EDN europe

FEBRUARY 2016
www.edn-europe.com

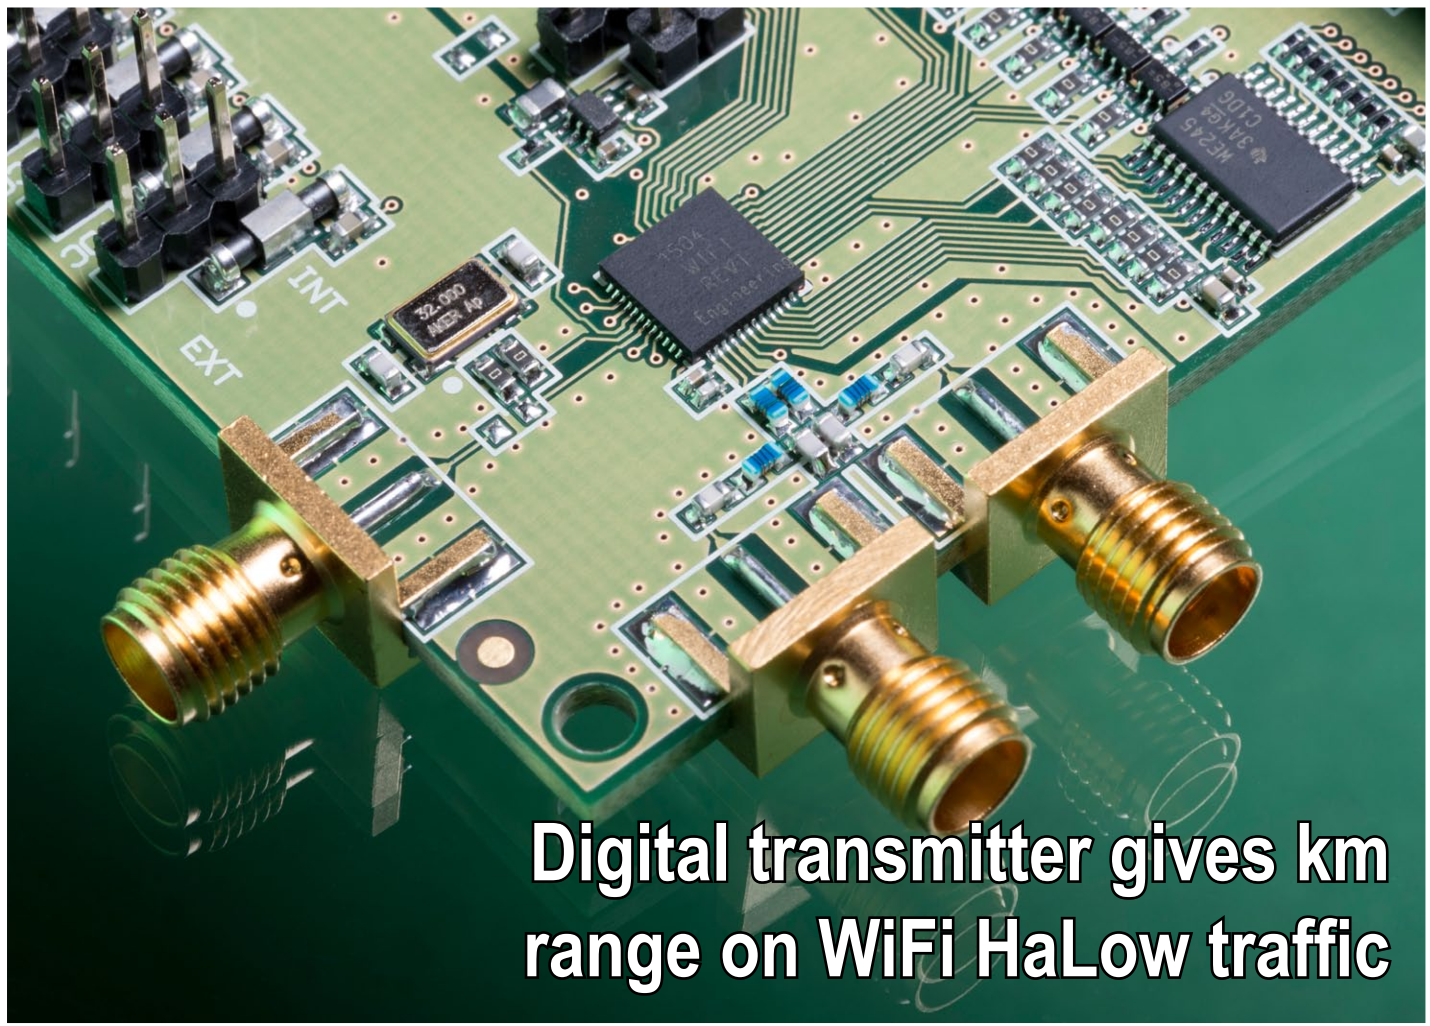

Digital transmitter gives km range on WiFi HaLow traffic

Best of Class:
Parts in Stock Available for Immediate Shipment
Source: Distributor Evaluation Study, UBM Tech, May 2015
51%
16%
5%
4%
3%
1%
1%
DIGI-KEY
DISTRIBUTOR B
DISTRIBUTOR C
DISTRIBUTOR D
DISTRIBUTOR E
DISTRIBUTOR F
DISTRIBUTOR G
1,250,000+ PRODUCTS IN STOCK
FREE SHIPPING
ON ORDERS OVER €65!*
4.75 MILLION PARTS ONLINE
OVER 650 INDUSTRY-LEADING SUPPLIERS
100% AUTHORIZED DISTRIBUTOR
Digi-Key ELECTRONICS
FIND CONTACT AND ORDERING INFORMATION FOR YOUR REGION AT
DIGIKEY.COM/EUROPE
*A shipping charge of €18.00 (£12.00) will be billed on all orders of less than €65.00 (£50.00). All orders are shipped via UPS for delivery within 1-3 days (dependent on final destination). No handling fees. All prices are in euro and British pound sterling. If excessive weight or unique circumstances require deviation from this charge, customers will be contacted prior to shipping order. Digi-Key is an authorized distributor for all supplier partners. New product added daily. © 2016 Digi-Key Electronics, 701 Brooks Ave. South, Thief River Falls, MN 56701, USA
ecia MEMBER
ecsn Electronic Components Supply Network member
CEDA

COVER

Low-power, long-range WiFi; imec shows all-digital HaLow transmitter

Nanotech researchers Imec (Leuven, Belgium) and the associated Holst Centre have implemented a Wi-Fi HaLow low-power fully digital polar transmitter for IoT applications. Wi-Fi HaLow is the designation created by the Wi-Fi Alliance for products incorporating IEEE 802.11ah technology. Wi-Fi HaLow operates in frequency bands below 1 GHz, and is configured to provide long range at low power for modest data rate traffic, and connectivity according to "Wi-Fi Certified" standards. It is intended for uses in areas such as smart home, connected car, and digital healthcare, as well as industrial, retail, agriculture, and smart city. Imec/Holst's presentation describes a transmitter with a ten-fold power reduction as compared to state-of-the-art OFDM transmitters; it is a 1.3nJ/b fully digital polar transmitter optimised for IoT applications and the IEEE 802.11ah Wi-Fi protocol, Compared to other IoT standards, says Imec, the IEEE 802.11ah standard's sub-GHz carrier frequency and mandatory modes with 1 MHz/2 MHz channel bandwidths allow devices to operate in a longer range with scalable data rates from 150 kb/s to 2.1 Mb/s. The standard uses OFDM to improve the link robustness against fading, which is important in urban environments, and to achieve a high spectral efficiency (data rate over a given bandwidth). Full Story.

FEATUREARTICLES

ONLINE THIS MONTH

EDN's columns

TAAS, THE NEW AUTOMOTIVE REALITY

A few weeks ago, our friends at distributor Premier Farnell sent me a summary of a survey it had conducted about attitudes to self-driving cars, and other aspects of automotive technology development. The headline read, "Global research shows drivers want improved safety and efficiency, not self-driving cars" - which didn't entirely reflect the picture painted by the figures in the study; *"Improving efficiency and passenger safety should be the top areas addressed by engineering and technology developments in the automotive sector, according to global consumer research from element14. In contrast, areas that are seeing significant funding and development, such as sensors, self-driving cars and electric vehicles, have considerably less consumer appeal... 56% of consumers feel technology in the automotive sector should be primarily focussed on efficiency and 53% said safety respectively. Only 7% agree the focus should be on driverless cars and 8% on sensors, for example for judging distances... The results are more positive when looking at overall enthusiasm for driverless cars, with 44% of survey respondents saying they are of interest, more than virtual reality (41%) or gesture control (43%)."* (You can read that research here.)

This might be somewhat missing the point. The outcome that today's drivers might not prioritise a shift to autonomous driving is not really a surprise; but to a certain extent (not quantified in any study I've come across) they may not be the target market for those vehicles when they do become available. The sector of whom those questions should be asked is the non-drivers; it is a documented fact that in many developed economies (even in the USA, according to research from 2014) the proportion of teenagers acquiring a driving licence is falling. Gaining that licence is no longer the automatic rite-of-passage experience it once was. This is (if you like) the Uber generation, the purchasers of TAAS – transport as a service. And its priorities may be different. Consider this quote from a newspaper column;

"...in 10 or 20 years the three biggest car makers in the world will not be Toyota, Volkswagen and General Motors. They will be Google, Apple and Uber. That's because they have come to understand something more important than what goes under the bonnet. They've come to understand that you don't need a car. But you do need one tomorrow morning, at about ten past eight. And on Thursday afternoon for a couple of hours. Think about it. This morning your car is parked outside your house doing nothing except costing money. You probably won't use it at all until tomorrow morning, when you'll be forced to drive through all the roadworks and all the jams so that you can leave it outside your office all day... Life will go on like this for months until one day it needs a service, or it breaks down, and that'll be a nuisance because then you'll have to make alternative travel arrangements, which will be a chore. You're paying thousands of pounds a year, then, for something that you use for — what? — 5% of the time? Two per cent?"

Who wrote that? An eco-warrior? In fact it is by a motoring journalist; Jeremy Clarkson writing in the Sunday Times (UK). You have to allow that Mr. Clarkson is writing for entertainment and effect but, nevertheless, when motoring journalists start to question the merits of "default" car ownership, something my have changed. Or be changing.

This matters, because it feeds directly back to decisions about where the focus of on-going research and development in automotive engineering should lie. Should we work towards a self-driving car that is like today's car, but with "autopilot", assuming some residual skills on the part of the operator? Or one that "contains no user-operable controls"? Do we continue to add sophistication in the infotainment and connected-services dimension? If the car becomes less of an aspirational purchase, perhaps that is less important: conversely, if the occupants are to have nothing else to do, perhaps it becomes of even more importance.

There is almost no area of automotive engineering, from basic energy source, to "V2X" communications, to entertainment bandwidth, that does not face dramatic changes, and on (for the auto industry) unprecedented short timescales. Having the profile of the buying market transform itself at the same time can only add to the fun.

pulse

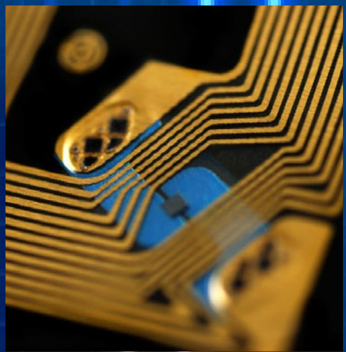

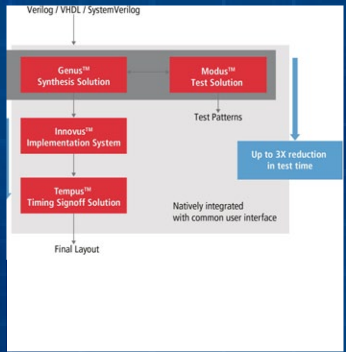
Verilog / VHDL / SystemVerilog
Genus™ Synthesis Solution
Modus™ Test Solution
Test Patterns
Innovus™ Implementation System
Up to 3X reduction in test time
Tempus™ Timing Signoff Solution
Natively integrated with common user interface
Final Layout

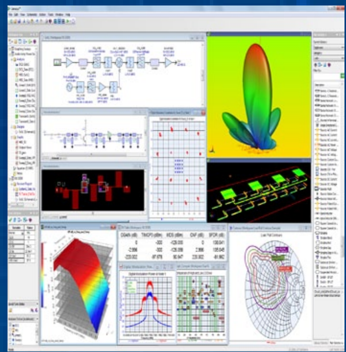

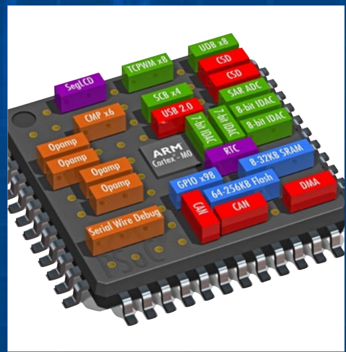
SegLCD
TCPWM x8
UDB x8
CSD
SCB x4
SAR ADC
USB 2.0
8-bit IDAC
CMP x6
Opamp
ARM Cortex-M0
RTC
8-32KB SRAM
GPIO x98
64-256KB Flash
DMA
CAN
Serial Wire Debug

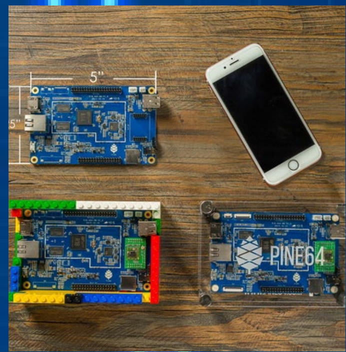
5"
PINE64

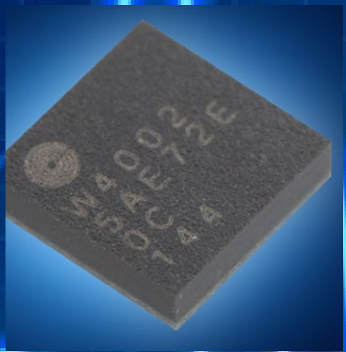

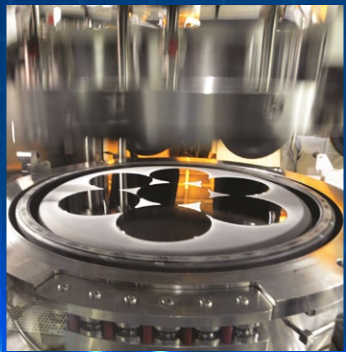

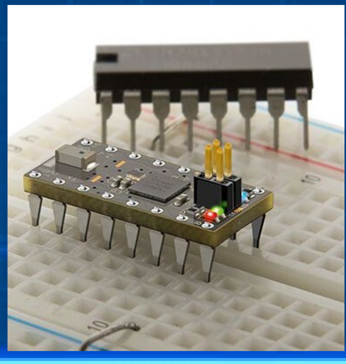

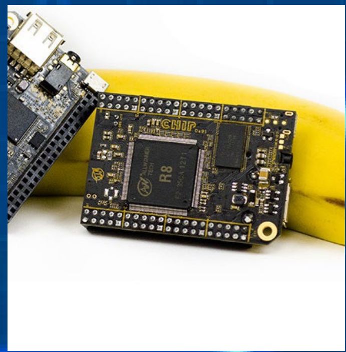
R8

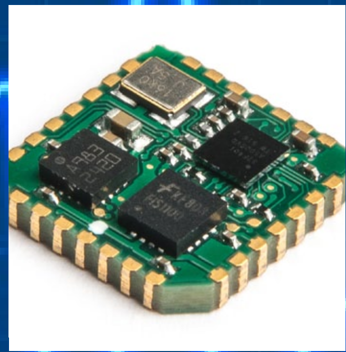

$9 single-board computer is set to make its debut

by Amy Norcross

What can you expect in a development board with a $9 price tag? Known as CHIP, which stands for "Computer Hardware in Products," the open source product features a 1-GHz R8 ARM processor, 512 MB of RAM, and 4 GB of onboard storage; has built-in Wi-Fi and Bluetooth connectivity; and is designed to work on any screen.

Smaller than a credit card, CHIP is the creation of Next Thing Co., which is based in Oakland, California, and was started by friends Dave Rauchwerk, Gustavo Huber, and Thomas Deckert. The group raised more than $2 million from over 39,500 backers through a May Kickstarter campaign, far exceeding its initial $50,000 goal.

"If you're wondering how CHIP could be this inexpensive, you can thank cheap Chinese tablets," said David Scheltema of *Make:*. "The system on chip (SoC) used in the development board is based on an A13 processor by Allwinner, a Shenzhen-based semiconductor company. As recently as 2013, Allwinner was the second-largest tablet manufacturer in the world, and the A13 was the most successful processor in Allwinner's lineup."

CHIP comes preinstalled with dozens of applications, tools, and games. The Chromium browser makes Web surfing possible. LibreOffice can be used to create word documents and presentations and also edit spreadsheets. Scratch is a free basic programming language that allows users to create interactive stories, games, and animations. In addition, CHIP can run thousands of free applications from the open source community. Several tutorials are available on topics that include powering on the device, flashing the board, and setting up Bluetooth.

The company is currently focused on fulfilling orders from the Kickstarter campaign, but the system is expected to be available for direct purchase at a later date; the venture's site says June 2016.

CHIP joins another Kickstarter offering from Next Thing — OTTO, a customisable "hackable GIF" camera powered by Raspberry Pi, was introduced in May 2014 and has a starting price of $199.

Complete article, here

oe-a
Messe München
Connecting Global Competence

Business meets technology

Join your community and benefit from
- **Networking—**Where manufacturers and end users come together.
- **Knowledge transfer—**Research that inspires product ideas.
- **Industry orientation—**State of the art technology that reveals market opportunities.

This is the place where printed electronics come to life.

Book your ticket now!
lopec.com/tickets

©: FlexEnable Ltd (top), Takao Someya Group, University of Tokyo (bottom left), Thinfilm, Norway (bottom right)

Exhibition: April 6–7, 2016
Conference: April 5–7, 2016

Messe München, Germany

LOPEC
International Exhibition and Conference for the Printed Electronics Industry

Adding a little FE memory makes RFID chips hack-proof

Researchers at MIT and Texas Instruments have developed a radio frequency identification (RFID) chip that is virtually impossible to hack, potentially defeating capture of data in scenarios such as theft of credit card number or key card information, or of high-value goods where inventory systems may be deceived by dummy tags.

The problem is tags that are vulnerable to side-channel attacks, which typically analyse patterns of memory access or fluctuations in power usage when a device is performing a cryptographic operation, in order to extract its cryptographic key.

One way to thwart side-channel attacks is to regularly change secret keys. Such a system would still, however, be vulnerable to a "power glitch" attack, in which the RFID chip's power would be repeatedly cut right before it changed its secret key. Power-glitch attacks have been used to circumvent limits on the number of incorrect password entries in password-protected devices, but RFID tags are particularly vulnerable to them, since they're charged by tag readers and have no on-board power supplies.

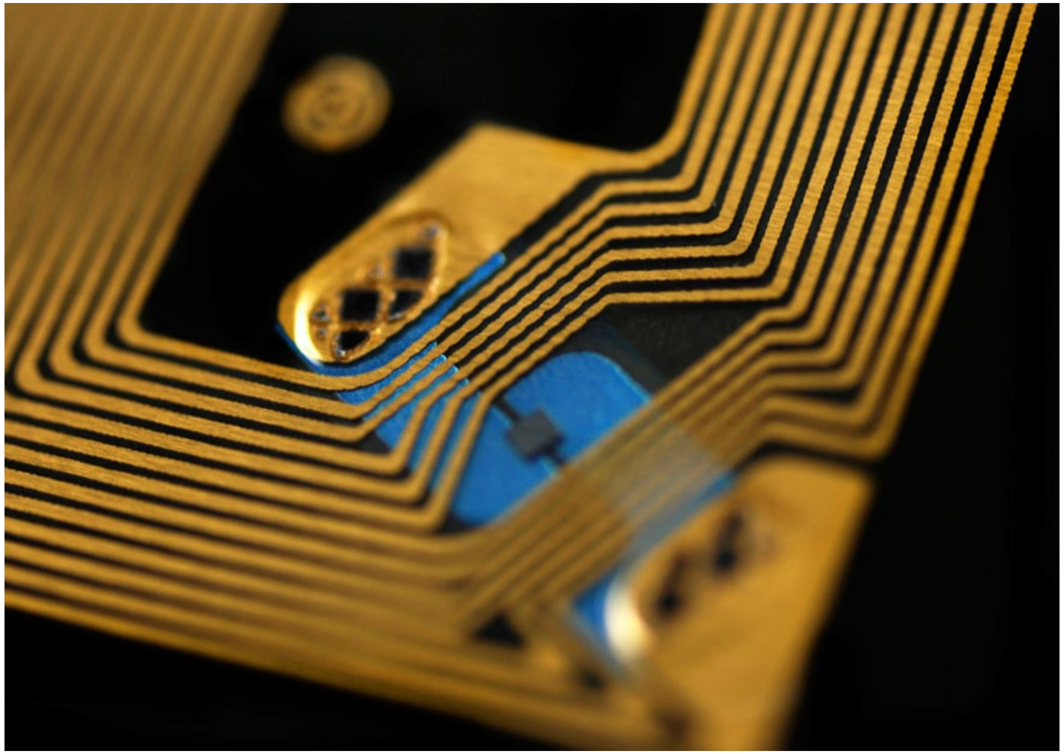

Two design innovations allow the MIT researchers' chip to thwart power-glitch attacks: One is an on-chip power supply whose connection to the chip circuitry would be virtually impossible to cut, and the other is a set of nonvolatile memory cells (NVM) that can store whatever data the chip is working on when it begins to lose power. Texas Instruments has built IC prototypes to the researchers' specifications, and for the NVM, used the ferroelectric memory technology (FRAM) that TI has in-house: the TI and MIT workers report the chips behave as expected.

Complete article, here

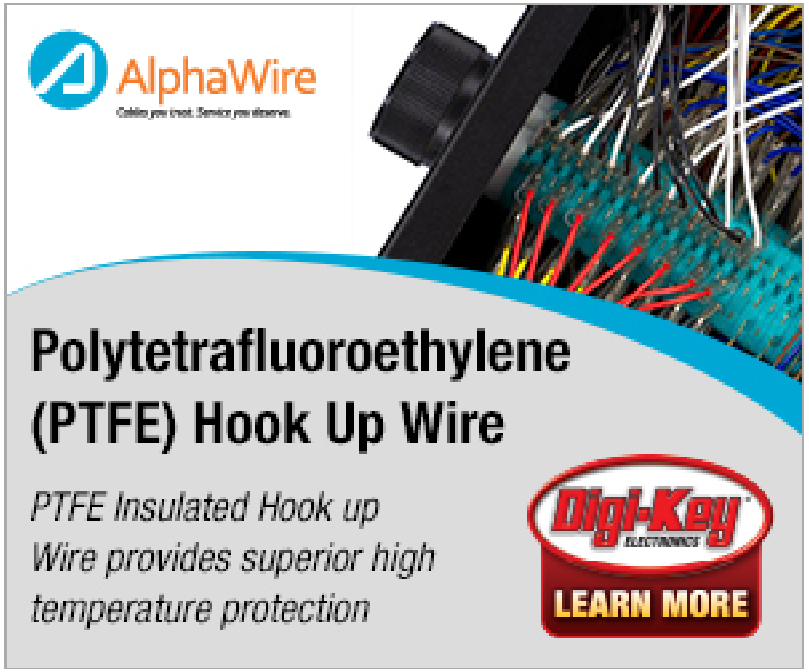

FPGA transceivers send 25 Gbps 5m over copper - Xilinx

The FPGA maker says that it has achieved a level of serial data performance that will, "bring greater cost efficiency to data centre interconnects," as its Virtex UltraScale devices are the "first FPGA to support 25 Gbps per lane copper cabling at five metres."

Xilinx's Virtex UltraScale devices have achieved compliance to the 25GE, 50GE and 100GE copper cable and backplane IEEE and related specifications which supports up to five metres of copper cabling in the data centre and up to one metre of backplane interconnect. These specifications include the IEEE 802.3bj 100GBASE-CR4/KR4, IEEE 802.3by 25GBASE-CR/CR-S/KR/KR-S, 25 Gigabit Ethernet Consortium 50GBASE-CR2/KR2. Data centre users can now make use of nx25G lanes of copper cabling versus optics for more cost- and power-optimised solutions to connect servers to top-of-rack switches using any off-the-shelf, specification-compliant vendor.

Along with integrated 100G Ethernet MAC IP, soft error-correction (RS-FEC) IP, and ASIC-class logic fabric, Virtex UltraScale FPGAs provide a complete, high performance and low-latency Ethernet solution for data centre workload acceleration. More on this story here.

16nm Finfet now shipping

Xilinx has also announced that it has made a "first customer shipment" of the Virtex UltraScale+ FPGA, which it describes as the first high-end FinFET FPGA built using TSMC's 16FF+ process. Ultrascale+ is the architecture behind Xilinx' latest generation of high-end FPGAs; the company says it has over 100 "engagements" on the technology stage; and has shipped devices and/or boards to over sixty of these customers.

Together with the Zynq UltraScale+ MPSoCs and Kintex UltraScale+ FPGAs, Xilinx says it now has availability of the three families in its 16nm portfolio. Virtex UltraScale+ features capabilities including 32G transceivers, PCIe Gen 4 integrated cores, and UltraRAM on-chip memory technology; Xilinx configured the parts for duties in areas such as next generation data centre, 400G and terabit wired communications, test and measurement, and aerospace and defence.

Complete article, here

SoC test software cuts test times, eases DFT burden

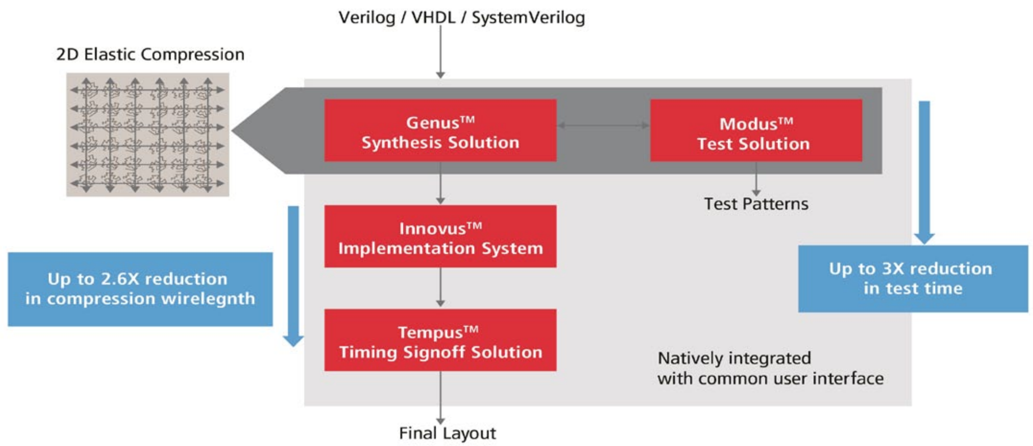

Cadence has released its Modus test package that enables a reduction of up to three times, in system-on-silicon-scale IC test times. It incorporates technology termed physically-aware 2D Elastic Compression architecture that can reduce test logic wirelength by up to 2.6-times, and enable compression ratios to scale beyond 400X without impacting design size.

SoC test engineers, the EDA vendor says, have become used

to an upper limit on test compression - the factor by which (in broad terms) test patterns prepared for production can be reduced compared to a one-at-a-time accessing of all of the scan chains formed when the target device is placed in test mode. This factor has been accepted at [around] 100; now, Cadence says its latest code can extent that to 400, shortening text times. Test compression and decompression has conventionally been a combinational process; Modus adds [some] registers and a degree of combinational logic. Although this extends the data injected into each scan chain, a higher level of compression is enabled, with overall benefit.
The software also brings a new approach to adding the wiring (which must be inserted into the chip's metallisation layers) that is needed to give the scan chain access in testing. Modus takes a more chip-wide view of the logic when creating that interconnect; and it uses a grid-based approach to laying those wired down. The result is a less congested, more manufacturable physical layout that has less impact on the DFT (design for text) insertion steps. Overall the impact of DFT at synthesis stage is reduced; chip time on the tester is cut; and/or the added test-structure wire length is reduced – all without impacting test coverage.

Complete article, here

Spice up your Pi for IoT development

by Richard Quinnell

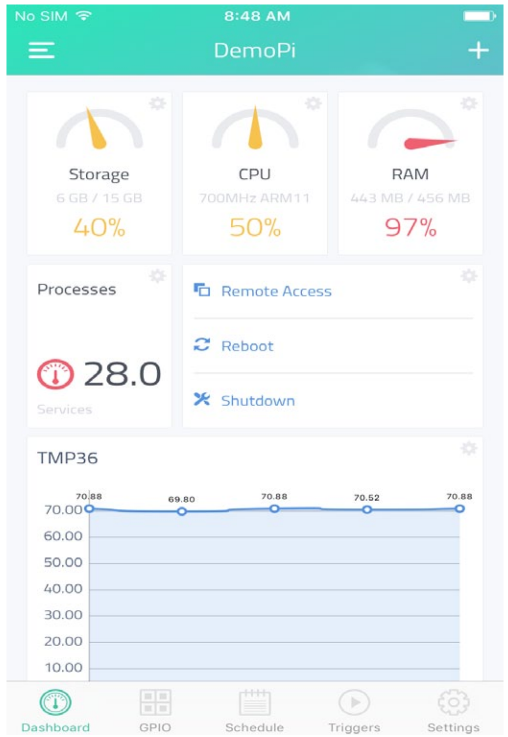

Developers seeking to enter the IoT market quickly discover that creating a nifty device is not enough. No matter how clever, the device is only one part of a larger system that includes such things as gateways, servers, analytics, web services, and even mobile apps. Something new has entered the market, which claims that it can solve those challenges for you, especially if your device design is based on the Raspberry Pi.
Platforms that IoT developers can leverage to support their device designs are not new. Companies like ThingWorx, Ayla Networks, Google, and Apple along with consortia like the Allseen Alliance are offering platforms that they hope will attract users and build an ecosystem of interoperable devices around their platform as a common standard. But nearly all of the platforms available require the IoT device be custom designed and programmed to operate with that platform.
A company called myDevices offers an alternative. Its platform, released in late 2015, is device-agnostic, supporting a wide variety of wireless technologies as well as standard data communications protocols such as CoAP, MQTT, and the REST API. The platform will adapt to the device rather than forcing the device to conform to it. It will also let devices using the platform talk with one another, serving as a translator.
myDevices has just released Cayenne, a tool for developers and makers to rapidly configure a complete IoT system, including mobile apps, rules engines, and analytics, starting with a widely-available development board – the Raspberry Pi. And you don't even need to work at configuring the Pi. Cayenne is designed to au-

tomatically discover any Pi connected to the same network as the host computer, then download its agents into the Pi. When the Pi reboots, it is ready to serve as the basis of an IoT device.
Cayenne provides developers with a host of capabilities for turning that Pi into a fully functioning IoT system node. Cayenne will automatically detect any sensors connected to the Pi and make their data available on a drag-and-drop configurable data display dashboard in a choice of numeric and graphical formats. Another dashboard gives developers full access to the Pi's GPIO resources for rapid configuration and sensing. Triggers allow events on one to invoke action on another device, or send messages via SMS or email. The tool also automatically configures a mobile or web-based app to replicate the dashboard.

Complete article, here

Design & simulation software speeds realisation of RF circuits

The latest release of Keysight Technology's RF simulation and synthesis software, Genesys 2015, features Keysight Sys-Parameters, that enables RF system simulation with off-the-shelf component datasheets; and comprehensive RF circuit synthesis to enable the fastest realisation of RF systems and circuits.

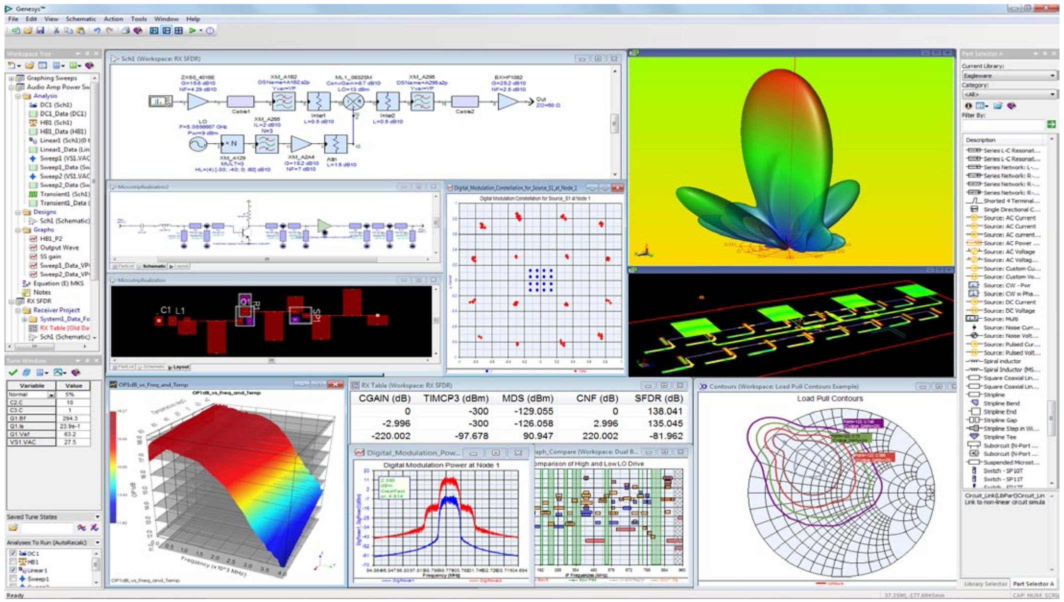

Keysight Sys-Parameters define RF system component parameters such as amplifier P1dB, IP3, gain and noise-figure over frequency, temperature and bias in a convenient spreadsheet format that designers can use directly in RF system simulation. This eliminates the tedious creation of non-standard data files or writing of custom programs to interpolate the data; addressing a long unfulfilled need of the RF industry to use component datasheet specifications in fast design simulations instead of wasteful prototyping iterations.
Genesys 2015 accepts multi-dimensional Sys-Parameter data in a regular spreadsheet format that a user creates from RF component datasheets. The software also provides Sys-Parameter libraries from Mini-Circuits and Analog Devices. Easy access to Sys-Parameters during design enables designers to accurately evaluate and select available RF system components before hardware implementation, thereby eliminating wasteful iterations.
Genesys 2015 includes tutorial videos for its comprehensive suite of automatic RF circuit synthesis capabilities. The tutorials are designed to teach users how to quickly design and realise circuits such as filters, matching networks, oscillators, mixers, couplers and transmission lines. Genesys 2015 also features a Matlab script debugger, multi-dimensional interpolator and interactive 3D graphing.

Complete article, here

Cortex-M0, programmable analogue & logic functions, in single-chip PSoC part

Cypress Semiconductor has extended its PSoC 4 programmable system-on-chip architecture; the PSoC 4 L-Series integrated single-chip solution has an ARM-Cortex-M0 core, up to 256 kB flash memory, 98 general purpose I/Os, 33 programmable analogue and digital blocks, a USB device controller, and a control area network (CAN) interface.

The PSoC 4 L-Series is intended for a broad range of industrial and consumer applications, applying the PSoC architecture to address multiple product variations, and Cypress's CapSense capacitive touch-sensing technology to implement reliable and elegant user interfaces.

Cypress adds, "The PSoC 4 L-Series introduces new capabilities such as dual-mutual CapSense blocks with up to 94 channels for large, capacitive-touch home appliance applications and USB and programmable digital blocks to create bit-perfect digital audio solutions. Additionally, it provides all the resources needed to create new products for the emerging USB Type-C market."

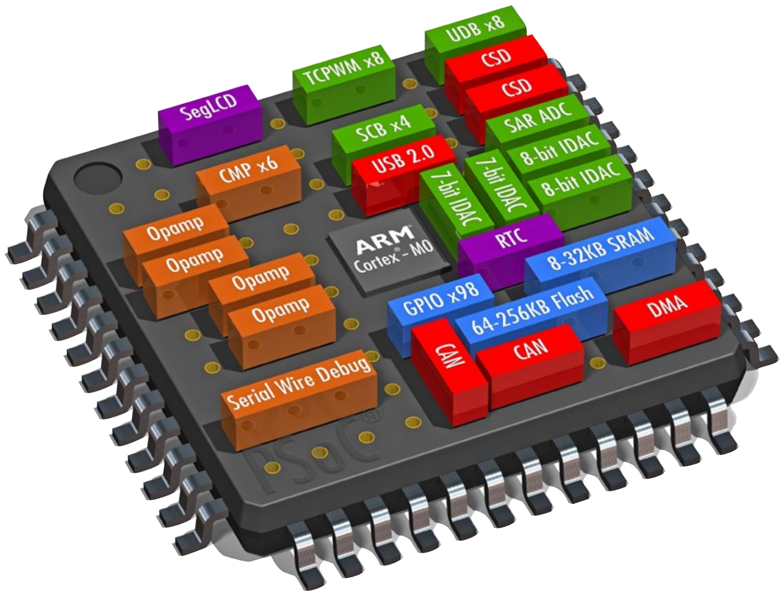

The PSoC 4 L-Series has up to 13 programmable analogue blocks including 4 high-performance opamps, 4 current-output digital-to-analogue converters (IDACs), 2 low-power comparators, a 12-bit SAR ADC and dual CapSense blocks with up to 94 capacitive-sensing channels. The programmable analogue blocks enable engineers to create on-chip, custom analogue front ends to support new end-product features, without increasing product costs, size or power consumption.

The devices host up to 20 programmable digital blocks including 8 timer/counter/PWM blocks, 4 serial communication blocks and 8 Universal Digital Blocks (UDBs) – these are programmable digital blocks that each contain two programmable logic devices, a programmable data path and status and control registers. UDBs can be configured as coprocessors to offload compute-intensive tasks from the ARM Cortex-M0 core. The blocks also enable engineers to implement custom digital peripherals, state machines or glue logic.

Complete article, here

Everything FPGA. ENCLUSTRA FPGA SOLUTIONS

LEDs in cubic GaN to overcome "green gap"

Plessey Semiconductors (Plymouth, UK), Anvil Semiconductors and the University of Cambridge have announced that they are working together to fabricate high efficiency LEDs in cubic GaN grown on Anvil's 3C-SiC-on-Si substrates, as part of a drive for more efficient LED light sources.
Cubic GaN has the potential to overcome the problems caused in conventional LEDs by the strong internal electric fields which impair carrier recombination and contribute to efficiency droop. This is particularly true for green LEDs where the internal electric fields are stronger and are believed to cause a rapid reduction in efficiency at green wavelengths known as "the green gap". The availability of cubic GaN from a readily commercialisable process on large diameter silicon wafers is seen as a key enabler for increasing the efficiency of green LEDs and reducing the cost of LED lighting.
The collaboration, which is partly funded by Innovate UK under the £14m Energy Catalyst Programme, follows on from work by Anvil Semiconductors and the Cambridge Centre for GaN at the University of Cambridge where they successfully grew cubic GaN on 3C-SiC (cubic silicon carbide)

on silicon wafers by MOCVD (metal oxide chemical vapour-phase deposition – picture shows a CVD chamber). The underlying 3C-SiC layers were produced by Anvil using its patented stress relief IP that enables growth of device quality silicon carbide on 100 mm diameter silicon wafers. The process is readily migrated onto 150 mm diameter wafers and potentially beyond without modification and is therefore suitable for large, industrial-scale applications. Plessey has started to commercialise LEDs produced in conventional (Hexagonal) GaN grown 150 mm silicon wafers using IP originally developed at The University of Cambridge. Anvil's high quality 3C-SiC on Silicon technology, which is being developed for SiC power devices, provides an effective substrate, to allow single phase cubic GaN epitaxy growth and provides a process which is compatible with Plessey's GaN on Si device technology.

Complete article, here

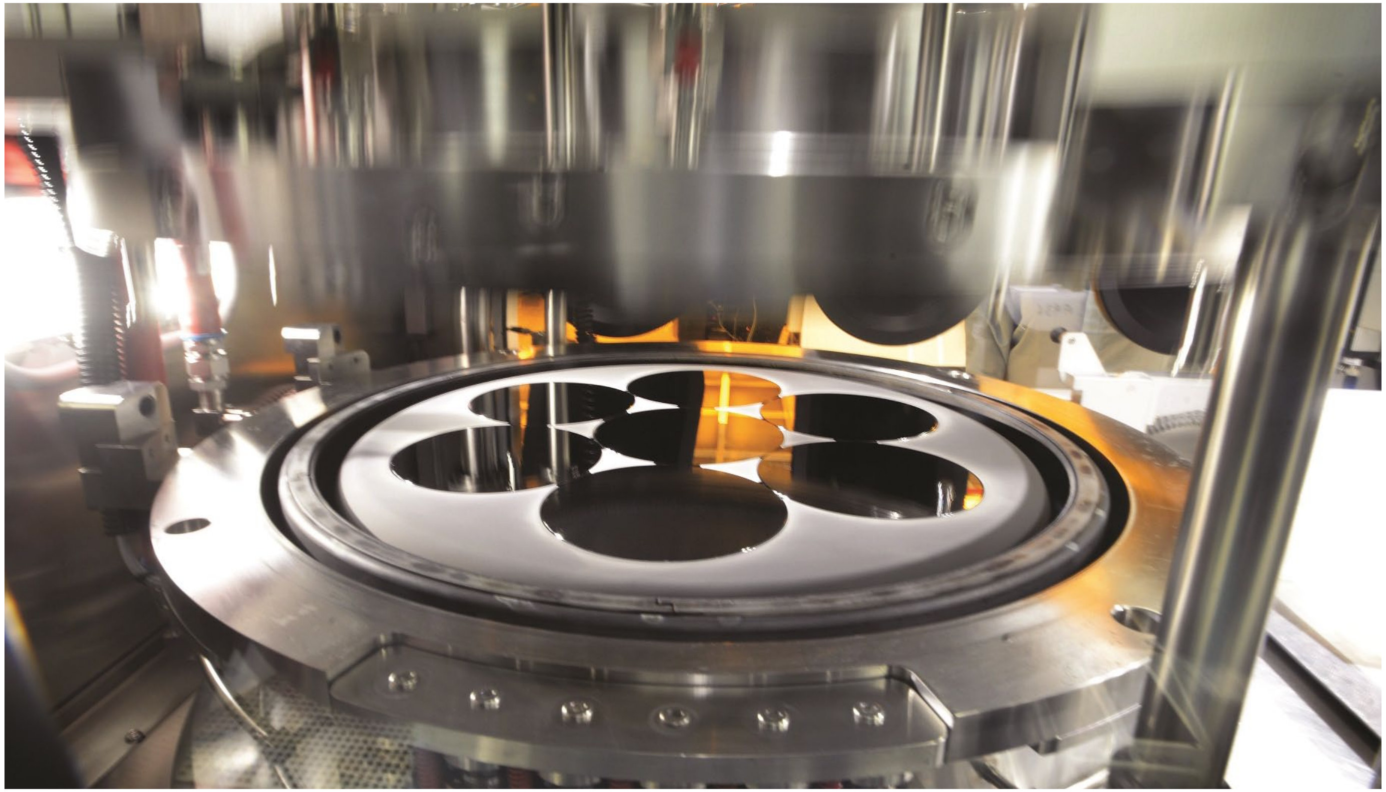

Toshiba to pare chip business down to NAND

by Junko Yoshida, EETimes

Toshiba Corp. has decided to exit most sectors of the company's semiconductor business—except for its NAND flash memory products, Japan's economic newspaper Nikkei reported on Saturday morning, 23rd February 2016.

The move is viewed as a necessary step by the Japanese giant to return to profitability, after expensive post-scandal restructuring. Under the new strategy, Toshiba will focus and consolidate its investment and business operations on "flash memory" and "nuclear energy" businesses. Toshiba has now positioned them as the two pillars to carry the company's future.

Toshiba, the world's second largest NAND supplier after Samsung Electronics, will keep its flash memory business but will sell everything else in its chip business. Those scheduled for sale include: analogue, LSIs, MCUs and other energy saving devices. These products have served the automotive and industrial markets, while also broadly designed into consumer appliances and home equipment.

Toshiba's overall semiconductor business (except for NAND), which will go on sale, is estimated to generate the annual sales of about 200 billion yen [approx. €1.5billion]. Toshiba hopes to devote proceeds of the sales to the company's Yokkaichi fab – the home of its NAND flash production.

Complete article, here

A craving for cheap 64-bit computing

by Julien Happich

Judging from the nearly 30,000 backers Californian startup PINE64 Inc. gathered around its $15 64-bit single board computer on KickStarter, 64-bit is the way to go for makers and electronics hobbyists of all horizons with the promise to run any complex task including graphics processing for multimedia streaming.

With only 48 hours to conclude its crowd-funding campaign, the company had secured $1,362,677, well over 40 times its initial goal. Based on Allwinner Technology's recently released 64-bit tablet processor Allwinner A64 (itself only costing $5), PINE A64 is described as a 64-bit high performance expandable single board computer (SBC). The Allwinner A64 features a quad-core ARM A53 64-bit processor that runs at 1.2 GHz and a dual-core MALI-400 MP2 running at 500 MHz for the GPU, capable of 1.1 Gpixel/s throughput. The GPU provides OpenGL ES 2.0, hardware-accelerated OpenVG, 4Kx2Kp30 H.265 decode, and 1080p60 H.264 high-profile encode and decode.

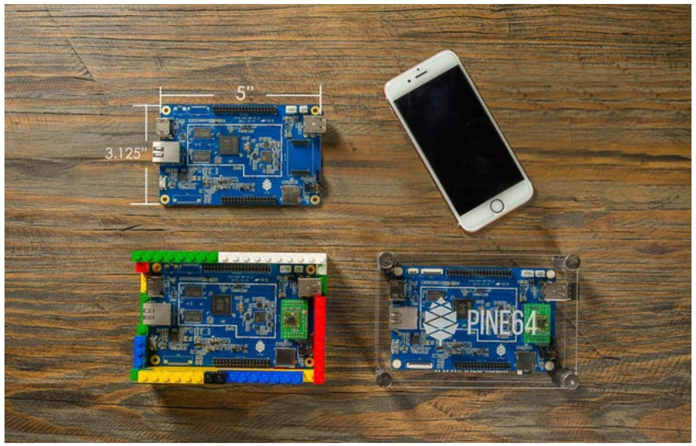

The founders claim their SBC delivers up to 20-30% better performance than other 32-bit open source boards (think Raspberry Pi or Arduino-type). While providing the same power of desktop computing, it only draws 2.5 to 3.5W, either from a 5V supply through MicroUSB or using a 3.7V battery through on-board power management.

Complete article, here

Chip-size Bluetooth module for wearables, in distribution

TDK's ultra-small SESUB-PAN-D14580 Bluetooth module is available from distributor Mouser Electronics; the Bluetooth 4.1 Low Energy (LE) micro module is claimed as the world's smallest module for Bluetooth Smart devices.

The 3.5 × 3.5 × 1 mm SESUB-PAN-D14580 module is based on TDK's Semiconductor Embedded in Substrate (SESUB) technology, reducing the size by 60% compared to modules using discrete components. The single-mode module's ultra-compact footprint and low current consumption suit it for battery-powered wearable devices.

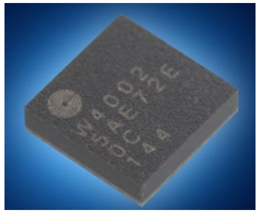

It integrates a Dialog Semiconductor DA14580 Bluetooth 4.1 chip, 32-bit ARM Cortex-M0 microcontroller, and DC-DC converter onto a thin substrate, along with all peripheral circuitry including a 16 MHz crystal, inductor, and capacitor. All inputs and outputs (I/O) from the substrate layers are routed to a ball grid array (BGA) footprint on the module's bottom surface. Interfaces include UART, SPI, and I^2C, helping to speed the hardware design process. The low-power module requires a voltage supply of 3.0V, and consumes only 5.0 mA when transmitting, 5.4 mA when receiving, and 0.8 μA in standby mode.

Complete article, here

Nordic silicon in Bluetooth Smart module packaged for the makers

Nordic Semiconductor says that a third-party open source module, 'OSHChip', supplied in a specially-developed 16-pin through-hole DIP style package can allow Bluetooth Smart and 2.4GHz proprietary wireless prototypes using Nordic chips to be built without soldering because the module plugs straight into any standard bread board, as used by makers and hobbyists.

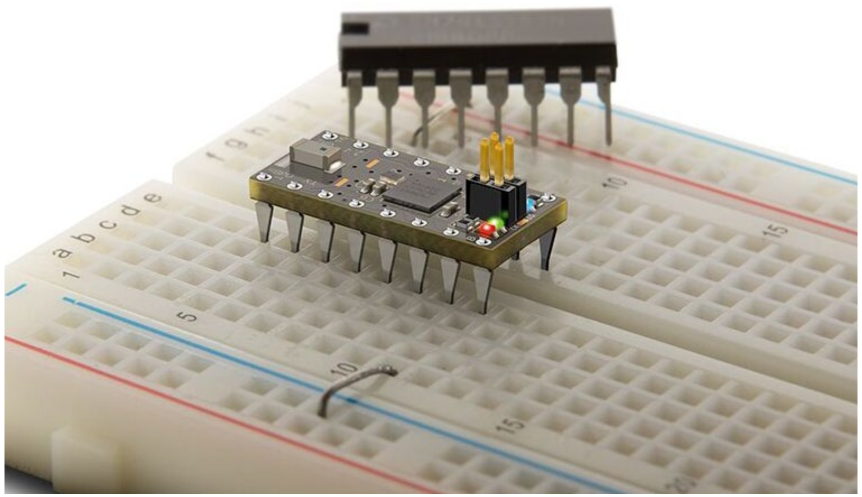

OSHChip (OSH standing for 'Open Source Hardware') from U.S. maker and hobbyist startup of the same name, employs a Nordic nRF51822 System-on-Chip (SoC) and in addition to eliminating the need for soldering, its 1.98 x 0.89 cm form factor is said to be less than 5%t the size of an Arduino UNO or 22% the size of an Arduino Nano.

To develop OSHChip OSH had to develop its own IC-style pins that would be compatible with the vast majority of off-the-shelf bread boards. OSH also says that were it not for the Nordic nRF51822's chip-scale package option it wouldn't have been able to fit his module into such a small form-factor. OSHChip can be programmed and de-bugged (Keil) from a partner OSHChip programmer board that supports USB drag-and-drop programming, SWD debugging, and bi-direction serial communications.

Complete article, here

Ford opts for Analog Devices' UTP-based Automotive Audio Bus for infotainment

Analog Devices has announced that the Ford Motor Company has selected the Automotive Audio Bus (A2B) as its primary infotainment network technology with in-vehicle deployments starting in 2016.

ADI's A2B technology is capable of distributing audio and control data together with clock and power over a single, unshielded twisted-pair wire (UTP), enabling advanced infotainment systems while reducing system costs in wiring-intensive automotive applications. Ford will use the AD2410 transceiver, the first product in the A2B portfolio, in four vehicle platforms moving to production in 2016. The A2B series also includes the AD2401 and AD2402 transceivers that have been specifically tuned for microphone connectivity, for applications such as active noise cancellation, hands-free, and in-car communications.

Complete article, here

Microchip signs to buy Atmel

by Peter Clarke

Microcontroller, analogue and flash memory IC vendor Microchip Technology(Chandler, Arizona) has agreed to buy Atmel (San Jose, Calif.) for about $3.56 billion and Dialog Semiconductor, previously set to buy Atmel, has been paid a $137.3 million termination fee.

Microchip's purchase is $8.15 per share made up of $7 per share in cash and $1.15 per share in Microchip common stock. Atmel has terminated its previous agreement to be acquired by Dialog Semiconductor plc (London, UK) and paid Dialog a termination fee.

The deal has been agreed by both boards of directors and is expected to close in the second quarter of 2016, subject to approval by Atmel's stockholders, regulatory approvals and other customary closing conditions. "As the semiconductor industry consolidates, Microchip continues to execute a highly successful consolidation strategy with a string of acquisitions that have helped to double our revenue growth rate compared to our organic revenue growth rate over the last few years," said Steve Sanghi, CEO of Microchip.

Complete article, here

Time-sensitive Ethernet ready for demo, says Innovasic

Specialists in deterministic Ethernet silicon, Innovasic (Albuquerque, New Mexico) demonstrated TSN-ready Deterministic Ethernet at the 2016 Automotive Ethernet Congress in Munich (3rd and 4th February 20126). Innovasic says it will demonstrate how time-critical control data can be effectively managed on an Ethernet network using technologies that are being adopted and evaluated by the IEEE standards association for 802.1 Time Sensitive Networking (TSN).

By using its fido5000 Real-time Ethernet Multiprotocol (REM) switch chip, Innovasic will demonstrate scheduled traffic (802.1Qbv), time synchronisation (802.1AS-Rev) and ingress policing (802.1Qci). These capabilities enable deterministic control of critical data packets, even in the presence of very high levels of non-critical traffic. The fido5000

chip has a programmable architecture with the necessary hardware in place to support the new TSN standards as well as existing AVB standards. The programmable architecture is important, the company says, as it allows users of the chip to adapt to specification changes as the TSN standards evolve and are finalised. Innovasic's device also supports pre-emption (802.1Qbu) and redundancy (802.1CB).
The use of the new IEEE 802.1 TSN standards, in conjunction with new 802.3 PHY standards, is expected to enable widespread adoption of Ethernet within the car both as a network backbone as well as to replace legacy serial connections. The widespread use of Ethernet in the car can greatly reduce the cost and weight of the wiring harness and also enable easier connectivity options for the connected car and autonomous vehicles.

Complete article, here

Motion tracking module incorporates sensor fusion

Intended to give product designers a simple way to add motion intelligence to their projects, and to accelerate time-to-market, Fairchild's FMT1000-series of highly accurate motion tracking modules enables the fast integration of motion intelligence into any system, including drones, autonomous vehicles, unmanned systems, heavy industry, construction, agriculture, VR headsets, and camera and platform stabilisation.
The FMT1000-series is a range of turnkey Motion Tracking modules outputting inertial data with a complementary offering of roll, pitch and (relative) heading corrections, resulting in the first inertial module, delivering accurate 3D orientation performance with industrial specifications.
The modules can be integrated with minimal hardware and software development, employ an easy-to-use communication protocol, open source drivers, and include source code and examples for embedded integration via I²C, SPI, ARM mbed, as well as a full software suite (for Windows and Linux).

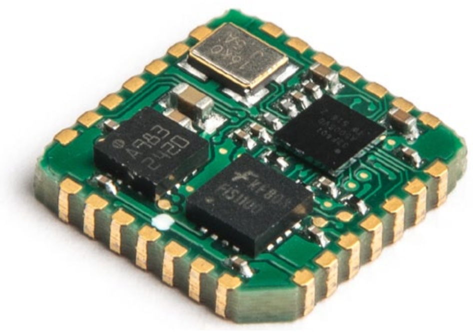

Built upon Fairchild's FIS1100 6-axis MEMS Inertial Measurement Unit (IMU), the FMT1000 modules are individually tested and calibrated over the full temperature range to reduce manufacturers' capital investment and development time. All modules are time-synced and run onboard advanced sensor fusion algorithms with auto-calibration and data formatting, relieving the host processor from calculations to save cost and power.

Complete article, here

Microchip opens PIC compiler licensing to low-cost rolling monthly subscription

With this addition to the business models on which it supports its development tools, Microchip says it is offering the ability to license PRO compiling optimisations on an ‘as-needed’ basis, providing flexibility and giving access to tools that enable best execution speed and code size for all ~1,400 PIC MCUs and dsPIC DSCs.

The flexible, low-cost, renewable monthly subscription license is available for the PRO editions of the MPLAB XC line of C compilers. The MPLAB XC8, XC16 and XC32 compilers offer reduced complexity, and are available with three optimisation levels: Free, Standard and PRO. By offering a monthly-subscription PRO license, Microchip is providing 8-, 16- and 32-bit MCU designers the ability to make use of the most features and best code optimisation - only when needed in the design cycle. Subscribers will also have the ability to receive updated versions of the MPLAB XC compiler type to which they subscribe, without the need for an active Microchip MPLAB XC High Priority Access (HPA) maintenance subscription.

Unlike most software subscriptions, this license is not cloud-based, providing designers the additional flexibility of using the compilers offline. Renewable monthly subscription licenses for the following compilers are priced at $29.95 per month, for each compiler.

Complete article, here

PRECISE BATTERY DIAGNOSIS MAKES MANOEUVRING XXL LOADS EASIER

By Athier J. Lafta, Product Manager Precision Measurement (VT-MT) Isabellenhütte

The shunt-based "IVT Modular" sensor module from Isabellenhütte measures battery current and voltage to within less than 0.1 percent. This extraordinary precision enables permanent energy monitoring and efficient utilisation, even on high-capacity battery systems. This is why KUKA Roboter GmbH fits the IVT sensor modules as standard in its mobile platforms for transporting heavy loads.

KUKA's omniMove platform stands out thanks to its ability to perfectly manoeuvre heavy loads in any direction and rotate them on the spot. "The omniMove makes it possible to move heavy parts very precisely in limited space, with 360° flexibility and a positioning accuracy of ± 2 millimetres", explains Paul Wyszynski, developer at KUKA Roboter GmbH. "In the lifting platform variant, the omniMove can move in the millimetre range, even when the platform is raised. This is used when painting airliners, for example." The platform is also called upon to transport gas turbines the size of a house, moving them between the individual production areas or to the testing area. The size, width and length of modular platform system are freely scalable, and it can be combined with additional vehicles to form tandem or tridem configurations. This makes it possible to transport loads the size of an entire aircraft fuselage. The platform can be operated via remote control, optical tracking, or fully autonomously with laser scanners.

Precisely measured battery capacity allows greater utilization

The omniMove's electrical drive is extremely efficient and quiet, even when used for long periods. The model KoM UTV-2 E375, for example, achieves a minimum operating duration of four hours. For technical and economic reasons, it is important that the available energy in a heavy-duty vehicle can always be used efficiently. This is why the lead gel batteries fitted as standard in the omniMove use current and voltage sensors based on sensor modules from Isabellenhütte, which are among the best in shunt-based current measurement technology.

Figure 1. *The omniMove mobile platforms from KUKA for extremely heavy loads of up to 100 tonnes, are used e.g. in aircraft construction.*

Only by determining the current and voltage as precisely as possible is it possible to get a reliable picture of battery function, such as the state of charge, state of health, or the ability of the battery to fulfil a certain requirement such as start capability (state of function). "These measurement values can then be used to control the flow of energy in the battery, improve charging times and cycles, and lengthen the service life of the battery", confirms Jens Hartmann, Sales Director ISAscale at Isabellenhütte. The high utilisation of the omniMove batteries is achieved by precisely measuring the battery capacity.

KUKA uses the "IVT Modular" sensor module from Isabellenhütte to measure the current and voltage values in the omniMove vehicles.

Download PDF of Full Article

Find Current sensor on EETsearch

FLEXIBLE COMMUNICATION FOR THE SMART GRID

By Emmanuel Gresset, CEVA

How DSP platforms help engineers address the smart energy network challenge; The communication and connectivity requirements of smart meters and other 'smart energy' equipment and devices present significant challenges for designers and developers. DSP platforms can offer a practical, simplified and cost-effective route to delivering the flexibility and future-proofing demanded by these applications.

Fuelled by legislative, environmental and economic factors, utilities are under pressure to improve efficiency and reliability and to better manage electricity production and distribution. Fundamental to achieving these goals is implementation of 'smart energy' information and communication technology (ICT) infrastructures. These infrastructures facilitate the automatic monitoring of consumer behaviour, deliver enhanced intelligence for suppliers that help them align supply with demand, and improve flexibility for customers in terms of how and when they consume energy.

Growth of the Smart Grid

While there is no formal definition of a smart grid the term is generally used to describe the application of automation, processing and two-way communications technology between the electricity generation and transmission infrastructure and the end user electric meter.

Driven largely by governments through legislation, mandated use and incentives, the global market for smart grid technologies is estimated to be growing at a rate of around 8% per year meaning it will surpass the $400 billion mark by 2020. A key element of the smart grid is the smart meter – an electronic device that interfaces between utility and consumer and combines usage measurement with in-built intelligence and communications capabilities needed to support enhanced monitoring and control. According to a 2014 report by the US International Trade Commission, the global market for smart meters will reach £20 billion in 2018.

Over the years smart meters have evolved from their original incarnation as AMR (automated meter reading) devices with basic one-way communication to AMI (advanced metering infrastructure) platforms incorporating two-way communications. As well as providing consumer data for billing and analysis, such meters allow commands to be sent to the consumer premises to implement initiatives such as 'time-of-use' pricing, distribution automation and demand-response actions. They also

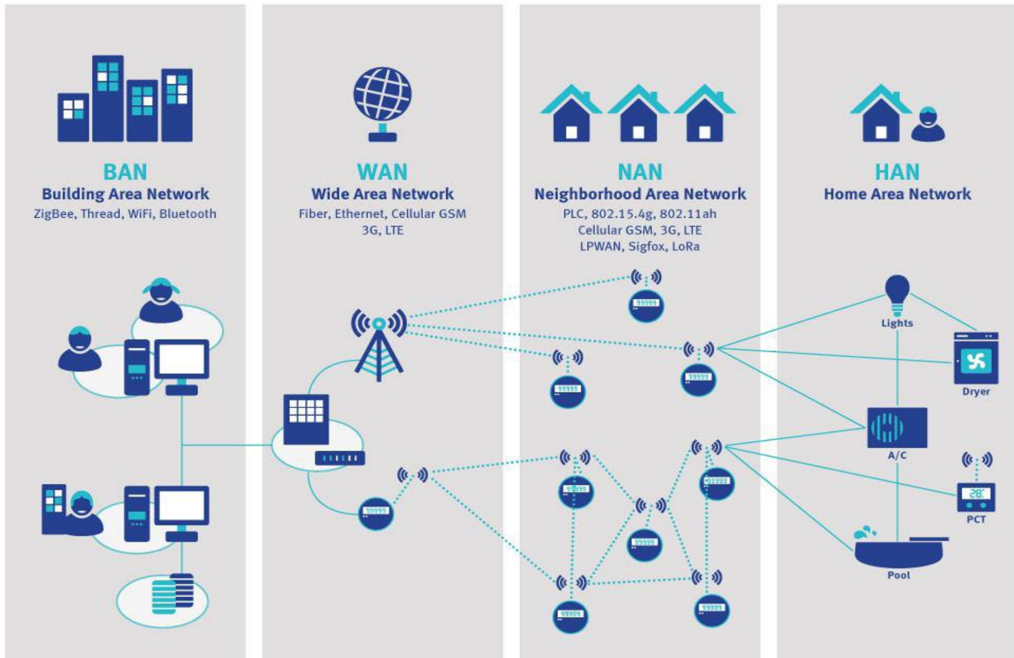

Figure 1. *Key elements of the Smart Grid*

enable remote switch-on and shut-off, provide a mechanism for assessing 'feed-in' payments, and give consumers the means to track their own energy usage via browsers or apps which enable direct connectivity to the smart meter.

Figure 1 depicts the key elements of the modern smart grid. The smart meter in the consumer premises is connected to monitoring and control consoles and other devices at the premises via a home area network (HAN) or building area network (BAN) and to a local 'concentrator' via a neighbourhood area network (NAN). Concentrators communicate data to and from the utilities using wide area networks (WANs).

Smart Meter requirements

Figure 2 shows the key elements of a smart meter, namely power supply and management; measurement circuitry for metrology; on-board 'intelligence' in the form of microcontroller and memory (and associated software); and the HAN and NAN communication blocks. A growing number of meters also incorporate GPS functionality to simplify location identification and to provide accurate date stamps for data.

Smart meter designs must address the need to deliver accuracy in line with legal requirements for billing purposes; keep power consumption to an absolute minimum (while traditional electromechanical meters may consume less than 1W, modern AMI meters can consume as much as 15W during operation and around 2W when idle); ensure security (including tamper protection and detection); and deliver reliable operation over long periods, preferably with options to update software remotely.

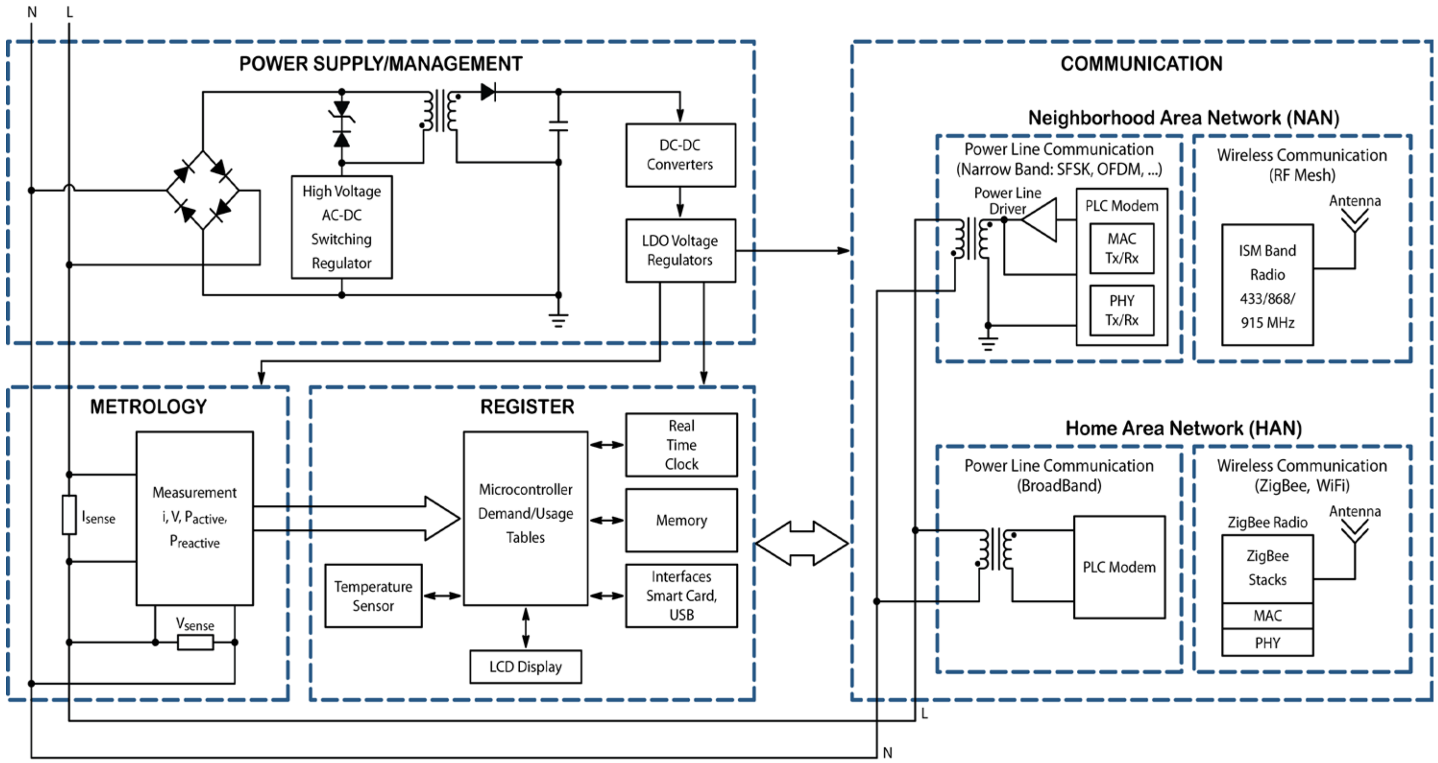

Figure 2. *Key elements of an AMI smart meter*

The article continues to propose an IP-based strategy to meet these needs; click for pdf

Download PDF of Article

Find IoT Communications on EETsearch

MAKING THE MOST OF AUTOMATED TOOLS FOR POWER SYSTEM DESIGN

By Rudy van Parijs, EBV Elektronik & Preet Sibia, Fairchild

How important is power design expertise? Power system design has often been considered a bit of a black art and hence frequently delegated to skilled power design engineers. This was fine in larger companies that could afford such dedicated resource and where a more traditional serial design process allowed the final power requirements to be determined once the system design was complete.

This is rarely the case today and, increasingly, designers need to embrace both the power supply design as well as the overall system design. Not only that, but time-to-market pressures mean that these design elements need to run in parallel, requiring a reasonably accurate anticipation of the power budget long before the system design is complete. This is challenging even for experienced hardware engineers – let alone someone new to developing a power system.

Regulatory requirements

Adding to the challenge is the need for power supply certification to ensure compliance with national and international regulations, such as UL for safety and FCC and CE regulations for emissions and interference. Here getting the design right first time is even more important because of the delays and added cost that are incurred if a design has to be reworked and re-certified.

So getting the power design right is vital.

Adapting to all skill levels

For the experienced power designer the ability to "tweak" all aspects of a power management system is even more important. This can reduce the system's energy consumption and provide a competitive advantage. It can mean longer operating time, lower operating costs or result in a smaller system. The ability to focus on particular areas such as transient response can provide benefits too.

There are many trade-offs throughout the design process, and while it is possible to look at these with a calculator and a list of equations, a spreadsheet or even an evaluation board, the biggest challenge is achieving a working design in a timely manner. Evaluation boards are optimised to demonstrate optimum performance under a fixed set of conditions and their performance rarely reflects real-world operating conditions.

Hardware engineers may turn to design tools for help, however, this often requires a range of tools, and switching from one to the other to get a complete picture, all of which can be confusing for novice power designers. It can be particularly difficult with different types of power design. Industrial LED lighting, motion controls, general switch mode power supplies and power trains all have different requirements that a general hardware engineer may not familiar with. Selecting the right components, from the MOSFETs to the power drivers and even the resistors, capacitors and inductors can take considerable time and effort.

Under these circumstances, Power Supply WebDesigner – a suite of tools from Fairchild for designing and optimising a power supply design – can save considerable time and effort. Device selection, design, analysis, and simulation takes only minutes because the models, calculations, and iterative steps of power supply design are built into the tools. The Automatic Design option provides all the mainstream settings and produces a full circuit diagram with a single click.

LED lighting supply

Consider the example of designing a dimmable 240Vac LED flyback power supply using the Fairchild FL6630 LED driver. Power Supply WebDesigner generates the simulation for the schematic in under three minutes and provides the transient analysis to show what is happening with the waveforms generated by the design. It rapidly produces a workable design, dramatically simplifying the challenge for the designer.

The FL6630 itself is packaged in an 8-pin small outline package (SOP). It's the basis for a single-stage, power factor corrected (PFC), offline LED driver circuit and uses Fairchild's TRUECURRENT technology to provide constant current control. This allows for the simplified circuit design for LED lighting applications that the design tool generates.

Using a single stage topology with primary-side regulation, an LED lighting board can be implemented with few external components and minimal cost, as it does not need an input bulk capacitor or feedback circuitry. The tool implements the power factor correction and ensures low total harmonic distortion, with discontinuous conduction mode (DCM) operation.

The constant current approach allows the operating frequency to change by the output voltage to guarantee Discontinuous Conduction Mode (DCM) operation. This provides higher efficiency and simpler design. The device provides protections such as open-LED, short-LED, and over-temperature, and the current-limit level is automatically reduced to minimise output current and protect external components in a short-LED condition.

However, the Auto-Complete feature avoids any of the optimisations that the experienced power designer would add to an implementation... *the conclusion of this short article notes that the design tool described also offers the experienced designer the option to make advanced refinements – click for pdf.*

Download PDF of Article

Find Automated PSU Design on EETsearch

smartsystems integration

International Conference and Exhibition
on Integration Issues of Miniaturized Systems
– MEMS, NEMS, ICs and Electronic Components

Co-organizer:
Fraunhofer ENAS
Fraunhofer IZM

Part of the activities of:
EPoSS

USING THE ADP2370 AS AN LED DRIVER

BY GLENN MORITA, ANALOG DEVICES

The fluorescent bulb in my battery powered fluorescent camp lantern failed. A new one was not available, so I decided to replace the defunct light bulb with several 1-W white LEDs. The best way to drive the LEDs was to use an off-the-shelf integrated buck converter.

Modern high-power LEDs are more efficient and have longer operational lifetimes than incandescent or compact fluorescent light bulbs. In addition, their lack of flicker helps to reduce eyestrain, making LEDs far superior to fluorescent lighting.

This analog tip shows how the ADP2370 adjustable-output-voltage buck converter can regulate the LED current from a variety of power sources to make a simple, robust, high-efficiency dimmable LED driver.

Regulated current LED driver

Buck converters such as the ADP2370 normally regulate a lower voltage from a higher voltage source, but the LED driver circuit shown in Figure 1 uses the ADP2370 to regulate the LED current instead of an output voltage. R7 and R8 function as both ballast and current sense resistors; R6 and R9 average the current sense voltages at the input of op amp U2 and equalise the LED currents. U2 amplifies the voltage across the current sense resistor, with the gain set by R2 and R3, and drives the FB input of the ADP2370, thereby "fooling" it into regulating the LED current. A gain of 5.32 sets the current in each LED to about 320 mA. R4, R5, R6, and R9 provide a dimming function. As R4 is decreased, the current-sense voltage is slightly offset above 0V, reducing the LED current. A photoconductive sensor could replace R4 to dim the LEDs as a function of the ambient light level. Over-temperature protection can be added by replacing or paralleling R4 with an NTC thermistor to reduce the drive current if the LEDs exceed a certain temperature. The thermistor must be in good thermal contact with the LEDs for accurate temperature sensing.

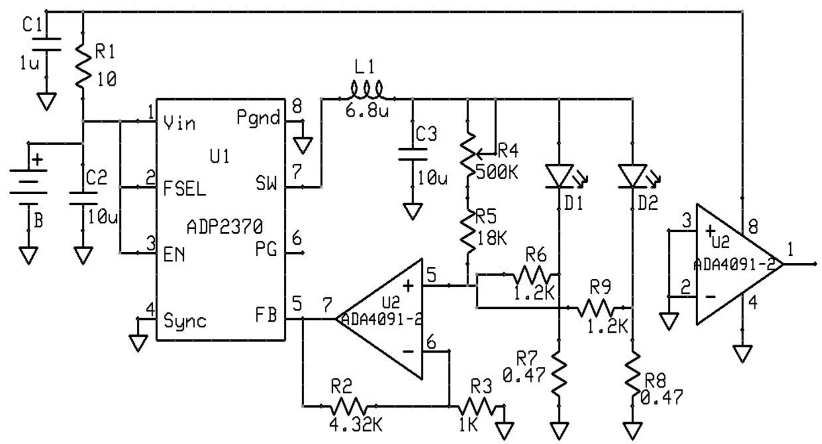

Figure 1. *Regulated LED driver using the ADP2370*

About the Author

Glenn Morita *[glenn.morita@analog.com] graduated from Washington State University with a BSEE in 1976. His first job out of school was at Texas Instruments, where he worked on the infrared spectrometer instrument for the Voyager space probe. Since then, Glenn has worked as a designer in the instrumentation, military and aerospace, and medical industries. In 2007, he joined ADI as an applications engineer with the Power Management Products Team in Bellevue, WA. He has over 25 years of linear and switch-mode power supply design experience at power levels ranging from microwatts to kilowatts. Glenn holds two patents for harvesting energy from body heat to power implantable cardio-defibrillators and an additional patent for extending battery life in external cardio-defibrillators.*

LOW POWER WIDE-AREA NETWORKING ALTERNATIVES FOR THE IOT

By Richard Quinnell, EDN

Wireless network technologies such as WiFi, ZigBee, and Bluetooth are fine for consumer applications of the Internet of Things (IoT), but many civic, industrial, and other IoT applications need to operate over vastly greater territory than these technologies can handle. Cellular and satellite machine-to-machine (M2M) technologies have traditionally filled the gap, but cost, power, and scalability concerns make these choices less appealing for the future. A number of low-power, wide-area networking (LP-WAN) alternatives have arisen that need careful consideration by developers looking to address these wide-ranging IoT applications.

The uses for wide-area IoT technology are legion. Civic infrastructure systems such as parking resources, traffic control, utilities monitoring and distribution control, and environmental monitoring are only a beginning. Agricultural uses such as monitoring of crop conditions and livestock movements need wide-area coverage. Asset monitoring and tracking, from taxicabs to refrigerated produce shipments need regional, national, or even worldwide coverage. Transportation infrastructures such as rail lines and roadways need wide-area monitoring. Even consumer applications such as health monitoring could benefit from having an alternative to cellphones for their wide-area connectivity.

Comparison of Low-Power WAN Alternatives

Name of Standard	Weightless -W	Weightless -N	Weightless -P	SigFox	LoRaWAN	LTE-Cat M	IEEE P802.11ah (low power WiFi)	Dash7 Alliance Protocol 1.0	Ingenu RPMA	nWave
Frequency Band	TV whitespace (400-800 MHz)	Sub-GHZ ISM	Sub-GHZ ISM	868 MHz/902 MHz ISM	433/868/780/915 MHz ISM	Cellular	License-exempt bands below 1 GHz, excluding the TV White Spaces	433, 868, 915 MHz ISM/SRD	2.4 GHz ISM	Sub-GHz ISM
Channel Width	5MHz	Ultra narrow band (200Hz)	12.5 kHz	Ultra narrow band	EU: 8x125kHz, US 64x125kHz/8x125kHz, Modulation: Chirp Spread Spectrum	1.4MHz	1/2/4/8/16 MHz	25 KHz or 200 KHz	1 MHz (40 channels available)	Ultra narrow band
Range	5km (urban)	3km (urban)	2km (urban)	30-50km (rural), 3-10km (urban), 1000km LoS	2-5k (urban), 15k (rural)	2.5- 5km	Up to 1Km (outdoor)	0 – 5 km	>500 km LoS	10km (urban), 20-30km (rural)
End Node Transmit Power	17 dBm	17 dBm	17 dBm	10µW to 100 mW	EU:<+14dBm, US:<+27dBm	100 mW	Dependent on Regional Regulations (from 1 mW to 1 W)	Depending on FCC/ETSI regulations	to 20 dBm	25-100 mW
Packet Size	10 byte min.	Up to 20 bytes	10 byte min.	12 bytes	Defined by User	~100 -~1000 bytes typical	Up to 7,991 Bytes (w/o Aggregation), up to 65,535 Bytes (with Aggregation)	256 bytes max / packet	Flexible (6 bytes to 10 kbytes)	12 byte header, 2-20 byte payload
Uplink Data Rate	1 kbps to 10 Mbps	100bps	200 bps to 100 kbps	100 bps to 140 messages/day	EU: 300 bps to 50 kbps, US:900-100kbps	~200kbps	150 Kbps ~ 346.666 Mbps	9.6 kb/s, 55.55 kbps or 166.667 kb/s	AP aggregates to 624 kbps per Sector (Assumes 8 channel Access Point)	100 bps
Downlink Data Rate	1 kbps to 10 Mbps	No downlink	200 bps to 100 kbps	Max 4 messages of 8 bytes/day	EU: 300 bps to 50 kbps, US:900-100kbps	~200kbps	150 Kbps ~ 346.666 Mbps	9.6 kb/s, 55.55 kbps or 166.667 kb/s	AP aggregates to 156 kbps per Sector (Assumes 8 channel Access Point)	--
Devices per Access Point	Unlimited	Unlimited	Unlimited	1M	Uplink:>1M, Downlink:<100k	20k+	8191	NA (connectionless communication)	Up to 384,000 per sector	1M
Topology	Star	Star	Star	Star	Star on Star	Star	Star, Tree	Node-to-node, Star, Tree	Typically Star. Tree supported with an RPMA extender	Star
End node roaming allowed	Yes	Yes	Yes	Yes	Yes	Yes	Allowed by other IEEE 802.11 amendments (e.g., IEEE 802.11r)	Yes	Yes	Yes
Governing Body	Weightless SIG			Sigfox	LoRa Alliance	3GPP	IEEE 802.11 working group	Dash7 Alliance	Ingenu (formerly OnRamp)	Weightless SIG
Status	Limited deployment awaiting spectrum availability	Deployment beginning	Standard in development. Scheduled release 4Q 2015	In deployment	Spec released June 2015, in deployment	Release 13 expected 2016	Targeting 2016 release	Released May 2015	In Deployment	In Deployment

Source: EDN.com - Copyright 2015 UBM Americas Rev. 9/15/15

Low-power WAN alternatives - expand page view for readability

LP-WAN essentials

While the applications are diverse, they have many common attributes on their network wish lists. These include:

Low cost – Most wide-area IoT applications anticipate a need for many hundreds or thousands of end-node devices for each installation. In some cases, such as city-wide parking space/meter monitoring, the numbers can get into the millions. With such high volumes, unit price is a major consideration in determining the return on investment (ROI) for the application.

Low energy consumption – Few of the applications for wide-area IoT have the luxury of a local power generator. Most will depend on batteries and some may even need to use energy harvesting. For those with batteries, replacing depleted batteries can represent a major logistical challenge as well as a substantial cost. The longer the battery life in the end node device, the better.

Extended range – All wireless networks connecting to the Internet need to work through an access point (AP) of one kind or another: gateway, concentrator, or the like. So an IoT design needs to consider both the endpoint cost and the cost of the access point infrastructure needed to support the application. The network's operating range, or allowable distance from an end node to its access point, can have a significant impact on that infrastructure cost. Range dictates the number and location of access points needed to cover the application's operating area, so in general the longer the range the lower the infrastructure cost.

Scalability – A given installation using a wide-area wireless IoT network may work well and the network may well have the capacity to handle any anticipated single user. But over time it's reasonable to expect that many different installations will be made in the same geographic area. If these different installations share common access points, like cellphones share towers, then the number of devices an access point can support can become a limiting factor and require increases infrastructure to overcome. Even if they don't share access points but do share the frequency spectrum, an increase in installations can erode the operational range of application through increased noise levels. In the worst cases, available channel capacity can fill and prevent new installations from operating at all.

Among the more established wireless networking technologies, only cellular and satellite communications offer the extended ranges that these applications require. Mesh networks such as ZigBee can potentially cover large areas but have limited scalability due to the need to forward traffic.

Unfortunately cellular and satellite communications technologies have limitations in the other attributes. Their radio requirements involve higher energy use and complex protocols that lower battery life and increase cost beyond what many applications can sustain. This arises in part from their history; they were originally designed to handle voice traffic. The networks are ill-adapted to handling short data messaging.

Still, some IoT applications and services – often called machine-to-machine (M2M) – did arise to exploit cellular and satellite communications networks. Many of them were based on the CDMA, or "2G" cellular technology. Unfortunately, those networks are now starting to be phased out by service providers in order to free spectrum for more advanced cellular technologies. However, the cellular community has made some strides toward improving the situation for M2M. The most recent specification for LTE (release 12) defined communications Category 0 designed around the needs of M2M traffic. Energy use and cost still remain concerns, however.

This situation has opened a door for alternative approaches to wide-area wireless networking for the IoT, approaches that focus on the low-power, low-cost requirements.

The article continues by considering the attributes required of low-power WAN schemes, and tabulating all of the competing alternatives – click for pdf.

Download PDF of Article

Find Low Power WAN for IoT on EETsearch

MEASURING 2 NV/√HZ NOISE AND 120 DB SUPPLY REJECTION IN LINEAR REGULATORS; THE QUEST FOR QUIET

By Todd Owen and Amit Patel, Linear Technology

There are numerous applications where linear regulators are required to provide quiet, often extremely quiet, power supply rails, but how does one ensure that the regulator performs as specified? Although we are in the domain of power supply design, this immediately presents a significant challenge in test & measurement

A quiet, well regulated supply is important for optimum performance in a number of circuit applications. Voltage controlled oscillators (VCOs) and precision voltage controlled crystal oscillators (VCXOs) respond to small changes in their supply very quickly. Phase-locked loops (PLLs) require a stable supply, as signal on the supply translates directly to phase noise in the output. RF amplifiers require quiet supplies, as they lack the ability to reject supply variations, and regulator variation will appear as unwanted side bands and lower the signal-to-noise ratio. Low noise amplifiers and analogue-to-digital converters (ADCs) do not have infinite supply rejection and the cleaner the regulator output is, the higher their performance.

Once fully built, one can determine if the supply has low enough noise for the application. Oscillator phase noise is measured and compared against results achieved with a known good supply, ADCs are checked to ensure they are getting the maximum number of bits. These are tricky, time consuming measurements and it would be better to make sure the noise levels are low enough for your needs without expensive trials.

"Measuring less than 1 µVRMS noise is not a trivial task"

In addition to noise, one must also consider the supply rejection capabilities of the linear regulator. Poor rejection from a linear regulator will bring switching regulator residue or other unwanted signals through, corrupting the hard work done to ensure a clean supply. Extremely low noise from the regulator is worthless if poor supply rejection brings enough signal through to swamp noise levels.

Measuring output voltage noise: being Quiet is Nothing New

The subject of noise has been broached before. Linear Technology Application Note 83, "Performance Verification of Low Noise, Low Dropout Regulators," published in March 2000, describes in detail a method for measuring output voltage noise of regulators down as low as 4 µVRMS with confidence. The amplifier circuit and filters in the Application Note gave 60 dB of gain across a 10 Hz to 100 kHz bandwidth. This is a good starting point to determine confidence in measurement of noise levels.

New linear regulators such as the LT3042 are now in production with much lower output voltage noise levels. While the family of regulators released around the publication of Application Note 83 operate with approximately 20 µVRMS noise in the 10 Hz to 100 kHz band, the LT3042 is now available with noise levels as low as 0.8 µVRMS across the same frequency band. Reviewing the circuit from Application Note 83 shows an input referred noise floor of 0.5 µVRMS, which provides less than 1% error when measuring noise levels as low as 4 µVRMS. With output noise levels of 0.8 µVRMS, this noise floor is now unacceptable; the regulator itself operates at noise levels only slightly above the measurement circuit. This translates to almost 20% error, making the measurement circuit too significant a factor to be able to measure signals with confidence.

Measuring less than 1 µVRMS noise is not a trivial task. Working backward from a 10 Hz to 100 kHz measurement band, this equates to a noise spectral density of 3.16 nV/√Hz (assuming white noise). This is equivalent to the Johnson noise of a 625Ω resistor! Measuring noise at these levels within 5% requires that instrumentation have an input referred noise of 1 nV/√Hz; measuring within 1% requires input referred noise of 450 pV/√Hz.

What measurement to make?

We now have an idea of the noise floor required by instrumentation, but there is a question as to what frequency range is critical and what instrument is to be used to measure the resultant noise. To measure noise spectral density, the regulator output can simply be fed through low noise gain stages [Ref. 1] and then fed into a spectrum analyser, blocking out unwanted frequencies from measurement. If peak-to-peak or RMS noise is desired, then band stops are warranted on the low noise gain stages to ensure that only signal in the desired bandwidth is measured.

A commonly used broadband noise measurement frequency range is 10 Hz to 100 kHz. This encompasses the audio frequency band and ensures minimal side bands for baseband data transmitted over RF. Low noise regulators used in phase-locked loops and high accuracy instrumentation require higher frequency measurements (up to 1 MHz and beyond), so we should not limit ourselves to only the 100 kHz range. Ideally, band stops would be absolute brick-wall filters at the desired frequency, but the realities of circuit design prevent us from achieving this. Higher order Butterworth filters are selected to maintain maximum flatness in the range of the frequencies of interest as well as their ability to give a better brick-wall approximation. The order of the filter is determined by the error introduced by their equivalent noise bandwidth (ENB): a second-order low pass Butterworth has an ENB of $1.11f_H$, which is too high an error. Fourth-order filters drop the ENB to $1.026f_H$, which gives error levels of approximately 1.3%. Higher order filters would add unnecessary complexity and cost while accomplishing minimal improvement in performance. Fourth-order filter error is coupled with errors introduced by the input referred noise, indicating that a measurement within 5% requires that input referred noise of the amplifier be targeted to contribute no more than 1% maximum error.

Circuit gain must be considered as well. If the gain is too low, noise of the measurement device will sum in and corrupt measurements the same as input noise of the amplifier. At the same time, instrumentation may not be sensitive enough to provide reliable results. For RMS noise measurements, an HP3400A RMS voltmeter has a bottom range of 1 mV, so 60 dB is an absolute minimum gain. Based on the noise floor of spectrum analysers currently commercially available (and available from the secondary market), it was decided that 80 dB would work best.

A block diagram of the noise measurement circuit is shown in Figure 1. - *this article continues to cover circuit design of the measurement setup; practical considerations of its construction to eliminate unwanted conducted, radiated and magnetic-coupling effects; and the instrumentation required.*

This article has appeared as 4-part series on the EDN Europe website, start here; or click below for a single pdf of the complete article.

Download PDF of Article

Find Linear Regulators on EETsearch

THE PHENOMENON OF TECHNICAL DEBT

By John Paliotta, Vector Software

The technical debt metaphor is gaining quite a bit of traction in the software development world. This term was coined by American Ward Cunningham in a report at the OOPSLA conference (Object-Oriented Programming, Systems, Languages & Applications, an annual ACM research conference) in 1992; he said, "Shipping first time code is like going into debt. A little debt speeds development, so long as it is paid back promptly with a rewrite…. The danger occurs when the debt is not repaid. Every minute spent on not-quite-right code counts as interest on that debt. Entire engineering organisations can be brought to a stand-still under the debt load of an unconsolidated implementation."

In simple terms, technical debt is the result of the trade-off's made during development: "we don't have time to do the right thing now; we'll do that in the next release". Just like financial debt, technical debt is not necessarily a bad thing if you're diligent about paying it back. The problem comes when you don't pay it back, or you incur it when you know you will never be able to pay it back.

Is all Technical Debt bad?

Technical debt is an unavoidable and inevitable part of building software. Our economy has an insatiable demand for new and improved electronic products, and software is fundamental to building these products. The business pressure to get products to market means that organisations will always incur some technical debt during development.

Some in the software development industry argue that there are actually different classifications of technical debt, some more dangerous than others as shown in the following illustration:

As depicted in Figure 1, 'Reckless' and 'Prudent' debt are both forms of acknowledged and deliberate debt, knowingly factored into the development of a software project. Inadvertent debt seems to be the most worrying kind as it implies a lack of visibility or understanding of the potential debt.

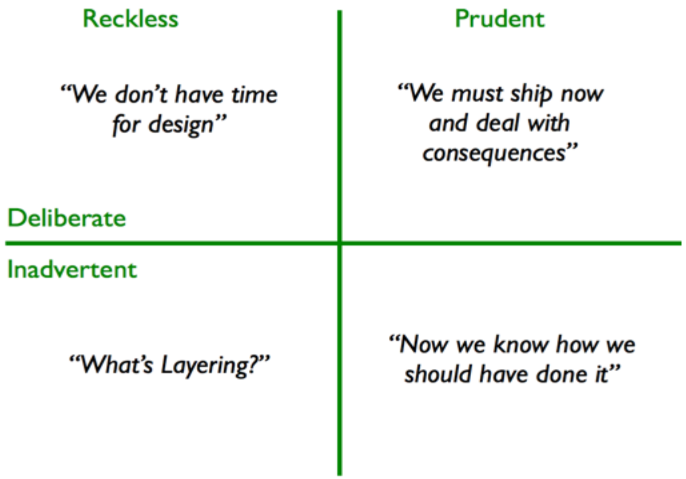

Figure 1. *The different aspects of 'technical debt', rendered as a 2-D space*

Where does Technical Debt come from?

Many things contribute to technical debt, but the top four coding problems are: architecture problems, overly complex code, incomplete testing, and lack of documentation.

Architecture problems are primarily the result of the initial design being built before a complete understanding of the problem exists. A common example is an application that does not scale well as the number of users increases. The initial design might have anticipated 1,000 simultaneous users. Two years later we are seeing 100,000 users.

Overly complex code results when new features are bolted onto existing algorithms. It is really tempting to extend something that does 'almost what we need' by adding some extra parameters and conditional code.

Incomplete testing most often results from pressure to ship version 1.0, or 5.x on a specific date. Since software development has historically been a linear process of: design, code, test - guess what part of the process gets skipped when a release date is looming?

Lack of code documentation is the result of lazy developers. Period. The most valuable software documentation is the comments inserted into the source as the code is developed; these comments should capture the assumptions that the developer has in mind as he/she writes the code, and are critical for software maintenance. If you are not enforcing code commenting, you should.

How to control technical debt

The biggest reason that technical debt is out of control in most applications is that no effort is being made to pay down the debt, and new debt is being added every day.

Technical Debt is a problem for even the largest corporations, with Bill Gates historically demanding that all Microsoft engineers stop writing new code for a full month to spend time focussing on fixing the bugs in the software that had already been built.

Just like financial debt, step 1 in dealing with Technical Debt is to get our spending under control. For Technical Debt, [the analog of] spending is when we increase debt by adding new code without any comments, or making a change that increases the complexity of a function. You might not have time to correct all of the sins of the past, but let's not make things worse with what we do today.

Step 2 is to allocate a portion of your development resources to refactoring existing code. Refactoring is an essential part in the battle against technical debt, but is often postponed due to time constraints and an element of fear; all of us have fragile modules in our applications that we are afraid to change.

Conclusion

Returning to Cunningham's quote from 1992, "a little debt speeds development so long as it is paid back promptly". We don't need to eliminate Technical Debt, but we do need to identify and control it; ensuring that the debt we have is Deliberate and Prudent. Allowing Inadvertent and Reckless debt to accrue will lead to bankruptcy; the inability to enhance an application and respond to market changes.

The software applications that we are building today will be in service for a long time. The vendors that will dominate are those that can quickly and efficiently enhance these applications to meet changing customer needs.

In a follow-up article, **John Paliotta** *moves on to considering how we can "pay down" the technical debt.*

Find Technical Debt on EETsearch

PATRICK MANNION, EDN

Teardown: Misfit Shine 2 – design win for sub-threshold silicon

At DesignCon 2016 in Santa Clara, we put the Misfit Shine 2 Fitness + Sleep monitor under the knife, and exposed the latest in low-power technology, processes, and power management techniques.

The Misfit Shine 2 is thinner, stronger, and smarter than, well, the original Shine. As per the maker's advertisements, it really is beautiful. Clad in T6061 T6 anodized aircraft-grade aluminium with glass-reinforced polycarbonate, the designers wisely opted for a minimalist aesthetic with a capacitive touch user-interface (UI) and visual feedback via 12 LEDs on the face. Tap the face, and the LEDs indicate your daily progress toward goal, followed by the time. You can set which comes first.

The Shine 2 is 8-mm thin, 30.5 mm in diameter, and weighs 8.5 g (with battery installed). It's water resistant up to 50m and uses a 3-axis accelerometer and magnetometer to sense movement.

A vibration alarm provides status feedback and notifications/alerts, which is useful for phone calls and texts from the paired Apple (iOS 7 or above) or Android (Jelly Bean) phone whenever they're in a bag or in mute mode.

Connectivity to the phone is achieved using a Bluetooth 4.1 radio and motion sensing is accomplished using a three-axis accelerometer and magnetometer working together. Also, by downloading the Misfit Link app you can connect using IFTTT recipes to control your mobile device to play music, take a selfie, or tap the Shine to control one of a plethora of home con-

Figure 1. *It's sleek design and ability to go 6 months on a single CR2032 3V coin cell battery made the Misfit Shine 2 an EDN teardown target at DesignCon.*

nect services, called IFTTT "channels," such as a WeMo switch, Nest thermostat, or even Twitter feeds.

Aside from aesthetics and ease of use, the Shine 2 stands out because the company addressed the common issue of battery life by opting for a coin cell that can last up to six months. This gets around the annoyance of having to recharge all too frequently and hoping it's done before you leave the house. Having used the first iterations of the Fitbit for a period of time, this was definitely a drawback.

Anyway, let's get inside and look at how the sleek design, low power consumption, and high functionality were achieved -- and who won and lost in the move from the original Shine to the Shine 2. In the full version of this article (click link) you'll find details of the disassembly – but skipping onwards...

The Bluetooth 4.1 interface came courtesy of Dialog Semiconductor's Smart-Bond DA14581, an optimised version of the DA14580 SoC... [and] it's the Ambiq Micro Apollo MCU that forms the brains -- and heart -- of both the overall processing and power management. So much so, that on talking with Keith Odland, director of marketing at Ambiq, it became clear that the implementation of Apollo on the Shine 2 was a case study in intelligent power management. read on... *click link to full article*

Complete article, here

designideas

- Precision programmable current source uses two ICs
- Circuit remotely adjusts PSU voltage

Precision programmable current source uses two ICs

By Stefano Salvatori & Pietro Oliva

This Design Idea mates a precision current source IC with precision difference amp chips to create a programmable current source or sink.

The resistor-programmable current source/sink in Figure 1 illustrates the basic topology, taking advantage of tightly matched on-chip resistor ratios instead of relying on absolute tolerances.

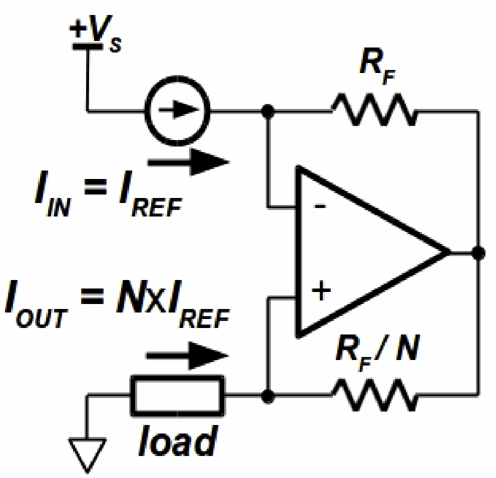

Figure 1. *I_{REF} can be increased or decreased to the load depending on N.*

The I_{IN} current source gives rise to a reference voltage across R_f at the op-amp inverting input. The same potential is found across the R_f /N resistor, so that an output current of N × I_{IN} will be generated. Despite a decreased compliance voltage due to the absence of direct connections to the op-amp inputs, the LT1991/5/6 is used as a single chip current divider.

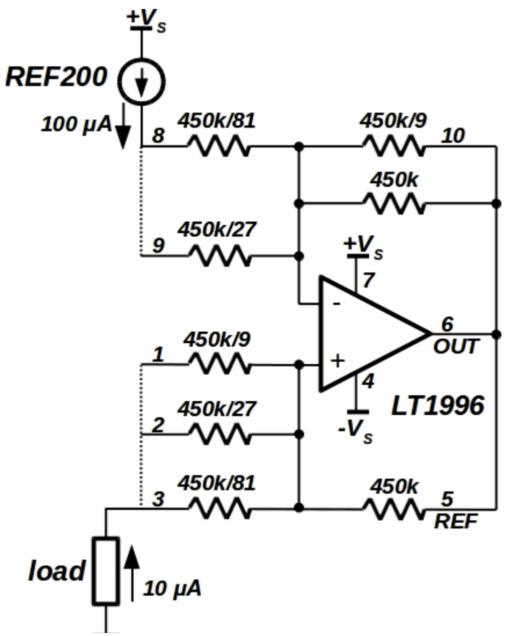

Figure 2. *LT1996 used as a pin-strap-programmable current sink/ source, in this case, dividing the reference current by 10.*

Figure 2 shows an example configuration, with a REF200 as the input current reference. Due to the high value of the internal feedback resistor connected to the op-amp inverting input (450 kΩ), a parallel connection with the 450/9 kΩ resistor is necessary to avoid op-amp output saturation induced by the injected I_{IN} current. The negative feedback resistance is thus equal to 450/10 kΩ, or 45 kΩ.

Given the 450 kΩ positive feedback resistor, N for Figure 2 is 0.1, yielding an output sink current of 10 µA. Different values of the output current can be obtained using the other available internal resistors. If an output source current is needed, reverse the connections to the REF200 IC, and connect it to $-V_S$.

The dotted lines in Figures 2 & 3 illustrate how unused internal resistors can be paralleled to reduce voltage drops in the reference and load paths.

Figure 3 shows a similar application based on the LT1995 chip and used to increase the output current, in this case, summing the currents from the two sections of the REF200 and multiplying by a factor of five.

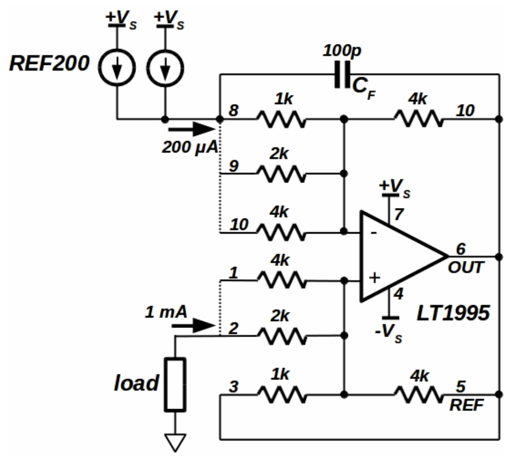

Figure 3. *LT1995 sinking five times the reference current*

The absence of an internal compensating capacitor in the LT1995 could induce instability. As ex-

I_{OUT} / I_{IN}	P1 pin 1	P2 pin 2	P4 pin 3	REF pin 5	C_F
2	-	OUT	load	-	33 pF
3	-	OUT	load	OUT	100 pF
4	-	load	OUT	-	33 pF
5	-	load	OUT	OUT	100 pF
6	load	OUT	OUT	-	100 pF
7	load	OUT	OUT	OUT	1 nF

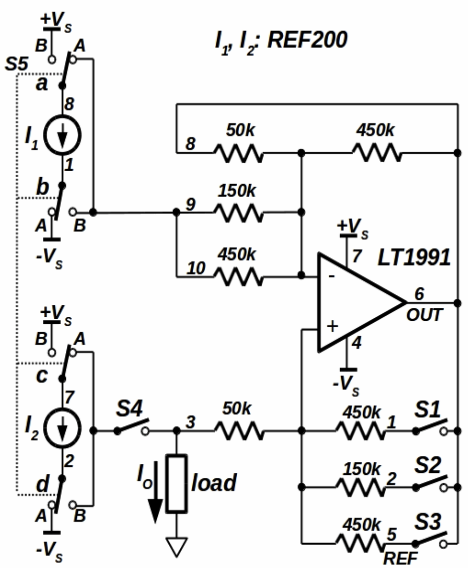

Figure 4. *An LT1991 is used to source/sink from 10-90 µA in 10 µA steps.*

perimentally verified, the circuit of Figure 3 will oscillate at a few megahertz if a 33-100 pF capacitor is not inserted as indicated.

Table 1 above summarises the connections for I_{OUT} / I_{IN} ratios between 2 and 7.

Figure 4 shows an LT1991-based programmable current source where internal resistors are ratioed 1:3:9.
The circuit can generate an output of 10 - 90 µA in steps of 10 µA. For currents up to 50 µA, the principle is the same as before. Above 50 µA, the second current reference in the REF200 is used to superimpose 100 µA onto the load to obtain the desired output current. Table 2 above illustrates switch settings for the various current levels.

I_{OUT} (µA)	S1	S2	S3	S4	S5	
					source	sink
10	ON	OFF	OFF	OFF	A	B
20	ON	OFF	ON	OFF	A	B
30	OFF	ON	OFF	OFF	A	B
40	ON	ON	OFF	OFF	A	B
50	ON	ON	ON	OFF	A	B
60	ON	ON	OFF	ON	B	A
70	OFF	ON	OFF	ON	B	A
80	ON	OFF	ON	ON	B	A
90	ON	OFF	OFF	ON	B	A

About the Authors

Stefano Salvatori *received in 1992 his Msc ("Laurea") in Electronic Engineering from the University of Rome "La Sapienza", discussing a thesis on the deposition and characterisation of amorphous silicon for Schottky Diode realisation. In 1998 he completed his PhD at the University "Roma Tre" defending a dissertation on CVD-diamond films on silicon for UV detection. He is currently Associate Professor with the Electronic Engineering Faculty at the University Niccolò Cusano of Rome, involved in the design, realisation and characterisation of CVD-diamond pixel-array detectors for X and UV radiation.*

Circuit remotely adjusts PSU voltage

By Chee Hua How

This Design Idea demonstrates a way to adjust the output voltage of a power supply, for OVP/UVP testing, margin-testing the load, voltage programmability, or any other reason.

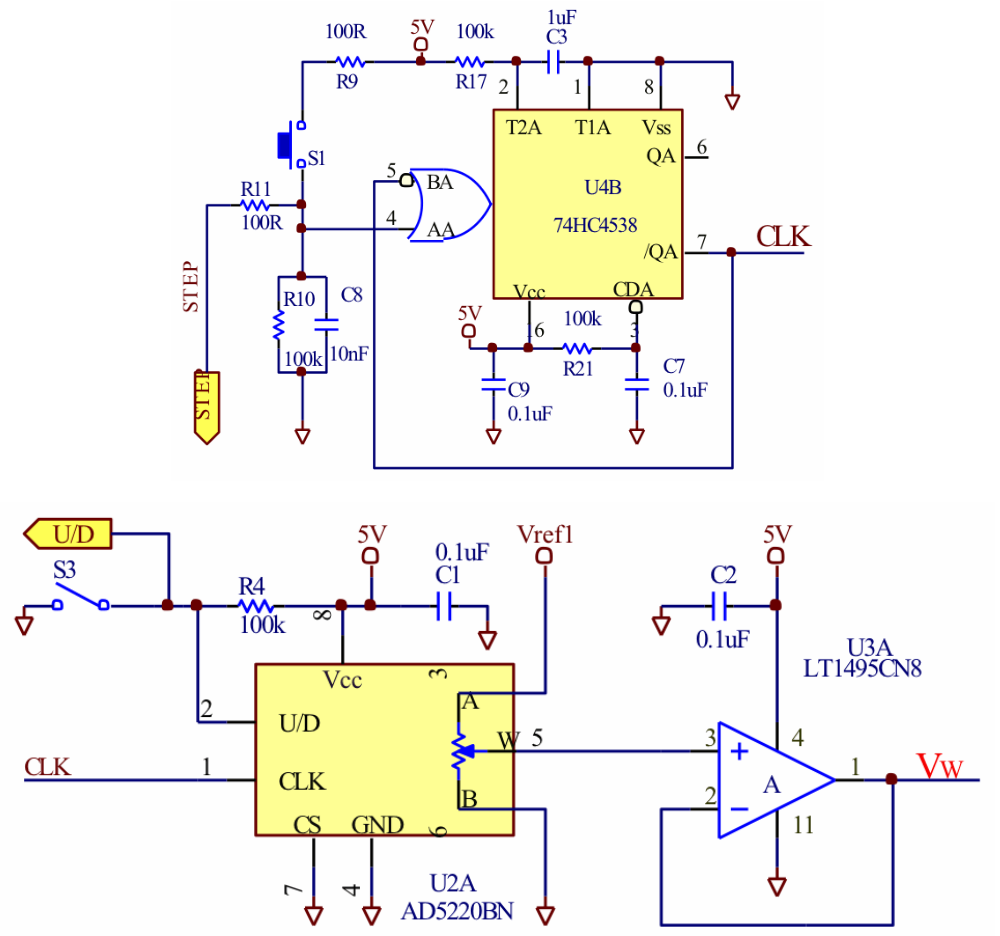

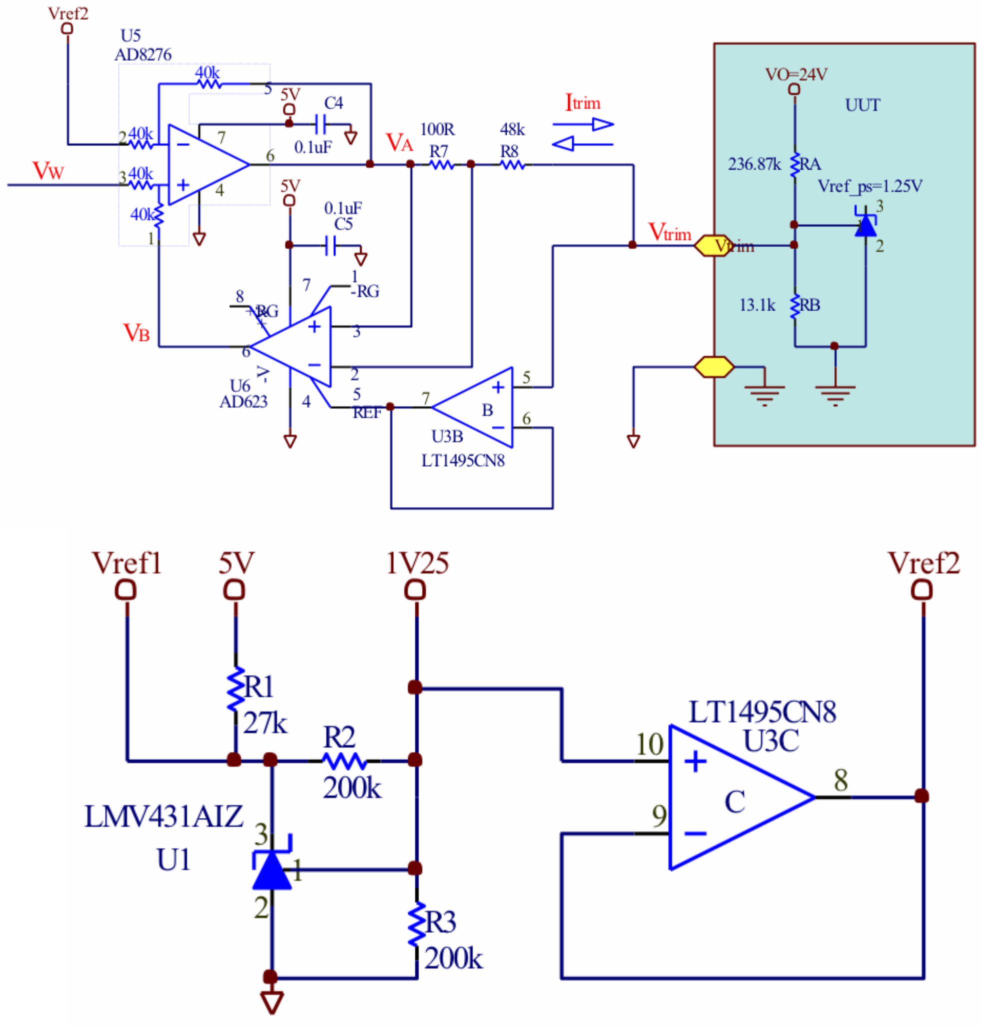

Figure 1. *Circuit for bidirectional trimming of PSU output voltage*

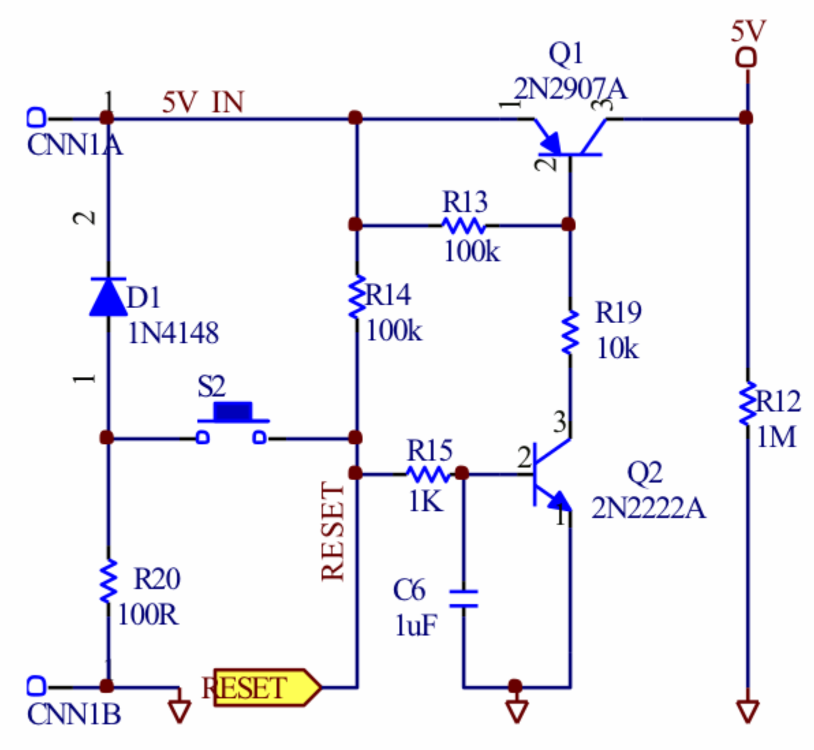

The circuit in Figure 1 is capable of bidirectional trimming of a power supply's output voltage by sourcing or sinking current into the feedback node. It can operate either manually with switches or digitally via three inputs: S1 (STEP), S2 (RESET), and S3 (U/D).

Every S1 rising edge increases or decreases V_O by one step (about 95mV in this design). S3 controls the adjustment direction (up/down), and S2 resets V_O to nominal.

One shot U4B ensures:
- One incremental step per press (debounces S1).
- Sufficient waiting period for the UUT's protection circuit to react.

The trim section comprising U5 & U6 features a voltage controlled current sink/source (VCCS). U3B & U3C shift up the VCCS's reference point so that:

- It's centred at the same level as V_{ref_ps}. Thus, in neutral condition (e.g., upon reset), $V_W = ½V_{ref1} \approx V_{ref2}$, $I_{trim} \approx 0$ & $V_O \approx V_{O(nominal)}$.
- The circuit can use a single supply while still capable of sinking and sourcing current.

Difference amplifier U5 generates I_{trim} to adjust V_O's magnitude, sourcing current into the V_{trim} node to lower V_O, or sinking current to raise V_O. U6 is a G=1 instrumentation amp that feeds the sensed current back to U5.
The U1/U3C circuit generates two reference voltages, V_{ref1} and V_{ref2}. V_{ref1} is a reference voltage for control signal V_w. V_{ref2} matches the UUT's 1.25V reference voltage.
Equations (1), (2), and (3) can be used to configure this circuit for different V_O settings:

$$V_{trim} = \frac{R_B}{(R_A+R_B)} \cdot V_O \mp I_{trim} \cdot \left(\frac{R_A \cdot R_B}{R_A+R_B}\right) \qquad (1)$$

$$I_{trim} = \frac{(V_W - V_{ref2})}{R_8} \qquad (2)$$

$$V_A = \left(\frac{R_7+R_8}{R_8}\right) \cdot (V_W - V_{ref2}) + V_{ref\_PS} \qquad (3)$$

The following example plugs Figure 1's parameters into the equations:
From (1) we have:

$$1.25 = V_O \cdot 52.4063 \cdot 10^{-3} \mp I_{trim} \cdot 12.4135 \cdot 10^3 \qquad (4)$$

So for V_O variations of ±25% (i.e., 30V to 18V), I_{trim} is in the range of -26µA to +25µA.
V_W is within 0V-2.5V. Putting this number into (2) gives R8≈48kΩ.

From (3):

$$V_A(V_W = 0) = \left[\left(\frac{R_7+R_8}{R_8}\right) \cdot (-1.25) + 1.25\right] \qquad (3(a))$$

$$V_A(V_W = 2.5) = \left[\left(\frac{R_7+R_8}{R_8}\right) \cdot (1.25) + 1.25\right] \qquad (3(b))$$

If R7=100Ω then $V_A(V_W=0) \approx 0V$ and $V_A(V_W=2.5V) \approx 2.5V$.
Given the 128 steps of U2, resolutions for V_W, I_{trim}, and V_O are 20 mV, 406 nA, and 95 mV respectively.

Figure 2 shows waveforms at a few key points. In the first phase, V_O (Ch3) is decreasing linearly from nominal with every clock pulse until it saturates around 18V. About half-way through, S2 is pressed to reset V_O to nominal and S3 is closed. V_O increases with each clock pulse to its upper limit of 29.5V.
Any mismatch between V_{ref2} and V_{ref_ps} introduces an offset on I_{trim} when U2's wiper is centred which shifts V_O from its nominal. This can be trimmed out if desired.
The circuit is powered by a 5V source, with current consumption under 2 mA. In some applications, V_O could be regulated and used for this.

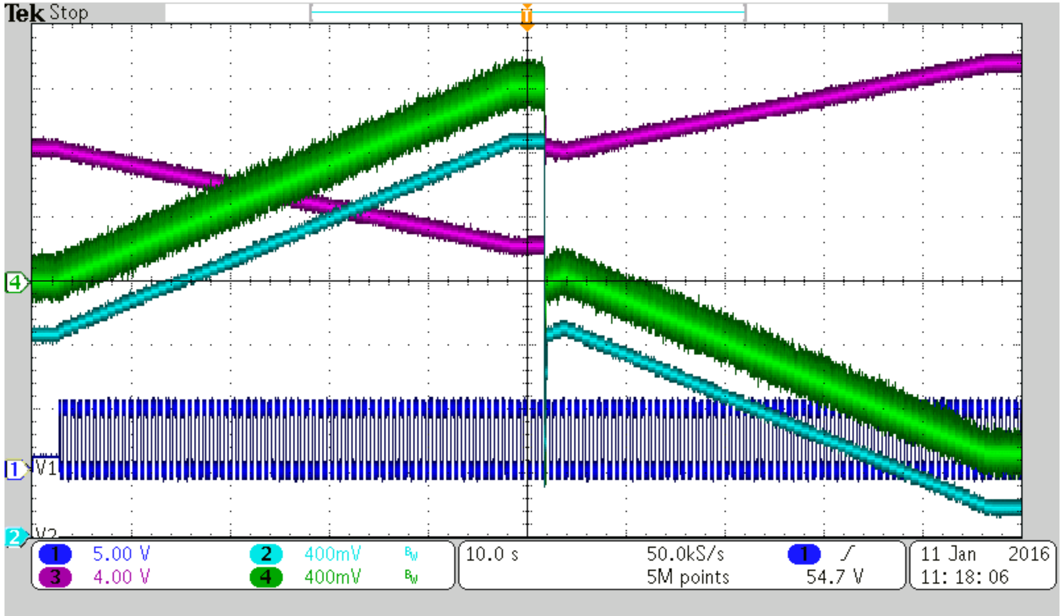

Figure 2. *Ch1:CLK; Ch2:VW; Ch3:VO; Ch4:V(R7+R8)*

About the Author

Chee H. How *is a Senior Power Electronics Engineer who graduated with a BSEE from Coventry University, UK and MSEE from Mutimedia University, Malaysia. He has more than 18 years of working experience in designing analogue & power electronics circuits and debugging production mysterious circuit failures. He can be reached at cheehhow@yahoo.com*

productroundup

productroundup

Add Ethernet ports to single port MCUs

Segger's TCP/IP stack now offers support for the "Tail Tagging" feature of Micrel/Microchip Ethernet switches. embOS/IP is the first embedded IP stack to support Tail Tagging. This enhancement establishes multiple virtual Ethernet ports when only one physical Ethernet port is available on the CPU - by choosing another PHY. Port addressing is done on a pure software basis and is transparent to the outside. The new feature allows every port to have its own assigned MAC-address so that they appear like different physical hosts in a network.

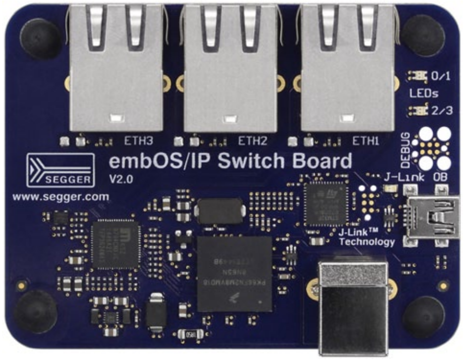

Complete article, here

High-voltage Li-ion batteries connect via 100-V high-side FET driver

bq76200 is presented as the first single-chip 100-V high-side FET driver for high-power lithium-ion battery applications, delivering advanced power protection and control. The high-voltage solution efficiently drives high-side N-channel charge and discharge FETs in batteries commonly used in energy storage systems and motor-driven applications, including drones, power tools, and e-bikes. Compared to typical 50-V low-side FET driver solutions, the 100-V high-side FET driver provides greater protection against possible inductive transient events in motor-driven applications.

Complete article, here

Explore digital power: development kit

Würth Elektronik (Waldenburg, Germany) and Infineon are launching the jointly developed "XMC Digital Power Explorer" evaluation kit. This synchronous step-down converter can be assembled with two different control cards (XMC1300 - ARM Cortex-M0 MCU and XMC4200 - ARM Cortex-M4F MCU). The kit is intended to assist developers of analogue power supplies, and embedded software programmers, to enter the world of digital power supply. The XMC Digital Power Explorer Kit is a complete solution with hardware, software and switchable resistance load bank.

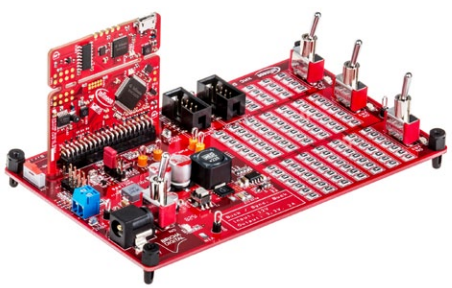

Complete article, here

STM32 MCU variant boosts graphics power

STMicroelectronics says it is bringing the power of the ARM Cortex-M7 to a broader range of applications with added graphics-centric STM32 microcontrollers; these devices are configured with up to 2 MB flash and 512 kB RAM. Latest members of the STM32 MCU family feature a 216MHz/462DMIPS/1082 EEMBC CoreMark Cortex-M7 core with double-precision floating-point unit and DSP instructions. Integrated alongside the core are up to 2 MB of dual-bank flash, ST's Chrom-ART Accelerator for graphics performance, a hardware JPEG accelerator, TFT-LCD controller, and MIPI-DSI host controller.

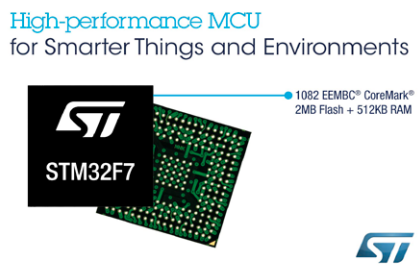

Complete article, here

productroundup

100V MOSFET H-bridge in a 22.5 mm² package, drive array-scale loads

DMHC10H170SFJ from Diodes Inc. is a 100V full H-bridge incorporating dual N-channel and P-channel MOSFETs into a DFN5045 package (5 x 4.5 mm). The device has a 100V drain-source breakdown voltage (BVDSS) providing sufficient headroom to support 48V telecom rails and industrial applications, while having a 5V gate voltage to simplify designs with a direct logic level interface to MCUs. An 11A peak pulse current rating also means the device can handle the in-rush current from energising a coil, which is usually greater than five times the typical operating current of a DC motor.

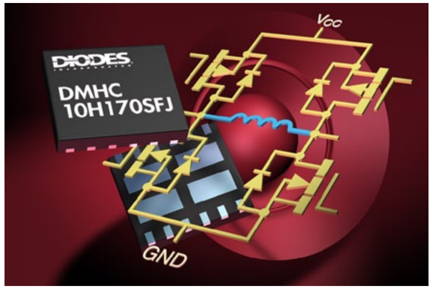

Complete article, here

Plastic fingerprint sensor detects underlying veins

FlexEnable and ISORG have jointly developed what they claim to be the world's first large area flexible fingerprint sensor on plastic designed for biometric applications. With an active area of 86 x 86 mm, a 84 μm pitch (78 μm pixel size with 6 μm spacing) and a 1048576 pixel resolution (1024 x 1024), the flexible sensor is only 0.3 mm thick and can operate in visible and near infra-red up to wavelengths of 900 nm. The technology is capable of measuring not only the fingerprint, but also the configuration of veins in the fingers, providing additional security versus that of a surface fingerprint alone.

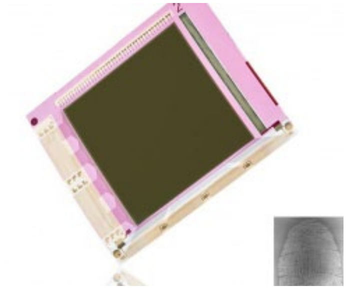

Complete article, here

Low power, inertial, pedestrian navigation IC

PNI Sensor has announced SENtrace, the first coprocessor for wearables providing accurate and ultra-low power pedestrian tracking indoors, as well as in urban canyons - anywhere that the global positioning systems (GPS) signal goes missing or is inadequate.SENtrace is a small custom ASIC that uses about 10x lower power than GPS demands. By building on PNI Sensor's proprietary embedded algorithms, it enables existing ultra-low power inertial sensors to track users when there is little or no GPS signal available. The chip provides tracking to one-metre accuracy over 100 metres travelled, supplying step-by-step data in lieu of extrapolations between two location points.

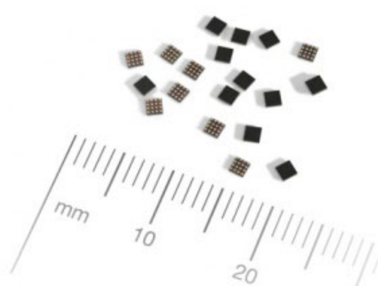

Complete article, here

Stepper motor drivers extend "AutoTune" setup

AutoTune technology, Texas Instruments asserts, dynamically tunes motors, and integrated current sensing, while saving 20% board space. TI's high-performance stepper motor-driver family gains three new devices for 24-V stepper motors. Two devices offer the AutoTune technology to eliminate stepper motor tuning; and two integrate current sensing to support differentiated 3D printers, robotics, factory automation equipment, currency-counting machines and more. AutoTune technology on the DRV8880 and DRV8881 devices eliminates the time-consuming, iterative process of manual tuning.

Complete article, here

productroundup

Pulse oximeter/heart rate sensor for wrist-bands

Providing an ultra-low power solution for wearable health and fitness applications, the MAX30102 pulse oximeter and heart rate integrated sensor module integrates red and IR LEDs to modulate LED pulses for oxygen saturation (SpO2) and heart rate measurements. It operates on a single 1.8V power supply and a separate 5V power supply for the internal LEDs. The device maintains a very small solution size without sacrificing optical or electrical performance; it integrates internal LEDs, photodetectors, optical elements, and low-noise electronics with ambient light rejection.

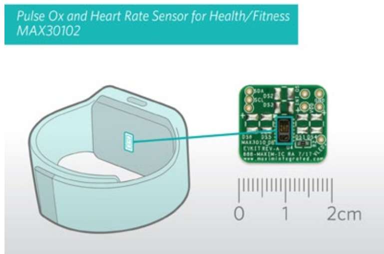

Complete article, here

LED lighting driver IC supports selection of circuit topologies

Power Integrations' LYTSwitch-5 single-stage LED driver IC family devices combine PFC and constant-current output, and support multiple common LED driver topologies. The flexible platform delivers high PF, low THD and highly accurate constant-current regulation. The LYTSwitch-5 family supports the safety-rated isolated flyback topology for ballast applications, and non-isolated topologies such as buck and buck-boost for low-cost lamps and bulbs, at over 90% efficiency, 0.9 PF (power factor), less than 5% THD and ±3% current regulation accuracy over line, load, production variance, temperature and multiple topologies.

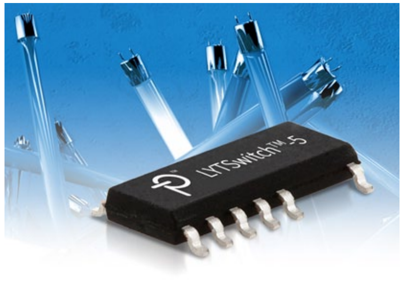

Complete article, here

Temperature & humidity sensor development demo

Distributor Mouser Electronics has the SHT31 Smart Gadget development kit from Sensirion, a reference design circuit board and development kit that demonstrates Sensirion's SHT3x series humidity and temperature sensors. Sensirion's Smart Gadget shows relative humidity and temperature values on the display, and can also communicate wirelessly with a Bluetooth Smart capable device, such as a smartphone or tablet. The Smart Gadget includes an onboard SHT31 humidity and temperature sensor, LCD driver and display, and Bluetooth module.

Complete article, here

Digital capture system adds to LeCroy scopes' analysis capabilities

An introduction by Teledyne LeCroy, the HDA125 High-speed Digital Analyser/ digital acquisition system, captures 18 channels of digital data at 12.5 Gsample/sec. The digital acquisition system interfaces to the company's oscilloscopes by means of the LBUS interface architecture to give what is claimed as the most flexible, highest-performance mixed-signal solution available. Also inroduced is the QuickLink probing system, featuring low-cost, high-fidelity probe tips that work with both the HDA125 and Teledyne LeCroy WaveLink probes using the Dxx30-QL QuickLink adapter.

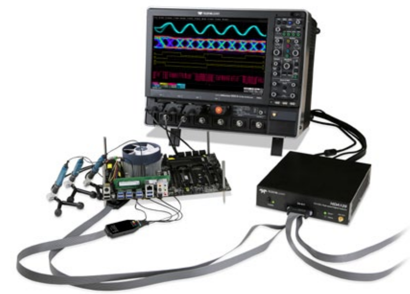

Complete article, here

productroundup

WiFi power amp delivers +22-dBm output

Optimised for 802.11a/n/ac WiFi infrastructure, the Anadigics AWL5911 5-GHz power amplifier provides 33 dB of linear power gain and low 1.8% EVM (error vector magnitude) with a 5-V supply and output power of +22 dBm. An integrated power detector permits accurate ±0.5-dB power control over varying load conditions (3:1 VSWR). With 50-Ω internally matched RF ports, the AWL5911 comes in a 4 × 4 × 0.80-mm surface-mount package. The device's high level of integration reduces component cost and saves board space for access points, media gateways, set-top boxes, and smart TVs.

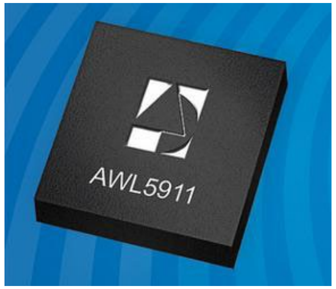

Complete article, here

Renesas configures MCU family for industrial/consumer

The RX230 group of microcontrollers extends the RX231 group of 32-bit MCUs, for cost-sensitive applications, combining digital signal processing (DSP) and floating point unit (FPU), and low power consumption. With 1.8V to 5.5V supply, the RX230 Group is aimed at applications that require high processing performance in environments with low current supply capacity, as well as the robustness needed in home appliance applications. RX230 is optimised with regard to embedded flash size and feature, for cost-sensitive industrial and consumer applications that do not require USB nor CAN interfaces, can fit their code in 256 kB flash or less, and need a maximum of 32 kB embedded RAM.

Complete article, here

NI updates its VirtualBench all-in-one instrument

National Instruments' latest version of VirtualBench – a software-based all-in-one instrument that combines a mixed-signal oscilloscope, function generator, digital multimeter, programmable DC power supply and digital I/O – now has 350 MHz of bandwidth, four analogue channels and Ethernet connectivity. An enhanced mixed-signal oscilloscope with protocol analysis delivers 350 MHz of bandwidth and four analogue channels for higher performance interactive test. A higher wattage programmable DC power supply outputs up to 3A for the 6V output channel and up to 1A for the 25V and -25V channels.

Complete article, here

PLL with VCO synthesiser ranges to 6.8GHz

Analog Devices has added a phased-locked loop (PLL) synthesiser with integrated voltage-controlled oscillator (VCO); the ADF4355 PLL with VCO synthesiser operates up to 6.8 GHz, a frequency band that allows significant margin to industry's current carrier frequencies. The integrated VCO has an output frequency ranging from 3.4 to 6.8 GHz. The VCO frequency is connected to divide-by-1, -2, -4, -8, -16, -32, or -64 circuits, which allows the designer to generate RF output frequencies as low as 54 MHz, thereby allowing the ADF4355 to cover a continuous range from 54 to 6800 MHz with no gaps. For applications that require noise isolation, the RF output stage features a mute function that is both pin and software controllable. The ADF4355 can operate in either fractional-N or integer-N mode, and with 38-bit modulus resolution offers exact mode frequency operation, independent of reference input frequency.

Complete article, here

DON'T SHIP MERELY-FUNCTIONAL SYSTEMS

BY JACOB BENINGO

Over the course of the last few years, I have noticed a trend among embedded system developers and teams that is quite disturbing. The trend consists of developing embedded systems that are functional (at best), but are not built or tested for a production environment. This trend leads towards disaster.

The primary cause of this "simply functional" trend appears to be due to three factors: making use of example code; a rushed development cycle; and a lack of understanding what it takes to build a production embedded system.

The first factor, building on example code, is actually a critical step toward jump-starting embedded software development. Example code helps in getting an embedded system up and running as well as in gaining important insights into the target hardware. Many microcontroller suppliers provide developers with much needed sample code [showing] how to setup peripherals and interact with the microcontroller.
But there are two issues that many developers often don't consider about this example code. First, example code is just an example; it is not intended for production. It is simply a guide on how to setup and interact with various peripherals. Yet developers will adopt the code, and once example code is brought into the system it usually stays in the system.
A close examination of example code from different microcontroller vendors often reveals a disclaimer that the provided code is not guaranteed fit for any purpose. It's not even warranted for any purpose but is simply provided "as is." Just reading the disclaimer should make embedded software developers squirm in their seats with discomfort when thinking about adopting that code. The software's producer isn't confident enough to stand behind their example so what makes anyone think that the example code product is production-ready?
A great example of functional sample code can often be seen when checking a hardware register flag. Figure 1 shows something similar to what one will usually find.

One problem with the code in Figure 1 is that the while loop assumes the operation will eventually complete successfully. Under ideal conditions this may be true, but what happens if hardware fails? Maybe an oscillator is drifting so that synchronisation cannot be achieved. Perhaps a write to flash fails. The hardware check could be on a communication

```
volatile uint8_t HARDWARE_FLAG  = 0;

/* Perform a hardware operation that requires monitoring a hardware flag */
OPERATION;

/* Wait for the ready flag to be set indicating that the operation is complete */
while(HARDWARE_FLAG != 1)
{
    HARDWARE_FLAG  = READ_HARDWARE_REGISTER_FLAG;
}
```

Figure 1. *Example code hardware register flag check*

bus flag when a faulty external sensor that has gone rogue has brought down the bus, making it impossible to complete a transmission. In these cases, the result of using the code in Figure 1 will be an infinite loop that will require intervention by an external force such as the watchdog timer. Even then, the watchdog timer would reset the system but there is no guarantee that the system wouldn't end up back in the loop, entering a perpetual cycle of constant resetting.

Software written for a production environment should accommodate the possibility of a failure. Some solutions to a scenario such as the while loop in Figure 1 might be to add a time-out to the loop based on the system tick or maybe establish a maximum number flag checks. These would prevent the system from entering infinite loops or perpetual reset cycles.

The example in Figure 2 demonstrates how additional conditions might be added to the while loop so that the system will exit the loop in the event of a failure. Rather than having the system hang in an infinite loop waiting for rescue, these additions generate error code that alerts the calling routine that the hardware flag of interest has timed-out. The system can then take corrective actions without invoking the last-resort watchdog.

The second factor leading to the trend of

```
/** Defines the number of milliseconds before a time-out occurs */
#define TIME_OUT      10

volatile uint8_t HARDWARE_FLAG  = 0;
uint32_t TickStart = 0;
uint32_t TickElapsed = 0;

/*
 * Get the start time from the system tick to determine whether the
 * loop has timed out.
 */
TickStart = SysTick_Get();

/*
 * Wait for the ready flag to be set indicating that the operation is complete
 * The operation should never time-out unless a failure has occurred.
 */
while((HARDWARE_FLAG != 1) &&
      (TickElapsed < TIME_OUT))
{
   HARDWARE_FLAG  = READ_HARDWARE_REGISTER_FLAG;

   TickElapsed = SysTick_Get() - TickStart;

   if(TickElapsed > TIME_OUT)
   {
      /* Set error code before exiting the loop */
   }
}
```

Figure 2. *Production code hardware register flag check*

building functional but not production-intent embedded systems is the rushed development cycle. Developing an embedded system can incur significant overhead costs to businesses, which makes the business want to get to market yesterday. Also, start-ups, small businesses, and sales teams are notorious for optimistically setting a production date without examining the real effort required to develop a robust and production-ready system. Many engineers either refuse to stand up to management in these circumstances: or if they do, they find their concerns falling on deaf ears. The end result is that corners get cut in an attempt to meet an unrealistic deadline, which results in a design containing merely-functional code that only works over a range of very controlled conditions.

The final factor that is contributing to the release of functional and not production-intent embedded systems is a lack of understanding about how to build a production-intent embedded system. Embedded software and systems engineers are in high demand and short supply. This situation has resulted in companies filling critical roles with either students right out of school or engineers from different disciplines, such as web or app development. The result is a knowledge gap in how to properly architect and implement robust embedded systems that don't require daily updates to patch bugs and fix security issues.
Green [inexperienced] and cross-discipline engineers aren't the entire story, though, leading to a lack of understanding what a production embedded system really is. Well-disciplined and experienced engineers will often be asked to develop a prototype or proof of concept. For demonstration to management, the engineers present a beautiful functional prototype based on functional, example code. The demonstration goes great, but that system is only working under controlled conditions. But because the demonstration went well, management wants to ship the system immediately, not understanding there is still a lot of work needed to make the system production ready.

Embedded systems are finding their way into every nook and cranny of our lives. The use of merely functional code may be fine for some devices operating under controlled conditions. But with the IoT and rapid progress towards an autonomous and smart society, the dangerous trend of shipping functional instead of production code is an accident waiting to happen.

Are you building production-intent systems or a functional house of cards teetering on the edge of collapse?

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

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
Copyright 2016 by European Business Press SA.
All rights reserved. P 304128

CONTACTS

PUBLISHER

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

EDITOR-IN-CHIEF

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

Patrick Mannion
Brand Director EDN Worldwide

CIRCULATION & FINANCE

Luc Desimpel
luc.desimpel@eetimes.be

ADVERTISING PRODUCTION & REPRINTS

Lydia Gijsegom
lydia.gijsegom@eetimes.be

ART MANAGER

Jean-Paul Speliers

ACCOUNTING

Ricardo Pinto Ferreira

european business press

SALES CONTACTS

Europe

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

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

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

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

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

Andrea Rancati
Italy
+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
MIshida@mx.itmedia.co.jp

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

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