩ total: R9 = 1 kΩ, XC = 1 kΩ.

$$Xc = 1k\Omega = \frac{1}{2\pi fC}; \; C = \frac{1}{2\pi*65kHz*1k\Omega} = 2.44\text{nF}$$

Output frequency is equal to:

$$f = \left(\frac{V(D)}{Vout}\right) * \frac{1}{2\tau}$$

$$f = \frac{1k\Omega}{1k\Omega + 1k\Omega} * \frac{1}{2*1k\Omega*1\mu F} = 250\text{ Hz}$$

Duty-cycle adjustment is made by adding a DC component to *V(D)* through R4. The usable adjustment range is about 10%-90%.

$$DutyCycle\ (\%) = Vout \pm \frac{DCAdjustV}{Vout}$$

***Arturo Rivera** is a professor at Universidad Nueva Esparta and owner of a company manufacturing DC-DC converters.*

productroundup

productroundup

Gigabit Ethernet FPGA module with TCP/IP offload engine

Orange Tree Technologies (Oxford, UK) has announced the ZestET2-J, an FPGA module with a very high performance TCP/IP Offload Engine (TOE) chip for Gigabit Ethernet interfacing, providing access to very fast data rates without having to integrate complex networking hardware and software. The module can be used as a programmable interface to external devices, for high speed processing of streaming data, and for data acquisition and control, in markets including industrial vision, radar, sonar and medical imaging.

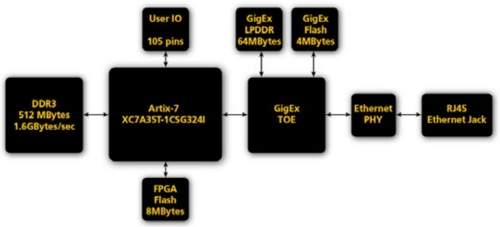

Complete article, here

Highly integrated power quality monitoring AFE

Analog Devices has released a polyphase analogue front end (AFE), with power quality analysis according to EN61000-4-30 Class S, designed to help extend the health and life of industrial equipment while saving developers significant time and cost over custom solutions. The ADE9000 AFE is an off-the-shelf solution with embedded power quality algorithms and integrated high-performance analogue-to-digital converters (ADCs) in one single device, to measure and analyze power pollution to prevent disruptions and long-term damage, primarily for three-phase power meters.

Complete article, here

Maxim's 1-Wire interface meets Arduino and mbed

Maxim Integrated has posted details of its reference design, MAXREFDES132#, which provides a platform for interfacing with the company's 1-Wire devices, 1-Wire Evaluation Kits (EV kits), and iButton devices using the DS2484 I²C-to-1-Wire master, or a bit-bang master on D2 of the Arduino form-factor pinout. Use the design with 1-Wire networks, data logging, and for rapid prototyping. Maxim says that its 1-Wire bus "continues to flourish in the era of IoT". This unique protocol provides communication and power along a single wire, at relatively long distances.

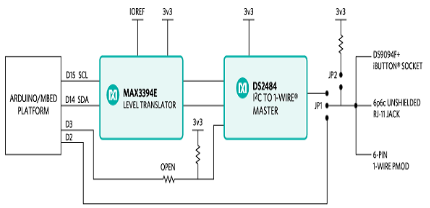

Complete article, here

'Heterogeneous inside & out' cores meet many-core needs

Imagination Technologies has added the MIPS Warrior I-class I6500 CPU to its IP offering; a multi-threaded, multi-core, multi-cluster design that promises increased levels of system efficiency and scalable computing for many-core heterogeneous designs. The I6500 provides a highly scalable solution which can coherently implement optimized configurations of CPU cores within a cluster – this is what Imagination tags 'Heterogeneous Inside' – as well as a variety of configurations of CPU clusters and GPU or accelerator clusters on a chip depending on the requirements of the system ('Heterogeneous Outside').

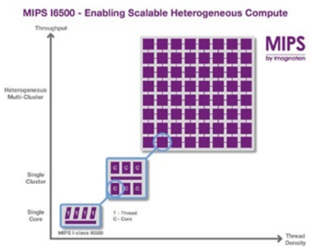

Complete article, here

productroundup

433-MHz band antenna offers high gain and simple integration

Antenova (Hatfield, UK), maker of antennas and RF antenna modules for M2M and IoT, has added Inca (part no SRFI028), to its flexiiANT FPC range. This antenna is for small devices in the 433MHz ISM band. The antenna measures 101.0 x 20.0 x 0.15 mm, and weighs just 0.5g, with a peak gain of 2.80 dBi. It is suitable for any small electronic device in the 433 MHz band, typically remote monitoring, robot control, smart meters, home automation and medical devices. It can also be used with low power wireless modules, as the 433 MHz band is increasingly being chosen for IoT applications that exploit the longer range that can be achieved with this frequency.

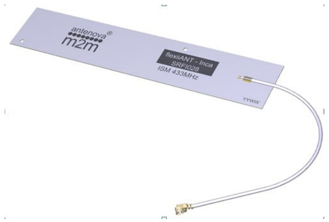

Complete article, here

Harmonise in-vehicle display panels with 14-ch gamma buffer

As vehicles are increasingly fitted with multiple displays, for both infotainment and for driver assistance (ADAS), the problem arises of inconsistent viewing experiences across those panels.

Intersil is addressing the problem with its ISL76534 14-channel gamma buffers IC, which it claims as the lowest-power part offering such functionality. The automotive-grade part ensures consistent brightness and colour matching across all LCD panels used in a vehicle. The ISL76534 is suitable for next generation LCDs designed for infotainment displays, ADAS, smart mirrors and instrument cluster displays.

Complete article, here

Intel samples Stratix 10 FPGAs on 14nm node

Intel has announced it is now sampling the Stratix 10 FPGAs, moving the FPGA fabric to its 14nm tri-gate process technology and combining it with a new architecture called HyperFlex. The new part sampling now claims 2X the core performance and over 5X the density compared to the previous generation and up to 70% lower power than Stratix V FPGAs for equivalent performance. The part delivers up to 10 TFLOPS of single-precision floating point DSP performance and up to 1TBps memory bandwidth with integrated High-Bandwidth Memory (HBM2) in-package. It embeds a quad-core 64-bit ARM Cortex-A53 processor.

Complete article, here

First 1000V silicon carbide MOSFET

Wolfspeed has introduced a 1000V MOSFET that enables a reduction in overall system cost, while improving system efficiency and decreasing system size, in applications such as efficient electric vehicle fast battery chargers. The company says that the FET offers system designers ultra-fast switching speeds with a fraction of a silicon MOSFET's switching losses. Designers can reduce component count by moving from silicon-based, three-level topologies to simpler two-level topologies made possible by the 1000 Vds rating of the SiC MOSFET. The increase in output power in a reduced footprint is realized by the ultra-low output capacitance, as low as 60 pF.

Complete article, here

productroundup

Authentication protects systems and securely controls IoT nodes

Maxim’s DeepCover Secure Authenticator, DS28C36, provides public-key and secret-key crypto functions that support new levels of embedded security protection; developers of industrial, medical, and IoT products now have an added level of IP and device integrity protection. Helping to integrate strong defences into products and systems, the DS28C36 provides a targeted set of asymmetric-key and symmetric-key cryptographic tools in a low cost and compact solution. Asymmetric public-key features are supported with the P256 based elliptic curve cryptography (ECC) algorithm and symmetric secret-key with SHA-256. The DS28C36 provides two authenticated GPIO pins with optional secure state control and level sensing as well as a function to enhance secure download/boot features. The device also integrates sophisticated countermeasures to protect against invasive and noninvasive security attacks.

Complete article, here

Zero-code-programmable mixed-signal IC in 4.4 mm²

Silego Technology has packaged its GreenPAK Configurable Mixed-signal IC into an ultra-small 2.0 x 2.2 mm 14-pin outline. Silego's SLG46169V and SLG46170V GreenPAK programmable mixed-signal matrix ICs allow engineers to design custom mixed-signal circuits using integrated, configurable digital logic and timing resources and, in the case of the SLG46169V, two analogue comparators. These devices are suited for integrating common functions such as system reset, power sequencing, glue logic, LED control and more, all with zero code.

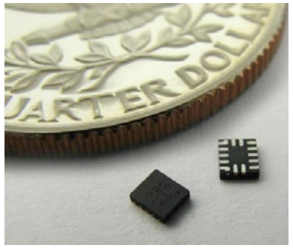

Complete article, here

24-pin MCUs with IOs to support up to 13 sensors

Renesas' RL78/G11 Group of microcontrollers is a group of small-package, low-power MCUs that support low-power sensor hubs and sensor application systems. The RL78/G11 MCUs support the largest number of sensors in an MCU with 24 pins: 11 analogue input channels and six serial communications channels. With the 24-pin compact package, the MCUs can support signal inputs from a maximum of 13 sensors. Processing of the signal inputs from the sensors by the built-in analogue functions and serial communications (reception) functions that operate even in standby mode allows low-power control, even in systems with multiple sensors.

Complete article, here

QM series; TDK’s revisit of modular PSU concept

TDK Lambda has added a new modular, configurable power supply series with the 1200W - 1500W modular QM range. It has been designed specifically to offer the lowest acoustic noise and to provide – for medical applications – comprehensive MoPP isolation. Product definition is made using an online configurator, giving an optimised module selection with a choice of signals, leakage current and standby voltages. With a wide range 90-264Vac, 47-440 Hz input the QM can deliver 1200W; and 1500W with a 150-264Vac input. Up to 16 outputs can be provided, with voltages ranging from 2.8V to 52.8V.

Complete article, here

productroundup

100 MHz to 40 GHz power detector replaces Schottkys

With 1 dB accuracy & 35 dB dynamic range (±1 dB), LTC5596 is a high frequency, wideband and high dynamic range RMS power detector that provides accurate, true power measurement of RF and microwave signals independent of modulation and waveforms. The LTC5596 responds in a log-linear 29 mV/dB scale to signal levels from –37 dBm to –2 dBm, at accuracy better than ±1 dB error over the full operating temperature range and RF frequency range, from 200 MHz to 30 GHz. In addition, the device's response has ±1 dB flatness within this frequency range.

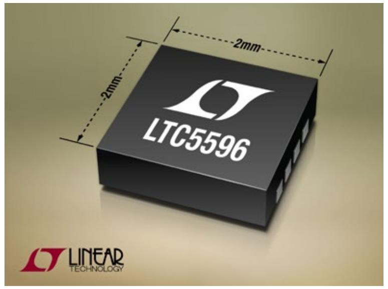

Complete article, here

Photorelays rated to 5A replace industrial E-M relays

Toshiba Electronics has added three DIP8 package photorelays that can replace mechanical relays in industrial applications. With 5A large drive current they can be used in applications ranging from heating, ventilation and air conditioning (HVAC) systems to inverters, PLCs and automated test equipment (ATE). A guaranteed pulsed ON-state current that is three times greater than that of the continuous ON-state current secures a margin for safety design. TLP3547 is a 60V product with 5A (max.) large drive current. The TLP3548 is a 400V product with a 0.4A (max) drive current that offers high-speed switching of 1 msec (max).

Complete article, here

AC/DC PSU outputs 40W from 4.5 sq in footprint

XP Power's ECF40 series of single-output 40W ultra compact AC-DC power supplies are open frame convection-cooled units capable of delivering their full output power without the need for any external forced air flow up to +50C. XP claims the units as the smallest 40W open frame power supply available, at 76.2 x 38.1 x 28.0 mm. XP intends the unit for a range of applications in the healthcare, industrial and technology sectors, rating it for an operating temperature range of -40 to +70C. The range comprises six models providing all the popular nominal output voltages from +12 to +48 VDC.

Complete article, here

RS-485 transmitters support 20 Mbps, draw under 0.6 mW

Exar's high-speed RS-485/RS-422 transmitters, XR33193, XR33194 and XR33195, operate from a 3.3V supply and draw only 180 μA (375 μA maximum) of supply current. All transmitters feature a shutdown mode that consumes less than 2 μA whenever the transmitters are disabled. The family is available in 3x3 mm 6-pin TSOT23 packages, ideal for high-speed point-to-point RS-485 applications where space is a concern. The XR33193/94/95 offer 250 kbps, 2.5 Mbps and 20 Mbps data rates; XR33193 and XR33194 have slew-limited outputs; XR33195 features 25 nsec propagation delay and 5 nsec driver-output skew.

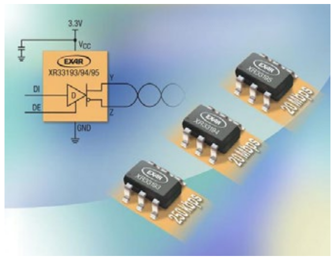

Complete article, here

LEN SHERMAN, MAXIM INTEGRATED

Inspiration from Chinatown

In the early 1990s, battery power products were taking off. We were getting lots of requests for power supply circuits that operated efficiently from AA batteries. In response to several customer inquiries, one of our engineers decided to run a contest within the company to see who could come up with the circuit that would wring the most run time from four AA alkaline batteries while generating 5V at 100 mA.

We ended up with about six circuits to test, but quickly realized that it would take a long time to test them in a manner that wouldn't hoard too much lab equipment or lab technician time, particularly if we wanted to run several tests to reduce the influence of battery variations. We did have an automated test rack at the time, but that was always busy doing "serious" work, and wasn't available for the amount of time we would need. We also didn't want to run the tests in series, one at a time, because that meant it would be weeks before we'd be able crown the winner. It would be much better (and more fun) if we could have a real race with all the candidates starting at the same time, a sort of "AA Marathon." But how could we do that without tying up too much lab equipment or making mind-numbing manual measurements?

So, what does all this have to do with Chinatown, the 1970s film noir private eye movie with Jack Nicholson? There is a brief scene where Nicholson is tailing a car. The target vehicle stops, the driver gets out, and enters a house. Jack doesn't want to wait (possibly all night) to see when the driver leaves, so he grabs an old wristwatch (out of a box of old watches he apparently keeps in his car's trunk just for this purpose), winds it up, sets it, and places it on the road, just under the tyre of the parked "target" car. The next morning Jack returns. The car has left the scene, but it's departure time is clearly preserved on his crushed watch, laying by the kerb.

That clever low-tech sort of "Lo-Jack/GPS/event timer" stuck with me and inspired a cheap and simple battery-monitoring alternative to data-acquisition systems, chart recorders, or even (at that time) boring meter reading and note taking. It also had added benefits in that it needed no programming skill, supplied an easily readable result, and even featured "nonvolatile" data retention!

All that was needed were low-power comparator/reference ICs (a single chip IC, even in 1993) to gate power to each of the six cheap $8 drug-store battery-powered analogue clocks. True to Chinatown, these had to be analogue clocks (with hands), not digital ones, since digital ones (unless they were mechanical) obviously would not retain their "hand position" once power was cut.

We wired up the comparators to open the clock's battery connection when the circuit-under-test's output voltage dropped under a set threshold. We stuck these on the output of each circuit in the race, put on a fixed resistor load (for 100 mA at 5V), and started them all. As each power supply output dropped below 4.5V (we could set this trip point with a pot), its clock would stop. The hands, from that point on, displayed the time when the circuit gave up. The clocks were all started at 12:00, so the hands showed running time in hours and minutes (within twelve hours of course).

This was a great triumph of "low tech." Just $60 worth of cheap clocks and a few ICs became an effective, and very cheap, substitute for a rack of HPIB hardware.

***Len Sherman** is a Senior Scientist at Maxim Integrated where he has defined dozens of power and battery management ICs, primarily for portable and handheld devices.*

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

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

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

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

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

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

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

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

USA & Canada

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

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

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

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

Asia

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

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

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