Electronic Specifier
design
July 2013 / Volume 3, Issue 7
www.electronicspecifier.com

A new world of treatment

How implantable pulse generators are stimulating the medical industry

In this issue:

contents


26
Amantys Power Insight in operation
QUALITY ASSURED
amantys
Power simplified, by design

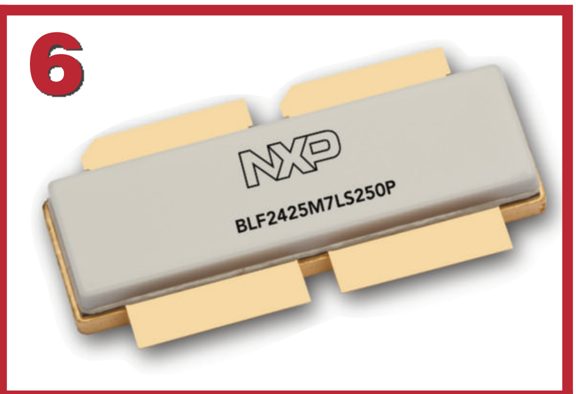
6
NXP
BLF2425M7LS250P

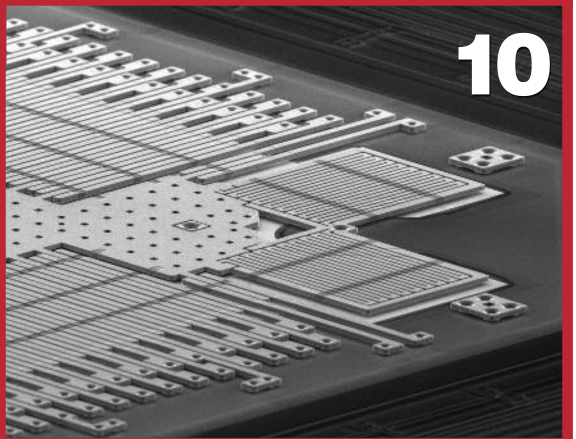
10


16


8
Tap Now

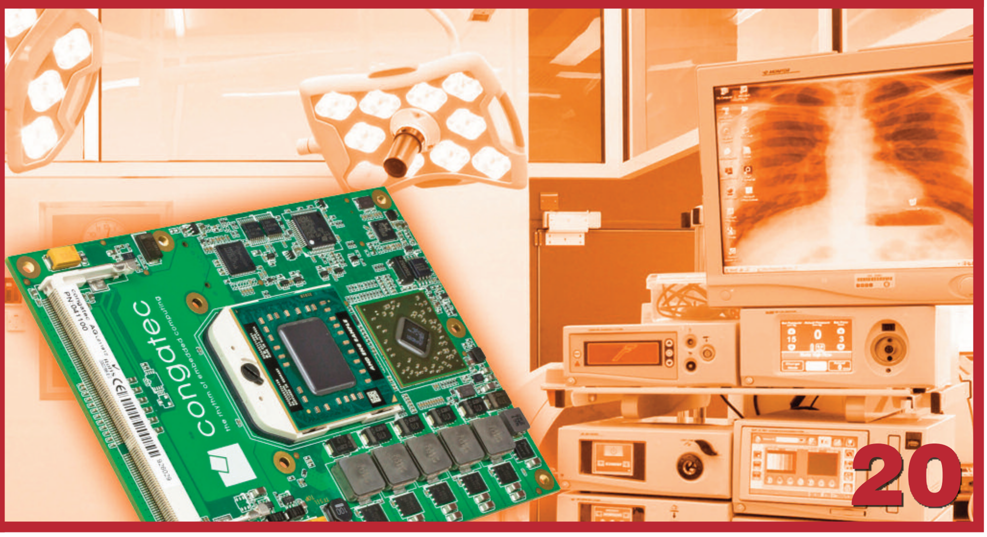
congatec
20

33


36
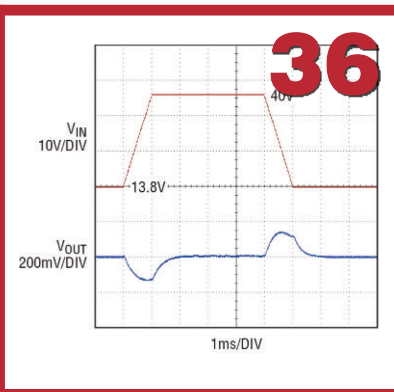

50


30


44
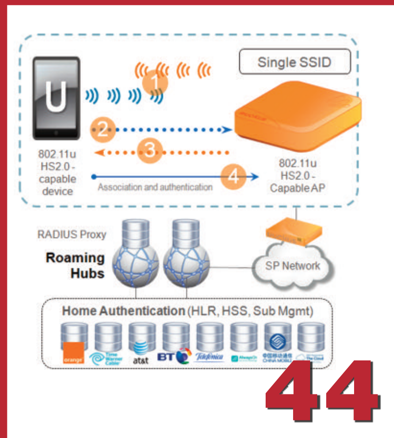

48
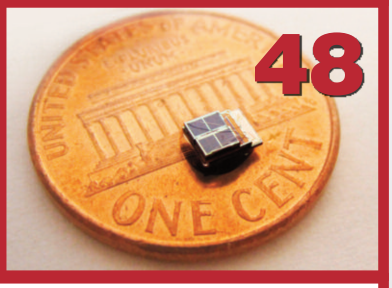

41
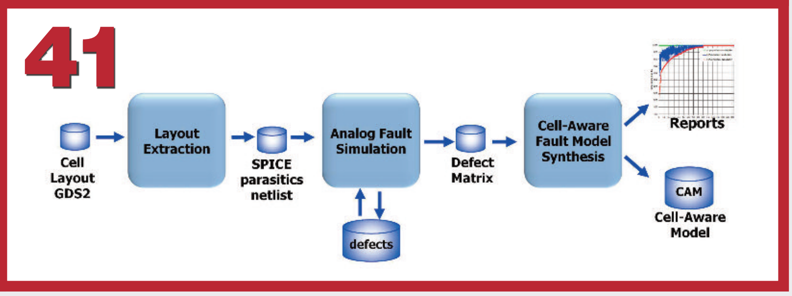

articles:

Heal thy self

It could still be described as 'burgeoning' but home-health treatment is undoubtedly picking up pace. What is apparent in this issue's Medical section is that electronics is also benefiting other areas of healthcare.

While home-health wil be enlisted to treat a growing number of patients with 'lifestyle' illnesses, there will always be a need for healthcare professionals to treat more serious conditions.

Technology is helping here, too, in the form of implantable devices and MEMS technology.

When coupled with self-diagnosis through connected technologies – or at least self-test and administer – the indications are that the patient could soon be their own primary health provider, or something close to it. Perhaps the only thing now missing is some form of augmented-reality assistance. 'Please state the nature of the medical emergency.'

Philip Ling, Editor

Electronic Specifier

ElectronicSpecifier Ltd.
Comice Place
Woodfalls Farm
Gravelly Ways
Laddingford
Kent ME18 6DA UK

Editor: Philip Ling
Email: phil.ling@electronicspecifier.com
Tel: +44(0)1622 679 221

Contributing Editors:
Sally Ward-Foxton
Vanessa Knivett

Advertising: Tim Anstee
Email: tim.anstee@electronicspecifier.com
Tel: +44(0)1622 871 944
Cell: +44(0)7771 628 990

Publisher: Steve Regnier
Email: steve.regnier@electronicspecifier.com
Tel: +44(0)1622 872 204
Fax: +44(0) 1622 870 279

Subscriptions
Subscribe to the Digital Edition of ElectronicSpecifier

Distribution Blog
Find out what's happening in the world of distribution

Vodafone targets M2M

The mobile network giant has announced its first all-in-one M2M package which will combine hardware, wireless network connectivity and application software. Vodafone's Remote Monitoring & Control Service (RMCS) aims to make it easier for companies to implement M2M.

According to Erik Brenneis, Director of M2M, Vodafone, the RMCS will provide a solution that can be implemented on a global scale, setting it apart from its competition.

Vodafone has also launched an annual 'M2M adoption barometer' report, which shows 50% of companies surveyed will have adopted M2M by 2015. The report identifies three important trends: the falling cost of M2M will accelerate uptake; smaller companies will overtake larger companies in their adoption of M2M, and; the manufacturing and consumer electronics market will grow faster than any other sector. The Asia Pacific region is also predicted to overtake the Americas in its M2M deployment in the next two years. Brenneis said: "The next two years are going to be a real game changer for the industry."

The report can be freely downloaded from: http://m2m.vodafone.com/barometer2013

Ada on ARM

Support for the Ada programming language running on an ARM processor is growing, as Green Hills Software introduces the AdaMULTI IDE for ARM, and Vector Software announced support for the AdaCore GNAT Pro compiler on ARM's Cortex family.

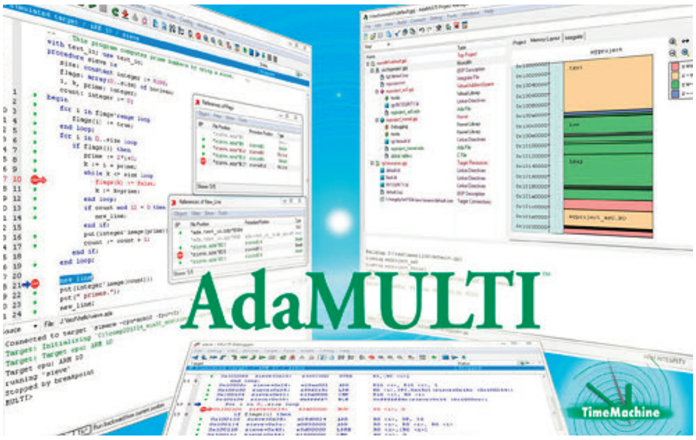

Ada is used widely in embedded systems where real-time and safety-critical operation are paramount. Both software tools vendors provide solutions for developing and testing safety-critical and secure software. The latest additions to their product portfolios recognise the growing presence of Ada being used instead of — or alongside — C/C++ languages.

Auto Alliance

Semiconductor manufacturers are moving deeper in to the automotive market, following a spate of collaborations between IDMs and car manufacturers.

Analog Devices has announced it is to collaborate with the Audi Group, while STMicroelectronics has established a joint laboratory with China's Great Wall Motor for the development of a new generation of automobiles.

With the semiconductor content in cars increasingly rapidly, the ripples of this mega-trend are now influencing many IDMs; Lattice Semiconductor has recently introduced automotive-grade versions of its ECP3 family of FPGAs which are qualified to AEC-Q100, while stating that future devices will be automotive-qualified at launch. Freescale, which has a well established automotive product range, has just released details of an integrated ADAS System for automotive radar solutions.

Solid state energy

Power delivered using radio frequencies could be the 'next big thing', according to NXP. It recently demonstrated a wide range of RF-powered applications at the IEEE International Microwave Symposium, using power transistors designed specifically for the 2.45GHz ISM frequency band.

They are claimed to allow RF energy to be used as a clean and highly efficient heat source and emphasised the point be demonstrating an RF spark plug, an RF plasma lamp and an solid-state cooker.

Engineers exploring RF power over the last 10 years have had to be content with 'brute force magnetrons' with 'next to no control', according to NXP's Mark Murphy, Director of Marketing, RF Power product line. This new range of solid-state RF power transistors have been optimised for 2.45GHz with a options for different power levels, control and cost.

We're now stocking over 22,000 TE Connectivity products.
TE connectivity
Authorized Distributor
Gold Cap
2.3V 3.3F
FIND IT.
DESIGN IT.
BUY IT.
Search TE Connectivity's extensive range by application at
www.rs-components.com/teconnectivity
RS

SoC for Multiband RF

A silicon-on-insulator process from STMicroelectronics, H9SOI_FEM, is claimed to allow RF front-end modules to cover more frequency bands with fewer components. Importantly, it could see the overall size of wireless devices being reduced even further.

Flavio Benetti, General Manage of the Mixed Process Division at ST, claimed: "The H9SOI_FEM dedicated process enables our customers to develop state of the art front-end modules that are half the size — or smaller — compared to today's front-end solutions."

The process is a 0.13um, dual-gate 1.2V and 2.5V MOSFET SOI process which supports multiple technologies, allowing it be used for the monolithic integration of all key functions in an RF front-end including RF switches, low noise amplifiers, power amplifiers, diplexers and antenna tuning, as well as RF energy management functions.

Sharper edge

A touch-screen tablet that doubles as a chopping board has secured an paid-for internship for the inventor with Sharp Laboratories. Siobhán Andrews, currently in her final year studying Sustainable Product Design at Falmouth University, beat more than 70 other entrants in the competition.

Andrews explained: "It is basically a screen that is behind a chopping board, which is interactive." As such, it allows a user to access recipes on the internet while preparing food.

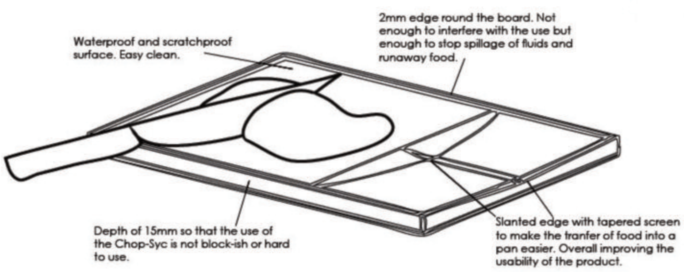

Energy My-cro

Silicon Labs has completed its acquisition of Energy Micro, with shareholders receiving an upfront payment of $115 million in cash and approximately $55m in deferred and earn-out consideration.

The strategic acquisition is expected to strengthen Silicon Labs' 32bit MCU and wireless SoC portfolio, adding nearly 250 devices which are expected to contribute around $7m in revenue to the second half of 2013.

Supply meets Demand

The UK's Electronics Technology Network (ETN) and technology trade association Intellect are to merge, potentially benefitting the ETN's SME membership and Intellect's relationship with large technology companies by better bridging the gap between supply and demand within the electronics industry.

In addition, Intellect's close relationship with Government is expected to help ETN members raise their international profile. Intellect will continue to operate the ETN as a free-to-join network.

NFC for pairing

With the proliferation of wireless-enabled devices, TI is hoping the complexity and inconvenience of pairing devices could present a market opportunity for NFC technology. It has launched a dynamic NFC transponder which combines an NFC and wired interface to connect a device to a host (classified as an NFC tag type 4 device); it also incorporates an ISO 14443B-compliant RF interface, allowing wireless access of NDEF messages.

Power plays

Following the recent announcement of Altera's acquisition of power management provider, Enpirion, Dialog Semiconductor has announced its intention to purchase another power management innovator, iWatt.

Dialog claims it will enhance its existing product portfolio in two high growth areas; AC/DC converters and LED solid state lighting. The acquisition will see iWatt receiving an up-front payment of approximately $310 million and will provide Dialog with around 110 patents based on iWatt's innovative power management technology.

The world's first digitally enhanced power analog controller

New analog-based power management controller with integrated microcontroller

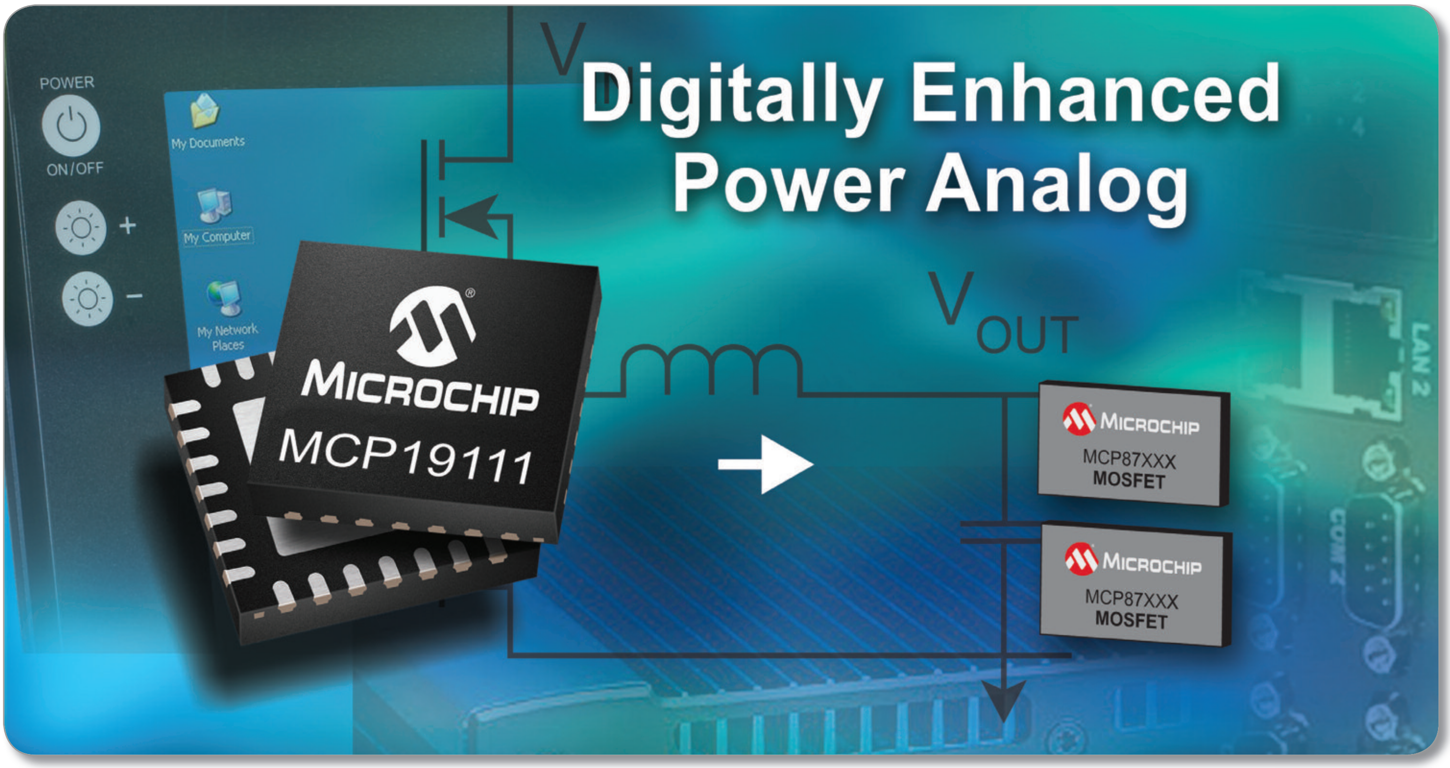

Combine the flexibility and I²C communication of digital DC/DC power conversion, with the speed, performance and resolution of analog-based control, by using Microchip's new MCP19111 power management controller.

The MCP19111 is a new hybrid, mixed-signal controller which combines analog and digital power management into a single chip. By integrating an analog-based PWM controller, a Flash-based 8-bit PIC® microcontroller, and MOSFET drivers for synchronous, step-down applications, the MCP19111 enables configurable, high-efficiency power conversion.

With transient performance of analog power conversion, the MCP19111 eliminates the need for a high-MIPS microcontroller or a high-speed A/D converter, minimising solution cost and power consumption.

To support even higher efficiency, the MCP19111 can be used to drive Microchip's latest MCP87xxx high-speed, SMPS-optimised MOSFETs. These logic-level 25V MOSFETs offer a low Figure of Merit (FOM) with on-state resistance from 1.8 mΩ to 13 mΩ to deliver higher DC/DC power conversion efficiency.

GET STARTED IN 3 EASY STEPS:

1. Evaluate the MCP19111 using the low-cost ADM00397 board
2. Combine with SMPS-optimised MCP87xxx MOSFETs
3. Customise DC/DC conversion to match your application

For more information, go to: www.microchip.com/get/eumcp19111

microchip DIRECT www.microchipdirect.com

MICROCHIP

Microcontrollers • Digital Signal Controllers • Analog • Memory • Wireless

Micro Motion

Understanding and utilising the fundamentals of MEMS motion sensors in medical devices. By Bob Scannell

The ability to detect, capture, and analyse motion with microelectromechanical systems (MEMS) has become a commonplace feature on consumer and mobile devices. Where technological advances have provided high-precision motion capture, applications have extended into industrial fields. Many potential medical diagnostic and instrumentation applications can benefit from merging the precision of industrial devices with the mobility and economies of consumer devices.

In some cases, the complexity of medical motion capture rivals that of high-end military systems. For instance, precision navigation, typically associated with applications developed for land, air, and sea vehicles, is increasingly being used in medical applications ranging from surgical instrumentation to robotics. Also, while the design requirements of a surgical navigation system share broad similarities with traditional vehicle navigation, there are distinct new challenges posed by the environment and the level of required performance.

This article first introduces some of the fundamentals of MEMS motion sensing, including key understandings needed for component selection. It also looks at the unique challenges of medical navigation applications, and explores possible solutions ranging from various sensor mechanisms, to necessary sensor processing, and the unique system characteristics and data processing required to provide optimal solutions. Critical sensor specifications will be reviewed and explained for their individual contributions and, more importantly, the potential error and drift mechanisms will be discussed to aid in sensor selection. Opportunities and approaches for sensor enhancement through integration, sensor fusion and sensor processing (such as Kalman filtering) will be highlighted as well.

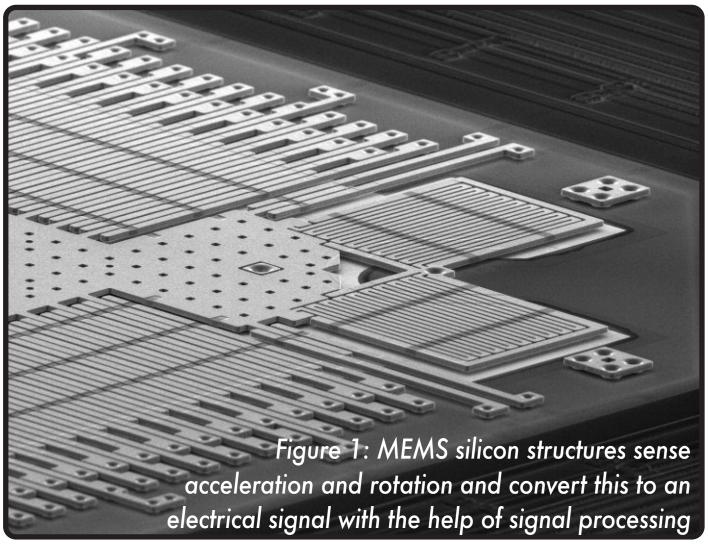

Figure 1: MEMS silicon structures sense acceleration and rotation and convert this to an electrical signal with the help of signal processing

Inertia

Silicon-based accelerometer and gyroscope sensors known as MEMS (Figure 1) are commonly found today in a wide range of devices. These inertial sensors detect and measure motion, with minimal power and size, and are valuable to nearly any application where movement is involved, and even those where lack of motion is critical. Table 1 outlines some of the basic pertinent medical applications by motion type. Later, more advanced applications where combinations of motion, in complex scenarios that present additional challenges, will be discussed.

Many medical applications such as accurately determining position and repetition rate in CPR, or the precise positioning of scanning equipment in relation to a patient's body, can benefit from relatively basic, yet still precise, motion information. In these cases, a single sensor type may be adequate, particularly if there are other sensor inputs, or at least fixed/known boundaries to the movement and use case.

Even with limited range of motion, or simpler motion dynamics, the individual sensors must have well understood and controlled drift factors, and it

SENSOR TYPE

APPLICATION	Acceleration/ Positon	Tilt	Angular Rate/Angle	Vibration	Shock	Sensor Fusion
	CPR Assist	Patient Down Monitors	Scanning Instruments	Tremor Control	High Value Equipment Warranty	Precision Surgical Navigation
	Activity Monitors	Bed-Patient Positioning/Aspiration	Basic Surgical Tools	Equipment Wear		Remote Diagnostics
	BioFeedback Monitors	Blood Pressure Monitors	Prosthetics			Rehabilitation
		Imaging Equipment				

Table 1: Inertial sensors accurately capture varied and complex motion to drive widespread medical applicability

is often desirable to have embedded compensation within the sensor, as well as the ability to tune it to the application via embedded filtering.

While simple motion detection – linear movement along one axis, for example – is valuable to a number of applications, such as detecting whether an elderly person has fallen, a majority of applications involve multiple types and multiple axes of motion. Being able to capture this complex, multi-dimensional motion can enable new benefits while maintaining accuracy in the most critical of environments.

In many cases, it is necessary to combine multiple sensor types – linear and rotational, for instance – in order to precisely determine the motion an object has experienced. As an example, accelerometers are sensitive to the Earth's gravity, so they can be used to determine inclination angle. As a MEMS accelerometer is rotated through a ±1-g field, (±90°), it is able to translate that motion into an angle representation. However, the accelerometer cannot distinguish static acceleration (gravity) from dynamic acceleration. In the later case, an accelerometer can be combined with a gyroscope, and post-processing of both devices can discern the linear acceleration from the tilt, based upon known motion dynamics models. This process of sensor fusion obviously becomes more complex as the system dynamics (number of axes of motion, types, and degrees of freedom of motion) increase.

It is also important to understand the environmental influences on sensor accuracy. Temperature is obviously a key concern, and can typically be corrected for; in fact higher precision pre-calibrated sensors will dynamically compensate themselves. A less obvious factor to consider is the potential for even slight vibrations to produce accuracy shifts in rotational rate sensors. These effects, known as linear acceleration and vibration rectification, can be significant depending on the quality of the gyroscope. Sensor fusion improves performance by using an accelerometer to detect linear acceleration and compensate for the gyroscope's linear acceleration sensitivity.

For many applications, particularly those requiring performance beyond basic pointing (up, down, left, right) or simple movement (in motion, or stationary), multiple degrees-of-freedom motion detection is required. For example, a six degree-

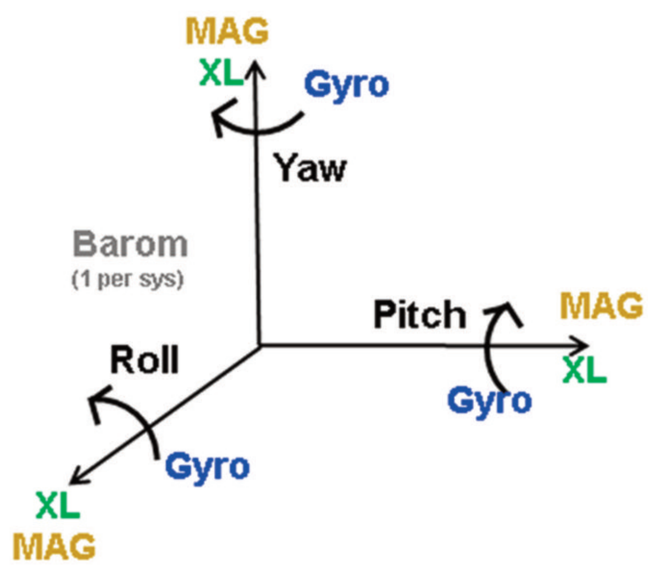

Figure 2: Linear X, Y and Z motion, plus rotational roll, pitch and yaw make up the six degrees of motion measurement required for full motion assessment; often augmented by both magnetometers and a barometer

Sensor Type	Major Advantage	Potential Limitations	Applicable to Medical Navigation?
GPS	Long Term Absolute Reference	Potential Blockages	No
Magnetic	No Required Infrastructure (except Earth)	Subject to Field Interference	Limited
Optical	Intuitive	Line of Sight Obstruction	Limited
Inertial	Self-Contained	Relative, not absolute reference	Yes

Table 2: An outline of various navigational sensors widely used in industry and their applicability to medical navigation

of-freedom inertial sensor has the ability to detect linear acceleration on each of three (X, Y, Z) axes, and rotational movement on the same three axis, also referred to as roll, pitch, and yaw; as depicted in Figure 2.

Basic navigation principles

The use of inertial sensors as a navigation aid has become prevalent in industry. Typically, they are used in conjunction with other navigation devices such as GPS. When GPS access is unreliable, inertial guidance fills the gap in coverage with what is called dead-reckoning. Other sensors, including optical and magnetic, may be added depending on the environment and the performance goals. Each sensor type has its own limitations. MEMS inertial sensors provide the potential to fully compensate for these other sensor inaccuracies since they are not affected by the same interferences and do not require external infrastructure: no satellite, magnetic field, or camera is needed; just inertia. The major navigational sensor approaches are outlined in Table 2, along with their strengths and potential limitations.

As with the potential for GPS blockage in vehicle navigation, the medical corollary is optical guidance and the potential for line-of-sight blockages. Inertial-based sensors perform dead-reckoning during the optical blockage, as well as enhancing system reliability by providing redundant sensing.

One medical application outlined earlier in Table 1 involves the use of inertial sensors in the operating room for more accurate alignment of artificial knee or hip joints with a patient's unique anatomical structure. The goal here is to improve joint alignment to less than 1° error from the patient's natural alignment axis, versus today's 3° or larger error using purely mechanical alignment approaches. Greater than 95% of total knee arthoplasty (TKA) procedures today are done with mechanical alignment. Computer assisted approaches using optical alignment have only slowly begun to replace some mechanical procedures, likely due to the equipment overhead required. Whether mechanical or optical alignment is used, approximately 30% of these procedures result in misalignment (defined as >3° error), which often leads to both discomfort and additional surgery.

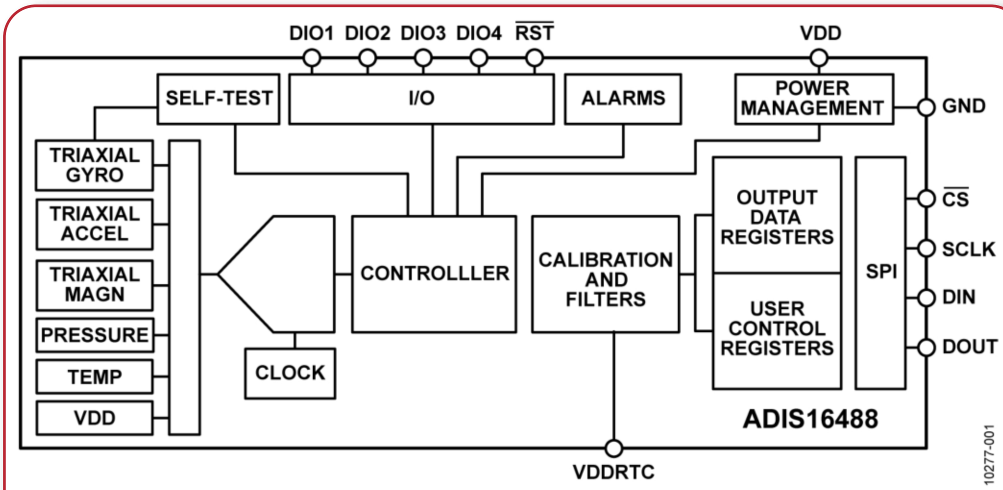

Figure 3: MEMS-based inertial measurement units provide precision six-degrees-of-motion measurement in compact form factors suitable to surgical instrumentation

Reducing misalignment has the potential of offering less invasive and shorter surgery time, increasing post-operative patient comfort, and producing longer lasting joint replacements. Inertial sensors in the form of a full multi-axis inertial measurement unit (IMU), as shown in Figure 3, have demonstrated substantially improved accuracy for TKA.

System Variable	Conditions/Considerations
Environment	Indoor/outdoor, temperature, shock/vibration, interference sources
Performance Rating/Goals	Accuracy, repeatability, speed, stability
Operator	Assisted or autonomous, trained or untrained
Safety	Life Critical, Inaccessible, Redundancy
Budget	Cost/Time to Implement, Risk

Table 3: Considerations in sensor selection

Sensor selection

There is a large variation in the performance levels of inertial sensors. Devices suitable for gaming are not able to address the high-performance navigation problem outlined here. The key MEMS specifications of interest are bias drift, vibration influence, sensitivity and noise. Precision industrial and medical navigation typically require performance levels that are an order of magnitude higher than is available from the MEMS sensors targeted for use in consumer devices. Table 3 outlines general system considerations, which through analysis can help focus the sensor selection.

Most systems will implement some form of Kalman filter to effectively merge multiple sensor types. The Kalman filter takes into account the system dynamics model, the relative sensor accuracies and other application specific control inputs to make the best determination of actual movement. Higher accuracy inertial sensors (low noise, low drift, and stability over temperature/time/vibration/supply-variance) reduce the complexity of the Kalman filter, the number of redundant sensors required and the number of limitations placed on allowable system operational scenarios.

The two primary challenges found in any high-performance motion capture implementation are the conversion of raw sensor data to calibrated and stable sensor data, and the translation of precision sensor data into actual position/tracking information. Overcoming the first hurdle involves integration of optimised sensor processing electronics coupled with motion calibration, which is based on intimate knowledge of motion dynamics. The second hurdle requires

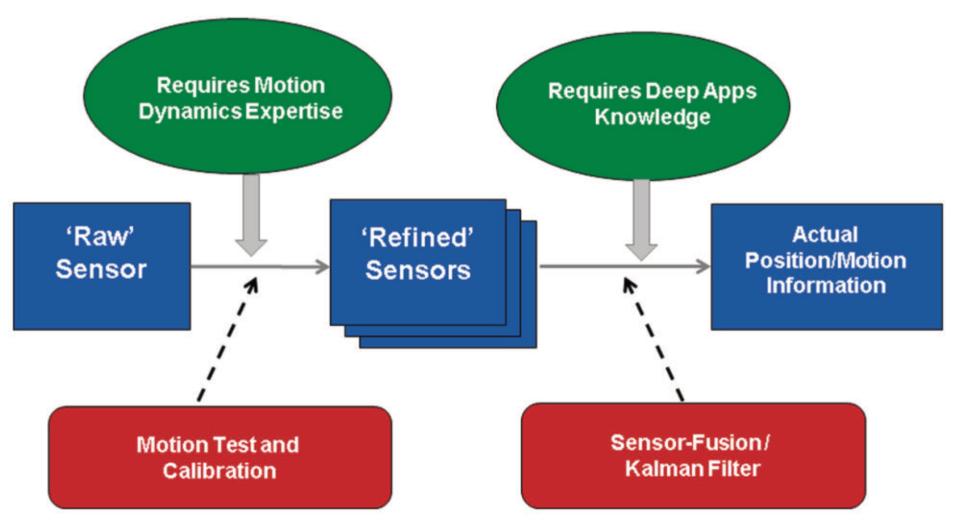

Figure 4: Precision motion detection begins with high-performance core sensors, coupled with optimised sensor processing and embedded application intelligence

merging an understanding of motion-dynamics with a deep knowledge of the peculiarities of the application at hand; as depicted in Figure 4.

MEMS inertial sensing is a highly mature technology in terms of both commercial viability and reliability. Beyond the well-known use cases in mobile devices and gaming, significantly more challenging needs exist in the medical and industrial fields. In these cases, substantially higher performance is required, along with much more complete integration and sensor processing. The complexity of motion involved in medical navigation, for instance, dictates the need for starting with highly stable inertial sensors as a foundation, then building on this with optimised integration, sensor processing and fusion.

The availability of highly accurate and environmentally robust sensor developments is driving a new surge in the adoption of MEMS inertial sensors within the medical field. These inertial MEMS devices are capable of offering advantages in precision, size, power, redundancy, and accessibility over existing measurement/sensing approaches.

Fortunately, many of the principles required for solving these next-generation medical challenges are based on proven approaches from classical industrial navigation problems, including sensor fusion and processing techniques.

Author profile: Bob Scannell is the Business Development Manager for MEMS Inertial Sensors at Analog Devices Inc.

More from ADI
Return to contents page.

A new world of treatment

How implantable pulse generators are helping to stimulate the medical device industry. By John Antalek & Mark Bartrum

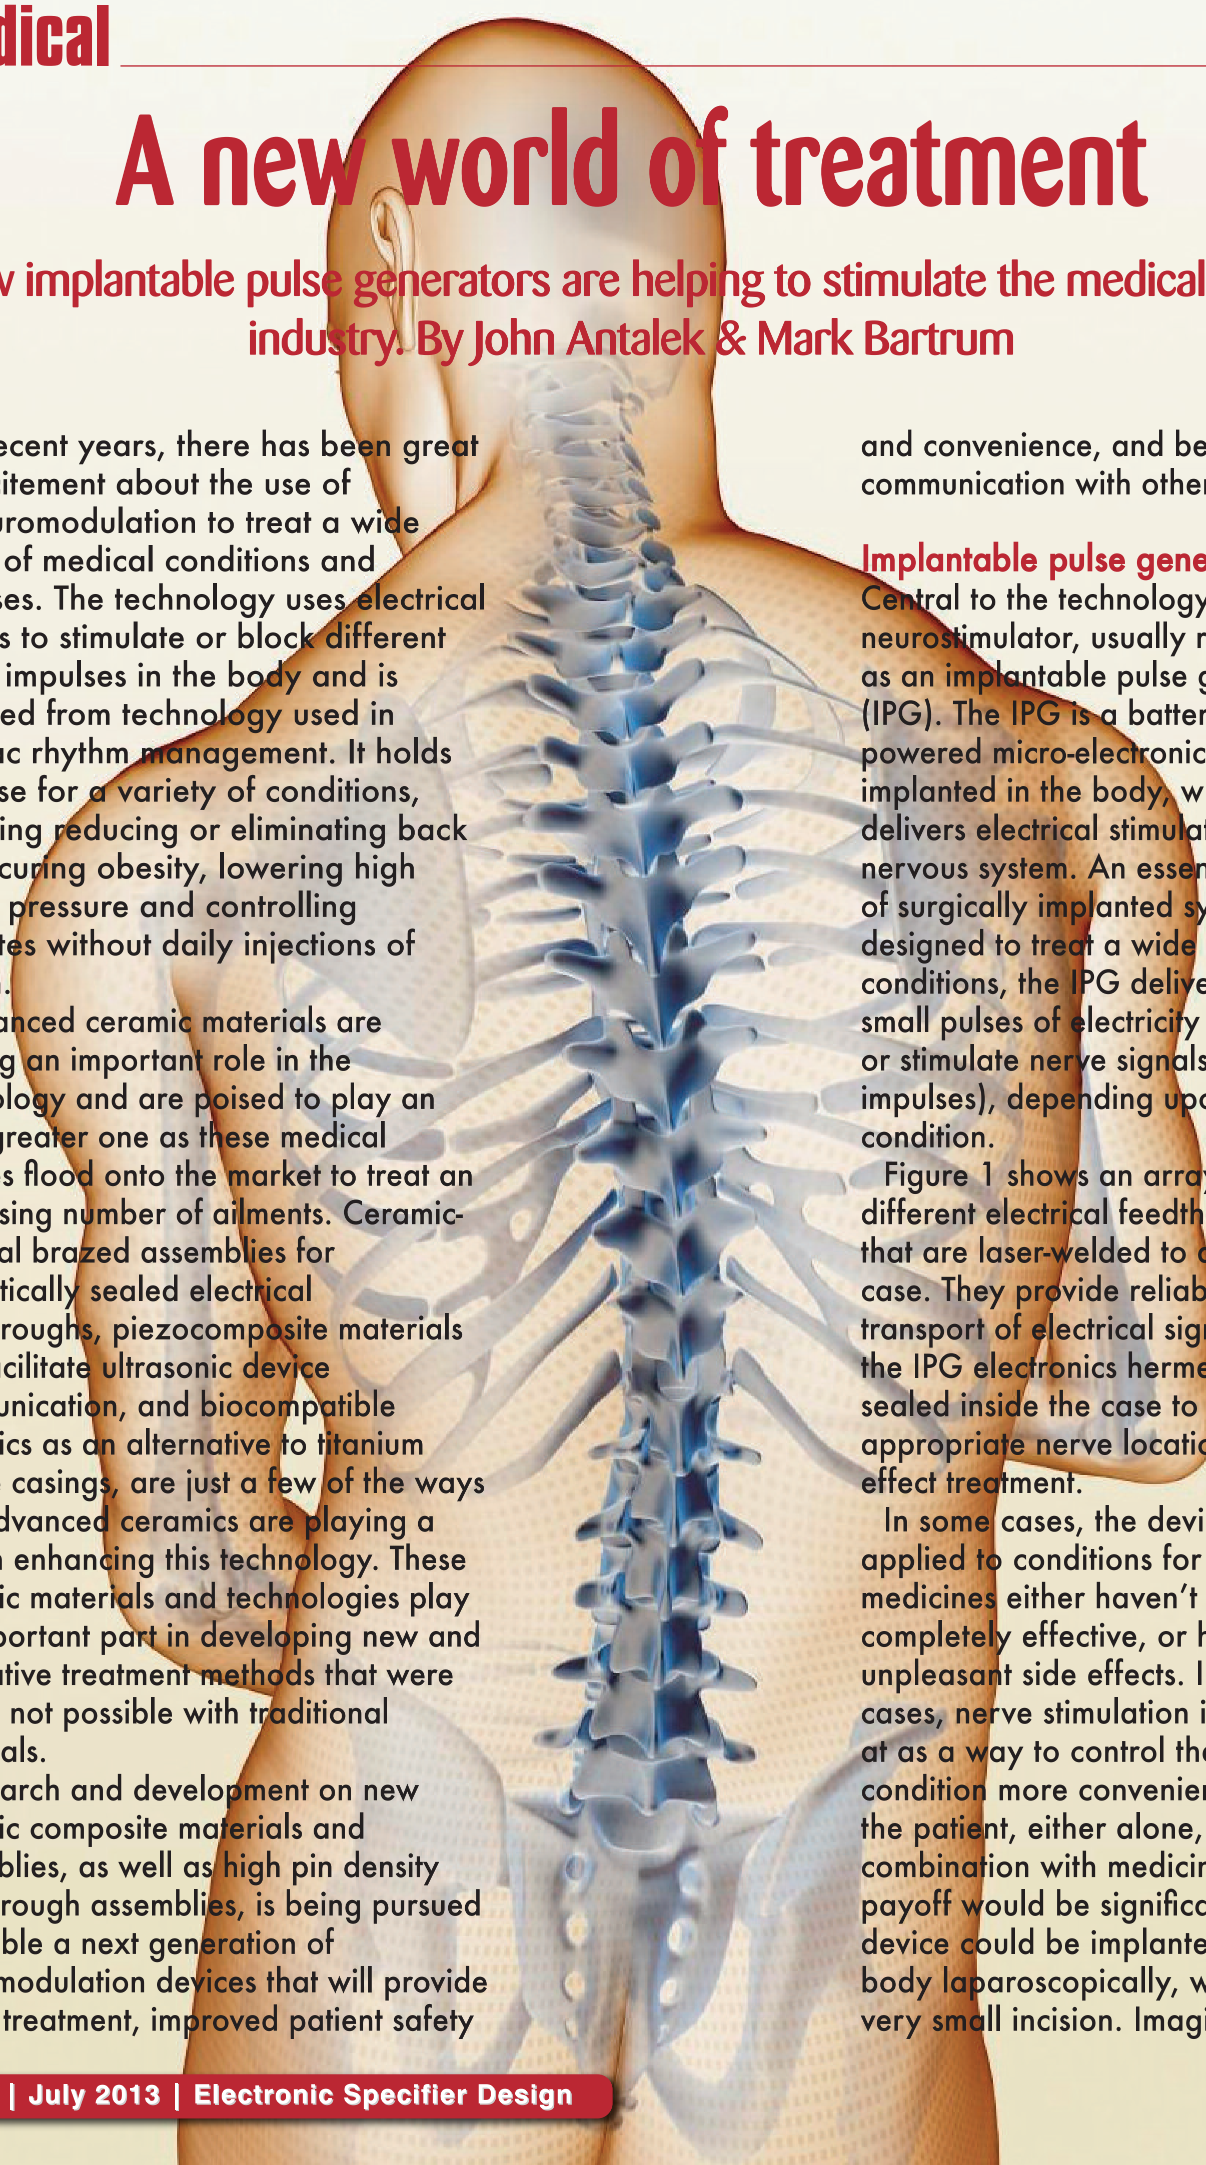

In recent years, there has been great excitement about the use of neuromodulation to treat a wide array of medical conditions and diseases. The technology uses electrical signals to stimulate or block different nerve impulses in the body and is adapted from technology used in cardiac rhythm management. It holds promise for a variety of conditions, including reducing or eliminating back pain, curing obesity, lowering high blood pressure and controlling diabetes without daily injections of insulin.

Advanced ceramic materials are playing an important role in the technology and are poised to play an even greater one as these medical devices flood onto the market to treat an increasing number of ailments. Ceramic-to-metal brazed assemblies for hermetically sealed electrical feedthroughs, piezocomposite materials that facilitate ultrasonic device communication, and biocompatible ceramics as an alternative to titanium device casings, are just a few of the ways that advanced ceramics are playing a part in enhancing this technology. These ceramic materials and technologies play an important part in developing new and innovative treatment methods that were simply not possible with traditional materials.

Research and development on new ceramic composite materials and assemblies, as well as high pin density feedthrough assemblies, is being pursued to enable a next generation of neuromodulation devices that will provide better treatment, improved patient safety and convenience, and better communication with other devices.

Implantable pulse generators

Central to the technology is a neurostimulator, usually referred to as an implantable pulse generator (IPG). The IPG is a battery-powered micro-electronic device, implanted in the body, which delivers electrical stimulation to the nervous system. An essential part of surgically implanted systems designed to treat a wide array of conditions, the IPG delivers very small pulses of electricity to block or stimulate nerve signals (or impulses), depending upon the condition.

Figure 1 shows an array of different electrical feedthroughs that are laser-welded to an IPG case. They provide reliable transport of electrical signals from the IPG electronics hermetically sealed inside the case to the appropriate nerve locations to effect treatment.

In some cases, the devices are applied to conditions for which medicines either haven't been completely effective, or have unpleasant side effects. In other cases, nerve stimulation is looked at as a way to control the condition more conveniently for the patient, either alone, or in combination with medicine. The payoff would be significant if a device could be implanted in the body laparoscopically, with only a very small incision. Imagine a twenty minute outpatient surgery, followed by years of 100% patient compliance, with no possibility of forgetting to take medication. The thought has doctors, patients and insurance companies very excited.

One key component of the IPG devices is the feedthrough, the mechanical structure that provides the electrical connection for the leads to get in and out of the device housing.

This tiny component performs several key functions. First, the feedthrough provides the conduit for communication of signals between the IPG and the body. Second, the feedthrough's hermetic seal keeps body fluids outside the IPG device and prevents electricity and battery materials from leaking out into the body. It must be completely and totally leak-free. It also has to be robust enough to withstand radio frequency and eliminate interference from MRI equipment and anti-theft scanners. This tiny piece of device real estate can contain as few as two leads, to around thirty.

Morgan Technical Ceramics - Alberox (MTC-Alberox) specialises in customisation of feedthroughs, providing design assistance and often working through several design iterations with the device designer, before arriving at the optimised feedthrough for the device.

The feedthrough technology is changing rapidly, as next generation devices get smaller and more compact and device designers seek to add more leads to improve the therapeutic value of the devices. Many device manufacturers often start with an off-the-shelf feedthrough to get their first generation device on the market. Their next generation devices are much smaller and more compact, which makes them more palatable to doctors and patients.

They may also have a lot more bells and whistles, so customers come to MTC to obtain a customised feedthrough that incorporates the additional features needed. Since MTC makes its ceramic components in house it has developed manufacturing processes capable of producing numerous sizes and shapes of feedthroughs to match the device design needs.

Stimulating from head to toe

New medical uses for IPG devices are patented frequently. Among the conditions for which the devices show the most promise are chronic back pain, hypertension and diabetes. Examples of devices focused on these conditions (either available now or under investigation) are provided below:

Chronic back pain

IPG devices deliver mild electrical pulses to the spinal cord, which interrupt or mask the transmission of pain signals to the brain. In this application, the IPG is implanted in the back, in close proximity to the nerve that doctors are trying to block.

Hypertension (high blood pressure)

Most patients with high blood pressure control the condition by using a regimen of anti-hypertensive drugs. However many studies have reported a persistence of refractive hypertension (elevated blood pressure despite using at least two anti-hypertensive drugs) in as many as 18% of the patient population. The IPG devices are being developed to provide a new and improved therapy for treating hypertension which is not only safe and effective, but avoids undesirable side effects of drug therapy.

The system includes an IPG, sensors and leads, external electronics for calibration, programming and periodic adjustment of parameters by the attending physician.

For example, one hypertension treatment clinical trial is investigating whether an implanted device can help control high blood pressure by stimulating pressure sensors called baroreceptors located on the carotid artery and in the carotid sinus. These sensors measure and report blood pressure to the brain, where it is compared to the needs of the body.

Another example is investigations on a device that stimulates the vagus nerve to control obesity; researchers testing the device have seen a dramatic drop in hypertension as an unexpected benefit of the therapy.

Diabetes

Several devices are being used or are under development, especially as an option for patients with diabetes that has proven unresponsive to drug therapy. An IPG is implanted and used to stimulate or inhibit the patient's vagus nerve to modulate its electrical activity to increase or decrease secretion of natural insulin by the patient's pancreas. The stimulator might be selectively activated in response to direct measurement of blood glucose or symptoms, or could be activated automatically for predetermined times or intervals. Alternatively, it could be automatically activated using an implanted sensor to detect the blood glucose concentrations.

One implantable system, originally tested to control obesity and other gastrointestinal disorders, is showing some promise in controlling diabetes. The device is designed to precisely control nerve and organ function using the vagal nerves, which regulate much of the activity of the stomach and the pancreas. The device being studied delivers high frequency, low energy electrical signals through laparoscopically implanted leads to block vagal nerve transmission. The delivery of energy to the nerves is intermittent and the effects of the therapy on the nerves and end organs are intended to be reversible. The system is designed to be precisely programmed and non-invasively adjusted to meet individual needs.

Headaches, obesity, epilepsy, depression, Parkinson's, syncope, sleep apnea and restless leg syndrome are a few of the many other conditions for which this treatment is being investigated.

Stimulation in the future

The general trend in implantable devices is miniaturisation. New devices tend to be smaller, but with a greater number of leads that get more signals into and out of a single device. It is clear that improving the feedthrough and expanding its capability is central to the next generation of IPGs. Future neurostimulator applications are currently looking at 100-200 leads, which will give device manufacturers opportunities to add further treatment to an implantable system.

Knowing where the IPG device is headed enables a proactive approach to providing improved feedthrough capabilities to help device manufacturers meet their needs; this is evident in how far technology has come with cardiac rhythm management. Pacemakers were the size of a Blackberry just a few years ago and had only two leads. Now the typical pacemaker is about the size of a lighter and can have as many as 10 leads, some of which allow better communication to the device, monitoring of other patient information and the ability to send information directly from the device to a doctor.

Additional leads could also build in intentional redundancy, which would reduce the device downtime and eliminate the need to remove the device if any of the leads fail. One of the most exciting avenues of research to increase the number of leads is the development of new high density feedthroughs that could contain ten times the number of leads, while keeping the current size and spacing. Today's feedthroughs are constructed by assembling many different parts, stacking them into complicated arrays with braze materials and putting them in a furnace for joining. However, researchers are now developing high density feedthroughs using cutting edge advanced ceramic materials and processing technologies that use miniaturisation techniques to pack many more wires together in a much tighter space.

Body communications

Another interesting development in the IPG arena is the development of 'body communications,' in which ultrasonic devices are placed into a medical device casing and used to remotely power and communicate with other devices in the body.

This next level of improvement has great advantages, because it could mean that no wires would have to be implanted. Implanting wires in the body can be a challenge. They may eventually fail and subsequent removal and replacement can be difficult. Also, using ultrasound, as opposed to radio frequency, means the communication stays within the body. This means one person's medical device is less likely to interfere with another person's device and it could be more readily protected against interference from MRI equipment, scanners, or other large electrical devices.

Using ultrasound to both power and interrogate remote sensors is a likely development for many implantable devices, but wiring would still be needed for neurostimulators, where leads are attached to the skull, brain, or spine. However, in the future, the main implant might be able to communicate with other devices implanted in the body or the external programmer via ultrasound, rather than radio frequency.

Another exciting development, the use of piezo ceramic components, is an outgrowth from the technology used in cochlear implants. The Bedford, Ohio plant has researched a custom assembly that uses piezo ceramic components to increase resolution for an annunciator that could be inserted into a main IPG and use intelligible speech to warn the user that an event is occurring.

Most IPGs are currently made of titanium, a strong and light metal that is lustrous and corrosion-resistant. However, along with efforts to improve the basic electrical feedthroughs so more leads can be added, research and development is being conducted using ceramic injection moulding (CIM) to develop a thin-walled ceramic case that could be smaller, simpler and provide more efficient communications to the device electronics. CIM enables production of small components with very high precision and without a secondary grinding process.

For example, MTC's Stourport UK CIM facility is currently in the early stages of developing an implantable housing for use in migraine and cluster headache treatment that is made using Zirconia injection moulding. Zirconia is the preferred material for this device because it has a high mechanical strength that allows the casing to be made with very thin walls. The ceramic casing would then be brazed and hermetically bonded to the feedthrough.

The Zirconia ceramic casing is stronger than titanium, allowing for a mechanically robust structure in a smaller sized housing. By comparison, an Alumina equivalent would be strong, but twice as thick. Both Zirconia and Alumina are transparent to radiofrequency energy, so signals could be passed through the wall for communication and possible charging. The hope is that developing a device casing out of Zirconia would mean that no feedthrough wires would be needed to connect to an antenna. Keeping the antenna within the device improves function for the patient. The Zirconia material is also inherently insulating, so electrical wires can be placed closer together than a metal flanged feedthrough brazed into a metal housing. Such a device would have no need for a battery, so it would not require replacement, as long as it retains its hermetic seal. Finally, the injection moulded Zirconia is a biocompatible material that complies with ISO 13356 implants for surgery.

Author profile: John Antalek is Medical Unit Business Manager for MTC-Alberox. Mark Bartrum is the Transducer Design Manager for the ElectroCeramics business of Morgan Technical Ceramics.

Figure 1: Examples of electrical feedthroughs that are laser-welded to an Implantable Pulse Generator (IPG) case

Next steps

The market for neuromodulation is estimated at more than $2 billion, with a compound annual growth rate (CAGR) estimated at 18 to 22% and a seemingly never-ending supply of new applications for the basic technology. With an increasing acceptance by the Food and Drug Administration and insurance companies, preference over some drug therapies and increasing device complexity to deliver more features and tailored effect, it is clear that developing the next generation of IPGs is critical to advancing neuromodulation technology.

What is also clear is that advanced ceramics will be a major part of that quest. Morgan Technical Ceramics' experience as a multi-national ceramic manufacturer gives them the ability to take technologies from other parts of the advanced ceramics business and use them in parts of the business where they haven't been used before. The same robust biocompatible materials already being used to make implantable drug delivery devices may now be used to develop a feedthrough that will be used in next generation IPGs.

Return to contents page.

Getting a view on medical imaging

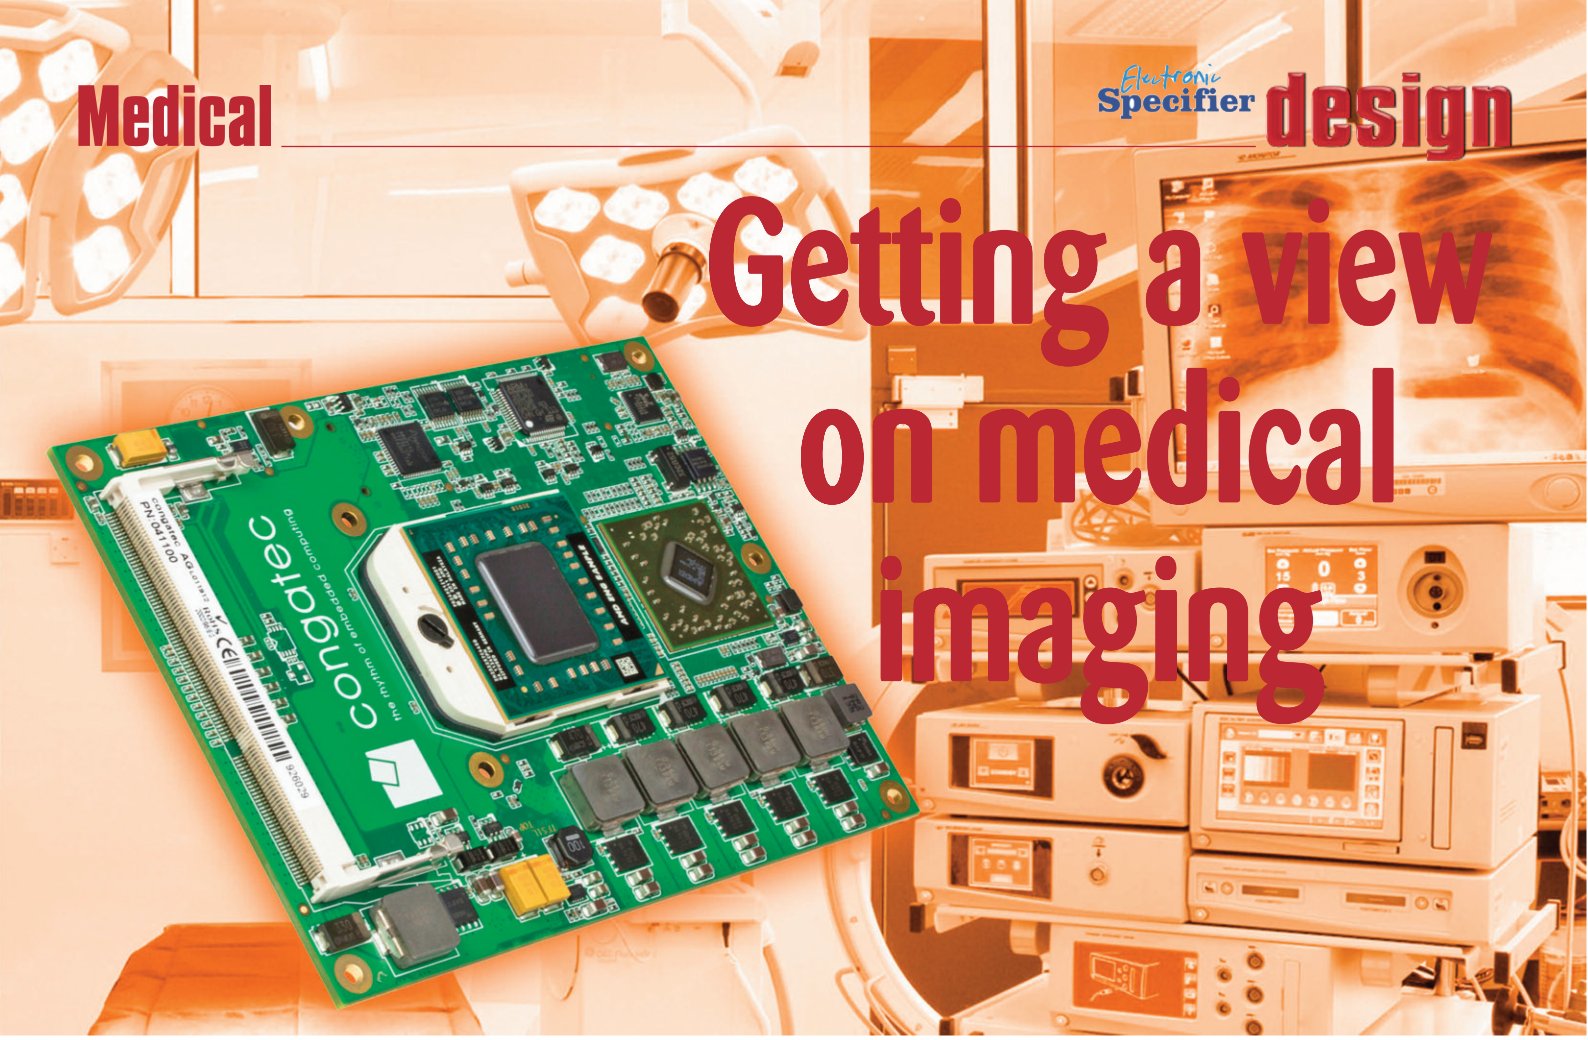

How modular-based embedded systems can meet the high-performance image display and processing requirements of medical applications.

By Martin Danzer

In addition to powerful image display, extremely fast image capture and image processing functions are key requirements for embedded graphic applications. Sensor data have to be processed to generate image data that can be displayed at very low power consumption and, ideally, in real-time. This requires the highest possible parallel processing power. While existing solutions based on DSPs or FPGAs are relatively powerful, they are proprietary and new developments incur great cost and effort.

The ideal solution for graphics-oriented applications therefore combines all requirements in one compact and energy-efficient system: High multifunctional computing power; high parallel processing power for data processing and imaging in real time; high graphics performance for visualisation; plus platform and hardware independence for high reusability.

That's exactly what a heterogeneous system architecture with Computer-on-Modules (COMs) based on the new Accelerated Processing Units (APUs) of the AMD R-Series offers. They integrate an efficient multi-core x86 CPU for classic PC tasks and scalable workloads with a programmable vector unit for parallel computing tasks and high-performance computer graphics on one silicon die. Thanks to tight integration and specialiszed processing units, the APUs are very compact and extremely energy efficient.

But which computing unit is responsible for the parallel tasks? The answer is simple; it's the integrated graphics processor. Fuelled by developments in the consumer sector – particularly computer gaming – graphics units have evolved to freely programmable specialists in parallel processing over the years. Modern GPUs now consist of several hundreds of processing units capable of performing complex calculations in parallel. They can do this with synthetically generated data from a computer game, but also with real data supplied by a wide range of sensors.

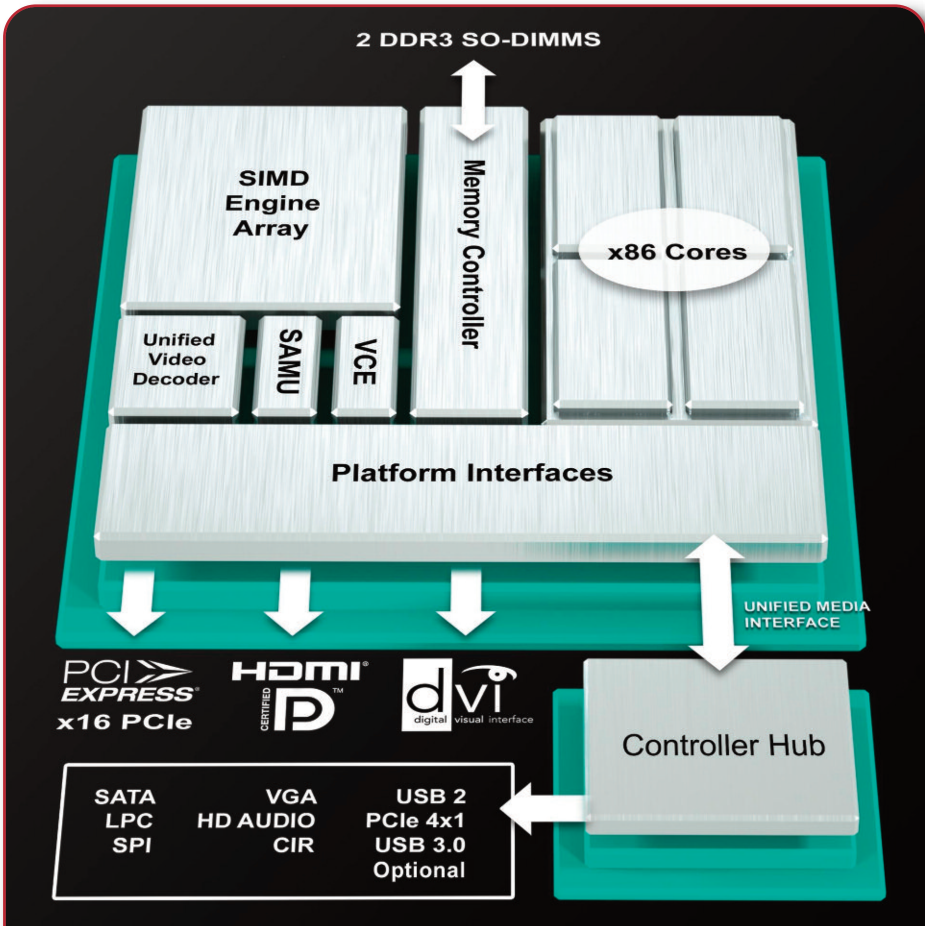

The innovative architecture of the AMD R-Series APU integrates all major system elements – including x86 cores, GPU vector (SIMD) engines and unified video decoder – in a space-saving two-chip solution.

High-speed parallel processing

The integrated AMD Radeon graphics unit of the 7000 family is extremely powerful, offering betwee 128 and 384 graphics cores with a clock speed of up to 686MHz. Achieving a 3Dmark Vantage "E" result of 13,066, the AMD R-464L APU far exceeds the performance class of previous integrated graphics units on the market. For parallel computing tasks, the AMD R-464L APU achieves a maximum of 576GFLOPS single precision performance. COMs which integrate these APUs are therefore an ideal platform for demanding medical imaging applications.

To enable developers to make the most of this parallel processing power, the conga-TFS COM with AMD Embedded R-Series APU supports the latest, cross-platform APIs such as OpenCL.

OpenCL is a powerful programming environment that allows computing tasks to be distributed and processed across hardware within heterogeneous processor systems. OpenCL is special in that multiple parallel execution is possible at each step (SIMD = Single Instruction Multiple Data), which means that classic parallel computing architectures are also supported. This is crucial because besides graphics display many analytical problems are also ideal candidates for parallelisation.

Using parallel computing, many highly complex and accurate calculations can be performed in just a few computational steps where classic, serial CPUs would require up to several thousand steps. This obviously reduces computation times and energy consumption for complex computational tasks drastically. Medical imaging technology, with its pronounced, often well parallelisable analytics, can draw particularly great benefit from this increased efficiency. An example from an image registration application where stable video images are key illustrates just how great the efficiency gains from using OpenCL in heterogeneous system architectures can be: The OpenCL based algorithm executes 120 to 130 times faster than a classical calculation on a x86 CPU.

COMs enable developers and OEMs to design these new features particularly efficiently into their medical devices and applications. COMs integrate the core computing functions of a system on a swappable module as pre-integrated components, while external interfaces and peripherals are implemented on an application-specific carrier board that is relatively simple to develop. The separation of carrier board and computing unit is a distinct advantage when developing medical devices as these must meet numerous specifications, such as EN6061 which specifies an extremely low leakage current through external I/Os. This requirement calls for specific I/O expertise that is easily met in the design of the carrier board without having to adapt the more complex computing unit.

The COM Express module conga-TFS currently supports three versions of AMD Embedded R-

Series APU ranging from the dual-core AMD R-272F APU to the quad-core AMD R-464L APU. The conga-TFS uses the AMD A70M Controller Hub and provides a powerful, compact two-chip solution with support for up to 16GB of dual-channel 1600MHz DDR3 memory.

The integrated graphics core supports DirectX 11 and OpenGL 4.2 for fast 2D and 3D imaging. A third generation hardware Universal Video Decoder provides seamless processing of H.264, VC-1, MPEG4 Part 2 and MPEG2 video streams. The choice of available graphics interfaces includes VGA and 18/24bit single/dual channel LVDS; three DisplayPort 1.2, one HDMI 1.4 and two single-link DVI for direct control of three independent displays are also available. Seven PCI Express 2.0 x1 lanes, one PCI Express 2.0 x8 link, four SuperSpeed USB 3.0 ports, four USB 2.0 ports, four SATA 6 Gb/s ports, a Gigabit Ethernet interface and high definition audio complete the comprehensive feature set.

Author profile: Martin Danzer is the Product Manager at congatec AG

More from congatec
Return to contents page.

Making technology better

How enclosures with antimicrobial properties can help sustain medical environments without compromising the equipment. By Paul Hoath

Hospital-acquired infections have become a serious problem in recent years. While improved housekeeping and new guidelines for staff have reduced the problem to a certain extent, the underlying problem of bacterial growth on surfaces remains. To help address the issue, manufacturers of plastic and metal enclosures and panels are adopting antimicrobial additives that can provide an effective means of reducing bacteria by up to 99.99% for the active life of the product.

Vero Technologies has entered into a technology partnership with Addmaster, a leading supplier of technically innovative additive formulations for plastics, textiles, papers, paints and coatings. The antimicrobial additive is sparingly soluble, providing a slow release of silver ions, safely inhibiting bacterial growth for an extended period, typically in excess of the product lifetime.

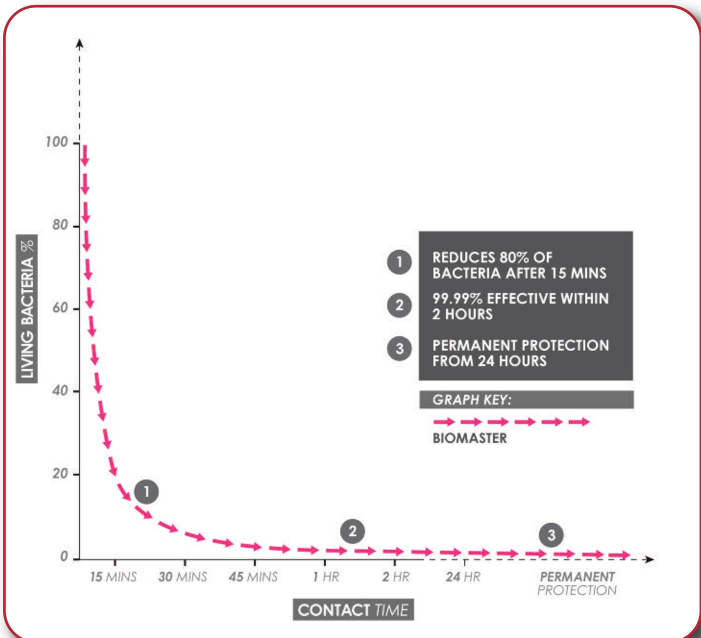

The key benefit is that the silver ion material is added to the raw ABS before the moulding process in the case of plastic enclosures. As the additive becomes an integral part of the material, the efficacy of the protection does not diminish over time, giving far superior protection to external coatings that will inevitably degrade through normal wear and tear, routine cleaning and accidental damage. For front panels and extruded aluminium enclosures where the antimicrobial material is added to the anodising process, trials are currently being conducted; preliminary results show similar reductions in bacterial growth rates to those found in plastic enclosures.

The additive, Biomaster, has been tested on over 50 different species of bacteria, including the commonly encountered MRSA, E.Coli, Listeria, Salmonella and Campylobacter species. Independent testing to ISO 22196:2011, using MRSA and E.Coli, has confirmed the effectiveness of the additive. In essence, the test procedure consists of inoculating test samples with a nutrient mix containing a known amount of bacteria. The samples are maintained at 37°C for 24 hours with the amount of viable bacteria colonies being measured.

The European test laboratory is the only one outside Japan to have SIAA (Society of Industrial technology for Antimicrobial Articles) accreditation. The SIAA is a cross industry-academic group for supervising antimicrobial technology.

Vero Technologies manufactures a wide range of plastic enclosures suitable for hand-held, portable instruments and desktop applications; all sizes and colours of the family can be supplied with the antimicrobial additive in the plastic, while the external appearance of the housings is identical to the normal versions. Also available are extruded aluminium cases and a large number of different front panel designs for cases and sub-racks. For OEMs involved in the medical equipment and instrumentation markets, specifying the antimicrobial versions of the products, which can be machined to suit the application without reducing the effectiveness of the antimicrobial protection, could also reduce time to market and offer a demonstrable benefit in terms of enhanced performance.

Author profile: Paul Hoath is the Sales and Marketing Manager for Vero Technologies

More from Vero Technologies
Return to contents page.

Board Mounted Power

- ✓ Low cost
- ✓ High reliability
- ✓ Short lead times
- ✓ Industry standard footprints
- ✓ Highest power density

20W DC-DC in 1"x1"

5W AC-DC in 1"x1"

AC-DC low power board mounted ranges

ECE series
- 5 to 40 Watts
- Ultra compact from 1" x 1"
- Encapsulated
- Single and dual outputs
- <0.3 Watts no load input power

ECL series
- 5 to 30 Watts
- Ultra compact size
- Single & multi outputs
- Open-frame & encapsulated
- <0.3 Watts no load input power

Industry standard DC-DC ranges

1 to 3 Watts
- SIP & DIP packages
- Regulated & unregulated versions
- ±10%, 2:1 & 4:1 inputs
- 500 V to 6 kV isolation

2 to 60 Watts
- 1 x 1", 1 x 2", 2 x 2" & DIP-24 footprints
- Regulated single, dual & triple outputs
- Wide 2:1 & 4:1 inputs
- Up to 3.5 kV isolation

Medical safety approvals
- 3 to 10 Watts
- Meets IEC60601-1 3rd edition
- 4000 VAC reinforced insulation
- 2µA patient leakage current

LED drivers
- 5 to 48 Watts
- Up to 1000 mA output
- Constant current output
- 95% efficiency

Surface mount power
- 1 & 2 Watts
- Regulated & unregulated versions
- ±10, 2:1 & 4:1 inputs
- Single & dual outputs

Download FREE
interactive guides to help select your AC-DC & DC-DC products at:
www.xppower.com/literature_downloads

Farnell — Also available from www.farnell.com

Visit our website to request a copy of our latest Power Supply Guide and see our complete line of power products.

GREEN•POWER

Selector App Available

XP Power
www.xppower.com

PCIM 2013, digested

Sally Ward-Foxton rounds up the most interesting MOSFETs, excellent IGBTs and the best of the rest from PCIM 2013.

Couldn't get to Nuremberg in time for this spring's PCIM (power control, intelligent motion) exhibition? Here's a digested summary of the most compelling new products and other interesting developments from the show.

IGBTs and IGBT drivers

British company Amantys used PCIM to launch its Power Insight technology last year, so it's no surprise that the company retained a big presence at the show this year. The company's booth had a live demo of their Power Insight IGBT monitoring software with live data coming from traction converters on board trains in the field in Portugal, thanks to its collaboration with railway maintenance giant EMEF (Figure 1). The software uses existing PWM and fault signals that are used to control the IGBT to monitor the device's health in real time and sends information to the train management system. Timely decisions can then be made on whether maintenance is required, streamlining maintenance efforts. Amantys is also partnering with various semiconductor manufacturers to incorporate Power Insight into their IGBT modules; Hitachi and Fuji Electric have been signed up so far.

CT-Concept previewed a dual-channel IGBT driver which can handle voltages of up to 4.5kV. The 2SC0635T's signal isolation is transformer based, making a cost saving when compared to fibre optic equivalents. As with all the latest IGBT drivers from the company, it's based on their proprietary ASIC which integrates driving, monitoring, status acknowledgement, DC-DC converters and isolation of all signals between the control and power sections.

Also launching at PCIM were a couple of new families of IGBTs; one for consumer appliances like rice cookers and microwave ovens, the other for solar inverters and UPSs. Designed to make these applications more energy efficient, according to manufacturer ON Semiconductor, the new Field Stop II IGBTs improve switching characteristics to reduce losses by up to 30% compared to the previous generation of products. Case temperature has also been reduced by 20%. The consumer products, NGTBxxN120IHRWG and NGTBxxN135IHRW, also have reduced conduction losses for induction heating and soft switching applications operating at frequencies of 15 to 30kHz. The NGTBxxN120FL2WG and NGTBxxN135FL2WG, for industrial use, have an operational junction

Figure 1. The Amantys and EMEF teams demonstrated live data from field trials in Portuguese trains at the show.

AC-DC Power Supplies with Battery Back Up

PQE75 & PQE100 Series

The PQE & PQP series are a range of AC-DC power supplies with universal AC input and either dual or triple DC outputs with battery back up in the event of an AC power failure. They are available in 75W or 100W power rating and can be either open frame construction (PQE Series) or fully enclosed (PQP Series). They are ideal for use in CCTV or Alarm systems applications or in fact anywhere where critical systems need to continue to operate in the event of an AC failure.

Features

- Universal 90-264VAC 47-64Hz input
- 12V or 24V Battery Back up versions
- DC output nominal 12V or 24V (depends on battery used)
- Auxiliary 5V 3A output on triple output versions
- Charger output 12V version is 13.8V 1.8A or 2.5A for 75W or 100W model
- Charger output 24V version is 27.6V 1A or 1.5A for 75W or 100W model
- Auto switching to battery when AC fails
- Short circuit, over load and over voltage protection
- Battery low and battery reverse polarity protection
- Brown out (AC input voltage low) protection
- Alarm signals for AC Power Fail & Battery Low
- Cooling by free air convection
- UL60950-1 and TUV/EN60950-1 approved
- EMI conforms to EN55022 Class B

PQP75 & PQP100 Series

POWER SOLVE

website: www.powersolve.co.uk
email: sales@powersolve.co.uk
Telephone: 44-1635-521858

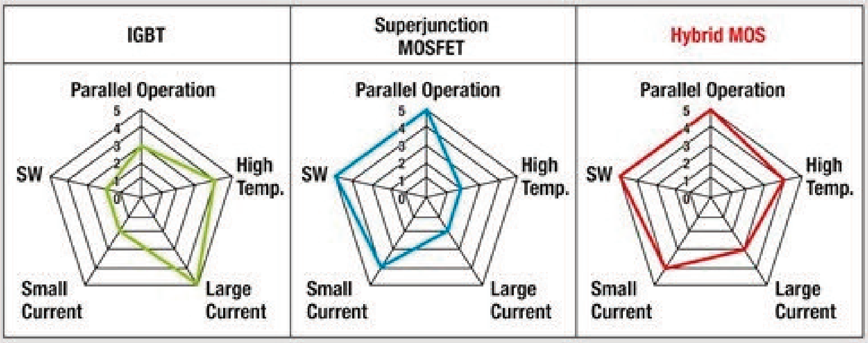

Figure 2. Rohm's Hybrid MOS devices, which combine superjunction MOSFET and IGBT structures, offer the best of both worlds in terms of characteristics.

temperature range of -55 to +175°C and increase in current ratings up to 100A in TO-247 packages.

MOSFETs

A newly-invented hybrid MOSFET-IGBT device was unveiled by Rohm at the show, which is designed to combine the best features of both device types into one part. Details of exactly how they have combined the structures of a superjunction MOSFET and an IGBT were a little sketchy, but the company did show off a chart of the expected features of the new devices (Figure 2). The idea is to combine the high switching speed of a superjunction MOSFET with the high voltage/high temperature capability of an IGBT. Samples will arrive this summer.

Toshiba has integrated high speed diodes into its superjunction MOSFET packages. High speed (fast recovery) diodes, essential for protecting the MOSFET in switch mode power supplies, can allow faster switching frequencies if their reverse recovery time, trr, is quick enough (Figure 3). Diodes in the three new devices offer t_{rr} of 100ns (TK16A60W5), 135ns (TK31J60W5) and 150ns (TK39J60W5). Additionally, a 30% reduction in $R_{DS(ON)}$*A (on-resistance x area) is down to Toshiba's latest DTMOS IV process, which means space could be made in the same TO-220SIS/TO-3P(N) packages to incorporate the diode. DTMOS IV uses only one epitaxial process to build the superjunction structure. The resulting structure gives the MOSFETs better high temperature characteristics too.

PCIM was also the venue for three MOSFET launches from Taiwanese manufacturer Advanced Power Electronics Corp (APEC). These devices, which are described as competitively priced, feature $R_{DS(ON)}$ as low as 0.19Ω. The AP11S60-HF-3, AP14S50-HF-3 and AP20S60-HF-3 high voltage MOSFETs also offer low gate charge and low conduction and switching losses. They provide a high blocking voltage to withstand voltage surges in demanding power systems, and can be used in PFC correction and PWM stages.

Keeping cool was the name of the game over at International Rectifier's booth, where the company introduced a new family into its COOLiRFET MOSFET range for automotive applications. Many applications run cooler using these devices than with other high performance MOSFETs, according to the company. In the new series are 22 40V N-channel MOSFETs, featuring IR's Gen12.7 trench technology. They are AEC-Q101 qualified, which requires that there is no more than a 20% change in $R_{DS(ON)}$ after 1,000 temperature cycles of testing. The benchmark D2Pak-7P AUIRFS8409-7P delivers $R_{DS(ON)}$ max as low as 0.75mΩ at 10VGS with a current rating up to 240A.

Power modules

Power IC supplier Vicor showed off new models in its Picor Cool-Power line, which are intended

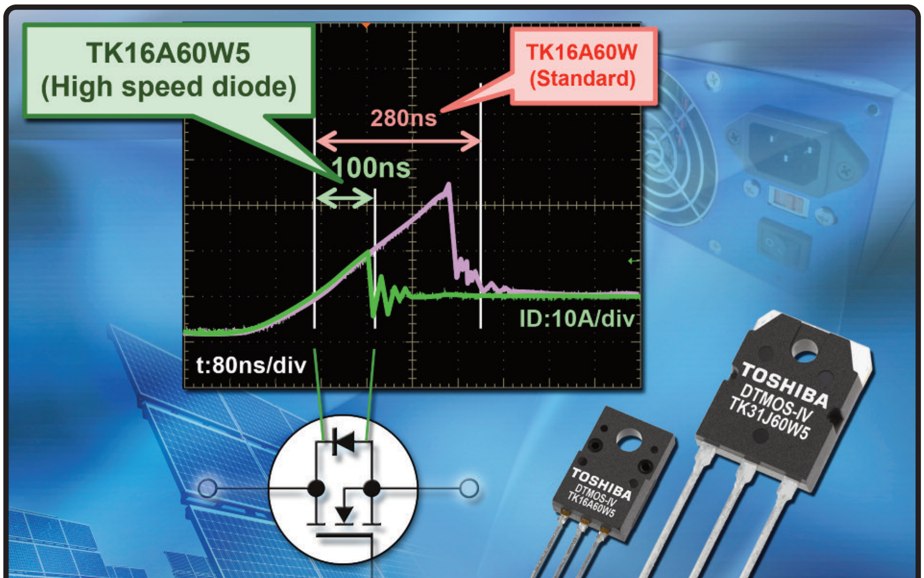

Figure 3. Toshiba's superjunction MOSFETs with integrated high speed diodes have decreased their reverse recovery times to 100ns.

for 24 and 28V applications, or those with demanding temperature requirements. These DC-DC converter modules are based on the ZVS (zero voltage switching) topology which allows higher frequency switching (900 kHz in this series) at higher input voltages without compromising efficiency. The PI31xx series provides up to 334W/in^3 and 2250V input to output isolation in a package smaller than a standard isolated 1/16th brick. Vicor recommends pairing these modules with its QuietPower EMI filters, which at less than 25% of the size of competing offerings, maximises density.

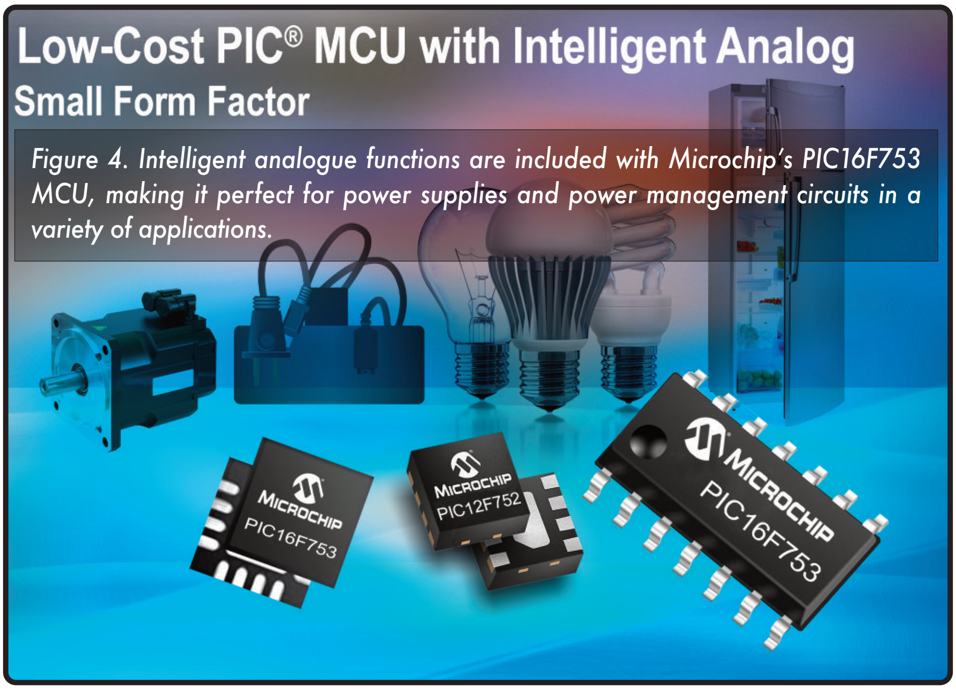

Figure 4. Intelligent analogue functions are included with Microchip's PIC16F753 MCU, making it perfect for power supplies and power management circuits in a variety of applications.

Meanwhile, Infineon chose PCIM to launch three EconoDUAL automotive qualified power modules for commercial, construction or agricultural vehicles which need extended reliability. These modules withstand more than three times better thermal cycling capability and better thermal shock capability by a factor of ten compared to the industry standard. They offer the highest power density available in this footprint; up to 600A/1200V. Copper wire bonding reduces internal lead resistance, while a new soft diode improves the modules' EMI behaviour.

Mitsubishi's offering in the automotive power module sector was the J1 series for driving inverters in hybrid and electric vehicles. This series currently comprises two parts; the 650V/600A CT600CJ1A060 and the 900V/400A CT400CJ1A090. Six CSTBT (carrier stored trench-gate bipolar transistor)-type IGBTs are in one package, saving space, and reducing collector-emitter saturation voltage to 1.4/1.7V. The package allows direct cooling via cooling fins, improving heat radiation by 40%, compared to previous products.

Best of the Rest

GaN systems was at the show hoping to prove that gallium nitride is ready for adoption in real applications and is a long way beyond the laboratory stage. GaN offers higher operating temperatures than silicon devices combined with high breakdown strength, so is well suited to power designs. The company has developed a new GaN transistor topology, called 'Island Technology', which is set to reduce size and cost of GaN devices by a factor of four.

MCU giant Microchip revealed an 8-bit micro with intelligent analogue and core-independent peripherals that it said would be perfect for power supplies and power management applications (Figure 4). The PIC16F753's features include a complementary output generator that produces complementary waveforms for inputs such as comparators and PWM peripherals, an op amp with 3MHz of gain bandwidth product, and a 9-bit DAC.

And finally: one of the more interesting new ideas from the show was Cymbet's solid state battery. These tiny batteries can replace coin cells, but since they're made using standard silicon processes, are essentially a battery IC with no toxic chemicals. The rechargeable battery die can be packaged with another IC, perhaps as a backup battery for a clock IC, easily. The company announced that these batteries had been proven safe for implantable medical devices at the show.

Return to contents page.

Forethought is a wonderful thing

Why design engineers need to think about circuit protection as part of the design process, rather than as an afterthought. By Bharat Shenoy

Smaller, faster, prettier; arguably these are the main criteria for today's electronics. A profitable lead for manufacturers, but also quite challenging for design engineers, who have to fit everything together in a much smaller environment.

Put safety on top of these requirements and the story gets even more complex; for both designers and manufactures, which is ultimately relevant to the end customer, who expects products to be not only fully functional, but also safe.

Imagine the following scenario: morning, jogging, wet weather... running shoes, hosiery, and nylon shirt. Taken together, for a brief moment, they can generate an electrostatic discharge (ESD) of 30,000V - a potential cardiac arrest for the smartphone or any other handheld device that plays your favourite music.

Imagine you walk up to your computer and suddenly everything resets. Your first thought is maybe its malware or a virus, or maybe the computer is just old and the hardware is failing. You might not think that the simple contact between you and the computer could triggered a brief but powerful electric incident.

Circuit protection is ever present but still, problems associated with over-voltage and over-current still remain an afterthought for many engineers. Why? Because, with their duties expanding and with design cycles compressing, most engineers relegate circuit protection to the end of the to-do list. Faster is the word of the day; engineers are expected and required to design the core functionality of their devices as quickly as possible. They have to get the form factor done, get the software done, get the prototype built, and prove out the concept. Only then, if at all, will they have time to think about circuit protection.

Saving development time

More than ever, though, that approach is creating problems for product designers. Cellphones, computers, and music players are getting smaller. Moreover, they're running on tiny voltages that are more susceptible to ESD, distant lightning strikes, motor switching and stray currents from process machinery. It's usually 10,000 or 15,000 volts, but experience nowadays is showing that it can get really high, even up to 30,000 volt parts.

The unfortunate result of leaving such matters to the last minute is that design functionality suffers. With devices getting smaller and smaller, and so many parts having to fit together, engineers can't find room anymore for circuit protection devices

on their printed circuit boards. The result: they end up re-spinning the boards and losing valuable development time; time that could have been saved had they thought about circuit protection from the very beginning. To make things worse, many times engineers do think of circuit protection, however they hurry up choosing whatever solution is at hand and the wrong protection device, resulting in functional failures, poor reliability, safety issues, shock, or even fire.

In order to guide the design engineering community to prevent such gloomy scenarios, here are some expert recommendations of engineers whose professional lives revolve around the subjects of over-current protection and shock immunity.

Think of circuit protection from the very ‘chip’ phase - Considering circuit protection too late in the project can really have some serious repercussions for any design engineer. One such scenario would be to put yourself in a situation where the space is not available for your ESD device. Or, even worse perhaps, you end up settling for a non-optimal location, where the device won’t function the way it’s supposed to.

To avoid these, the best time to start thinking about such matters is after the very moment you’ve picked out the chip set and begun laying out the circuit board. Doing it in that way, ESD ratings are available and designers know how robust or how sensitive the chips are. Some of these chips are running at 1.5V and you don’t have to do a lot to upset them. Experience tells us that the circuitry is more complicated and more sensitive than we sometimes realise.

Understand the threats - Fuses are simple, everybody understands them, but over-voltage may not be so obvious, and people might not realise the consequences. However, consequences do exist, even if they are not as catastrophic as those of over-current. Did you know, for instance that over-voltage has even incapacitated the Hubble Telescope, shut down refineries, killed smartphones, and stopped roller coasters mid-ride? In some portable medical devices, over-voltages can even be life threatening.

There are many possible sources of excessive current and voltage, starting from the common lightning, ESD, motors, arc welders, as well as the aforementioned running shoes and hosiery, among others. Take lighting, for instance; people understand lightning but they may not know it travels across the ground and can create huge glitches in power lines a mile away. And this is just one example for why, knowing and understanding the possible threats is so important.

Define the needs - To accurately predict a product’s circuit protection needs, the design engineer must first be able to imagine how it will be used. In other words, you have to know where the product might end up, understand its environment and what might be adjacent to it. For instance, a device will be more susceptible to a factory setting than to an office.

Once the designer understands the environment, he or she can begin making accommodations. The point to start is the connection points. The over-voltage device should not be five centimetres away, but as close as possible to the connector. Understanding the target environment of the product and the possible threats that might come up with that location will be crucial in selecting the right circuit protection solution.

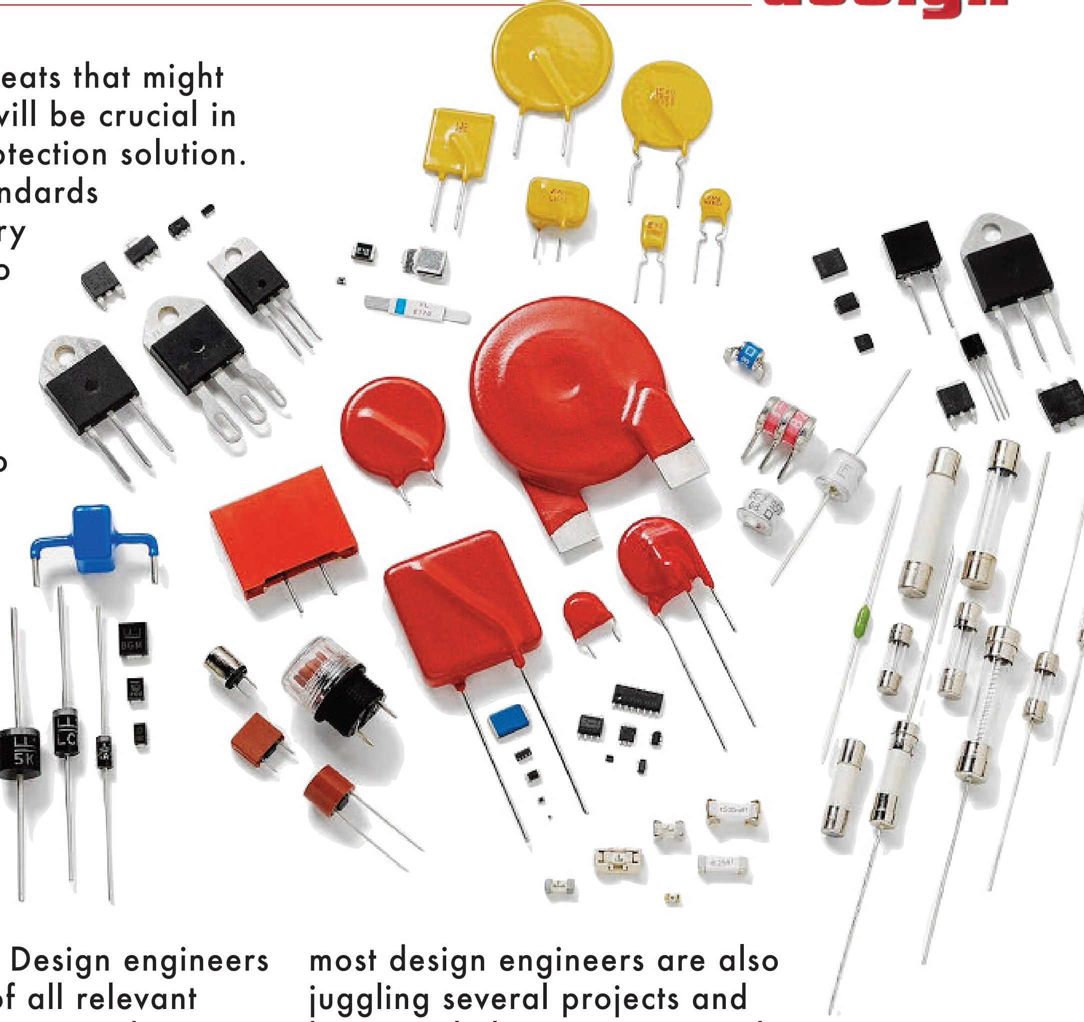

Know your standards - Standards determine the design of every product, all the way down to the circuit protection. Moreover, the list of standards that designers need to be aware of is seemingly endless. They also differ internationally, regionally or even locally. For circuit protection, standards include those from Underwriters Labs, Energy Star, NEMA, ATCA, CSA Group, IEEE, and standards bodies in Canada, South America, Japan, Korea, and Europe, among others.

Standards are a big part of every project or industry. Design engineers have the duty to be aware of all relevant standards before even commencing the project. Just like the components themselves, knowing and understanding the regulations that govern certain products locally or internationally, has to be part of the very first design phase as well. In other words, one can not create the product and then go look for the standards that would enable that product to work or not in a particular market.

Keep yourself informed - Because circuit protection isn't taught in universities, most engineers are more well versed in the complexities of product design than in the issues of over-current and shock immunity. That problem is compounded by the fact that most design engineers are also juggling several projects and have too little time to research the topic. Industry whitepapers, product specs, and case studies are always helpful as a source of information and education but also as examples of what worked in certain situations, possibly similar to yours.

Overall, although not a topic officially taught in schools, circuit protection is a topic relevant to the every day life of the end consumer, to the manufacturer whose brand reputation relies on delivering functionality and safety in its products and ultimately to the design engineer whose success is measured by the his development time and quality. It's something that definitely needs to be learned through experience and share of know-how across the industry. Littelfuse, for example, offers its speed2design site, which is geared toward helping time-challenged designers find circuit protection solutions.

Author profile: Bharat Shenoy is the Director of Technical Marketing, Electronics Business Unit, Littlefuse, and is responsible for market segment research, sales channel support and applications engineering for the Electronics Business Unit. Shenoy began his career as a US Naval officer in the nuclear power submarine fleet and has held positions in Tyco Electronics Circuit Protection as a Field Applications Engineer, Regional Sales Manager, and Sales Director.

More from Littlefuse
Return to contents page.

Redefining the midrange?

Broadly speaking, FPGAs are designed for either 'high-performance' applications or 'everything else' and as a result most FPGA vendors compete in the 'everything else' category, Philip Ling takes a look behind the latest announcements to see if we can expect a redistribution of wealth.

It's no secret that Altera and Xilinx compete, ostensibly, for the same design sockets within 'power' applications (predominantly in the telecommunications market) and that the rest of the FPGA vendors are 'ok with that'; after all, it requires a lot of money, effort and risk to develop on the bleeding edge, which is invariably where the top-end FPGAs have to be in order to win those design sockets.

While there are a limited number of design slots for high-end FPGAs, the 'everything else' market for FPGAs is vast and reportedly increasing. According to Microsemi's Vice President of Marketing, Paul Ekas, the market for FPGAs with up to 150,000 logic elements (LEs) is worth in excess of $1 billion (excluding the consumer market). Even with their higher selling price, this means the 'performance' end of the market may not be the most lucrative.

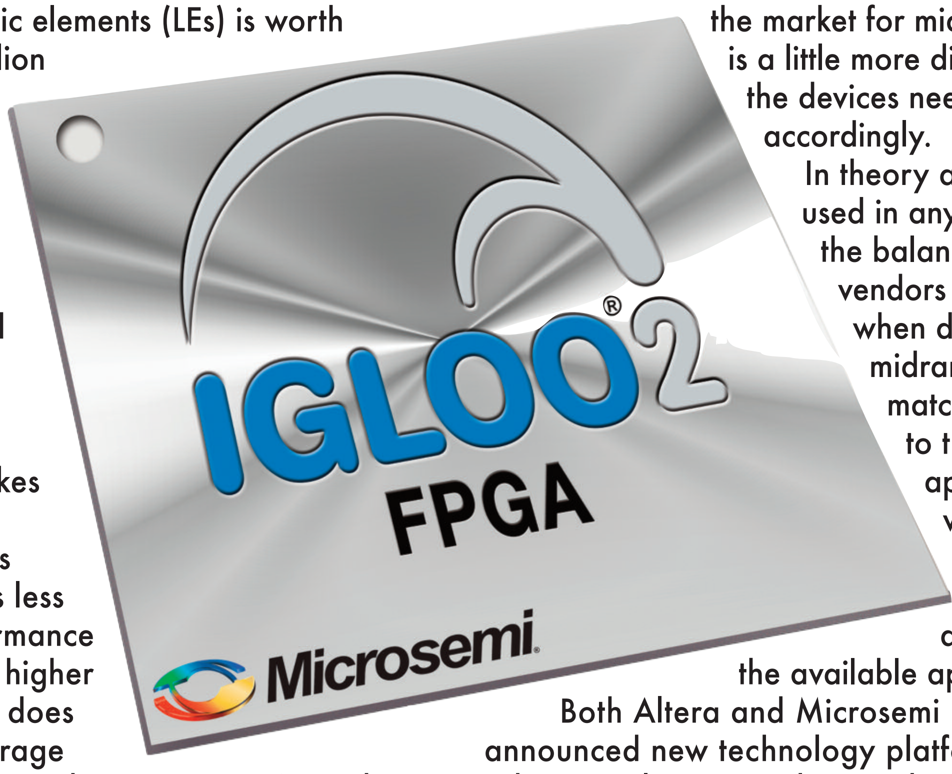

On paper, it makes sense to target a market that carries lower risk and has less demanding performance requirements with higher volumes, even if it does carry a lower average selling price. It's something most FPGA vendors seem to have acknowledged and are developing new devices targeting this part of the market. That means competition in the midrange isn't restricted to a subset of FPGA vendors, and as technology advances trickle down the product tree it also means that midrange devices are increasingly capable of taking on more demanding applications.

Midrange spread

The term midrange is probably easier to define in terms of applications, because from an FPGA vendor's point of view a midrange device is a relative term; each new generation of FPGA pushes performance up. Unlike the 'power' end of the market – customers who will make use of the fastest, widest and largest FPGA – the market for midrange devices is a little more discerning and so the devices need to be 'tuned' accordingly.

In theory any FPGA can be used in any application, so the balance all FPGA vendors need to strike when developing a new midrange device is to match performance to the (existing) application space, while pushing the boundaries enough in an attempt to extend the available application space.

Both Altera and Microsemi have just announced new technology platforms and, subsequently, new midrange devices. Arguably, they have both chosen the option of delivering more than 'just enough' performance, but not too much.

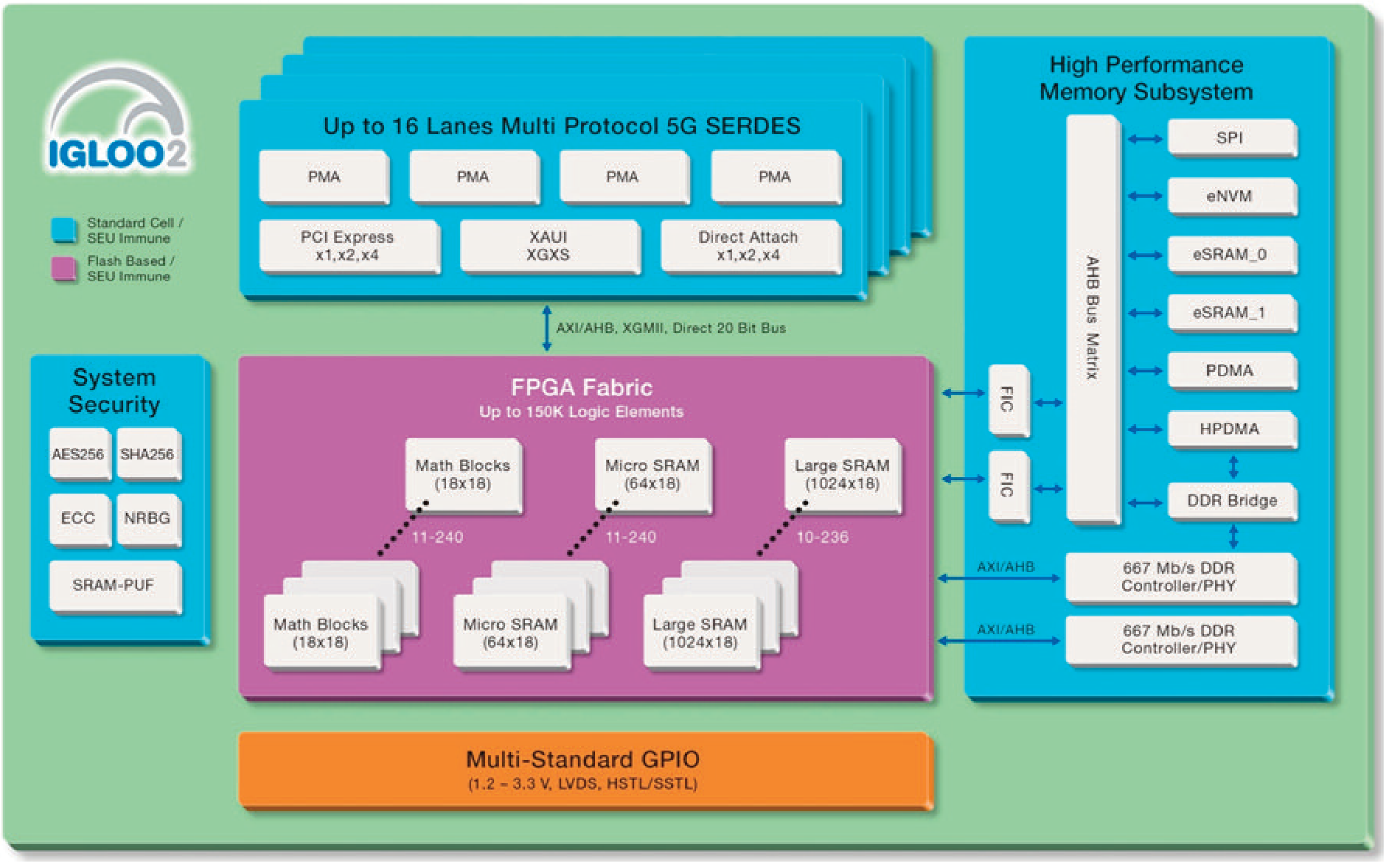

Microsemi's Igloo2 family is being marketed as a 'cost optimised' soluiton, with up to 150k LUTs

Microsemi's Igloo 2 family spans 6,000 to 150k logic elements and is designed on an all-new 4LUT (look-up table) fabric; the existing Igloo family uses a 3LUT fabric. According to Ekas, the choice to limit the family to 150k LEs was commercial but the fabric is optimised for this density, which means if Microsemi chooses to increase density it will probably need to redesign the fabric.

The move from 3LUT to 4LUT has pushed up the baseline performance of the fabric; it is reportedly five times faster than the Igloo range. Measuring FPGA performance is largely application-dependent, but Ekas stated that Igloo 2's performance is on a par with midrange devices offered by other FPGA vendors.

Going in hard

In its 30th year since incorporation, Altera has just released the first details of its 10th generation of FPGA, which it has chosen to call Generation 10. This sees both the Arria and Stratix families getting a technology upgrade and perhaps of most interest is the news that the flagship family, the Stratix, will be manufactured using Intel's 14nm Tri-Gate process (known elsewhere in the industry as finFET). This will enable developers to choose – in terms of performance or power efficiency, depending on the application's needs – how to best make use of that new technology.

For example, in comparison to the Stratix V, if maximum performance is needed the Stratix 10 can deliver a 2-fold improvement for 'only' 30% more power. Conversely, if the performance offered by the Stratix V is adequate, engineers can reduce power by as much as 70%. This power/performance tuning comes in both the FPGA and SoC versions of Stratix 10, with the SoC variant also featuring a new ARM Cortex processor sub-system. Currently the Stratix V SoC uses the 'first generation' of processor sub-system, based on a dual-core Cortex-A9 configuration. The configuration for the Stratic 10 hasn't been announced but Altera is differentiating it from the sub-system its integrating in to the Arria 10 SoC (which it has dubbed '2nd Generation') by referring to it as 'ultra high-performance' 3rd Generation. This could suggest the core being used in the Stratix 10 SoC isn't the A9 but something more advanced, such as the A12; heralded as the A9 successor thanks to delivering higher performance at the same power.

What is known already is that the 2nd Generation processor sub-system, as being integrated in to the Arria 10 SoC, will run at 1.5GHz (which is almost twice as fast as the 1st Generation sub-system), and is (still) the Cortex-A9 MPCore.

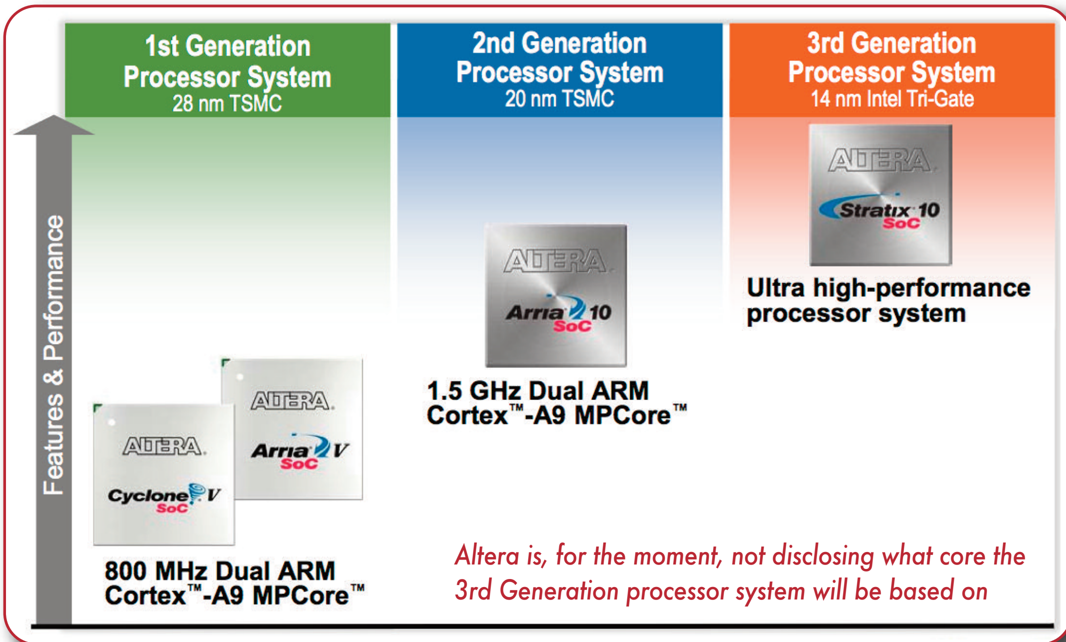

Altera is, for the moment, not disclosing what core the 3rd Generation processor system will be based on

Altera admits to being 'surprised' by the level of interest and uptake of the ARM-based SoC members of its FPGA families, stating that demand could overtake demand for non-SoC devices.

One of the target applications for this tiny FPGA is sensor hubs in devices that fall under the heading of 'the Internet of Things'. This is still an emerging mega-trend, of course, but the industry expects to see the number of connected devices rise in to the multiple billions within the next decade (the actual figures vary depending on the source of the research). The majority of these devices will be sensor-based, relaying data through the internet. Given the lack of standardisation in sensor interfaces, a small FPGA may be the way to go.

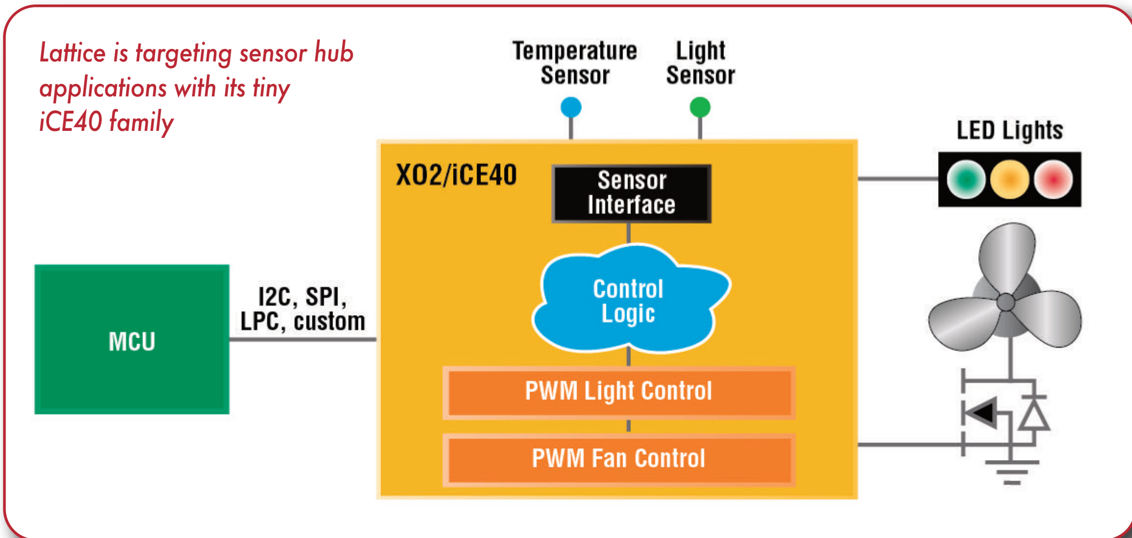

Lattice is targeting sensor hub applications with its tiny iCE40 family

FPGAs consistently extend their usefulness, either by integrating more hardwired functionality or extending the performance of the underlying fabric. Following the standard cost model, they also continue to fall in price, promoting their position as the alternative to ASIC and, increasingly, ASSPs.

Low end

If the midrange is difficult to quantify then the definition of a low-end application must be equally vague, particularly given that the cost of entry-level FPGAs is tumbling. The latest indication of this came from Lattice Semiconductor recently, when it introduced what it claims to be the industry's smallest FPGA; the iCE40 LP384 (which offers 384 LUTs). The roadmap for the family includes a 2mm^2 package at a price of $0.50 in multi-million volumes.

Return to contents page.

Flexibility through design

Expanding input capabilities and regulating seamlessly through automotive cold-crank and load-dump transients. By John Canfield

Handheld devices, industrial instruments and automotive electronics all demand power supply solutions that can support an expansive range of input voltages resulting from automotive input voltage transients, resistive line drops and a wide variety of power sources. As a further design challenge, applications often require a variety of regulated voltage rails, including some that fall within the input voltage range.

The proliferation of electronic subsystems in automobiles has created demand for small size, high reliability power supplies that can operate under the stringent conditions presented by the automotive environment. Figure 1 shows a 5V automotive supply ideal for use in engine control units and other critical functions including safety, fuel system and drive train subsystems where processors must remain powered without glitch during even the most severe input voltage transients. This application uses a 2MHz switching frequency to minimise its footprint and eliminate interference with the AM broadcast band.

The V_{CC} rail provides power to the internal circuitry of the LTC3115-1 (see sidebar) including the power device gate drivers and is ordinarily powered from the input rail via an internal linear regulator. In this application, diode D1 bypasses the internal linear regulator and delivers power to the V_{CC} rail directly from the regulated output to improve efficiency and output current capability. This is particularly advantageous in applications with higher switching frequencies, given that the increased gate drive current is provided more

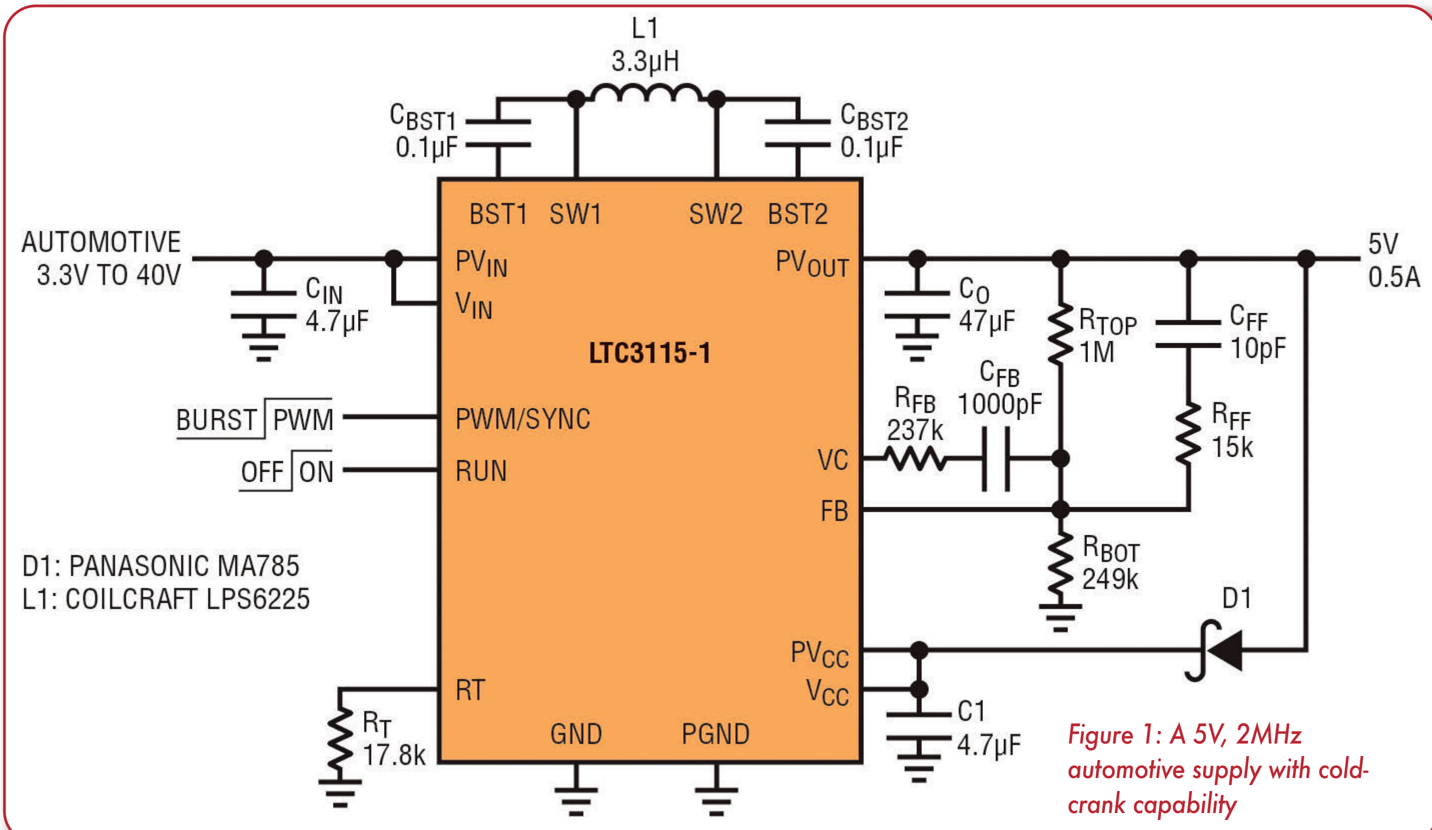

Figure 1: A 5V, 2MHz automotive supply with cold-crank capability

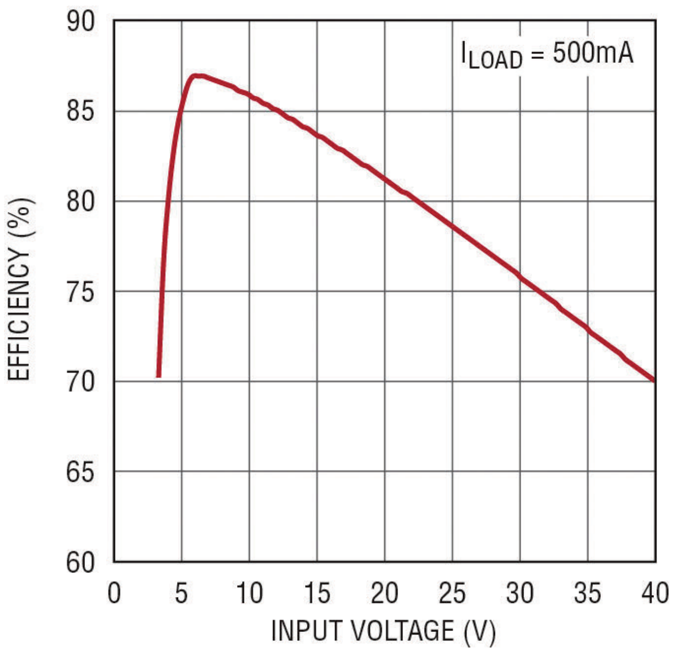

Figure 2: 5V, 2MHz automotive supply efficiency versus V_{IN}

efficiently from the converter's output rail than through the internal linear regulator. Figure 2 shows the efficiency of this application circuit with a 500mA load for input voltages from 3.3V to 40V.

Riding through line transients

Of commonly utilised power sources, the automotive supply rail presents one of the most challenging inputs to a power supply. Its nominal voltage varies from 10.6V to 15V depending on the state of charge of the battery, the ambient temperature and whether the alternator is charging or idle. In addition to the variability in its nominal voltage, the automotive power rail is also subject to a wide range of dynamic disturbances induced by changes in engine RPM, transitioning loads such as power windows, wipers and air conditioning, and inductive transients in the wiring harness.

However, the most extreme conditions occur during a load-dump transient which can produce voltages in excess of 120V for a duration of hundreds of milliseconds. A load-dump transient occurs when the alternator is charging the vehicle's battery and an electrical open-circuit causes a momentary disconnection of the battery from the alternator. Until the voltage regulator can respond, the full alternator charging current is applied directly to the automotive power bus, raising its voltage to potentially dangerous levels. Such a transient could be caused through a physical disconnection of the battery by a mechanic working on the vehicle, but could also result from a faulty connection in the battery cable or corrosion at the battery terminals.

Automotive electronics must also be designed to survive a double-battery jump start, where they are subjected to 24V for extended durations as the vehicle is jump started using a series-connected second battery or from a commercial vehicle with a dual battery electrical system. An additional overvoltage condition on the automotive bus is caused by alternator voltage regulator failure and is often included in the battery of tests conducted by automotive electronics OEMs. Such a malfunction can result in full application of the alternator charge current to the battery and an overvoltage of approximately 18V for extended durations.

The automotive power rail is also polluted with short duration overvoltage transients due to rapid load changes produced by switching high power loads such as power doors, fans and cooling fan motors interacting with the significant inductance in the vehicle's wiring harness.

In most vehicles a passive protection network consisting of a lowpass LC filter and transient voltage suppression (TVS) array is used as a first line of defence to clamp the peak excursions of the power bus. Typically,

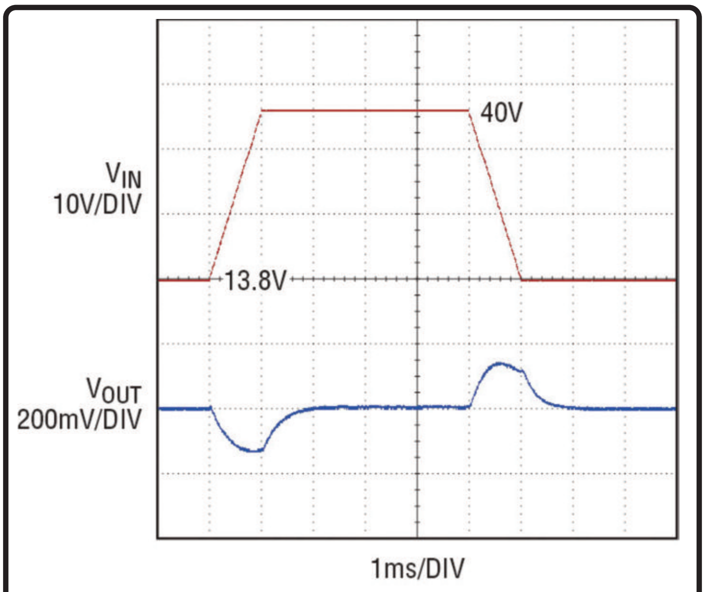

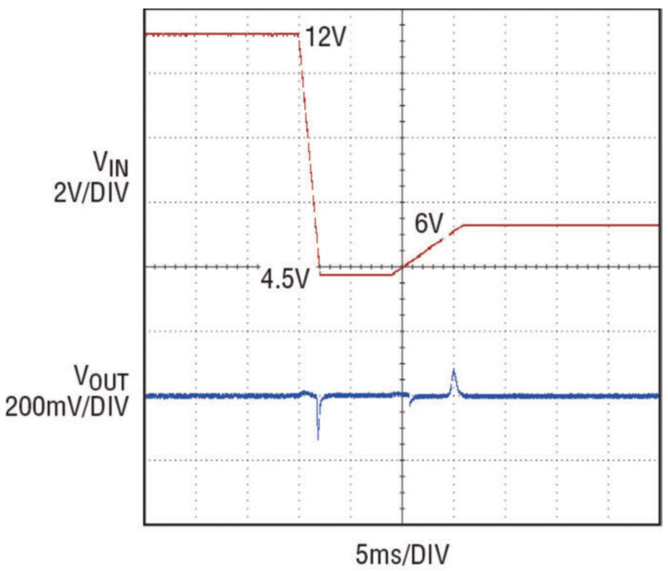

Figure 4: A 12V to 4.5V cold-crank line transient

automotive electronics located downstream from the protection network must survive up to a 40V transient without damage. Critical systems must not only survive, but must also function seamlessly through such transients without interruption. Figure 3 illustrates the ability of the LTC3115-1 to maintain uninterrupted regulation of a 5V supply rail through a 13.8 to 40V momentary line transient with 1ms rise and fall times.

Cold-crank transients

High voltage transients are a problem on the automotive power bus, but perhaps the more challenging problem is undervoltage transients. The most severe of these is known as cold-crank, which occurs when the engine is initially started.

A typical cold-crank voltage waveform is shown in Figure 4. The initial low voltage plateau is the most extreme and is caused when the starter motor begins turning over the engine from a dead stop. During this phase, the vehicle's bus voltage can fall below 4V. Colder temperatures exacerbate the situation since the higher viscosity of the engine oil results in a higher required torque from the starter motor. The first plateau is followed by a second somewhat higher voltage plateau, typically near half the nominal battery voltage, as the starter maintains the engine rotation. Once the engine starts, the battery recovers to its nominal voltage.

Safety devices and engine critical components such as the engine control unit and fuel injection system are required to remain operational throughout a cold-crank transient. As shown in Figure 4, the LTC3115-1's buck-boost architecture enables it to maintain output regulation through even the most severe cold-crank transients by automatically and seamlessly switching to boost mode operation during the undervoltage event.

Cold-crank capability for automotive electronics has expanded in importance as cars now include automated fuel-saving, on-demand engine start/stop, whereby the vehicle's engine is turned off during momentary vehicle stops at stoplights or in traffic. Vehicles equipped with on-demand starting are subjected to frequent cranking undervoltage events. As a result, auxiliary electrical components that previously had no need to function through the occasional cold-crank event in a traditional vehicle must now operate through such transients to eliminate any disturbance to infotainment, navigation, dashboard electronics and lighting systems.

The LTC3115-1 features a low noise forced PWM mode where both switch pins operate at constant frequency for all loads, producing a low noise spectrum, independent of operating conditions. The predictable spectrum and

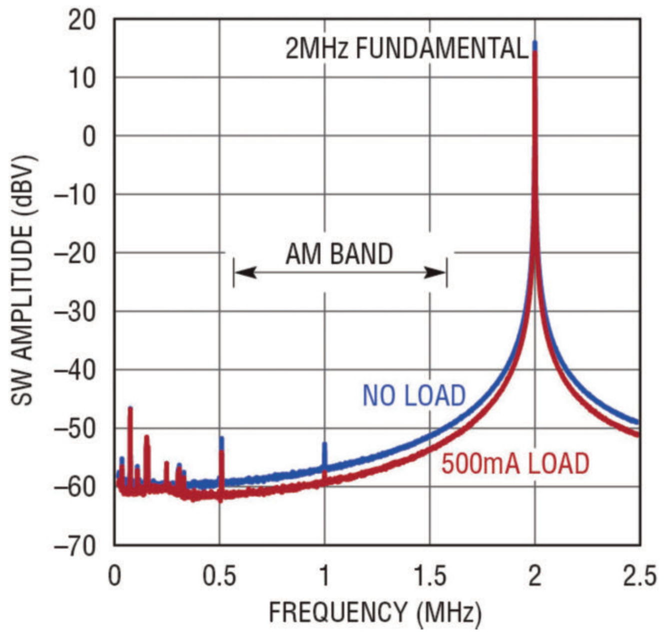

Figure 5: Fixed frequency low noise PWM minimizes emissions across the AM band

minimal subharmonic emissions help reduce interference and aid in compliance with strict automotive EMI standards.

The LTC3115-1 supports switching frequencies up to 2MHz so that the fundamental switching frequency component, and all of its harmonics, can be located above the AM frequency band to minimise interference with radio reception. Figure 5 shows the spectral emission of the LTC3115-1 over the AM band for the automotive application circuit of Figure 1 operating at no load and with a 500mA load. In both cases the entire range of frequencies within the AM broadcast band is free from any significant spectral emission.

Multiple power sources

To increase flexibility and enhance the user's experience, many portable electronic devices are being designed to work from various power sources. These power sources can vary widely in voltage, especially when accounting for connector and cable drops.

Under USB 3.0, the nominal supplied voltage is 5V ±5%, but a fully compliant powered device must be able to operate down to 4V when accounting for allowable cable and connector voltage drops. In addition, a downstream USB power rail is permitted to drop as low as 3.67V under transient conditions such as when additional devices are plugged into the host or powered hub.

The newly approved USB PD (power delivery) specification allows for higher power delivery over USB with support for supply voltages up to 20V. Firewire ports deliver an unregulated power rail with a voltage that varies over a wide range, typically 9V to 26V depending on the class of the power provider.

The ubiquitous wall adapter remains perhaps the most common source of power for portable devices. A typical wall adapter is simply a transformer followed by a bridge rectifier, offering no active regulation. That task is left to the end device to avoid the effects of cable drop. Unregulated wall adapters are designed to provide rated current at the specified typical output voltage. Being unregulated, the output voltage is a load line function, increasing substantially at lighter loads and decreasing

The LTC3115-1 provides uninterrupted operation through load dump transients and even the harshest cold-crank conditions. Its programmable switching frequency optimises efficiency and supports operation at 2MHz to ensure that switching noise and harmonics are located above the AM broadcast band. It employs a proprietary low noise PWM control algorithm that minimises electromagnetic emissions over all operating conditions even during transitions between the step-up and step-down modes of operation and over the full range of load current. An internal phase-locked loop allows switching edges to be synchronised with an external clock for further control of EMI in noise-sensitive applications.

An accurate RUN pin provides a programmable input undervoltage lockout threshold with independent control of hysteresis. By consuming only 30μA of quiescent current in Burst Mode operation and 3μA in shutdown, the LTC3115-1 reduces standby current drain on automobile batteries to negligible levels.

The LTC3115-1 is also well suited for handheld devices, which are required to interface to an expanding array of power sources. While it was once common for portable devices to be powered by a dedicated AC adapter or a single power source, many must now be compatible with a variety of inputs including automotive, USB, Firewire and unregulated wall adapters. Next generation military radios and support electronics are an extreme example, requiring the capability to operate from all available power sources for emergency use and to minimise the number of battery varieties carried in the field.

under heavy load. In addition, the AC line voltage is permitted to vary between 105V and 125V, adding an additional 10% variability in the unregulated wall adapter's output. It is not uncommon for a 12V unregulated wall adapter to produce an output voltage of 17V or greater at light load.

The LTC3115-1 operates directly from all of these portable power sources as well as from a variety of battery chemistries including lithium (single cell or series connected), sealed lead acid, three or more series alkaline cells and even a bank of supercapacitors for backup applications. Multiple power sources can be combined through a Schottky diode-OR circuit.

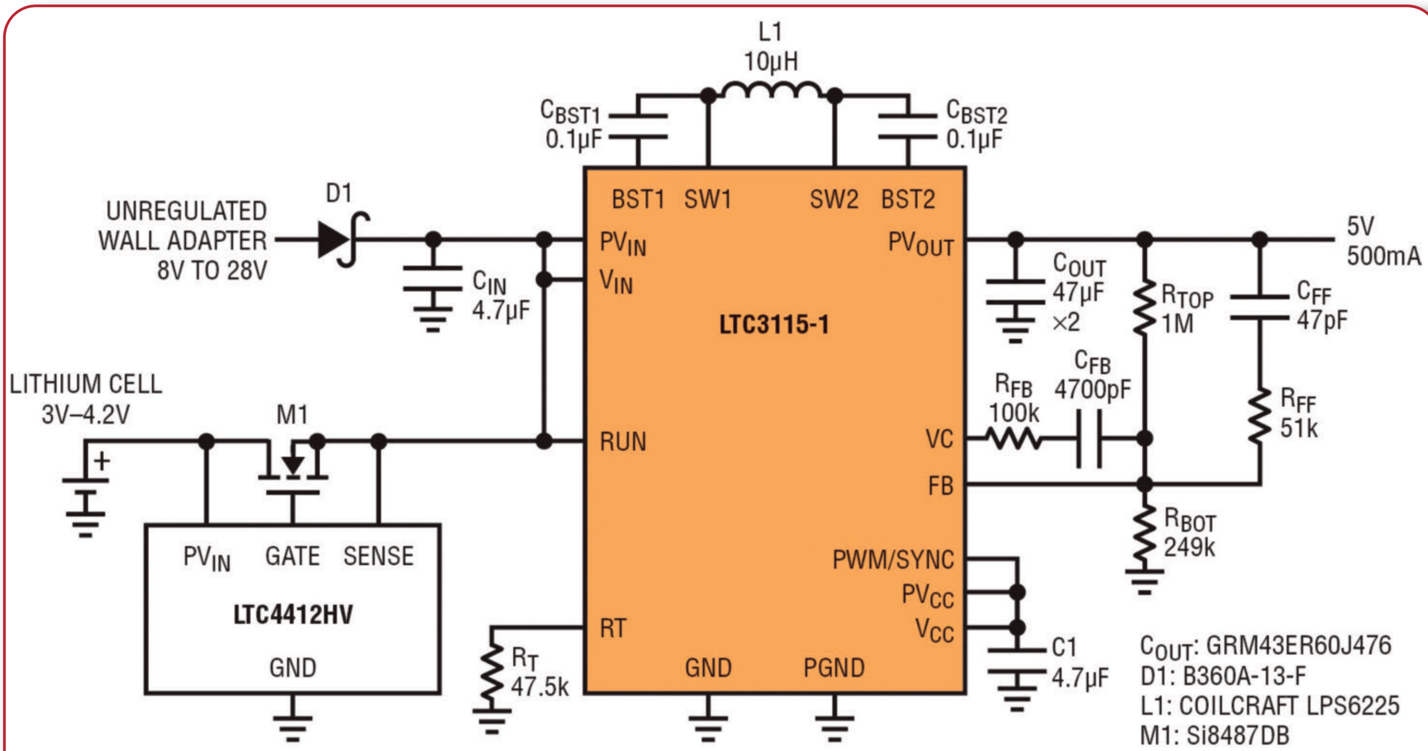

Figure 6: For high efficiency, this dual input 5V supply uses a LTC4412 low loss PowerPath controller and a P-channel MOSFET in the battery path instead of a Schottky diode. An inexpensive Schottky diode is used on the higher voltage input where its voltage drop is insignificant.

For higher efficiency, the LTC3115-1 can be combined with an ideal diode PowerPath controller to provide automatic switchover between multiple power sources using the low voltage drop of a power P-channel MOSFET to replace the Schottky diode. Figure 6 shows how the LTC3115-1 can be combined with the LTC4412HV to obtain a dual input – single lithium and unregulated wall adapter – 5V supply. In this case, a series PMOS is used on the lower voltage lithium input while an inexpensive Schottky diode is used on the higher voltage input where its voltage drop is insignificant. The overall efficiency of this supply including the converter and PowerPath is given in Figure 7 for each power input.

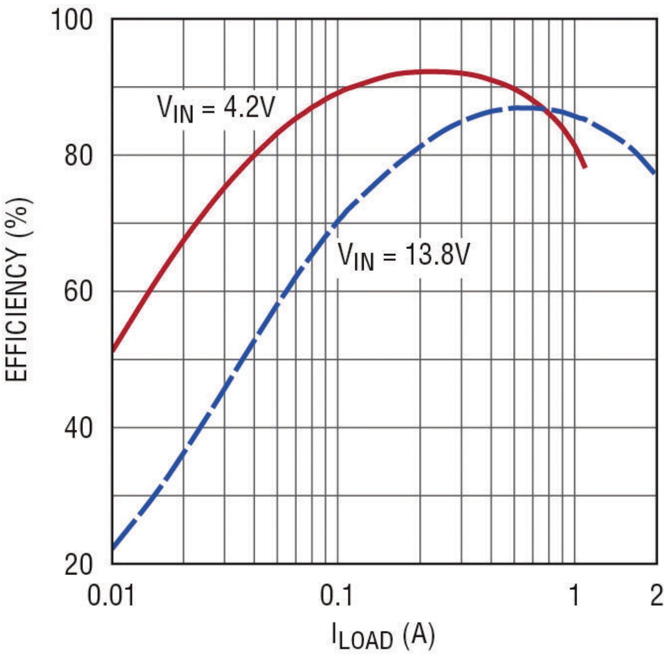

Figure 7: Overall efficiency of the PowerPath and LTC3115-1

The flexibility and high efficiency of the LTC3115-1 make it suitable for many next generation automotive electronics and portable devices, especially those operated from multiple power sources. Its internal power switches and programmable switching frequency minimise the power solution footprint, supporting the increasing demand for miniaturisation of electronic devices in the portable and automotive arenas.

Author profile: John Canfield is a Design Engineer with Linear Technology Corp,.

More from Linear Technology
Return to contents page.

Electronic Specifier design

Cell-Aware ATPG

Finding, identifying and fixing manufacturing defects and systemic yield limiters within library cells at 90nm and beyond. By Stephen Pateras

The goal of silicon test is to find defective parts before they are shipped to the customer. The widely used methodology is to add scan test structures to the design, then deliver test patterns through the test structures that reveal defects when the chip responses are observed. Test patterns are generated from fault models that represent potential defects with simple properties. These fault models are based on observations of silicon defect behaviour, and are developed by the electronic software design automation (EDA) vendors.

This provides high fault detection, but mostly for defects at the gate (or cell) boundary or between library cells. However, today's fault models like stuck-at, transition- and path-delay are no longer sufficient for detecting all defects within cells. TSMC has stated that, "for 90nm and beyond, a significant number of manufacturing defects and systematic yield limiters lie inside library cells." With more recent fabrication technologies, the population of defects occurring within cells is significant, perhaps amounting to roughly 50% of all defects.

The traditional approach finds some of the defects within cells, but many cell-internal defects require a unique set of stimulus to excite and observe the defect. Particularly for zero-defect designs, or those using emerging technologies like FinFETs, or to simply get a quality edge over your competition, you need to use these "cell-aware" fault models that specifically target the defects within cells. A new methodology to do this, called cell-aware test pattern generation (ATPG) is available to define and use fault models that target these cell-internal defects.

The cell-aware ATPG methodology has been shown to increase defect coverage by over 1% compared to what can be achieved with traditional test patterns. Multiple studies have been performed that show tests generated using cell-aware ATPG will find defects that the normal test methods miss. Published reports have shown improvements over 800 DPM (defects-per-million) with cell-aware testing (Figure 1), which is very significant for many production environments.

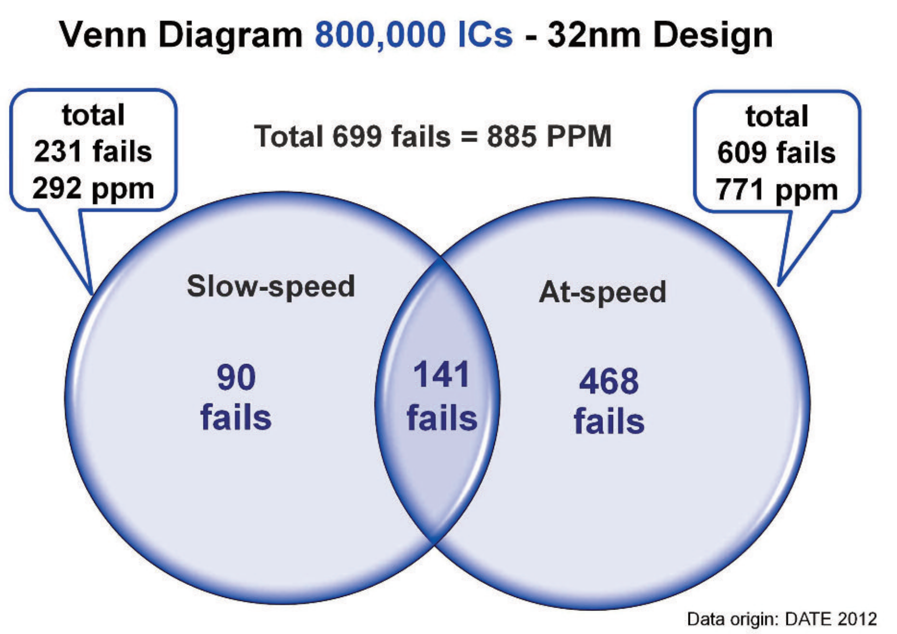

Figure 1: A summary of the part-per million (PPM) defective part reduction results from an industrial 32nm design. Slow-speed cell-aware patterns detected 231 defective devices and at-speed cell-aware tests found 609 defective devices that were missed by state-of-the-art transition and stuck-at patterns.

Why cell-aware?

During normal ATPG, thousands of patterns are produced. So even if a cell-internal defect was not targeted by a traditional

Figure 2: Generating cell-aware fault models through library characterisation

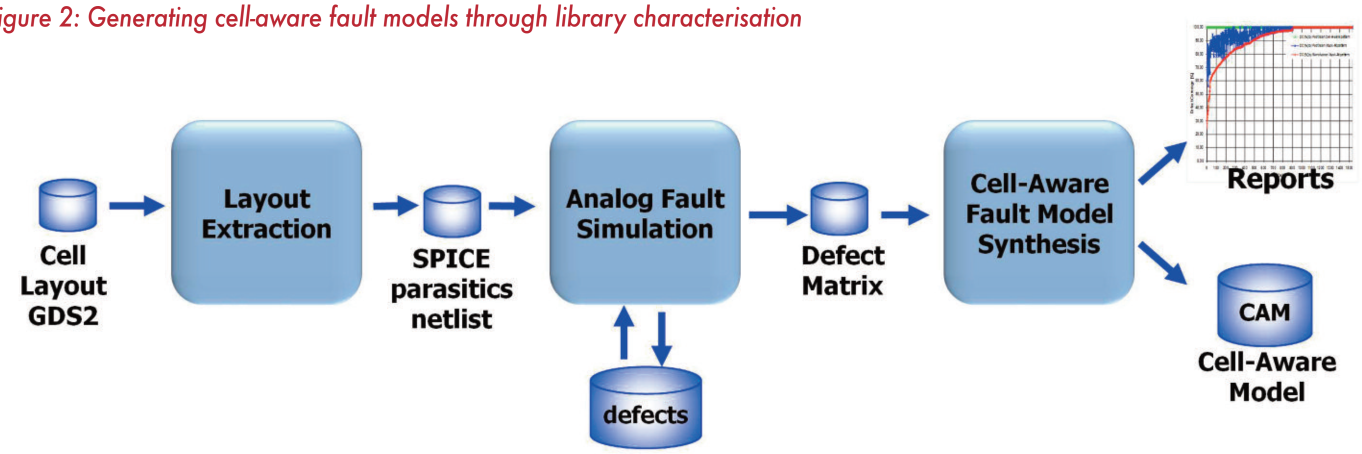

fault model, it could be fortuitously detected (detected by random luck). However, when considering millions of gates in a design, it is not effective to rely on luck to detect these potential cell-internal defects at every gate. One option would be to apply every possible combination of inputs at every gate. This fault model is referred to as a gate-exhaustive fault model. It would be effective in detecting many cell-internal defects since it would apply every possible combination. For example, for an 8 input cell, gate-exhaustive testing would apply all possible 28 (256) input combinations. It is easy to see that to apply such an exhaustive set of patterns quickly becomes impractical.

Furthermore, many defects inside cells can manifest themselves as timing-related defects and are not detectable using static tests. A two pattern test is necessary to detect such defects. So for our 8 input cell example, two cycle gate exhaustive testing would require the application of 28 x 28 = 216 patterns. A much more efficient method is clearly necessary to target static and dynamic cell-internal defects for low DPM products. This is the objective of cell-aware ATPG.

Cell-aware ATPG starts with an automated cell library characterisation process, shown in Figure 2. The goal is to convert potential physical defects that can manifest themselves in each technology cell layout into a fault model that ATPG can process.

Each fabrication process has a set of technology libraries used to describe the logic behaviour and physical layout of the lowest level component in the netlist. Usually, the ATPG and time-based simulation processes are based on a logical library composed of logic models. A corresponding physical library exists for the IC layout. The cell-aware characterisation process starts with an extraction of the physical library, represented in GDSII.

Each extracted cell results in a transistor-level design with parasitic resistance and capacitance. A resistance location represents a conductive path with the potential for an open defect and capacitance identifies locations with the potential for a bridge defect. Each potential defect is SPICE simulated against an exhaustive set of stimuli to determine if there are sets of cell inputs that produce an output different from the defect-free result.

The simplest case is to simulate each capacitive location with a 1Ω resistance representing a hard bridge. There are many possibilities to model other defects using analog simulations. For example, various resistances can have different effects and might detect defects that aren't detectable with hard shorts (1Ω bridges). In addition, simulating over multiple cycles is also useful to detect bridges or opens that are only observed as dynamic defects.

The final process in cell-aware characterisation is to convert the list of input combinations into a set of the necessary input values for each fault within each cell. Because this information is defined at the cell inputs as logic values, it is basically a logic fault model representation of the analog defect simulation. This set of stimulus for each cell represents the cell-aware fault model file for ATPG. Within this file, a simulated defect (now a fault) can have one or more input combinations.

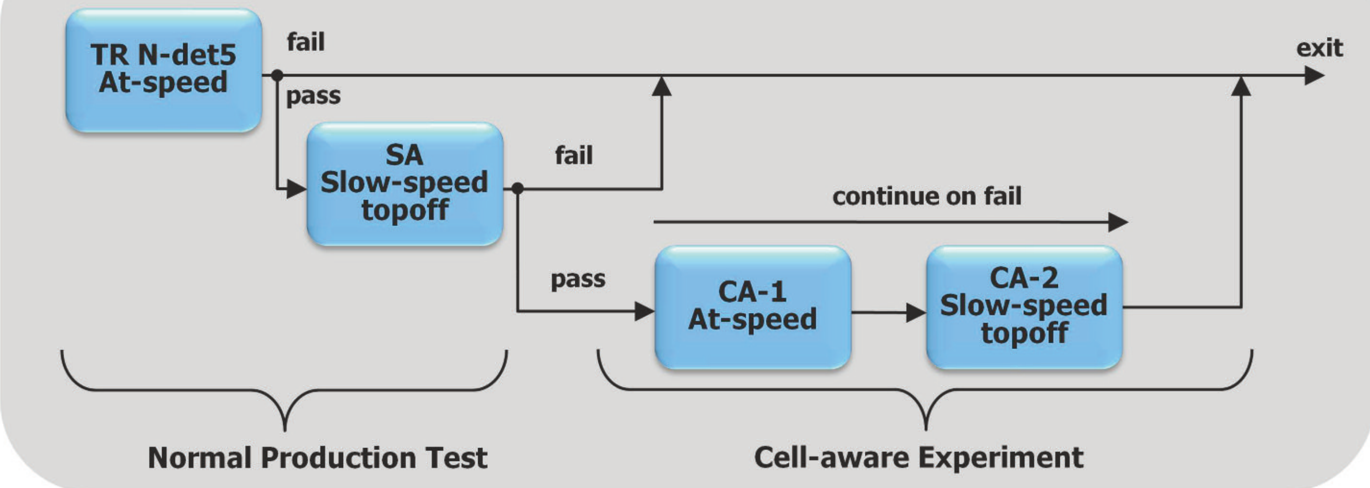

Figure 3: Production test flow for measuring cell-aware test effectiveness

Note that because the cell characterisation process is performed for all cells within a technology library, any design using that technology can read in the same cell-aware fault model file. Characterisation only needs to occur once and then can be applied to any design on that technology node.

Measuring effectiveness

When evaluating the effectiveness of a new fault model it is important to determine if it uniquely identifies defects the other fault models do not and the impact of the fault model on test time and pattern count.

Several manufacturing test experiments have been performed to measure the above parameters. In these experiments, the tester is typically setup to be in data collection mode, also known as 'continue on fail'. A typical flow is illustrated in Figure 3. The production test flow is to perform at-speed transition tests. For all dies that pass this test, top-off stuck-at testing is performed. For all dies that pass the stuck-at top-off tests, both the cell-aware at-speed and cell-aware slow-speed top-off patterns are run. Any die that fails either or both of these tests are logged. With sufficient volume, this process allows to calculate the reduction in DPM that is achieved by applying the cell-aware patterns. The greater than 800 DPM result mentioned earlier is one such value. Many other experiments have shown similar strong results.

Traditional fault models fully test the periphery of standard cells and the interconnections between them, but they can miss some bridging or open defects internal to the cells. Advanced process technologies, notably the use of FinFETs, introduce a variety of new failure modes that need to be addressed during IC testing.

The cell-aware characterisation process can create a set of fault models based on the simulated behaviour of defects within the cell layout. As a result, a higher quality pattern set can be produced by ATPG enhanced to support the new fault model. Silicon results have already shown significant additional defect detection beyond standard stuck-at and transition patterns when using cell-aware ATPG.

Author profile: Stephen Pateras is product marketing director for Mentor Graphics Silicon Test products. He was previously the VP of Marketing at LogicVision. While at LogicVision Pateras also held senior management positions in engineering, and was instrumental in defining and bringing to market several generations of LogicVision's semiconductor test products. From 1991 to 1995, Pateras held various engineering lead and management positions within IBM's mainframe test group. He received his Ph.D. in Electrical Engineering from McGill University in Montreal, Canada.

More from Mentor Graphics
Return to contents page.

Roaming seamlessly

With the demand for data only likely to increase, service providers are looking towards new technologies like HS2.0 to help quench consumers' thirst for data and greatly increase the capacity of their networks. By David Stephenson

Today's cellular networks are being overwhelmed with data traffic, much of it being generated by the rapid proliferation of smartphones. Operators have realised that 4G will not alleviate the capacity demands for data; many operators understand that their data needs cannot be met by enhancements to their cellular networks and are turning to Wi-Fi as a solution. Working alongside cellular networks, it is the perfect technology to stabilise the data capacity issue. Wi-Fi is seen as an excellent option as it has access to upwards of 600 MHz of spectrum, supports dense Access Point (AP) deployments, is available on all data-centric devices, and is available in all locations where people congregate.

However, capacity and ease of deployment are only the first steps in enabling a carrier-class solution. The industry is now focussed on improving the Wi-Fi user experience while roaming, providing a service that is akin to that of cellular networks. This work is known as Hotspot 2.0 (HS2.0) and is being driven by the Wi-Fi Alliance (WFA), which also certifies HS2.0 interoperability as part of their Passpoint program. The Wireless Broadband Alliance (WBA) is also very much involved in this process through their Next Generation Hotspot initiative. It comes as no surprise that many see it as 'the next big thing' when it comes to wireless technologies.

Revolutionising Wi-Fi

HS2.0 is the new technology getting carriers excited, because it takes away many of today's manual tasks – like authentication – and automates them; it also enables users to roam securely without the hassle of searching for Service Set Identifications (SSIDs).

The key enabling protocols are IEEE 802.11u, along with IEEE 802.1X, selected Extensible Authentication Protocol (EAP) methods, and IEEE 802.11i. The latter three are part of the WPA2-Enterprise certification program in the Wi-Fi Alliance, and are standard on

Current activity within the HS2.0 sphere

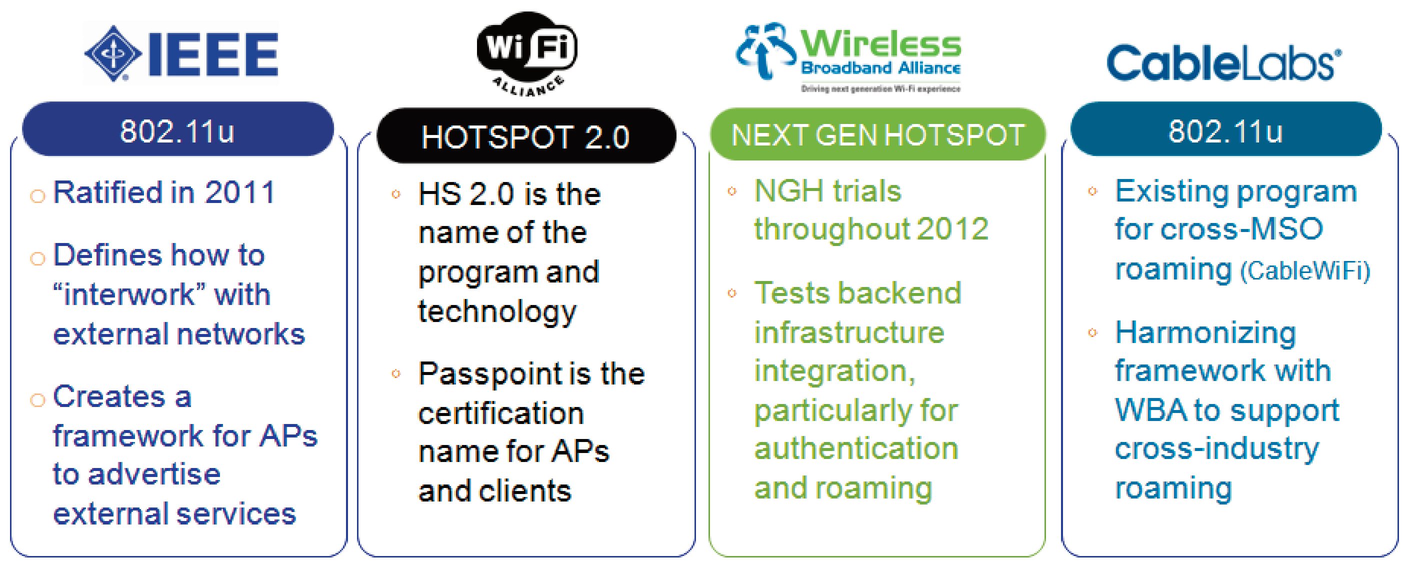

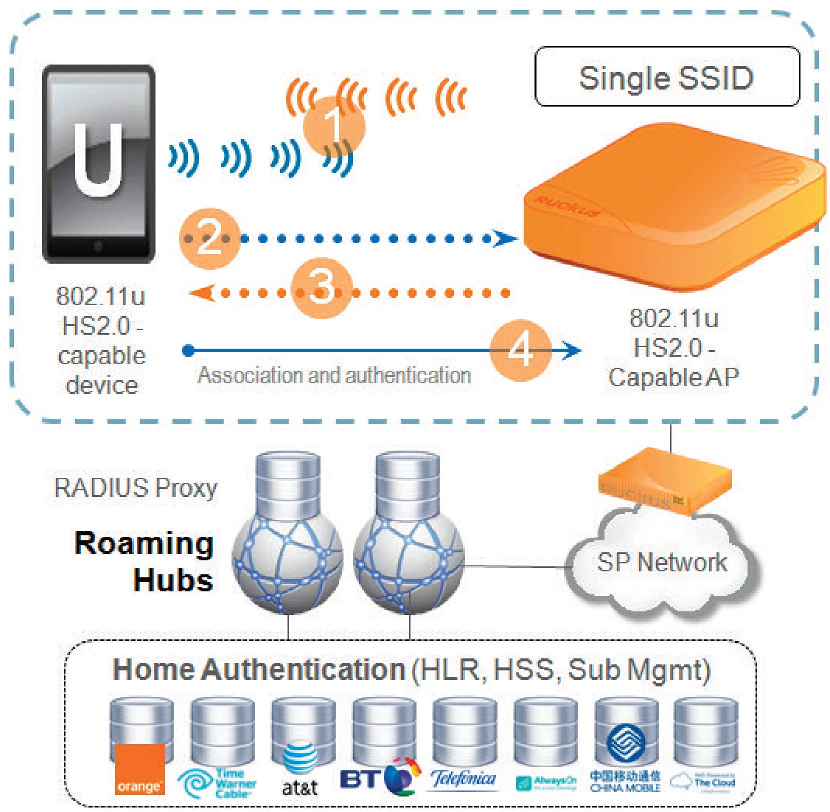

1. 802.11u -capable AP broadcasts its HS2.0 support, while HS2.0-enabled devices look for support
2. Device processes information contained in ANQP payload - such as 'reachable' authenticators, and capabilities of the hotspot
3. (Optional) Device requests full authenticators list, the AP responds to ANQP query with requested information
4. Device compares provisioned network-selection policy with HS2.0 data from APs and associates itself to the best BSSID

all smartphones. While the certification is called 'WPA2-Enterprise', the end result is a process that is every bit as secure and easy to use as what exists in the cellular world.

The IEEE 802.11u protocol enables a mobile device to have a dialog with a Wi-Fi AP 'pre-association' to determine the capabilities that the network can support. The two protocols that 802.11u uses to make this happen are the generic advertisement service (GAS) and the access network query protocol (ANQP). These protocols run on top of 802.11 and enable the Hotspot 2.0 experience.

When a user with an HS2.0 capable mobile device comes within range of an HS2.0 capable AP, it will automatically open up a dialog with that AP to determine its capabilities. This is done using ANQP packets that are carried at layer 2 by the GAS service (Note: the device has not yet attached and does not yet have an IP address). It is the exchange of ANQP packets that allows the mobile device to automatically learn the capabilities of an AP.

If the AP is part of the user's home network then no roaming is required and the user can move straight to authentication. However, if the AP is not on the user's home network, then roaming is required. In this situation, the list of roaming partners that are supported by the AP must be passed down to the mobile device via the ANQP protocol. This can be provided in the form of a PLMN (Public Land Mobile Network), ID, realm, or the organisational Identifier. Once the mobile device learns the roaming partners and the identity of the AP operator, it invokes some basic, built-in network selection policies to determine which AP to join. This means a user can roam on a Wi-Fi network even if a mobile device does not see an AP belonging to its home network provider.

This is the beginning of a new era where users no longer have to think about SSIDs, authentication or fumble around with passwords. It will make it possible to link a huge network of effectively random Wi-Fi access points through a web of interconnections, so that users can enjoy a seamless experience as they move between Wi-Fi networks from almost any location. This is very much like the cellular experience that we all enjoy when getting off an airplane just about anywhere in the world. Wi-Fi roaming would apply anytime a mobile device does not see an AP belonging to its home network provider. A user could roam on a Wi-Fi network that is across town or on the other side of the world. Roaming partners could now potentially include MSOs, MNOs, wireline operators, public venues, enterprises, and basically any other entity that has Wi-Fi assets.

New opportunities

HS2.0 has been developed and promoted predominately by carriers and equipment suppliers, but it is very possible that it will have its greatest impact and appeal within the enterprise; that's what will make it a game-changer. HS2.0 will be about much more than the technology enablement of a better mobile user experience; it will shift relationships between carriers and building owners who may already have a Wi-Fi

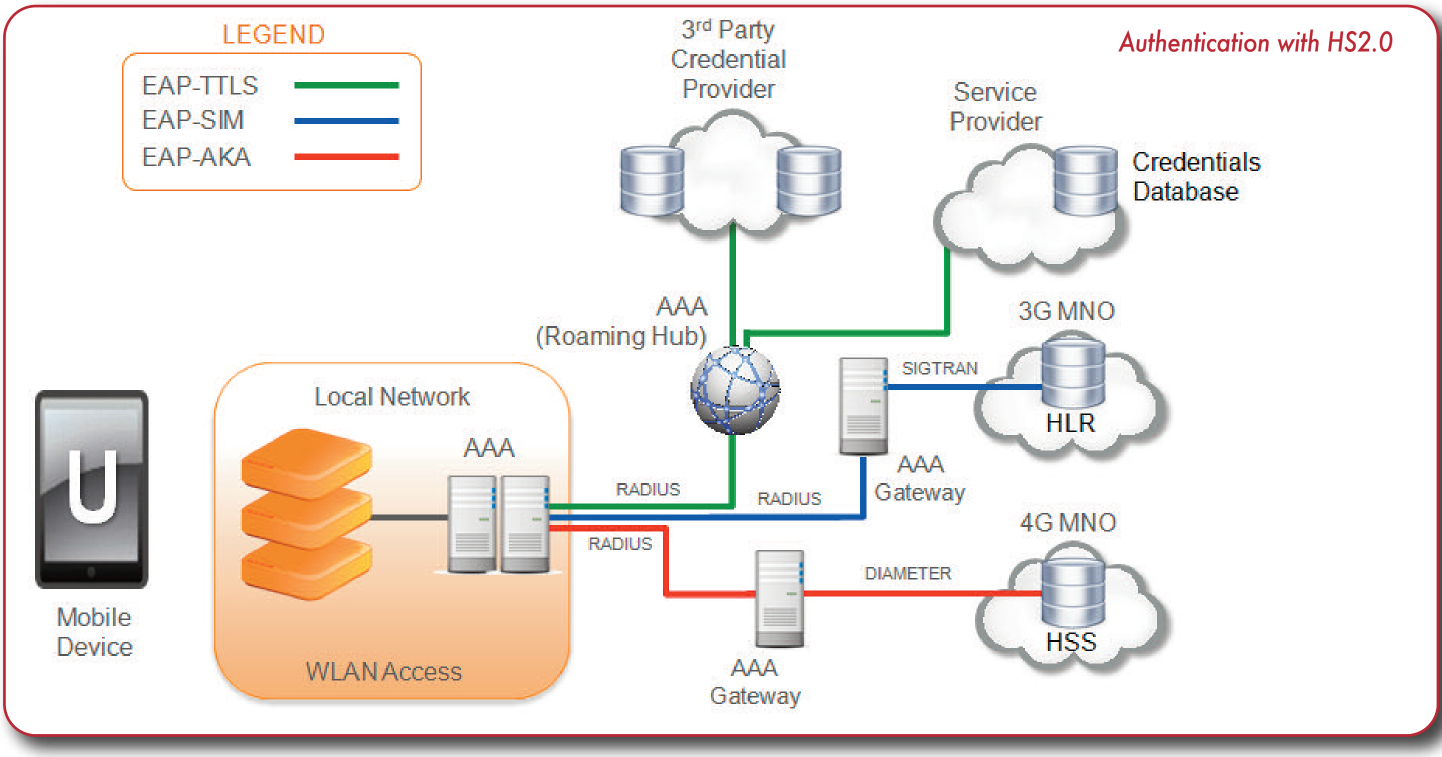

footprint deployed – those who want to provide the uninterrupted service as part of their continued strategy to deliver better subscriber experiences, and those who own the locations essential for providing that continuity of service. This commercial collision course between the worlds of the Service Provider and the Enterprise, combined with an important leap forward technologically, is going to give HS2.0 its place in the Wi-Fi hall of fame.

Two principle parties are interested in the provision of Wi-Fi services: the owner of the venue or building, and the service provider. Now their interests are coinciding. The widespread and growing use of Wi-Fi across public venues such as hotels, schools, shopping centres, retail outlets, public transport, sports venues – in fact, anywhere where people gather and expect to use their mobile devices without encountering any problems – is an incredible opportunity for venue owners, and for the enterprise. These are usually the owners of the network infrastructure. Since operators want Wi-Fi network access, the real opportunity will emerge for any enterprise or venue owner to wholesale their existing wireless LAN capacity to operators; charging them recurring fees for that access.

Mobile service providers want to automatically connect their subscribers to their own 'branded broadband' service, through the venue's available high-speed Wi-Fi network. It is this connection that HS2.0 will make possible; giving the Wi-Fi network an interconnection with the subscribers' 'home' service provider, so the device just carries on functioning in the way we all expect it to in the 21st century. These back-end connections might be direct, but more likely will be indirectly provided through third-party hubbing services. Either way, from a user perspective, the connections will be seamless and automatic, with no loss in service.

Becoming truly global

One of the key challenges for Hotspot 2.0 in becoming a truly global technology is its compatibility with mobile phone handsets and other Wi-Fi enabled devices. Passpoint is the Wi-Fi Alliance's designation for devices and network equipment that supports the Hotspot 2.0 standard. The first signs of its uptake come from the fact that the new Samsung Galaxy S4 is Passpoint enabled. With Samsung starting to ship Passpoint on the Galaxy S4, we can expect the other major smartphone vendors to follow suit in short order.

However, challenges don't just lie in the compatibility of devices, it is also vital that Wi-Fi access points are upgraded to become compatible with Hotspot 2.0, in order to provide the smooth transition of Wi-Fi across disparate locations. Only then will people be able to walk around with a continuous Wi-Fi connection without risk of drop out.

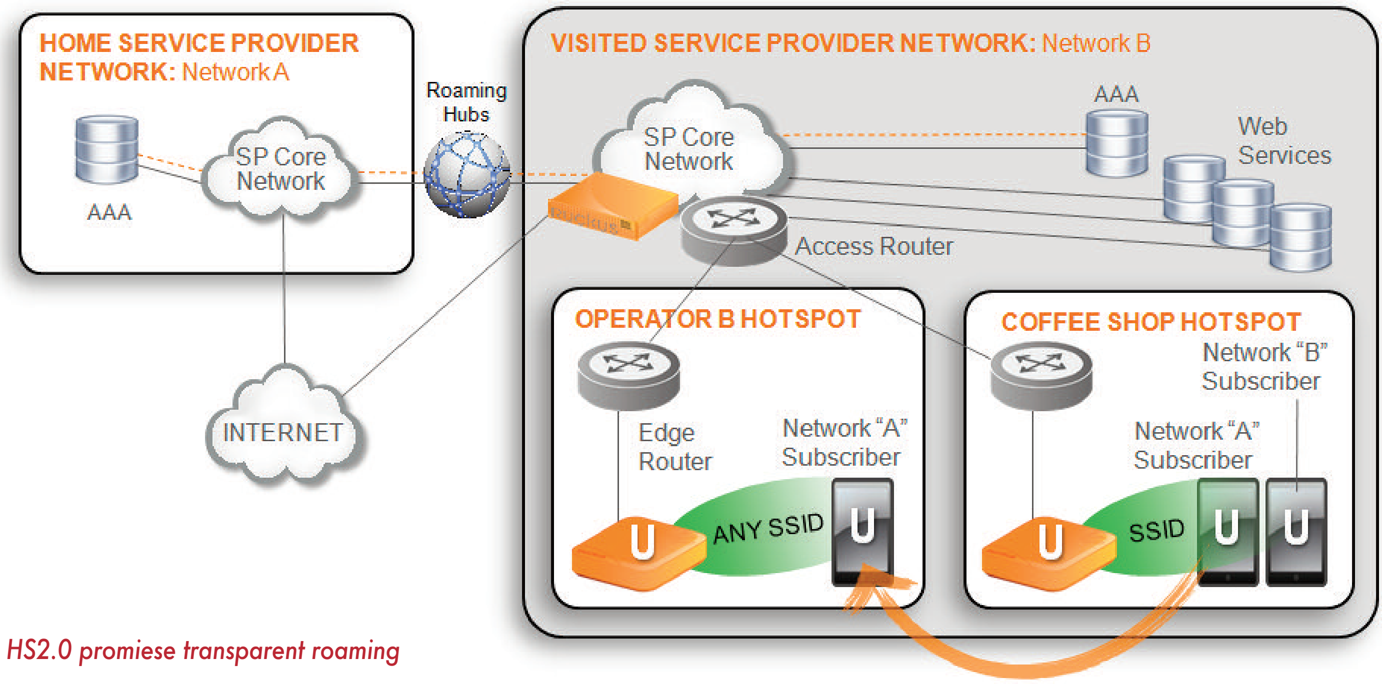

HS2.0 promiese transparent roaming

In addition, HS2.0 enabled APs must meet the critical requirements that carriers demand for a seamless Wi-Fi experience. In order to provide a carrier grade solution, APs must have a strong radio performance, both for access and backhaul, in order to face the unprecedented subscriber growth, high interference, and challenging high density deployments. Without addressing these issues, the APs will provide a poor user experience, with only a limited number of people able to connect to a single AP, poor signal due to intense RF environments and unreliable service in high density environments. Enterprises and carriers alike must ensure they are using the smartest Wi-Fi technology in order to provide the service that consumers crave.

The impact of Hotspot 2.0

HS2.0's impact on the industry will be enormous. Mobile operators are already seeing their networks overloaded by data traffic and are looking at all available options to increase densification. At the top of their list are technologies like Wi-Fi and LTE small cells. Cable and wireline operators are taking advantage of their backhaul capabilities to rapidly build-out an extensive Wi-Fi footprint. This technology has also been extensively deployed in public venues like hotels, airports, exhibition centres, stadiums, hospitals, etc. With HS2.0, it will now be possible to link together this huge footprint of Wi-Fi APs through a web of roaming arrangements. Users will be able to seamlessly roam onto Wi-Fi networks from almost any location.

The net result for mobile operators is much greater network capacity than could be achieved by building out a network of APs on their own and a much better experience for the subscriber. Users no longer need to know or care about SSIDs and authentication protocols. Instead, they get an always best connected experience.

Venue owners, enterprises and operators can begin to better monetize their Wi-Fi network investments through these roaming arrangements and the settlements they will entail. A mobile operator that deploys a Wi-Fi network in a stadium can now monetize that asset by allowing subscribers of other operators to roam onto that network. Enterprises can likewise allow subscribers of all different mobile operators to roam onto their in-building Wi-Fi networks, turning it into a profit centre rather than a capital expenditure cost.

In short, Hotspot 2.0 technology will radically transform the wireless industry, and it is set to emerge in 2013 in a very big way.

Author profile: David Stephenson is a Senior Principal Engineer with Ruckus Wireless.

More from Ruckus Wireless
Return to contents page.

Thinking networks

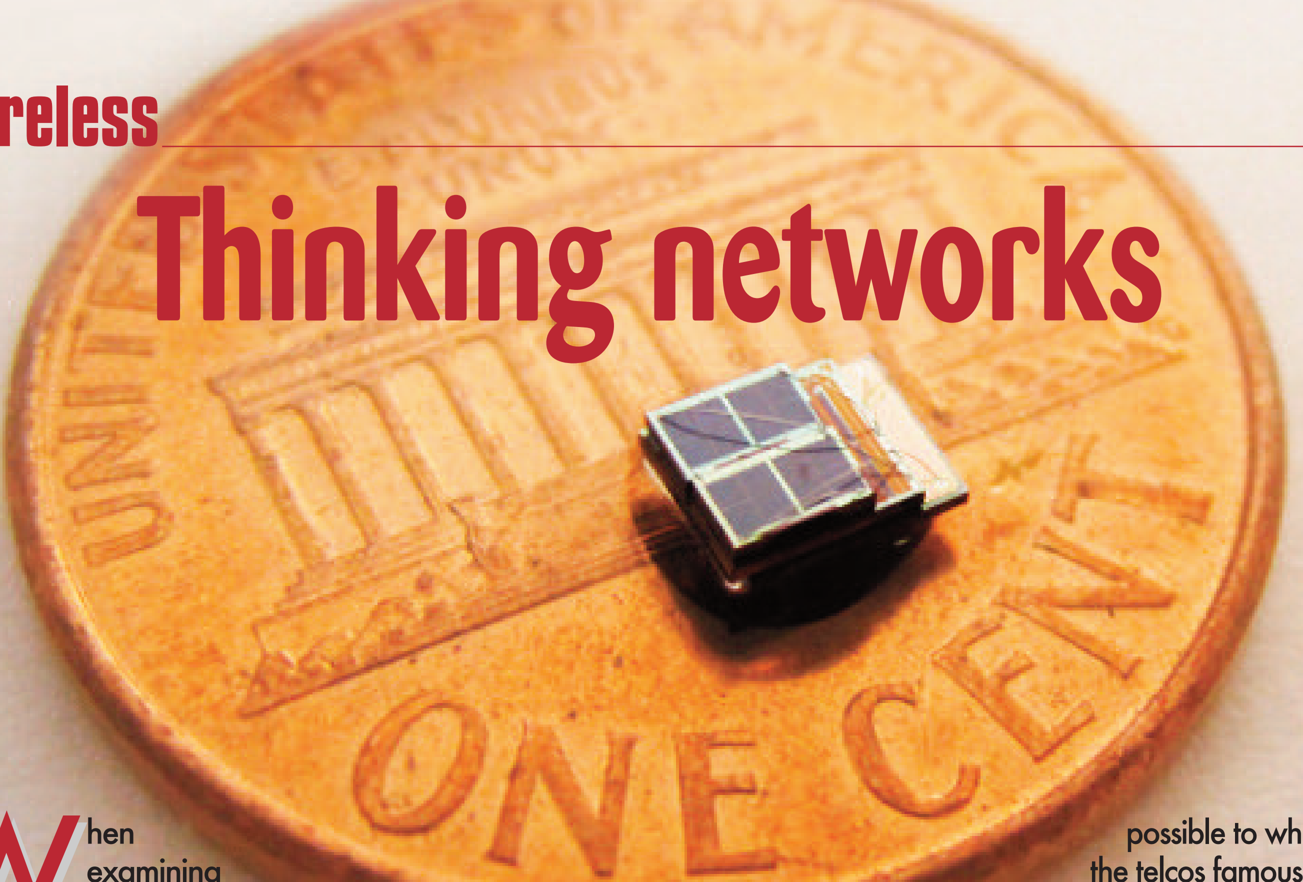

Photo Credit: University of Michigan low-power, smart sensor Phoenix system: Daeyeon Kim (http://www.engin.umich.edu/newscenter/feature/smallsensor/)

When examining Flame, the most sophisticated malware that has appeared to date, investigators discovered an interesting feature: Flame can steal and transmit data from computers that have no Internet connections.

It does this by using unsuspecting humans for bi-directional data transport. The process begins by copying itself to every digital storage device that it encounters, including USB sticks and external hard drives. When humans hand-carry portable digital storage devices to unconnected PCs and desktops, Flame copies itself into the new computers and starts stealing data. At this point, traditional spyware would have to upload the loot to a remote Internet server, but Flame is a bit more sophisticated; if there's no Ethernet connection Flame just waits for the next external storage device to come along. When one does, Flame copies itself again and brings along a copy of the stolen data. It repeats the process until it finds a computer with an Internet connection, and then it starts transmitting.

Looked at in a certain way, you could say that Flame uses its human data mules as high-latency Ethernet connections. You could also argue that it's an early example of a promising new tool in data communications: Disruption Tolerant Networking (DTN).

New approach to networking

The old network model assumed that the bulk of a network's intelligence was located in controlling computers. If connections were lost, data would be lost as well. So network designers tried to get as close as possible to what the telcos famously refer to as 'Five Nines' (99.999 percent) uptime. The closer you got to perfection, the harder it was to achieve the next incremental improvement and the more expensive things became.

That model is changing; downstream devices aren't necessarily 'dumb' anymore. As integrated circuits become steadily smaller and more powerful, and as software becomes more sophisticated, it is becoming easy and cost-effective to distribute localised intelligence all across the network.

The University of Michigan, for example, has developed a low-power, smart sensor system that demonstrates many of the key principles that will be employed in the smarter, 'thinking' networks of the future. At just nine cubic millimeters it's the size of a Vitamin C tablet, but is solar powered, has an internal battery and radio, and is equipped with its own processor called the Phoenix (photo). The processor employs a unique power gating architecture and an extreme sleep mode to achieve ultra-low power consumption.

Smart network nodes like the Phoenix system will provide network designers with opportunities that haven't been available in the past. Smart nodes will require less bandwidth, for example. They'll be equipped with situational awareness programming that considers parameters like power, network availability and the status of surrounding nodes to make independent decisions about whether there is any need to log in and report. Smart network nodes will also be able to collect data, time stamp it, log it, and – like

Flame – report the data whenever a network connection becomes available. Even if the network is down for minutes or hours, this Disruption Tolerant Networking model will ensure that data is not lost, thus freeing designers from the need to pursue 'Five Nines' uptime.

With their increased efficiency and intelligence, and their ability to hibernate when they have no useful information to report, smart network nodes will require far less power than their 'dumb' predecessors. When combined with advances in power harvesting, like the tiny solar panel demonstrated on the Phoenix sensor system, this will give thinking networks the ability to extend the network edge to include locations and applications that would previously have been completely inaccessible. The thinking networks of the future will include network nodes that are completely independent of the power grid.

Hardware for thinking networks

Like Flame, with its human data mules, thinking networks will use multiple techniques to transmit and deliver information. Single vendor solutions and proprietary data communications protocols have already become a thing of the past. Users are demanding network-wide interoperability and the ability to make use of any data communications option that may be available. Legacy serial devices are being Ethernet-enabled with Wi-Fi connections using both embeddable and external Wi-Fi Access Points. Where a fibre optic buildout would be impractical, designers are deploying cellular routers that can establish network nodes anywhere there's cellular telephone coverage. Thinking networks will continue the trend, using various combinations of copper cable, fibre optics, wireless and cellular data transmission. And as the network nodes get smarter, the methods used to transmit data will become increasingly irrelevant. Like Flame, thinking networks will find their way around obstacles, take advantage of whatever connection options are available, see to it that the data gets where it needs to go, and ensure that it arrives intact.

None of this will make life any easier for network designers. Why? More and more network nodes will become independent of cable connections and the power grid. So even though network designers will be able to invest less time and energy in pursuing perfect uptime, they'll be forced to start thinking about power budgets. Recommendations for things like 30-day battery replacement cycles will be unacceptable. Network designs will have to use power efficiently enough to ensure that remote network nodes and devices will never go dark. Traditional 'always on' communications schemes won't get the job done.

Pros & Cons

Distributed network intelligence and integrated communications infrastructures will be engines for increased efficiency and productivity. It's estimated that 50 billion devices will be network-enabled by 2020. Many of them will be M2M devices that communicate with one another to resolve problems with no need for human intervention, much the way 'smart' metering relieved utilities of the need to send employees out to visit the meters in person. Thinking networks will ultimately be able to deliver just about any data, just about anywhere, and the transmission methods involved will be completely transparent to the end user, whether that end user is a machine or a human being.

But the same hardware and software that help data travel across thinking networks will greatly complicate network security issues. Flame malware has already found a way to access the Internet from locations in which no Internet connection exists. We will soon live in a world where Internet connections are virtually everywhere.

Author profile: Mike Fahrion is the Director of Product Management of B&B Electronics

More from B&B Electronics
Return to contents page.

It's time to appcessorise!

Photo courtesy of: Wahoo Fitness

Adhering to a few simple guidelines and the careful selection of components are key to achieving long battery life in wireless applications. By Jay Tyzzer

Ultra low power (ULP) wireless is on the up. Proprietary 2.4GHz technologies have steadily carved a growing niche in applications as diverse as wireless desktops, remote controls and fitness monitoring sensors such as heart rate belts. But it is the introduction of interoperable technologies, allowing OEMs a choice of chips from a multi-vendor environment that are guaranteed to work together, that's really going to push this technology into the mainstream.

Technologies such as ANT, from Canadian developer Dynastream Innovations, and Bluetooth low energy (a part of the latest Bluetooth v4.0 specification) are receiving support from major consumer electronics manufacturers such as Apple, Microsoft and Sony-Ericsson. Devices from these manufacturers incorporate software and hardware that make it easy for developers to design ultra low power wireless peripherals that connect seamlessly with the companies' PCs, tablets and smartphones.

These peripherals plus an associated downloadable app are fuelling a new market for so-called 'appcessories'. An example is a blood glucose monitor incorporating Bluetooth low energy paired with a Microsoft Surface tablet computer. The diabetic's blood glucose readings could be automatically transferred from the meter to the tablet via Bluetooth v4.0 and an app could indicate trends and perhaps highlight the risk of a hypoglycaemic (low blood glucose) episode before it occurs. Better yet, on their next visit to the specialist, the diabetic's monitor could transfer all its data to the physician's Windows 8 PC, allowing him to make a rapid and accurate assessment of the condition. But this application represents just a drop in a huge ocean of potential.

Analysts have a bit of a reputation for exaggerating the potential of new technologies, but even if they are only partially correct about the potential of the appcessories market it will still be huge. According to consultants Juniper Research, for example, the market for wireless accessories linking to smartphones and tablets will reach 110 million units a year by 2017.

Longer life

The key to ULP RF technology's success is that it extends wireless connectivity to devices that – for cost, size, weight or accessibility reasons – are unable to support the batteries needed for technologies such as classic Bluetooth (versions prior to v4.0) or Wi-Fi. Even 'low power' wireless technologies like ZigBee and low-power Wi-Fi consume too much energy for typical ULP applications.

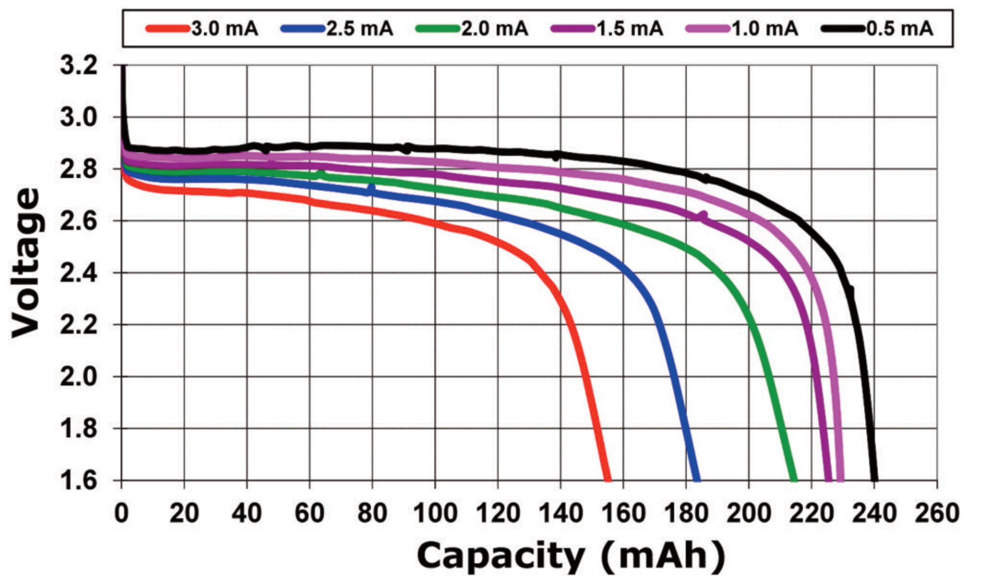

Figure 1: CR2032 coin cell continuous discharge battery voltage curves

A typical power source for ULP wireless applications is a 3V, 220mAhr CR2032 coin cell battery. While these devices feature impressive energy density, there is not a huge power budget to play with. And yet, consumers will not tolerate changing the battery in their heart rate belt or wireless mouse frequently. The design needs to cater for months of service or even up to a year.

Assuming a device such as a heart rate belt is used for an hour a day, the average current draw on the battery will need to be 220mAhr/365 x 1hr = 600μA for it to last a year. But there is much more to it than that; the pattern of demand influences how much of that initial 220mA will actually be available to the application and it's likely to be quite a bit less than the data sheet would have you believe.

ULP wireless applications are characterised by short periods of activity when the radio wakes up to transmit or receive data followed by much longer periods of inactivity. Duty cycles of just a few percent are typical. Of course, it depends on the application, but a 1ms burst of activity at a frequency of 40Hz is a good example. So the loading on the battery is a pulse pattern with a peak current draw dictated by the radio, microprocessor and any other component that needs to be powered while communication is underway. Unfortunately, drawing power from a coin cell in this way, particularly if the peak current is high, is a sure fire way to limit its capacity; even if the average current is in the microampere range that might otherwise be assumed to give good battery life.

Maximising accessible capacity

As is the case with all batteries, a lithium coin battery's available capacity (measured in milliamp hours (mAh)) is dependent on the rate of discharge. Figure 1 shows the voltage level of a CR2032 coin cell battery as it is drained with continuous currents ranging from 500μA to 3.0mA. When discharging the battery at a rate close to the rate stated in battery data sheets (< 500μA), it can be seen that discharge follows the familiar curve and that the full 230-to-240mAh capacity can be accessed before the voltage drops too low. However, as the rate of discharge increases, the effective capacity in the battery drops.

The accessible capacity – referred to as the Functional End Point (FEP) – will also be dependent on the voltage level. This is the minimum voltage level the electronic circuitry can operate at and is therefore a parameter you decide when choosing active components for your application. Let's assume that your application has an FEP of 2.0V. As seen in Figure 1, at 500μA a CR2032 can deliver its full capacity of 240mAh before reaching this FEP. But if the battery is drained at 2.5mA, the FEP is reached after only 175mAh of the battery capacity is used.

So much for continuous currents; ULP wireless application are characterised by pulse loads as

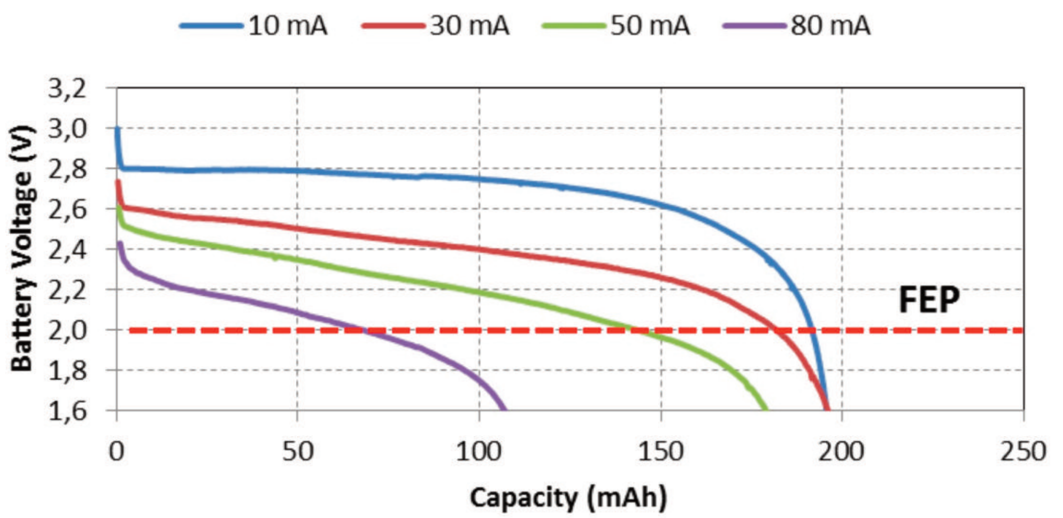

Figure 2: Peak current magnitude impact on battery capacity for CR2032 with a 25ms pulse cycle

the radio switches on to receive or transmit and then goes back to sleep. The pulses vary depending on the application and radio technology but for the purposes of illustration let's consider a fairly typical pulse period of 25ms (1ms 'on'/24ms 'off').

Figure 2 shows the drain curves when subjecting the battery to 10, 30, 50, and 80mA peak currents under this 25ms pulse cycle. At 80mA it can be seen that both the high peak current and the high average current challenge the battery. The 2.0V Functional End Point (FEP) is reached after approximately 70mAh. As the peak current drops, the 2.0V FEP creeps closer to 200mAh, but even at 10mA the peak drain is still too high for the battery to be able to deliver the full 240mAh capacity specified in the data sheet.

Figure 2 shows that, as long as the application can be designed with a FEP of 2V or lower, and peak currents lower than 30mA, the battery capacity loss is 'manageable', providing around 175 to 185mAh. But if the FEP is higher (> 2.0V), the impact from the higher peak current rapidly increases. For example, if the application has a FEP = 2.4V the battery can still deliver 175mAh when the peak current is restricted to around 10mA, but the capacity drops to 100mAh if the peak current is around 30mA.

Work fast and go back to sleep

Sometimes the answer to conserving energy can be counter-intuitive. Selecting a more powerful processor, for example, would seem to be a way to burn more power. But it depends on what you do with that processor.

Nordic Semiconductor earned its reputation designing ULP 2.4GHz transceivers. When the company designed its new range of wireless connectivity solutions, the nRF51 Series, the engineers wanted to improve on the competitive power consumption performance of the previous generation products but take advantage of the computational power of a 32-bit microprocessor instead of the 8-bit device used in the previous generation.

The processor core chosen for the nRF51 Series was the 16MHz, 32-bit ARM Cortex-M0 processor. Although the ARM core consumes about 4.4mA when executing code from Flash – a similar amount to the enhanced 8-bit 8051 core found in the previous generation chips – its average power consumption is lower than the older processor because a fast start-up (up to 100 times quicker at just 2.5μs) and rapid execution of complex code (the ARM core has 10 times the processing power of the 8-bit device) allows the ARM core to return to a low power sleep state very quickly. And, as we've shown, average power consumption – along with peak currents – is a key factor in determining battery life.

Apart from using an ARM core, the nRF51 Series included some other tricks to keep the system power down. For example, Nordic implemented a bus technology called Programmable Peripheral Interconnect (PPI) that allows the processor to initiate a peripheral function yet return to a low power sleep state while the peripheral devices complete the operation autonomously.

Further small power savings come from the new 2.4GHz radio that offers a 9.5dB improvement in radio link budget while consuming less than 10mA RX/TX peak current (using the on-chip DC-to-DC converter with a 3V supply). That compares with the 13mA peak current of the previous series.

Three power supply options are available: an on-chip DC-to-DC buck converter operating from a 2.1 to 3.6V supply; an LDO for 1.8 to 3.6V operation, and an LDO-bypass mode for operation between 1.75 to 1.95V, maximising the available capacity of the CR2032 coin cell. The end result is a chip that exhibits up to a 50 percent lower average current consumption than the company's previous generation of products (see Figure 3).

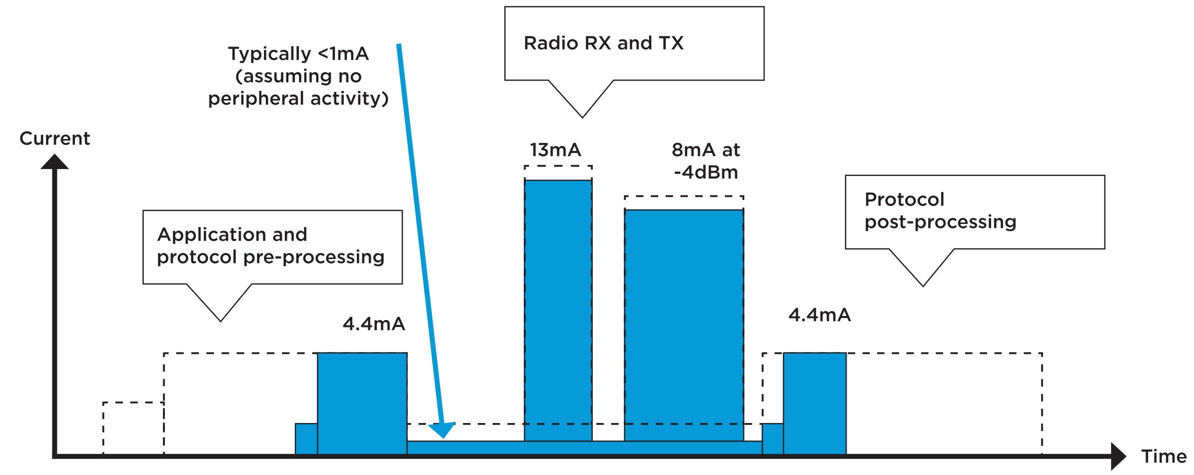

Figure 3: Power savings for the nRF51 Series compared with Nordic's nRF24L Series. Faster processor core, more efficient radio, and power-optimised architecture add up significant average current savings

Designing for ultra low power

If a designer is limited to drawing power from a modest coin cell, the design must work in the most efficient manner possible and so requires careful component choices. If the design is based round a 2.4GHz radio chip, it is essential to check the data sheet for both peak and average currents in applications similar to the target use case.

The typical pulse load of a ULP wireless application plays havoc with the battery life if the power peaks push up to 20 or 30mA. Also bear in mind that radio chips with a raw data rate of 250kbit/s have to be on air for much longer than devices with megabit bandwidths to transmit equivalent payloads, which pushes up the average current.

Beyond the radio chip consider maximising battery life by increasing the FEP margin by selecting active devices (ICs) that have as low a minimum supply voltage as possible, minimising the current drain by selecting devices with the lowest peak current, and designing the power management to prevent multiple high drain devices (for example RF chip, LEDs and LCDs) loading the battery at the same time.

Then identify the worst-case drain in the application by first ascertaining in which mode of operation the activity is highest and therefore the battery demand greatest. Search the IC data sheets for power consumption during normal operation and periods of RF activity.

Next, estimate your corrected battery capacity (which is likely to differ significantly from the nominal capacity in the data sheet); use this corrected battery capacity and the application's highest average current consumption to estimate the worst case battery lifetime using the formula:

$$\text{Battery life (h)} = \frac{\text{Corrected battery capacity (mAh)}}{I_{\text{AVERAGE MAX}}\ \text{(mA)}}$$

Note that as long as your application can stay in a mode with a lower average current consumption than that used in the calculation, it will be able to operate longer, but depending on how much battery capacity is left, it may fail once it again enters a mode with higher average current consumption.

Finally, remember that these rules apply for any high peak but pulsed battery load, not only wireless/RF circuitry. The high peak pulsed loads could also be from any combination of LEDs, vibrating motors, piezoelectric buzzers, LCD backlights and so on.

Author profile: Jay Tyzzer is a Senior Field Application Engineer with specialist 2.4GHz ultra-low power transceiver maker Nordic Semiconductor. He is responsible for engineering support in the Americas. Tyzzer has also worked with OEMs such as Space Vector, Pertec and Tandon Magnetics and tier one distributors, Hall-Mark, Hamilton Avnet, and Symmetry Electronics. He received his engineering degree from Piece College and is a member of IEEE.

More from Nordic Semiconductor

Return to contents page.