October 2012
Volume 118
Issue 1918
£5.10

www.electronicsworld.co.uk

Electronics WORLD

THE ESSENTIAL ELECTRONICS ENGINEERING MAGAZINE

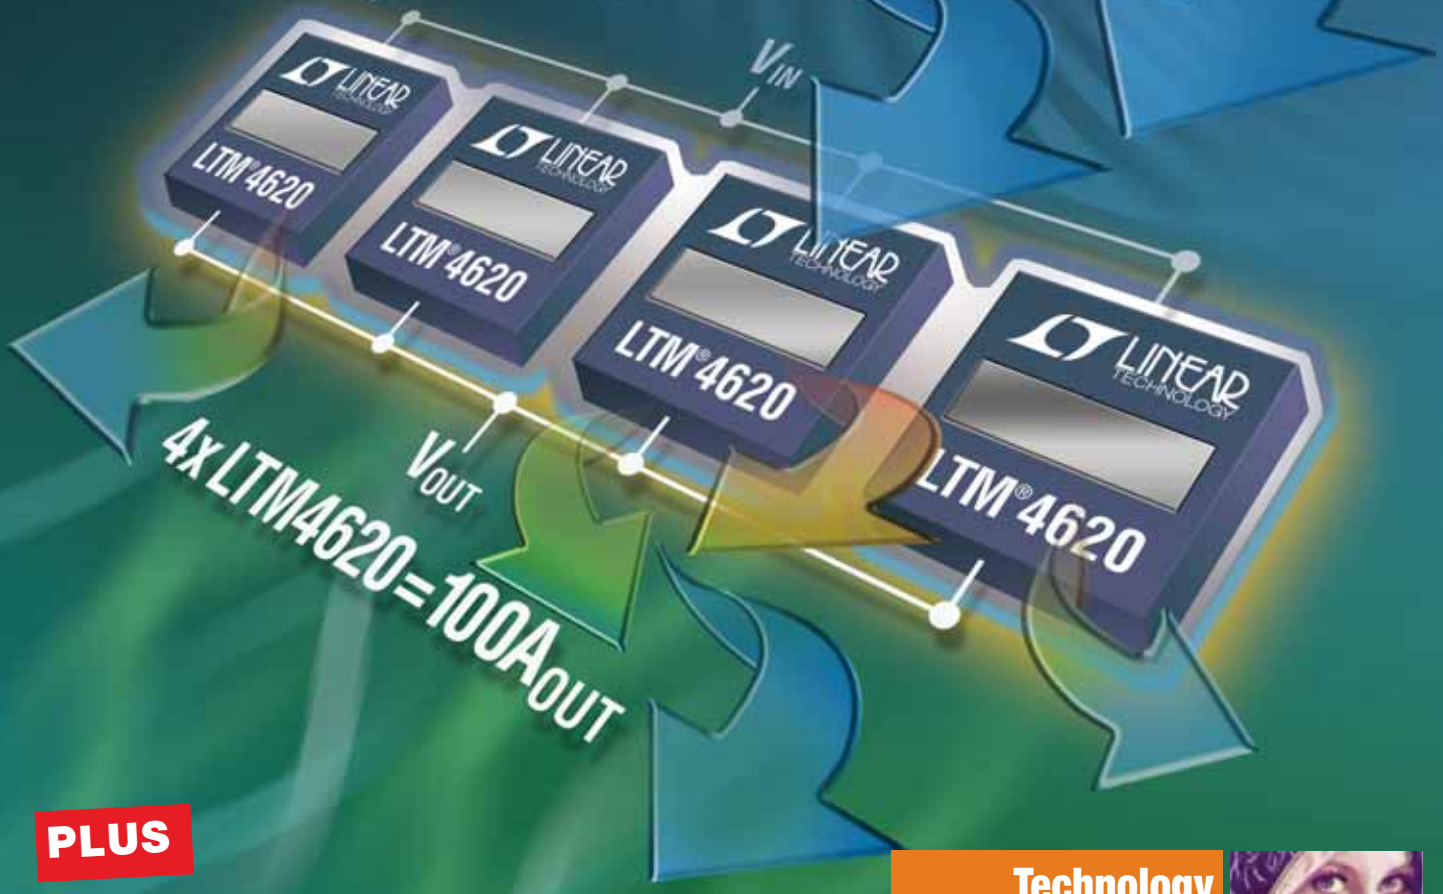

PLUS

MICROCONTROLLERS REPORT

Technology
First full colour 100,000dpi images created

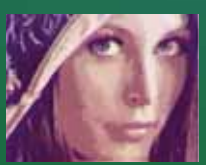

Processors
APUs geared for performance

Event
electronica 2012

Be the WaveAce

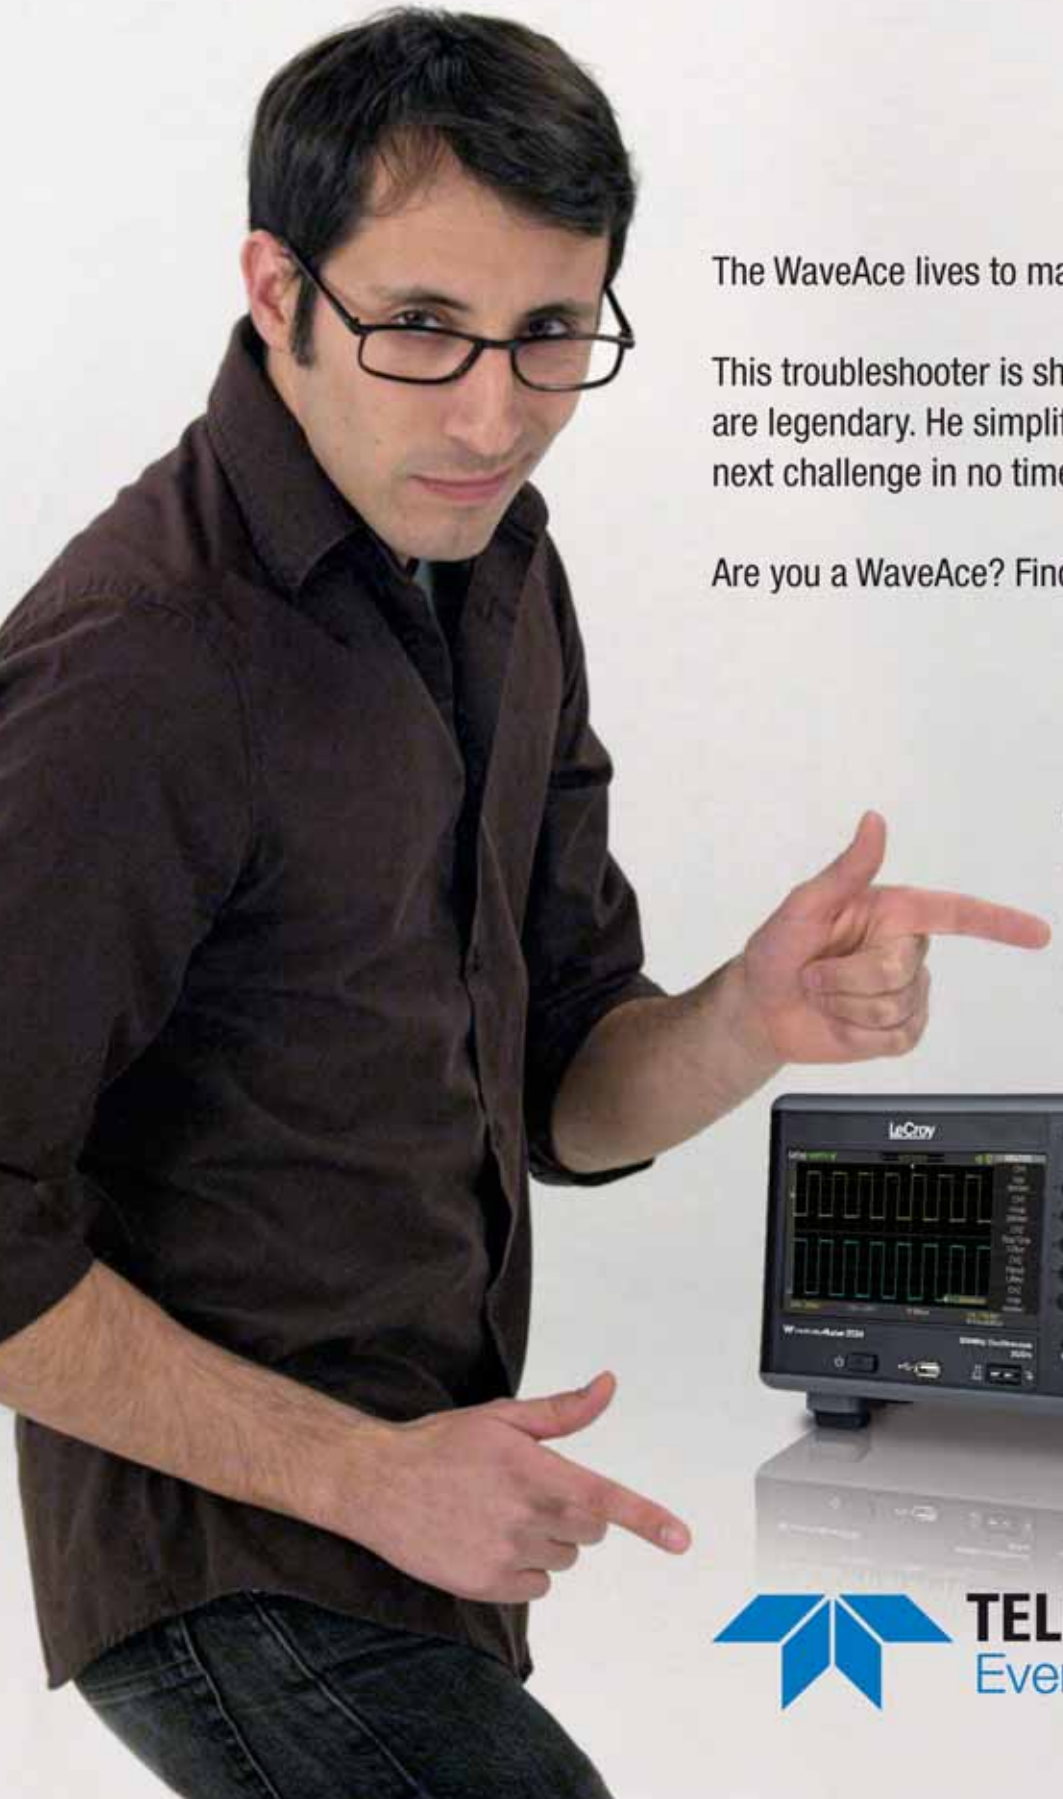

The WaveAce lives to make things work.

This troubleshooter is sharp. His measurement tools are legendary. He simplifies debug and is on to his next challenge in no time.

Are you a WaveAce? Find out at **bethewaveace.com**.

WaveAce 1000 & 2000 Oscilloscopes

- *Bandwidths from 40 MHz to 300 MHz*
- *Sample rates up to 2 GS/s*
- *Waveform memory up to 2 Mpts*
- *Starting from EUR 710*

TELEDYNE LECROY
Everywhereyoulook™

Cover prepared by LINEAR TECHNOLOGY
More on pages 7-9

REGULARS

10
Next Generation Networks

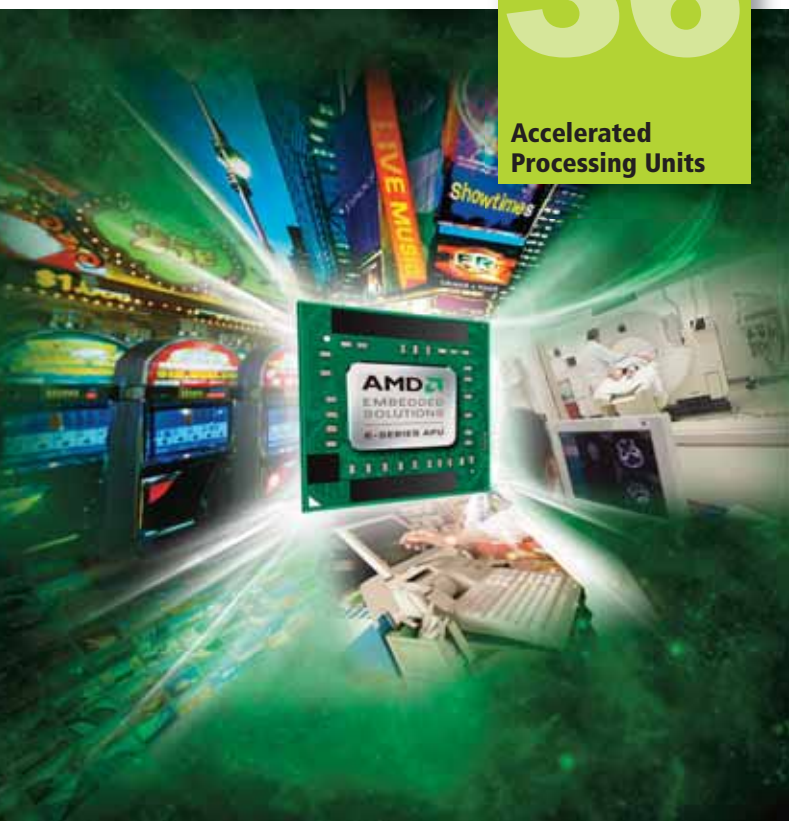

36
Accelerated Processing Units

FEATURES

electronica 2012 inside tomorrow
Hall A4, Booth 538
13-16 November 2012

Isolated LED Current Control with Active PFC

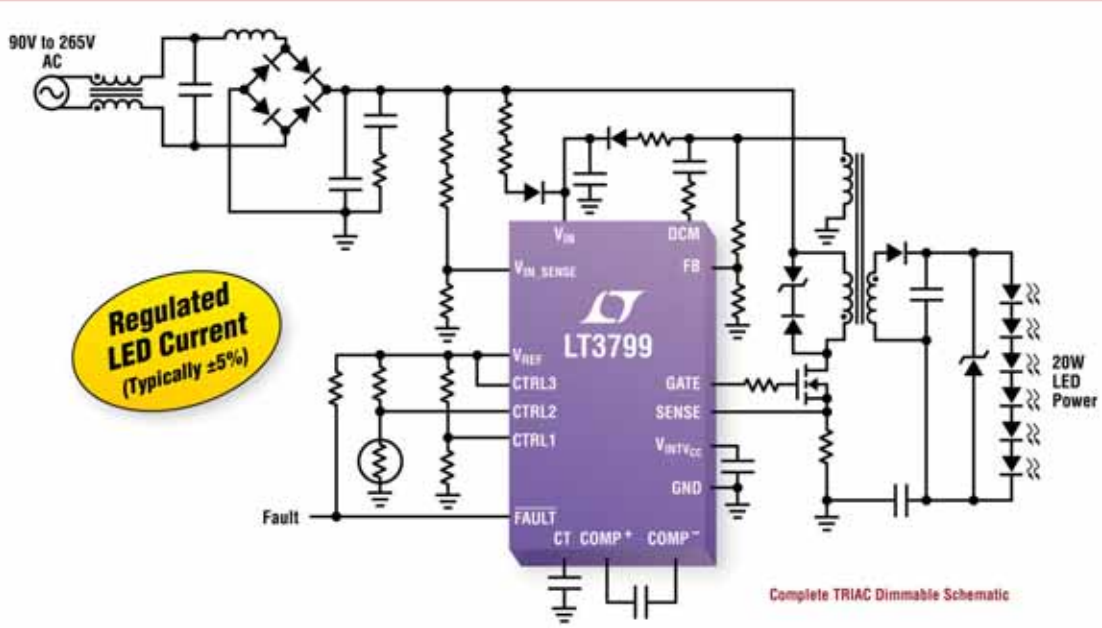

Complete TRIAC Dimmable Schematic

TRIAC Dimmable LED Driver Needs No Opto-coupler

Our LT®3799 isolated LED controller with active power factor correction (PFC) is specifically designed for driving LEDs over a wide input range of 24V to 480V+. It is ideal for LED applications requiring 4W to over 100W of LED power and is compatible with standard TRIAC in-wall dimmers. The LT3799's unique current sensing scheme delivers a well regulated current to the secondary side with no opto-coupler, enabling it to provide ±5% LED current accuracy. It also offers low harmonic distortion while delivering efficiencies as high as 90%. Open and short LED protection ensures long term reliability and a simple, compact solution footprint addresses a wide range of applications.

LED Current vs TRIAC Angle

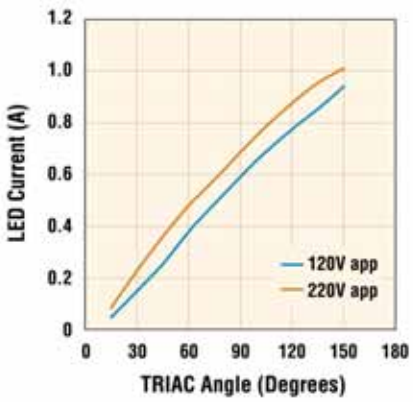

LT3799 Demo Board (25W)

Info and Purchase

www.linear.com/product/LT3799

Free Power Management for LEDs Brochure

www.linear.com/ledsolutions

Linear Technology (UK) Ltd.,
3 The Listons, Liston Road,
Marlow, Buckinghamshire,
SL7 1FD, United Kingdom.

Phone: 01628 477066
Fax: 01628 478153
Email: uksales@linear.com

WEARABLE TECHNOLOGY MARKET TO EXCEED $6BN BY 2016

Increasing demand for actionable, real-time data in a range of applications is driving strong demand for wearable technology. Some 14 million wearable devices were shipped in 2011; by 2016, wearable technology will represent a minimum revenue opportunity of $6bn, according to '*World Market for Wearable Technology – A Quantitative Market Assessment – 2012*', the new report from market analysis house IMS Research.

Current wearable devices are concentrated around a few products mainly in the healthcare and medical sectors, and fitness and wellness application areas. In these areas, there is a greater case for wearable technology in order to transmit data such as vital signs, and track user performance. Among the most prevalent wearable products are continuous glucose monitors – such as from Abbott and Medtronic, activity monitors – such as Fitbit, Adidas miCoach and Nike Fuelband, and fitness and heart-rate monitors – such as from Garmin, Polar and Suunto.

Increasingly self-aware consumers seek more data on their health and fitness, leading to an even more rapid expansion in the market for wearable technology

By 2016, the penetration of wearable technology in a number of current product categories is projected to rapidly increase as the range of new wearable devices widens. IMS Research envisages smart watches, smart glasses, sleep sensors, industrial and military head-up displays and hand-worn terminals to be most affected by wearable technologies.

"A $6bn market in 2016 is our most conservative forecast, which assumes that the adoption of wearable technology will be limited by factors including lack of suitable technology, poor user compliance and lack of an overall enhanced experience from devices that are wearable as compared to non-wearable products," said Theo Ahadome, senior analyst at IMS Research.

"In our mid-range and upside projections, product introductions such as Google's smart glasses and the rumoured Apple smart watch come to fruition and are successful. In addition, increasingly self-aware consumers seek more data on their health and fitness, leading to even more rapid expansion in the market for wearable technology", added Ahadome.

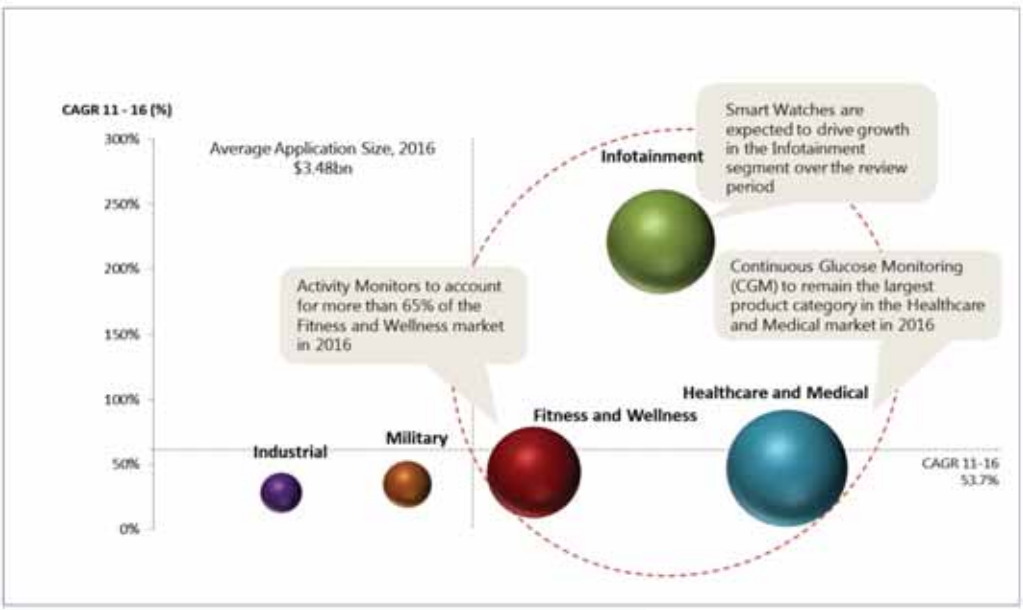

Figure 1: Wearable technology by application (the size of the bubble indicates relative market size)

IMS Research is a supplier of market research and consultancy to over 2500 clients worldwide. It was established in the UK in 1989 but was recently acquired by IHS, a source of information, insight and analytics for businesses and governments. IHS has been in business since 1959 and now employs over 6,000 people in more than 30 countries. www.ihs.com

EDITOR: Svetlana Josifovska
+44 (0)173 883392
Email: svetlanaj@stjohnpatrick.com

DESIGN: Tania King
Email: taniak@stjohnpatrick.com

DISPLAY SALES: John Steward
Tel: +44 (0)20 7933 8974
Email: johns@stjohnpatrick.com

SALES EXECUTIVE: Orla Cullen
Tel: +44 (0)20 7933 8999
Email: orlac@stjohnpatrick.com

PUBLISHER: Wayne Darroch

ISSN: 1365-4675

PRINTER: Pensord Magazines & Periodicals

SUBSCRIPTIONS:
Tel/Fax +44 (0)1635 879361/868594
Email: electronicsworld@circdata.com
SUBSCRIPTION RATES:
1 year: £51 (UK); £74.50 (worldwide)

John Patrick Publishers

Follow us on Twitter @electrowo

Join us on LinkedIn http://linkd.in/xH2HNx

NOVEL METHOD USED TO CREATE THE FIRST FULL COLOUR 100,000 DPI RESOLUTION IMAGES

Researchers from Singapore have demonstrated an innovative method for producing sharp, full-spectrum colour images at 100,000 dots per inch (dpi) suitable for reflective colour displays, anti-counterfeiting and high-density optical data recording.

The researchers from A*STAR's Institute of Materials Research and Engineering (IMRE) used metal-laced nanometer-sized structures, without the need for inks or dyes. This method contrasts current industrial printers such as inkjet and laserjet printers that achieve only up to 10,000dpi, with research grade methods only able to dispense dyes for single colour images.

The novel breakthrough allows colouring to be treated not as an inking matter but as a lithographic matter, which can potentially revolutionise the way images are printed and developed, such as for use in high-resolution reflective colour displays and high density optical data storage.

The inspiration for the research comes from stained glass, which is traditionally made by mixing tiny fragments of metal into the glass. It was found that nanoparticles from these metal fragments scattered light passing through the glass to give stained glass its colours. Using a similar concept with the help of modern nanotechnology tools, the researchers precisely patterned metal nanostructures and designed the surface to reflect the light to achieve the colour images.

"The resolution of printed colour images very much depends on the size and spacing between individual 'nanodots' of colour," explained Dr Karthik Kumar, one of the team. "The closer the dots are together – and because of their small size, the higher the resolution of the image. With the ability to accurately position these extremely small colour dots, we were able to demonstrate the highest theoretical print colour resolution of 100,000dpi."

Dr Joel Yang, project leader, explained further the process: "Instead of using different dyes for different colours, we encoded colour information into the size and position of tiny metal disks. These disks then interacted with light through the phenomenon of plasmon resonances. The team built a database of colour that corresponded to a specific nanostructure pattern, size and spacing.

"These nanostructures were then positioned accordingly. Similar to a child's colouring-by-numbers image, the sizes and positions of these nanostructures defined the 'numbers'. But, instead of sequentially colouring each area with a different ink, an ultrathin and uniform metal film was deposited across the entire image causing the 'encoded' colours to appear all at once, almost like magic!" he added.

The IMRE researchers are now looking into further collaborations and to licensing the technology.

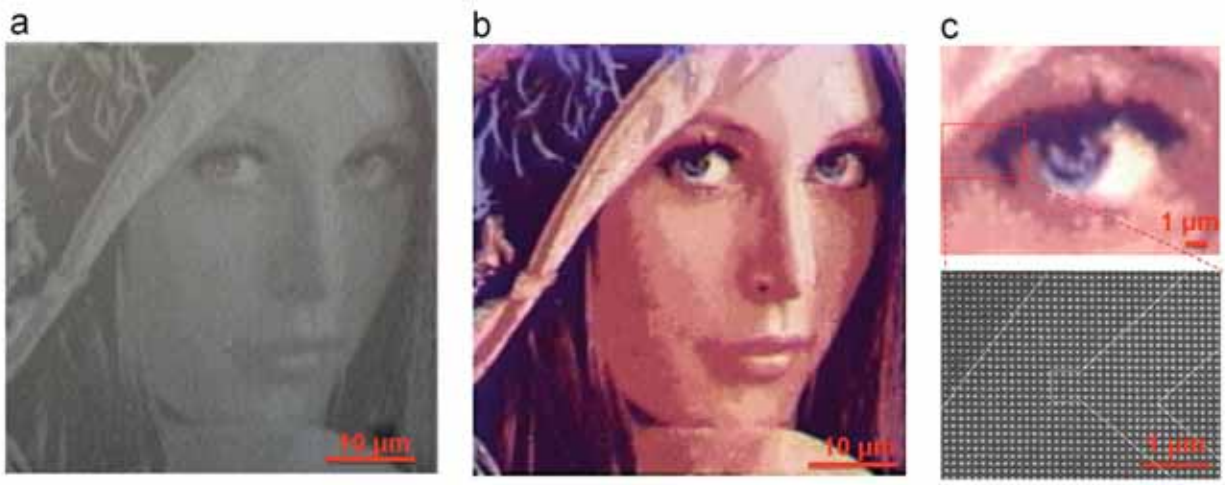

A coloured nanoscale rendition of a standard test image used in image processing experiments: (a) Before the addition of metal in the nanostructures, the image has only grayscale tones as observed under an optical microscope; (b) Colours are observed using the same optical microscope after addition of the metal layers to the nanostrucutres and in specific patterns; (c) Zooming into the image with the same setup, the specular reflection at the corner of the eye is observed showing the refined colour detail that the new method enables

NEWS IN BRIEF

Fujitsu Semiconductor Europe (FSEU) has demonstrated the transmission of > 100Gbps over a single CEI-28G-VSR channel, effectively quadrupling the data rate throughput defined by the Optical Internetworking Forum (OIF) for this chip-to-chip electrical interface.

This serves as a benchmark for what can be achieved over short-reach electrical channels using proven CMOS converter technology deployed in long-haul optical transport systems today.

Technology companies could be allowed to list on the London Stock Exchange for as little as 10% of their business, under new proposals being considered by the UK government.

Currently, tech firms need to float at least 25% of their business on the LSE. But with more firms looking outside the UK for investment, new measures to reduce the minimum float are gaining support.

HIGHER POWER, SMALLER & COOLER POINT-OF-LOAD DC/DC REGULATION

Innovative SoC Packaging with Integrated Heat Sink

By **Eddie Beville**, Afshin Odabaee & Mike Stokowski, Linear Technology Corp.

Each generation of high end processors, FPGAs and ASICs burdens power supplies with heavier loads, but system designers rarely allocate precious additional system board space to correspond to the power inflation. The squeeze on power supplies is compounded by the widespread requirement for greater numbers of dedicated board-mount power supplies, which provide POL (point-of-load) regulation for multiple voltage rails. Individual rails must increasingly support from tens to over a hundred amps at low voltages, under ≤1V, requiring an initial accuracy of ~1% and superb load transient deviation of less than a few percent. So the challenge is to find power supply solutions that are accurate, can deliver high load currents at low voltages, while taking little system board space.

Once a suitably powerful regulator solution is found, it must be evaluated for power loss and thermal resistance. These two parameters can break an otherwise good regulator solution if it can't meet system heat requirements, especially when the system must operate in an elevated ambient temperature environment. Obviously, conversion efficiency must be high in order to limit power loss, and the package design must feature low internal thermal resistance and a low thermal resistance connection to the ambient environment. As solutions shrink, the thermal resistance between the regulator and the board decreases in area, making it increasingly difficult to keep the board cool because the power regulator usually conducts most of the power loss back into the system board, significantly increasing the internal system temperature.

THE REAL ISSUE: HEAT & COST OF COOLING

System and thermal engineers spend a lot of time modeling and evaluating these complex electronic systems in order to design solutions that remove power loss in the form of heat. Air flow and heat sinks are typically used to remove this unwanted heat. The real issue is that modern processors, FPGAs, and custom ASICs usually dissipate significantly more power as the internal system temperature increases. This unfortunately requires more power from the power regulators, and will increase their internal power loss thus increasing system temperature even further. So eliminating power loss and heat is very important, and high density power solutions must limit power loss and remove heat effectively. But most compact packaged power solutions either dissipate too much power or cannot effectively remove the heat and therefore cannot operate at elevated temperature without significant de-rating. A reasonable solution is needed to help alleviate the real issue.

It's no surprise that to keep the temperature of a high power design to reasonable levels attention to cooling methods is crucial. Installation of fans, cold plates, heat sinks and sometime submerging the system in special liquids are examples of approaches that some designers are

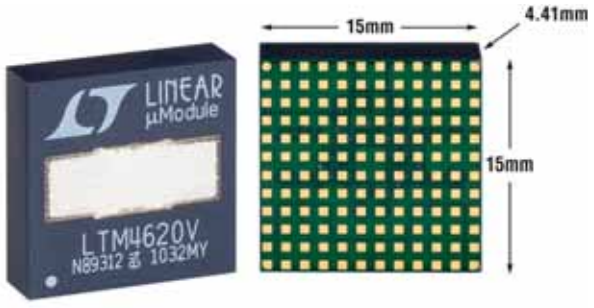

Figure 1: LTM4620 Package: 15mm x 15mm x 4.41mm LGA

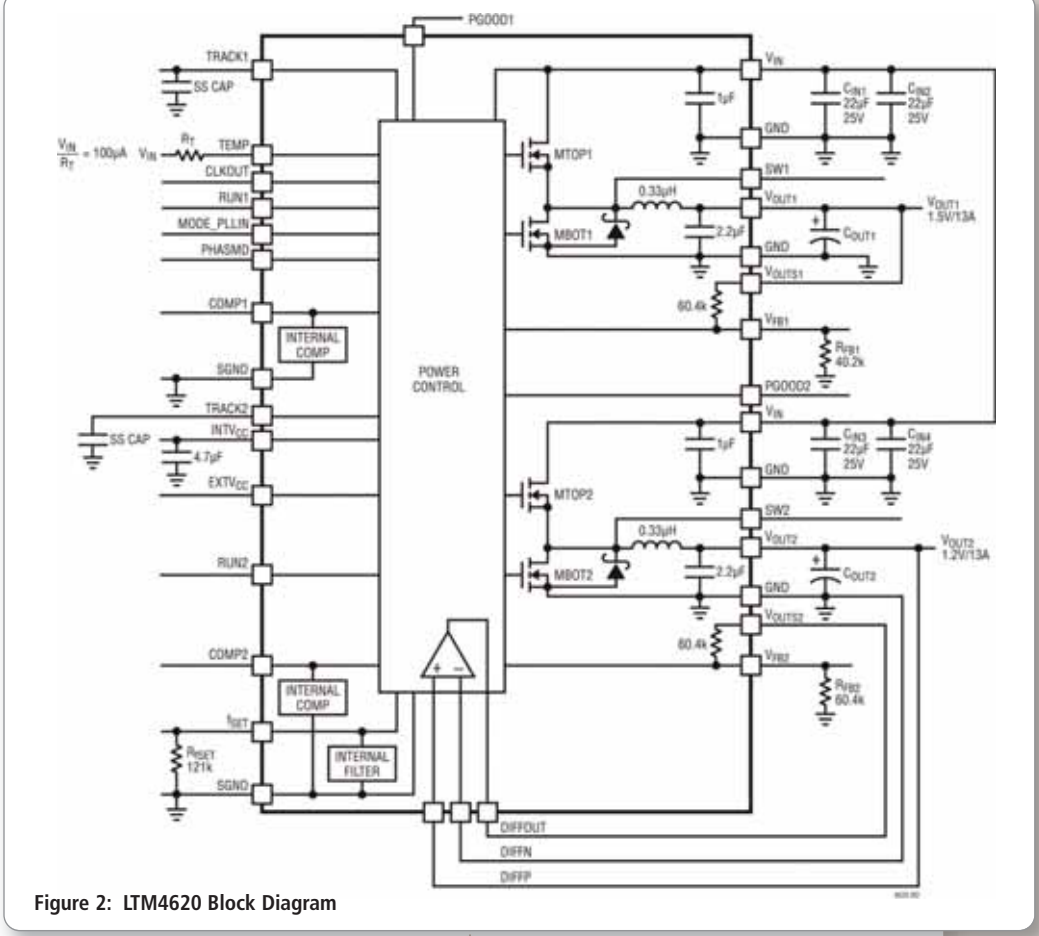

Figure 2: LTM4620 Block Diagram

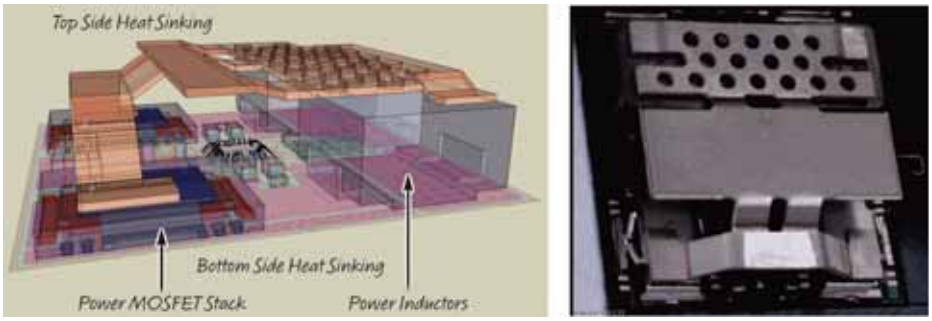

Figure 3: LTM4620 Side View Rendering & an Unmolded LTM4620

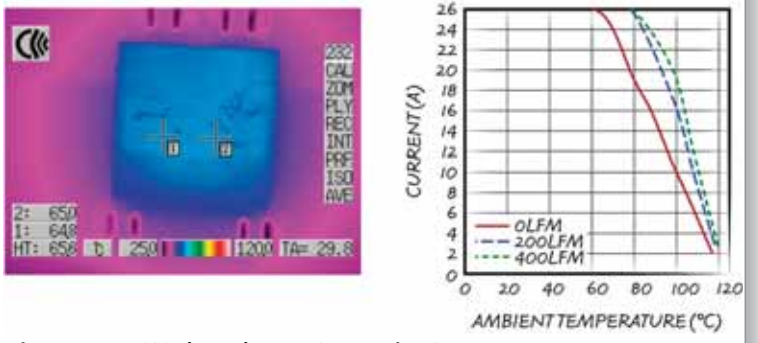

Figure 4: LTM4620 Thermal Image & De-Rating Curve

forced to implement. All are costly, but necessary. However, if a high power point-of-load regulator could deliver the required power while dissipating heat evenly and efficiently, the requirements for cooling that portion of the circuit will be reduced saving on cooling size, weight, maintenance and cost.

POWER DENSITY IS MISLEADING

The topic of high power density DC/DC regulators is misleading because it does not address the behavior of the temperature of the device. System designers should be educated to seek more information from the device's data sheet once they decide on a product that meets the system's electrical, physical and power requirements for a DC/DC regulator. Here is an example: if a DC/DC regulator in 2cm x 1cm delivers 54W to a load, its power density is rated as 27W/square cm. This number may impress a few designers and satisfy their search: desired power, desired size and desired price. However, what's forgotten is heat which finally translates into temperature. The key piece of information is to study the DC/DC regulator's thermal impedance looking for values for the package's junction to case, junction to air and junction to PCB.

Continuing with the above example, the device has another attractive attribute. It operates at an impressive efficiency of 90%. It dissipates 6W while delivering 54W output in a package with 20°C/W junction-to-air thermal impedance. Multiply 6W by 20°C/W and the result is 120°C rise on ambient temperature. At 45°C ambient temperature, junction temperature of the package of this seemingly impressive DC/DC regulator is calculated at 165°C. 165°C is not an impressive value for two reasons: a) its above maximum temperature of most silicon ICs which is roughly 120°C and b) it requires special attention to keep the junction temperature at a safer value below 120°C.

Above simple calculation is sometimes ignored. A DC/DC regulator that seemed to address all the electrical and power requirements failed to meet thermal guidelines of the system or proved too costly to use due to additional measures to operate at a safe temperature environment. It's important to remember to study thermal performance of a DC/DC regulator as one first becomes involved in evaluating attributes such volts, amps and centimeters.

This article will describe a new high density and scalable LTM4620 µModule® regulator. The discussion will comprise of the electrical, mechanical/package and thermal performance along with different scalable power designs. The goal is to show a new high density scalable power regulator that has excellent electrical performance, low power loss and a unique thermally enhanced package design to help resolve high power density challenges.

THE LTM4620 DUAL 13A OR SINGLE 26A µMODULE REGULATOR

Figure 1 shows a photo of the LTM4620 µModule regulator. The SIP (System-In-Package) is a 15mm x 15mm x 4.41mm LGA device. It is capable of delivering two independent outputs at 13A, or a single output at 26A. The package supports both top and bottom heat sinking for excellent thermal management.

Figure 2 shows a block diagram of the LTM4620 µmodule regulator. The LTM4620 consist of two high performance synchronous buck regulators. The input voltage range is 4.5V to 16V, and the output voltage range is 0.6V to 2.5V and 0.6V to 5.5V for the LTM4620A. The LTM4620 electrical features are ±1.5% total output accuracy, 100% tested accurate current sharing, fast transient response, multiphase parallel operation with self clocking and programmable phase shift, frequency synchronization and an accurate remote sense amplifier.

The protection features are output overvoltage protection feedback referred, fold back overcurrent protection and internal temperature diode monitoring.

THE LTM4620 UNIQUE PACKAGE DESIGN

Figure 3 shows the side view rendering and a top view photo of an unmolded LTM4620. The package design consists of a highly thermal conductive BT substrate with adequate copper layers for current carrying capacity and low thermal resistance to the system board. A proprietary lead frame power MOSFET stack is utilized to provide high power density, low interconnect resistance, a high thermal conductivity to both the top and bottom of the device. The proprietary heat sink design attaches to the power MOSFET stacks and the power inductors to provide an effective top side heat sinking. An external heat sink can be applied to the topside exposed metal to remove heat with air flow. Airflow alone with no heat sink removes heat from the topside due to construction of the heat sink and the mold encapsulation.

Figure 4 shows a LTM4620 thermal image and a de-rating curve for a 12V to 1V at 26Amp design. The temperature rise is only 35°C above ambient temperature with 200LFM of air flow, and the de-rating curve shows that the maximum load current requires no de-rating out to ~80°C. Figure 4 reveals the thermal data that shows the real merits of a thermally enhanced high density power regulator solution. The unique package design keeps the power loss as low as possible for the small size, and effectively removes heat as a function of the power loss.

THE LTM4620 ELECTRICAL PERFORMANCE

Figure 5 shows the LTM4620 operating in the dual output current sharing operation. This configuration provides a very high density 1.5V

at 26A solution. The RUN, TRACK, COMP, VFB, PGOOD and VOUT pins are tied together to implement the parallel operation. The design shows one way of monitoring the LTM4620 internal temperature diode using a LTC2997 temperature sensor monitor. The temperature sense diode can be monitored by many different devices that monitor a diode connected transistor. Figure 6 shows the 1.5V efficiency for the 2 phase parallel output and the two channel current sharing performance. The 86% efficiency is very respectable for such high density solution, and as shown in the Figure 4 thermal data, the temperature rise is well controlled due to a low ΘJA thermal resistance after board mount. The effective top and bottom heat sinking enables the LTM4620 to operate at full power with low temperature rise. Figure 6 shows both VOUT1 and VOUT2 current sharing performance. The LTM4620 internal controller is accurately trimmed and tested for output current sharing. This makes the LTM4620 an excellent choice for high density scalable power solutions. The high efficiency and fast transient response current mode architecture fits well with the low voltage core power supply requirements needed for high performance processors, FPGAs, and custom ASICS. The outstanding output voltage initial accuracy and the differential remote sensing provide proper DC voltage regulation at the load point. The unique thermal capabilities and the excellent current sharing allows for scaling the output current capability up to 100+ Amps. No external phase shifted clock sources are needed for setting multiphase operation for each regulator channel. Each LTM4620 has a Clock In pin and a Clock Out pin with internal programmable phase shifting for clocking the paralleled channels. Either external frequency synchronization or internal on board clocking can be selected. These clocking features further enable the power scaling concept.

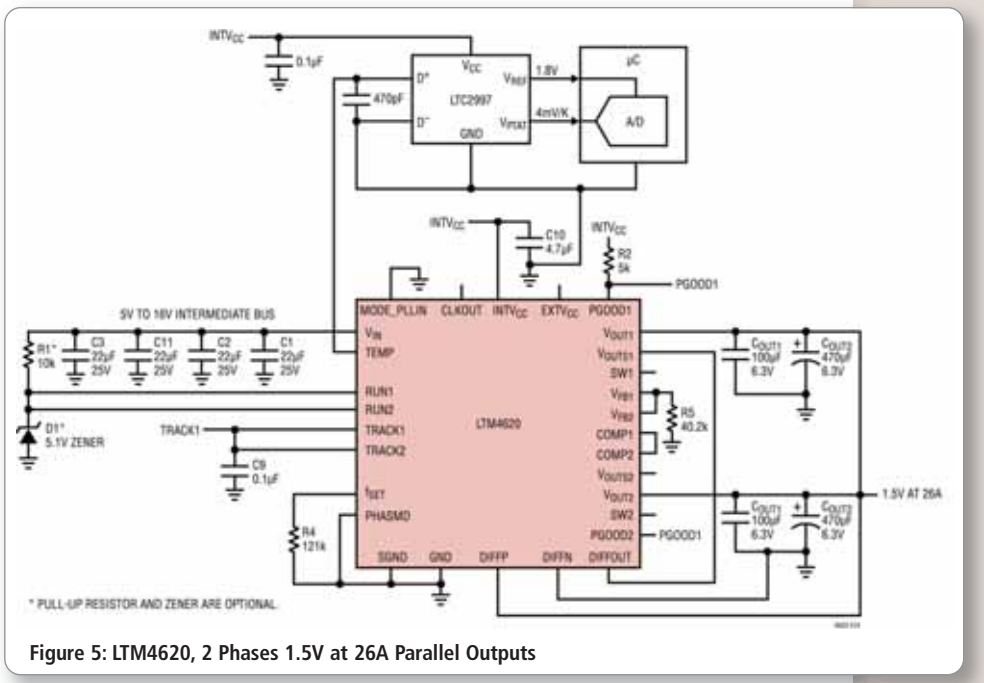

Figure 5: LTM4620, 2 Phases 1.5V at 26A Parallel Outputs

Figure 7 shows both a snapshot of an 8-Phase, 4 µModule regulator 100A design and the current sharing graph for all four regulators. All eight phases are clocked phase and tied together to implement the current sharing scaled 100A implementation. As noted in Figure 7, the actual µModule regulator board space is about 1.95 square inches of space to support the 100A power solution. This provides an excellent high density power solution for these high currents. A heat sink can be applied across all four modules to remove power loss with air flow. This keeps a lot of power loss from being dissipated into the system board.

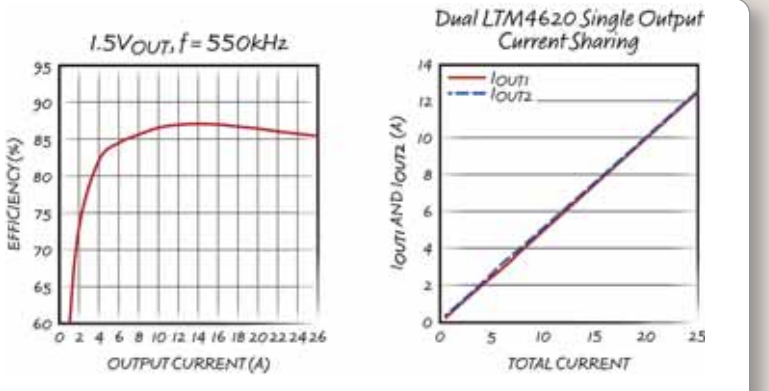

Figure 6: 2-Phase 1.5V Efficiency & Current Sharing Graphs

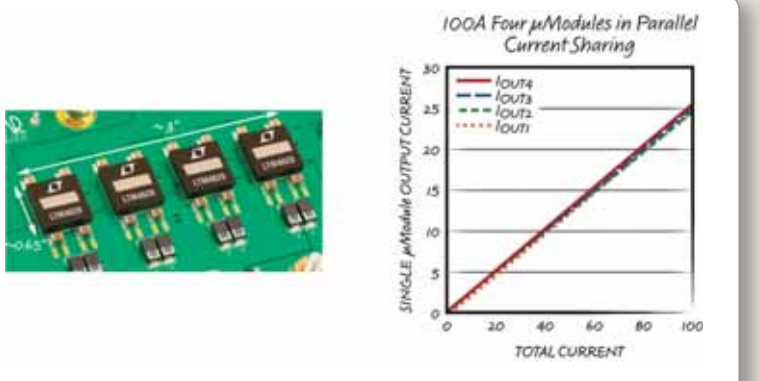

Figure 7: 8-Phase, 4 µModule Regulators Scaled to 100A Design

PROOF OF PERFORMANCE

To validate the performance of the LTM4620, four quick TechClip videos are provided, demonstrating the setup and measurements. These TechClip videos cover topics of short circuit protection, thermal behavior and temperature rise at 26A and 100A, heat sink attachment and precision current sharing at startup, steady state and shutdown. View these videos at http://video.linear.com/p4634-126.

CONCLUSION

The LTM4620 µModule regulator brings a new concept to high density power solutions. The high performance regulator housed inside a superior thermally designed package enables high power designs to be implemented in a very small form factor. The multiphase clocking features with the accurate current sharing enable scalable designs in 25 Amps, 50 Amps and 100+ Amps. The LTM4620's unique thermal properties provide full power operation at elevated ambient temperatures. High current designs can be implemented while controlling power loss and temperature to acceptable levels.

Linear Technology (UK) ltd
Tel 01628 477066 • Email: uksales@linear.com
www.linear.com

Building the Bridge to Next Generation Networks

PRAKASH KANTHI, SENIOR MARKETING MANAGER, TEST BUSINESS, TEKTRONIX COMMUNICATIONS, EXPLAINS THE SHIFT TO NEXT GENERATION NETWORKS AS THE NEED FOR TEST SOLUTIONS THAT SUPPORT BOTH LEGACY AND NEXT GENERATION ARE NECESSARY IN UNDERPINNING THE TRANSFORMATION

Global LTE network rollouts are accelerating, and expenditure on network improvements will top $250bn by 2015, with LTE accounting for 46% of spend over this period, according to Infonetics. This transformation is set to be a gradual process, driven by business needs and technology upgrades, leading to the emergence of so-called 'heterogeneous networks' – a mixture of legacy and next generation technologies. However, it is also creating a series of challenges for operators.

> *Some operators – for economic, business or technical reasons – are not yet ready to transition away from legacy networks*

Simultaneously, network components are becoming functionally more complex, having to operate in environments that involve multiple technologies, vendors and standards. This has meant that niche testing solutions are becoming insufficient, as they only consider a fragment of the overall problem.

The mobile user experience is being revolutionising by heterogeneous networks

The Move to Next Generation

Legacy mobile network technology has developed from standard voice calls, through SMS messaging to the more advanced mobile data services we now see. The fabric of these networks is constructed with multiple technologies. Each technology enables an individual service or application – such as SMS and voice.

Designated technical and engineering teams were assigned for each separate technology and service, which promoted a silo-type service model in legacy networks. Such networks initially relied heavily on circuit switching and, over time, evolved to packet IP switching. However, this evolution was slow – with some portions of the network transi-tioning to IP faster than others. Some operators, for economic, business or technical reasons, are not yet ready to transition away from legacy networks. Carriers are forced to cope with multiple network topologies and their associated underlying legacy technologies.

The 3G and HSPA+ technologies, combined with the proliferation of smart devices and tablets, have seen an entirely new class of mobile broadband services enter the market. However, these services are creating traffic volumes that are unsustainable over legacy networks, which were primarily designed for voice and low-speed data. Contemporary operators need networks

that can deliver low latency, high bandwidth and greater flexibility, in order to provide efficient and cost-effective mobile broadband services.

Hence the adoption of LTE (Long Term Evolution), which is IP-based and addresses many of the legacy challenges that operators are encountering. However, the LTE journey is complex as operators need to adopt an all-IP mindset, and yet not lose sight of where they have come from.

As well as LTE, operators are turning towards IMS (IP Multimedia Subsystem) as one of the key enabling technologies to leverage and deliver services across the network. To this degree, IMS defines a complete architecture and framework that enables the merging of voice, video and data over an IP-based infrastructure. With IMS at its core, next-generation networks are becoming less siloed and moving towards flatter, more efficient topologies.

In spite of these challenges, the migration to next-generation all-IP networks continues. Operators favour all-IP networks as they can reduce network complexity, capex investment and overall service and maintenance. As operators have experienced the continuing inconsistency between traffic growth and revenue, next-generation networks, based on all-IP technologies, provide the required platform and economies of scale for operators to leverage their position in the overall communications ecosystem.

IMS defines a complete architecture and framework that enables the merging of voice, video and data over an IP-based infrastructure. With IMS, next-generation networks are moving towards flatter, more efficient topologies

Negotiating the Chasm

At present, operators are attempting to bridge the gap between validating their all-IP network solutions for conformance, performance and interoperability, and in some cases, retaining a first-class service for some of their legacy infrastructure. Negotiating this divide will enable operators to make the most of their current network assets, safe in the knowledge that their network will perform equally well under legacy and next-generation conditions. However, in order to achieve this, equipment manufacturers and operators will need to fully distinguish and test their networks, subsystems and new all-IP elements to make sure they are operating properly, under loaded conditions that simulate real-world traffic in a mixed environment. This will result in operators being able to ensure end-to-end quality of the network and all the services being offered through it.

New test strategies and solutions must be flexible, robust and innovative to link the void between legacy and next generation. Mobile networks spanning legacy and next-generation contain a collection of technologies and standards, numerous protocols and equipment made by differing vendors. These individual network components, forming the complete architecture, must function correctly and as part of a larger network. To ensure this, a test and measurement solution must have the ability to test specific protocols and standards, over multiple technologies, to inspect if legacy and next generation components are working individually, but also communicating efficiently with each other.

One answer is a single-platform test and measurement tool, with a single user interface, that can perform legacy and

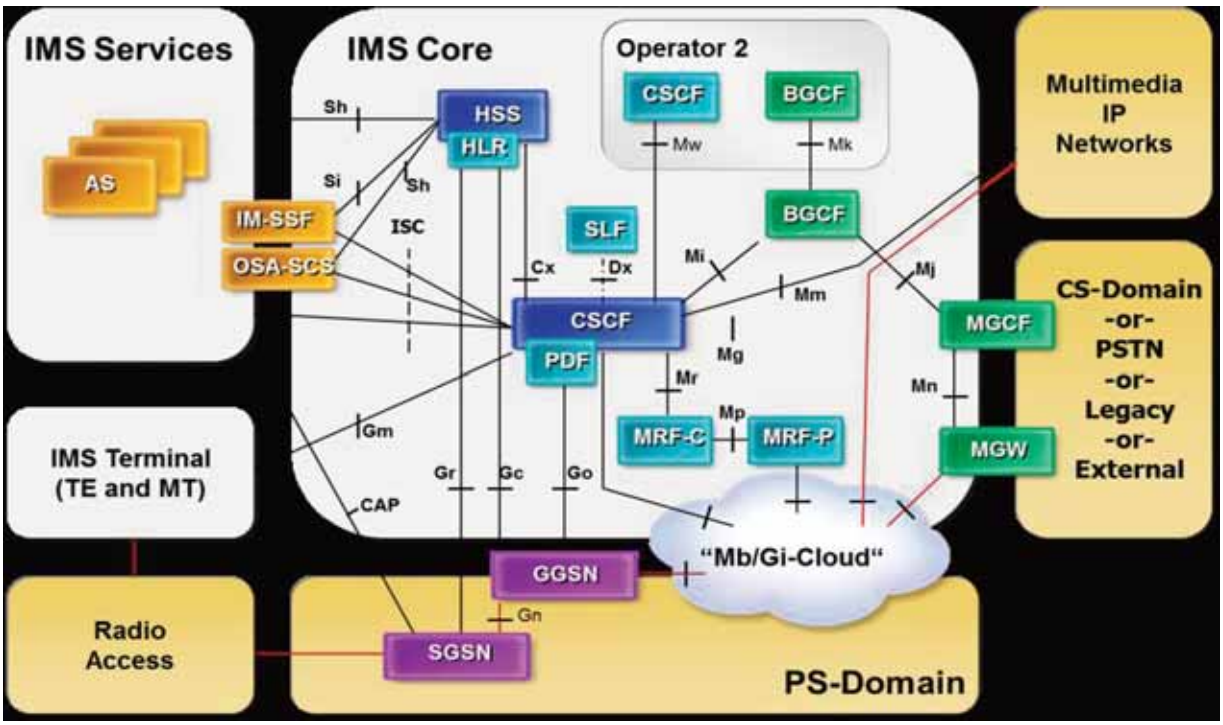

The complexity of the current mobile network

IP testing at the same time, and monitor legacy media conversion to and from IP.

Due to the level of equipment in mobile networks from multiple vendors – it is also important that this test and measurement system is vendor agnostic (capable of monitoring any vendors' equipment). Having a single, functional and flexible test system like this can help operators make the transition to an all-IP network easier, while retaining the optimum performance of their legacy networks.

Sealing the Gap

Both legacy and IP components can be served by a single test and measurement solution, that will also provide more accurate and repeatable measurements, as there is no need to rely on multiple pieces of test equipment which may produce different results.

As networks undergo the transition to next generation technologies, the amount of equipment and overall complexity at a high level is decreasing. However, the drawback is that individual components are now becoming more integrated with higher levels of functionality, as technologies become more standard and off-the-shelf (OTT), and chipsets become more common. As a result, the network equipment vendor market has become more competitive and is attracting new market entrants. This abundance of vendor solutions with greater functionality increases the complexity of testing network components, thus it heightens the need for a test and measurement solution that is vendor-agnostic.

Looking Forward to a 4G Future

The industry is undergoing a major transformation. Next generation technology is ensuring telecommunications moves from a capital-intensive model focused on individual technologies, to a subscriber-centric service delivery model that is high-quality and cost-efficient. This has been driven by competitive pressures within the traditional market, and the disruptive business models of OTT players impacting operators' traffic and revenues.

As the transformation to all-IP networks continues, the barriers to entry for equipment manufacturers are dropping, thus increasing overall competitiveness. Advances in smart devices and tablets, combined with new technologies such as Voice-Over-LTE and the GSM Association endorsed Rich Communication Suite, are revolutionising the mobile users' experience. As mobile services become more integrated and more personalised, operators need to manage and test the underlying technologies delivering these services.

To successfully manage the evolving range of technologies, protocols and equipment required for next-generation networks, the industry needs flexible, thorough and open T&M solutions to ensure that network components and subsystems are operating properly in both legacy and next-generation technologies. ●

Fog

MYK DORMER IS SENIOR RF DESIGN ENGINEER AT RADIOMETRIX LTD
WWW.RADIOMETRIX.COM

The "figure of merit" in comparing low power wireless modules is frequently the unit's range, so when a unit in the hands of the customer fails to achieve the claimed, predicted or calculated range, complaints begin to fly.

Complaints can cost money, in support time and lost sales, so understandably a lot has been written on this topic already. Most of this concentrates on getting the best performance out of the receiver (optimising data rates and modulation formats, minimising noise figure, maximising sensitivity), the transmitter (best possible efficiency to utilise the very maximum band-legal power output) and the aerial (best type of installation, optimum location, and minimum feeder loss).

The point when things get really nasty is when all this work has been done, all the guidelines are being followed, but the customer's system is still getting (for example) half the range you think it should (based on known performance of the radios used). "Nasty", because this is the point at which you probably have to start accepting that the problem might not be the radio; might not be the way the radio is being used, but might actually be inherent to the customer's installation. It's the point at which you have to point the finger at the rest of the hardware and say: "It's noise! Digital noise from your (whatever, but probably a PC) is de-sensitising the radio".

This is something that no RF engineer wants to do. It takes the responsibility away from the radio design – which is something we think we understand and have actual control over – and drops it into a questionable half-way house of accusation and counter-accusation, incomplete tests and inconsistent and hard-to-replicate results. Your customer will think you are incompetent or fraudulent, and will keep blaming the radio hardware, while you will have less and less respect for the customer's technical abilities.

At this point we need to call a truce: there is always a fair possibility that the installation of the radio equipment has some problems, or that the radios supplied are defective. These issues need to be addressed first but, after these obvious faults have been ruled out, the most likely cause of poor range, in my experience, is still de-sensitisation of the receiver by digital noise emanating from co-sited computer equipment, or (as it is sometimes termed) "electric fog".

Any electronic system using non-steady state signals will, by definition, radiate

After these obvious faults have been ruled out, the most likely cause of poor range is still de-sensitisation of the receiver by digital noise emanating from co-sited computer equipment, or as it is sometimes called "electric fog"

some form of electromagnetic energy. Modern computing equipment, which uses a multitude of fast, sharp-edged, square waves, is very prone to radiating troublesome interference. Occasionally this will manifest as visible, discrete interferers – harmonics or sub-harmonics of a system clock for example – which can be seen on a spectrum analyser, but more insidiously it can appear as a zone of wide-band noise, covering tens or hundreds of MHz. This "radio frequency effluent" is usually within the legal RFI limits for the device, leaving no recourse to complain to its supplier or the regulatory bodies, but is still more than sufficient to noticeably de-sense any radio link located within a couple of meters of the device.

Theorising about the cause of a range reduction is not enough, and to the customer it looks like buck-passing. It is necessary to identify and, if possible alleviate, this problem.

Identifying this form of interference is quite easy, once you know what to look for. It is rarely possible to see such signals on spectrum analyser/wide-band aerial setups so elimination tests are the best way forward: A typical narrowband receiver might have a working sensitivity of -120dBm. Assuming a 10dB co-channel rejection and a halving (approximately 12dB link margin reduction) of the range, then the interfering signal will be at about -118dBm, and it's unlikely to be a discrete carrier.

1. If it is possible to test the system with all associated computer equipment turned off, possibly using a radio-modem's internal test routines, then this can easily show the problem in a comparative range-trial. Unfortunately this cannot always be arranged.
2. Assemble a "range comparison trial" prototype, using the same radio link, data structures and aerial but without the customer's digital hardware. Then test this link over the same (or comparable) terrain.
3. Re-locate the aerial (and radio hardware if possible) at least 3m further from the suspected digital interference source, keeping the elevation above ground as near to the same as you can. Conduct a comparative range trial.
4. Look at the RSSI (received signal strength indicator). The receiver's analogue signal strength meter output is one of your most valuable tools. Any indication of a signal above the usual

A SIMPLE DEMONSTRATION OF "DIGITAL FOG"

TRY THIS SIMPLE TEST: Assemble a narrowband 433MHz band receiver module, such as an LMR2-433; a whip aerial; a battery; and a voltmeter connected to the module RSSI output – a lashup on a piece of matrix board will do.

In the absence of any transmitters on channel – and in an open space – observe the signal strength meter reading. This roughly corresponds to the "noise floor".

Now approach the screen of your desktop PC, with an eye on the meter.

At 50cm a change will be noticeable.

At 10cm it's frankly appalling.

Remember: if you are seeing a difference between the usual noise floor reading and the reading in proximity to a computer, the radio link will probably be seeing a reduction in maximum range.

no-carrier noise floor at times when the link transmitter is off can be an indicator of noise de-sense. It can be useful to bench calibrate the equipment before trials, to determine the RSSI voltage corresponding to a range of input signal levels. This is not foolproof, but in some cases can be very revealing. On one troublesome radio system I have seen an RSSI reading from digital noise alone corresponding to a -108dBm carrier. The radio in question had a calling sensitivity of -120dBm.

5. Listen to the receiver AF output and look at it on a 'scope. This does not always yield results; sometimes digital noise just looks like more noise on channel, but sometimes distinct "patterns" or "notes" can be heard in the audio (the human ear is very sensitive to this) or "quieting" in the absence of a wanted signal seen on a 'scope.

Of course, identifying a digital noise source is one thing, but sorting it out is another. It's not enough to tell the customer: "That PC/PDA/etc that you use the radio to communicate with and on which your entire system depends – you can't use it. Turn it off". (Not if you want to keep the customer, anyhow).

Computing equipment is a part of the 21st century world, so we RF engineers need to find coping strategies including:

1. Where possible (and unfortunately, unless the system design has anticipated this requirement from the outset, it is unlikely to be feasible) time slice the periods of radio communication and periods of heavy computational effort, so they do not overlap. Using "sleep" or partial "hibernate" modes may be effective.
2. Screen the computing system as best you can (using metal or conductive plastic housings), paying attention to RFI filtering on all cables into and out of the screened volume, or they may become aerials in themselves.
   Do not assume the basic "EMI compliance" screening will be enough.
3. Remember peripheral devices. Ethernet cables are notorious radiators, and the viewing surface of typical LCD displays lack adequate shielding, which can be added with conductive transparent coatings or fine metal meshes. Disable unused functions, such as Bluetooth links or extra I/O ports. Minimise additional external hardware.
4. Use the lowest power, "greenest" PC or PDA you can. In general the amount of wide-band digital "mush" generated will be proportional to the power consumed, although there are occasional exceptions. Be prepared to compare several types of unit, if design constraints and time allow. Not all manufacturers are equal.
5. Distance is the best screen. Of all "cures" I know this is the most likely to succeed: separate the aerial – and the rest of the radio hardware if possible – from the PC (or whatever) by the maximum distance that mechanical and usage constraints allow. A computer in the basement is unlikely to degrade the performance of a radio modem in the attic. On vehicles, locating the antenna at the extreme rear with the rest of the digital hardware at the front of the passenger compartment will help, especially as it gains some extra screening from the metal body panels.

Even a yard can make a difference. If you need convincing then see the 'Demo' box above).

FINDING THE RIGHT SIZE MICROCONTROLLER

KEITH CURTIS, TECHNICAL STAFF ENGINEER IN THE SECURITY, MICROCONTROLLER AND TECHNOLOGY DEVELOPMENT DIVISION AT MICROCHIP TECHNOLOGY GIVES A PRACTICAL GUIDE TO THE CONSIDERATIONS THAT HAVE TO BE TAKEN INTO ACCOUNT WHEN CHOOSING AN MCU

The choice of which microcontroller to use is made at the beginning of an embedded-control project. The options are to select an 8-bit, 16-bit or 32-bit MCU and either start from a low-cost micro and migrate up, or downsize from a high-end micro.

There are no hard and fast rules: a number of factors will influence the decision, including the level of control and processing power the project will require, the power limitations and so on – the list is almost endless. Requirements such as operation in a harsh environment or the man-machine interface can be as vital as considering how fast the product has to respond to changes. The end result can be selection 'paralysis', as an engineer weighs conflicting requirements and demands.

System Considerations

The solution can be found in a requirements document which allows an engineer the opportunity to weigh the alternatives. The first section should tackle basic functionality:

1. What tasks will the system perform?
2. What are the inputs and outputs for it?
3. How much data storage will it require?
4. How quickly must the system perform its tasks and respond to events?

Table 1 shows an example list for a simple thermostat control.

Design constraints should feature in part two:

1. What is the material and assembly cost target?
2. What is the power requirement of the system?
3. What is the physical-size limitation of the system?
4. What will be the operating environment for it?

Table 2 shows a sample list for the same simple thermostat control.

Table 1: Initial project requirements list

Intelligent Home Thermostat

Function list

1. Heater control
2. AC control
3. Serial output to PC
4. Serial input from PC
5. LCD display driver
6. Capacitive-touch input
7. Temperature control
8. Command decoder
9. Real-time clock

Data storage

- ☐ 32 bytes: Serial transmit buffer
- ☐ 32 bytes: Serial receive buffer
- ☐ 20 bytes: Command parser
- ☐ 10 bytes: Real-time clock
- ☐ 20 bytes: LCD buffer
- ☐ 6 bytes: Tone generator
- ☐ 45 bytes: Temporary variables

165 bytes total

System inputs/outputs

1. RS-232 serial input
2. Cap-touch sensor 4
3. Heat/cool switch
4. Temperature sensor
5. Heat/cool output drive
6. RS-232 output
7. Segmented LCD display
8. Piezo buzzer
9. Battery Voltage

Speed

- ☐ 2 kHz: Touch interface scan rate
- ☐ 60 Hz: LCD update rate
- ☐ 4 kHz: Tone generator
- ☐ 5 Hz: Heat/cool drive
- ☐ 1 Hz: Battery check
- ☐ 1 kHz: Serial port character rate (6900 baud)

Once these requirements are clearly defined, a preliminary list of resource requirements for the system can be drawn up, as shown in the following example:

1. Data memory – the size of the RAM required?
   165 bytes
2. Flash memory – the size of the program space required?
   2300 words
3. Peripherals – list what on-chip peripherals are a 'must have' and what on-chip peripherals are 'nice to have'.
   Must Have:
   a) LCD peripheral
   b) USART
   c) ADC
   Nice to Have:
   a) Cap-touch peripheral
   b) Real-time clock peripheral
4. External circuitry – what other signal-conditioning or control circuitry are required?
   a. Temperature sensor
   b. Watchdog timer
   c. Open collector drivers for heat/cool
   d. Voltage regulator
5. Processing speed – how many MIPS are required for the job?
   500KIPS to 1MIPS

At this point, accuracy isn't essential; the aim is to get a rough order of magnitude to establish a numerical basis for the trade-off analysis that will follow. Figure 1 shows some of the alternatives that need to be weighed as part of the design.

The chart shows a range of 8-bit through 32-bit; each item listed is on a continuum, ranging from one extreme to the other. This doesn't mean that full functionality of all elements at the extreme of the chart are impossible, it just means that they are easier to implement. For example, Real-Time Response is listed at the 8-bit end of

Table 2: Initial project design constraints

Intelligent Home Thermostat
Cost Targets
☐ Material cost < $4
☐ Mechanical cost < $2
☐ Assembly cost < 10
System Power
☐ Battery backup < 20 µA
☐ Line power < 5 mA
Physical Size
☐ 3" x 4" x 1"
Operating Environment
☐ Temperature 0C to 30C operating -20C to 75C storage
☐ Electrical noise 3V Common mode noise (no affect) 1 KV ESD (human body model) Cell phone 2' proximity

Figure 1: System alternatives

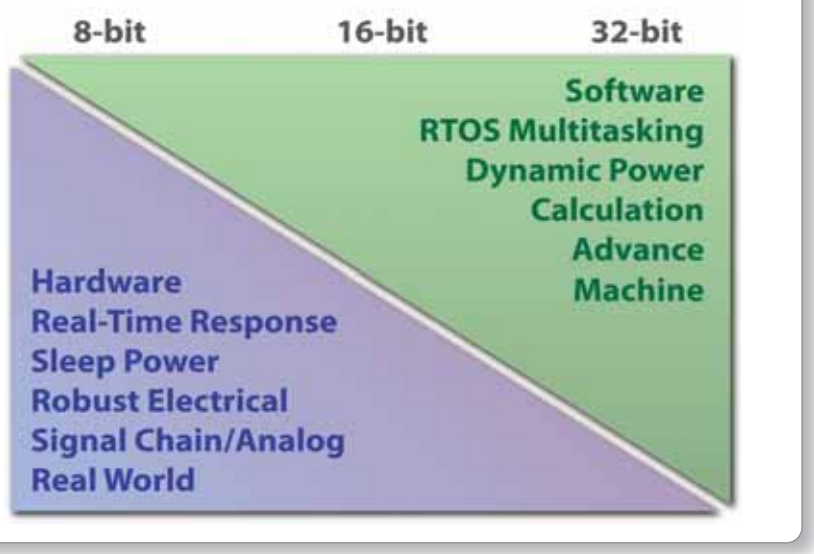

the chart, while RTOS is listed at the 32-bit end (even though an RTOS solution for 8-bit microcontrollers is possible, it's just that RTOSs are more prevalent at the 16- and 32-bit end, and their memory footprints are proportionally smaller in 16- and 32-bit MCUs).

Hardware or Software?

When it comes to digital peripherals, the first question relates to how to implement the functionality: software or hardware. This has far-reaching implications for the design. Processing power must be traded off against hardware complexity. It is the difference, for example, between implementing a software-based stepper motor controller and a controller chip. The selection will affect processor speed and program memory requirements, system cost, board size and possibly current consumption.

Some limitations may be hidden. For example, a USART function can be built in software, as both transmit and receive functions can be emulated in this way. The problem is the receive function must continuously poll for the falling edge of the start bit to accurately synchronise its receive timing. This can put a considerable drain on the processing power. Even if an interrupt-on-change function is available, the latency time of the interrupt service routine may make accurate timing of the edge problematic. At this stage, the best approach is to be aware of the possible limitation but initially just note the potential options and their associated costs.

Real-Time Response or RTOS?

Decisions must be made on how to multi-task in the design and making the choice between using a real-time operating system (RTOS) and building a system out of interleaved state machines. The RTOS handles all of the switching between states and simplifies the software design, while a real-time-response system requires a smaller memory footprint and makes real-time-response control easier to implement. Both options will impact the data/programming memory requirements and the processing power required for the system.

RTOSs fall into two basic categories: pre-emptive and cooperative. Both allow switching between multiple tasks, but differ in the trigger for the switch, which in turn has an impact on the requirements for variable storage and peripherals and should be factored into the system requirements.

Sleep Power or Dynamic Power Control?

The options are a control that has all-or-nothing current consumption, or a more graduated system that allows sections of the design to be turned off. Both have their merits; the decision will be heavily affected by the choices made earlier regarding the hardware and software. The question centres on whether some of the system should be left awake and drawing current to handle the system idea tasks, or employing an automated hardware system that can handle simple tasks while the processor is asleep and drawing much smaller quiescent current.

Robust Electrical Design

Although in an ideal world all designs should be electrically robust, there may be compromises concerning how robustness is achieved. This relates to both power and the choice of hardware and software.

For example, a power-conversion function could be delivered using a more software-centric design, which uses a software function to generate the feedback control. The alternative is to employ an op-amp based loop-filter with a simpler analog feedback PWM. Both systems work and are used by the industry. The differences are that the software-based system must remain active while the power converter is operating, which may require a higher operating current and, in the event that the program counter in the microcontroller is corrupted, the hardware-based system is unaffected by the microcontroller recovery function, making it more electrically robust in a noisy environment.

Analog Signal Chain or Calculation (DSP)?

Going back to examining the question of hardware versus software, the compromise here concerns system controls and whether to use an analog signal chain to condition input signals and generate outputs, or to use software. The answer has a significant impact on the microcontroller choice.

Realistically, the microcontroller must have a couple of on-chip features to implement a Proportional-Integral-Derivative (PID) function or a digital-filter design. The first requirement is a Hardware Multiply function that is capable of keeping up with the system requirements for feedback and control, and the second is a high-speed, high-resolution Analog-to-Digital Converter (ADC).

Simple feedback controls implemented using micros without these features tend to

Table 3: Example of a microcontroller selection matrix

Product Family	Architecture	5K $ Pricing	Flash (KB)	EEPROM (Bytes)	RAM (KB)	CPU Speed (MHz, MIPS)	LowPower	Comparators	ADC Channels	ADC Bits	Total UART	SPI	I2C	USB	Ethernet	LIN	CAN	Total Timers	Input Capture	PWM Channels	Parallel Port	Segment LCD	Supply Voltage
PIC24FJ64GB0	16	3.07	64	0	8.00	[32,16]	XLP	3	13	10	2	2	2	Devi		Yes		5	5	5	PMP	0	2 to 3.6
PIC24FJ64GA0	16	3.07	64	0	8.00	[32,16]	Yes	2	16	10	2	2	2			Yes		5	5	5	PMP	0	2 to 3.6
dsPIC33FJ32G	16	3.07	32	0	4.00	[120,4(	Yes		16	10	2	2	2			Yes		5	4	4	GPIC	0	3 to 3.6
PIC18F2331	8	3.08	8	256	0.75	[40,10]	Yes		5	10	1	1	1			Yes		4	2	2		0	2 to 5.5
PIC18F67K90	8	3.08	128	1024	3.70	[64,16]	XLP	3	16	12	2	2	2					11	10	10		132	1.8 to 5.5
PIC18F87J50	8	3.09	128	0	3.80	[48,12]	Yes	2	12	10	2	2	2	Devi		Yes		5	5	5	PMP	0	2 to 3.6
PIC32MX320F(	32	3.09	32	0	8.00	[40,40]		2	16	10	2	2	2			Yes		6	5	5	PMP	0	2.3 to 3.6
PIC16F84A	8	3.11	1.75	64	0.06	[20,5]			0			0	0					1	0	0		0	2 to 6
PIC18F86K90	8	3.11	64	1024	3.70	[64,16]	XLP	3	24	12	2	2	2					11	10	10		192	1.8 to 5.5
PIC18F87J93	8	3.12	128	0	3.80	[48,12]	Yes	2	12	12	2	1	1					4	2	2		192	2 to 3.6
dsPIC33FJ64G	16	3.12	64	0	8.00	[80,40]	Yes	2	10	12	2	2	1					7	4	4	PMP	0	3 to 3.6
dsPIC33FJ32M	16	3.12	32	0	4.00	[80,40]	Yes	2	9	10	2	2	1					7	4	4	PMP	0	3 to 3.6
PIC24HJ64GP2	16	3.12	64	0	4.00	[80,40]	Yes	2	10	12	2	2	1					7	4	4	PMP	0	3 to 3.6
PIC16F737	8	3.13	7	0	0.36	[20,5]	Yes	2	11	10	1	1	1					3	3	3		0	2 to 5.5
PIC18F8390	8	3.15	8	0	0.75	[40,10]	Yes	2	12	10	2	1	1			Yes		4	2	2		192	2 to 5.5
PIC18F4510	8	3.18	32	0	1.50	[40,10]	Yes	2	13	10	1	1	1			Yes		4	2	2	PSP	0	2 to 5.5
PIC18F2520	8	3.18	32	256	1.50	[40,10]	Yes	2	10	10	1	1	1			Yes		4	2	2		0	2 to 5.5
PIC24FJ96GA0	16	3.19	96	0	8.00	[32,16]	Yes	2	16	10	2	2	2			Yes		5	5	5	PMP	0	2 to 3.6
PIC18F66J65	8	3.19	96	0	3.70	[42,10.	Yes	2	11	10	1	1	1		10 B	Yes		5	5	5		0	2 to 3.6
PIC24FJ64GA1	16	3.20	64	0	16.0	[32,16]	Yes	3	16	10	4	3	3			Yes		5	18	9	PMP	0	2 to 3.6
PIC18F87K22	8	3.21	128	1024	3.80	[64,16]	XLP	3	24	12	2	2	2					11	10	10	PSP	0	1.8 to 5.5

be slow and typically use up all of the available processing power to accomplish the task. The simpler option is to opt for an analog system to accomplish the necessary functions with the microcontroller performing a supervisory role. Some functions will require greater speed or are more difficult – even impossible – to implement in analog terms, such as non-linear feedback controls and Infinite Impulse Response (IIR) filters. In these cases the best option is to use a software solution alongside any necessary analog fault-detection circuitry.

Real World or Advanced Machine?

The final question relates to usage and if the system has a small automated control with a limited user interface, or a more advanced system with networking and a more complex user interface. Whilst there would be no need to use a QVGA with a touch screen to implement a $9.99 home thermostat control, such a user interface may be needed for a system that is networked into a large building environmental control system. The answer lies in the size of the required bandwidth for the user interface and if any connectivity is needed. The aesthetic "Look and Feel" trends in the marketplace may impact the choice as well as more practical requirements for operation by a user wearing gloves. All of these requirements should be weighed in turn to determine the appropriate complexity for the system.

Once the list of system resource requirements has been adjusted to allow for the acceptable compromises, the designer can start to look for the appropriate microcontroller to meet those requirements. Table 3 shows an example microcontroller selection matrix which reflects the examination of the processing power, memory (flash and RAM) and peripheral mix required.

An exact match is unlikely; the designer will need to revisit some of the compromises and make some hard choices. Having already highlighted the peripherals which are essential and those which fall in to the category of a "nice to have" proves valuable when making the final decision. It is recommended to employ a microcontroller that is part of a family with upward and downward-compatible siblings and which provides options in the event that memory and/or peripherals need to be added or subtracted to add features or cut costs.

There is silicon- and a test-cost associated with each function, and they soon start to add up. Some functions have more incremental cost than others, so when the final product selection options are narrowed down, it's good to refer back to the "must have" vs "nice to have" lists. Also, the cost and complexity of the required design tools should be considered as this will impact the support, testing and production staff making the final product.

Additional Considerations

The availability of pre-built stacks and library functions is another important design consideration. Some manufacturers make a pre-built code available for specific functions, such as graphical display drivers, serial-communication functions and capacitive-touch functions. These standard blocks can save a considerable amount of time and testing, but those blocks may not be compatible with the design. For example, some blocks may monopolise one or more peripherals, preventing other functions from accessing them. There may also be timing requirements that conflict with other functions in the system, so it is important to ask specific questions before committing to the use of prefab code.

Finally, it's worth considering code reuse as part of the final decision as there is no reason to keep reinventing the wheel in each new design. Finding a manufacturer that produces a broad range of microcontrollers and sticking with them will allow the reuse in future designs, and this will cut development and testing time. ●

New 8-bit Microcontrollers with integrated configurable logic in 6- to 20-pin packages

Microchip's new PIC10F/LF32X and PIC12/16F/LF150X 8-bit microcontrollers (MCUs) let you add functionality, reduce size, and cut the cost and power consumption in your designs for low-cost or disposable products, with on-board Configurable Logic Cells (CLCs), Complementary Waveform Generator (CWG) and Numerically Controlled Oscillator (NCO).

The Configurable Logic Cells (CLCs) give you software control of combinational and sequential logic, to let you add functionality, cut your external component count and save code space. Then the Complementary Waveform Generator (CWG) helps you to improve switching efficiencies across multiple peripherals; whilst the Numerically Controlled Oscillator (NCO) provides linear frequency control and higher resolution for applications like tone generators and ballast control.

PIC10F/LF32X and PIC12/16F/LF150X MCUs combine low current consumption, with an on-board 16 MHz internal oscillator, ADC, temperature-indicator module, and up to four PWM peripherals. All packed into compact 6- to 20-pin packages.

Go to www.microchip.com/get/eunew8bit to find out more about low pin-count PIC® MCUs with next-generation peripherals

FAST-START DEVELOPMENT TOOLS

PICDEM ™ Lab Development Kit - DM163045

PIC16F193X 'F1' Evaluation Platform - DM164130-1

PICkit™ Low Pin Count Demo Board - DM164120-1

Free CLC Configuration Tool:
www.microchip.com/get/euclctool

MICROCHIP

Microcontrollers • Digital Signal Controllers • Analog • Memory • Wireless

PUSHING PERFORMANCE LIMITATIONS IN MICROCONTROLLERS

TRUE GAINS IN EFFICIENCY OF MICROCONTROLLERS MUST BE BASED ON ARCHITECTURAL INNOVATIONS, SAYS **KRISTIAN SAETHER**, PRODUCT MARKETING MANAGER, AVR, AT ATMEL

Microcontrollers continue to evolve into more complex systems that are expected to perform a wider range of tasks at a lower cost and power than previous generations. To operate at the high data rates and frequencies required for real-time embedded processing, these microcontrollers must achieve higher processing efficiency. Today's microcontrollers need to perform a wide range of tasks, including managing real-time control algorithms, decode high-speed communications protocols and process signals from high-frequency sensors. Polling methods, such as checking interface ports to see if new data has arrived, consume too many CPU cycles and often have too high a maximum response time to reliably service I/O and peripherals.

For most embedded applications, developers rely upon interrupts to meet their real-time requirements for managing peripherals. Interrupts, however, can only determine when a real-time event has occurred. Developers must still directly involve the CPU to read I/O and the peripherals before data is lost.

Handling an interrupt requires potentially interrupting other latency-sensitive tasks; it also incurs context switching overhead and introduces a wide range of esoteric challenges, such as managing latency when multiple interrupts occur concurrently. All of these reduce predictability and processor efficiency.

To be able handle the high data rates and frequencies of real-time I/O and peripherals, microcontrollers must achieve higher processing efficiency. This efficiency, however, needs to be founded not on increased clock frequency, which comes at the expense of higher power consumption, but through internal changes in microcontroller architectures. Specifically, microcontrollers have begun to integrate coprocessors that offload specific task blocks, multi-channel DMA controllers for facilitating penalty-free memory access and integrated event systems which route signals between internal subsystems to offload I/O and peripheral management.

More Than One Way to Offload a CPU

Integrated coprocessors have become widespread in a range of embedded microcontrollers. Among the more commonly recognized coprocessors are encryption and TCP/IP offload engines. Effectively, coprocessors offload entire tasks or assist in the more computer-intensive portions of complex algorithms. For example, an encryption engine reduces AES computations on the CPU down from thousands of cycles to hundreds of cycles per operation, while a TCP/IP offload engine makes it possible to terminate an Ethernet connection with little CPU overhead. In addition, offload engines simplify implementation of these tasks, eliminating the need for extensive driver code creation by allowing developers to access this advanced functionality through the use of simple APIs.

Microcontrollers have begun to integrate coprocessors that offload specific task blocks, multi-channel DMA controllers for facilitating penalty-free memory access and integrated event systems which route signals between internal subsystems to offload I/O and peripheral management

DMA and event system technologies are less familiar to developers and often not used for this reason. DMA controllers offload management of data movement from the CPU by performing data accesses, such as peripheral registers to internal or external SRAM in the background. For example, a developer can configure the DMA controller to preload a block of data into on-chip RAM so that it is available for fast access before the CPU needs it, thus eliminating wait states and dependency delays. Alternatively, a DMA controller can assume most of the burden of managing communication peripherals (see Table 1).

The savings in terms of cycles from using a DMA controller can be significant: many embedded developers have found themselves unable to fit an application within the resource limits of a microcontroller, only to be introduced to a DMA controller and find extra cycles – sometimes between 30-50% across the system. It is only when they face a processing 'wall' that many developers discover the untapped potential that's been available to them from the start.

Even fewer developers are familiar with event systems that work in conjunction with DMA controllers to further offload CPU cycles, as well as reduce overall power consumption. An event system is a bus that can connect internal signals from one peripheral on a microcontroller to another. When an event occurs at the peripheral, it can trigger action to be taken in other peripherals, without involving the CPU and within a two-cycle latency, much the way the human body processes reflexes like pulling a person's hand out of a fire without having to first consult the brain.

Specifically, an event system routes signals throughout the microcontroller using a dedicated network connecting the CPU, data

bus, peripherals and DMA controller. Normally, peripherals must interrupt the CPU to initiate any action, including reading the peripheral. By routing events directly between peripherals, the event system in effect offloads these interrupts from the CPU.

Developers have the flexibility to configure peripherals to follow different event channels, thus defining the particular event routing required to meet the specific needs of the application.

Flexible Offloading

The combination of DMA working with the event system enables developers to offload entire tasks, much like a coprocessor does. One key difference is that coprocessors are not programmable. They implement a well-defined task in hardware and at best are configurable.

The programmability of a DMA controller and event system makes them appropriate for a variety of tasks, from the most simple to the very complex. In the case of using DMA with an event system, the DMA manages the transfer of data throughout the microcontroller's architecture, while the event system controls when these transfers occur with low latency and a high degree of accuracy. Put another way, the event system makes sure the values managed by the DMA are sampled or output at the right time/frequency.

Figure 1 shows a block diagram of how an event system and DMA work together. The ADC is connected to a sensor and will collect samples. An internal counter is set to match the sampling frequency, providing regular and accurate intervals. Rather than interrupt the CPU to sample the ADC, the event system directly initiates sampling of the ADC. This results in the sampling frequency being extremely accurate relative to the microcontroller's clock. When the ADC settles and the conversion is completed, the ADC then triggers the DMA to store the value through the event system.

A DMA controller and event system work together to offload peripheral processing from the CPU. An internal counter sets the sampling frequency, providing regular and accurate intervals, or an input signal (event 1) could trigger an ADC to sample (event 2) and save the value to DMA (event 3) until the DMA buffer is full (event 4). In this configuration, the CPU is only interrupted once there is a full buffer of data for it to process.

Event management can be extended to include multiple events, connecting several peripherals to create more complex configurations. For example, an input signal (event 1) could trigger an ADC to sample (event 2) and save the value to DMA (event 3) until the DMA buffer is full (event 4). In this configuration, the CPU is only interrupted once there is a full buffer of data for it to process.

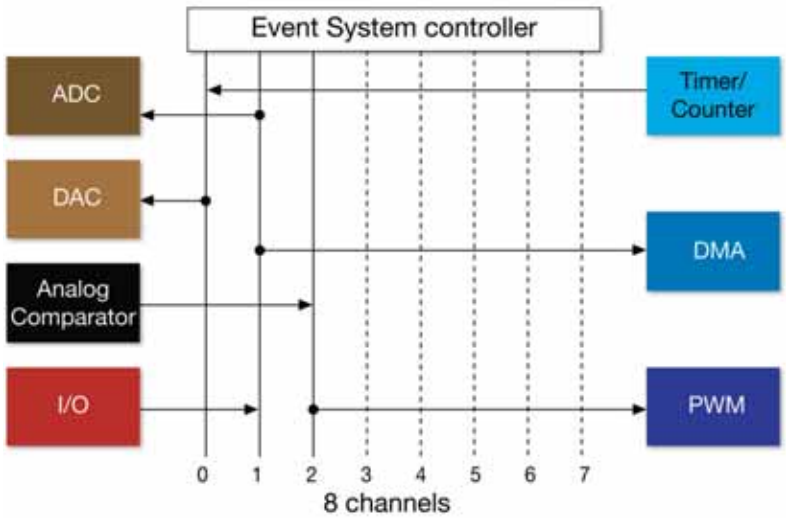

Figure 1: Block diagram of how an event system and DMA work together

Both DMA controllers and event systems also support multiple channels. This allows developers to configure an interconnected fabric that operates in parallel to the main CPU. As a result, multiple concurrent real-time tasks can be coordinated in a deterministic fashion.

Determinism and Latency

Determinism plays a key role in limiting latency and managing the responsiveness of real-time embedded systems. The more deterministic the system, the more consistent its responsiveness will be.

A primary factor affecting determinism is how many interrupts a system must handle concurrently. In general, increasing the number of interrupts in the system will erode its determinism.

Consider a system with a single interrupt that completes within 50 cycles. Latency for such an interrupt, then, is consistently on the order of 50 cycles. Note that even the simplest interrupts take on the order of 50 cycles by the time the microcontroller saves the context for a limited number of registers, accesses a peripheral, saves the data, restores the context and performs a pipeline flush.

However, it isn't the prospect of handling a single interrupt that creates the most problems for developers in terms of determinism and latency. Rather, the hard challenge in meeting real-time deadlines arises when many interrupts happen at the same time. For example, if a higher priority interrupt is introduced to the system that completes within 75 cycles, the latency of the first interrupt is affected since the higher priority task can interrupt. Latency now ranges from 50 to 125 cycles for the lower priority task.

As more interrupts are introduced, the latency for lower priority interrupts increases

UART Transfer Speed	CPU usage with	CPU usage without
9.6	0.01	0.26
19.2	0.01	0.52
38.4	0.03	1.04
57.6	0.04	1.57
115.2	0.08	3.14
1200	0.85	34.15
3500	5.17	99.5

Table 1: DMA controllers can assume most of the burden of managing communication peripherals

		Active		Idle Mode		Power	save	
Atmel XMEGA	Cycles	Time	Current (mA)	Time	Current (mA)	Time	Current (mA)	Average current
No DMA, No Event System	1200000	10%	3.80	0%	1.90	90%	0.55	380.495
DMA on	960000	6%	3.80	4%	1.90	90%	0.55	304.495
DMA and Event System	840000	5%	3.80	5%	1.90	90%	0.55	285.495

Table 2: Power figures for an application requiring 1.2 Mcycles/s

as determinism drops. A 50-cycle task might be interrupted multiple times and require on the order of 100s to 1000s of cycles to complete. This factor is important because not all interrupts can be high priority, relative to each other.

Determinism directly affects responsiveness, reliability and accuracy. If a developer knows latency is fixed at 50 or 500 cycles, this can be taken into account during processing. However, when latency can vary from 50 to 500 cycles, the best developers can do is assume a typical latency such as 200 cycles and count any variation as error. Additionally, worst-case latency may begin to approach real-time deadline limitations and threaten system reliability.

Reducing the potential number of concurrent interrupts – even low frequency ones – through a DMA controller and event system can substantially increase system determinism as well as lower latency. Higher determinism also results in important factors such as higher accuracy.

Achieving Higher Accuracy

To understand how determinism affects accuracy, consider the implementation of a power supervisory task to maximize AC power efficiency when driving a heavy load such as a motor. As the most energy is available when the voltage is at its peak and in phase with the current, this is when the system should draw the most current. Conversely, the closer the voltage gets to zero (i.e. the zero-crossing point), the less power is available and so the less efficient the current draw. By implementing Power Factor Correction (PFC), power efficiency is improved by switching in and out large capacitors that will adjust the load to keep the AC current and voltage in phase.

Typically a comparator is used to detect the zero-crossing. When the voltage drops below or rises above a set threshold, the comparator toggles. Instead of the comparator triggering an interrupt and forcing the CPU to switch the capacitors, the event system can route the comparator event directly to the Timer/Counter output controlling the switch without CPU intervention.

Interrupt latency for a low priority task like PFC could run into thousands of cycles, depending upon how many higher priority interrupts occur concurrently. Higher latency means the capacitors are switched later than the optimal moment, reducing overall efficiency by a significant amount. Latency from event routing, in comparison, is at most two cycles.

Consider these numbers over the microcontroller's clock rate. If the microcontroller is clocked at 32MHz, two-cycle latency introduces negligible error (2/32M). On the other hand, latency in the thousands of cycles could materially affect the accuracy of high-frequency tasks which themselves must be processed every few thousand cycles. Note that this latency could be reduced to the order of 50 cycles if the interrupt were made a higher priority task. However, this leads to assigning priority based on accuracy requirements, rather than importance of the function to the system. It also merely shifts the inaccuracies from lack of determinism to other tasks.

> Coprocessors offload well-defined, compute-intensive tasks from the CPU; DMA controllers relieve the CPU from moving data throughout the system; and event systems eliminate bottlenecks associated with multiple and frequently-triggered interrupts

Higher accuracy also plays a key role when generating signals, not just sampling them. Consider creating a 100kHz waveform. Using interrupts, the accuracy of the waveform will be affected by the variable latency relative to the signal rate, slightly slower or faster, based on context switching and the other interrupts that have piled up. Note that while the waveform will be accurate on average, what matters in many cases is the relative difference between two consecutive samples.

High-Frequency Signal Processing

Generating signals is becoming a more common task in a wide range of embedded applications. Signals are used to generate sound, manage a voltage converter regulator, control actuators in industrial applications and serve myriad other functions. The higher the frequency of the signal, the higher the load on the CPU when interrupts are employed, and the greater the potential of increased latency for other tasks.

For events with a higher frequency of occurrence, CPU load becomes a major consideration. For example, a high-speed sensor needs to have samples collected before the next sample is ready in order to prevent loss of data. A flow meter, multi-axis positioning system, or instrumentation system collecting 2MS/s for fast and accurate measurements will consume tens to hundreds of millions of cycles each second just to collect the samples. With an event system and DMA controller, all these cycles are offloaded from the CPU to actually process the samples, not simply buffer them. Even assuming a simple task that requires only 50 cycles to complete the task with context switching overhead, this results in offloading 100 Mcycles from the CPU. For this reason, many systems may use a separate microcontroller to manage individual high-frequency sensors or motors.

For higher frequency tasks, an event system and DMA controller also enable:

- **Accurate time-stamping:** Time-stamping samples enable developers to better synchronize signals to external events. With two-cycle latency, time stamps are far more accurate than those marked by interrupts that have latency on the order of thousands of cycles.

- **Oversampling:** One way to increase sensor resolution is to oversample. For example, dividing a counter by 16 will result in 16 times the number of samples collected, increasing the overall accuracy of the sensor. Because the CPU is not directly involved in taking and storing samples, it becomes possible to oversample without much penalty.
- **Dynamic frequency:** Certain applications require higher sensing accuracy only at certain times or under specific operating conditions. For example, a water meter could sample faster when the flow rate is changing quickly, and scale back when the flow is turned off or consistent. Sampling frequency is easily adjusted without impacting real-time responsiveness.
- **Reduced stack size:** An additional effect of reducing the number of concurrent interrupts is the ability to maintain a smaller stack. As each interrupt must perform a context save by pushing potentially dozens of registers to the stack, eliminating several layers of context significantly reduces the size of the stack required. This could lead to using less RAM for the application.
- **Immunity to scaling:** Given that different microcontrollers support a different number of peripherals, the number of interrupts can vary in an application across its price line. Even though the same microcontroller family is in use, a higher-end system that supports more functions will have more interrupts, degrading overall determinism. Thus, migration of a design to a more integrated microcontroller could impact signal latency and consequently accuracy, both for sampling and outputting.
- **Easy software changes:** By having event handling that eliminates CPU intervention, software changes can be done without affecting real-time responses. Even if more CPU time is needed to handle additional functions, the event handling and response time will remain exactly the same. Without this it can be difficult to implement changes during a product's lifetime for real-time applications.

Autonomous Control

There are many tasks an embedded microcontroller could perform to reduce power consumption, increase accuracy and improve the user experience. Many such tasks are but simple monitors, checking only a single value. For example, a battery monitor watches until the voltage drops below a certain level. Then the system triggers a shutdown to save application data while there is still enough power to do so.

Improving the user experience is often a key differentiator in many consumer products. For example, an event system enables faster responsiveness to a wake-up keystroke or peripheral input, enabling reaction within two cycles. Compare this to the responsiveness of using an interrupt. Interrupts also require the system to return to active mode, lowering power efficiency. For this reason, developers often extend timer intervals, reducing responsiveness.

Using interrupts, the cost of implementing such tasks has been too high in terms of CPU processing required, added latency and decreased determinism. With an event system and DMA controller, developers have the ability to implement such features while effectively bypassing the CPU. This not only reduces the number of interrupts the system must manage but also simplifies task implementation and management.

For example, consider an application that will play a warning message to a user under specific operating conditions. The pre-set sound file can be stored in a buffer and fed to the speaker through the appropriate peripheral using DMA. The event system, using a timer, will ensure that the data is applied at a rate of exactly 44,056kHz. As a side benefit, because the frequency is accurate and consistent, sound fidelity is increased as well. From a performance standpoint, once the DMA and event system have been configured, the CPU is completely uninvolved in the playback task.

To say such tasks are "free" would be an overstatement. Implementing them in this fashion, however, makes them feasible in a much wider array of applications. The combination of coprocessors, a DMA controller and an event system frees up a controller to handle just the processing of a signal rather than have the majority of its resources consumed in the signal's cycle-intensive acquisition. As a result, the CPU retains much of its processing capacity for signal processing. In this way, it becomes possible for a single controller to manage several high-frequency tasks rather than just one. This also simplifies system design, permits more tasks to be implemented at a lower price on a single microcontroller, enables easier correlation between multiple signals and improves power efficiency.

For many applications, the ability to support multiple tasks can lead to important product differentiation. For example, a motor control application utilizing a DMA controller and event system could free up enough resources on the microcontroller that developers could implement advanced features such as PFC without increasing the system BOM.

In addition to increasing the performance capacity of a microcontroller by offloading interrupts, the event system can also reduce power consumption by a factor of up to 7X, depending upon the application. Table 2 shows the power figures for an application requiring 1.2 Mcycles/s, using an Atmel XMEGA MCU as an example device. At 12MHz the microcontroller is in active mode 10% of the time and standby mode the rest. Implementing a DMA controller and event system offloads the number of cycles the CPU must execute each second, enabling the microcontroller to drop into idle sleep mode. Given that the current draw in active mode is substantially greater than when in idle sleep mode, even a few percentage changes in active mode results in significant power savings.

An event system and DMA controller not only increase CPU capacity and performance, they can also significantly reduce power consumption, depending upon the application, by enabling the microcontroller to drop into idle or sleep mode more often. Given that the current draw in active mode is substantially greater than when in idle or sleep mode, even a small percentage change in active mode results in significant power savings.

Continued Evolution

Embedded microcontrollers continue to improve performance through architectural changes that improve overall CPU capacity. Coprocessors offload well-defined, compute-intensive tasks from the CPU; DMA controllers relieve the CPU from moving data throughout the system; and event systems eliminate bottlenecks associated with multiple and frequently-triggered interrupts.

By reducing the number of concurrent interrupts the system must handle, developers can increase system determinism, leading to lower latency, improved signal resolution and accuracy, higher consistency and predictability and greater system reliability. As a result, a single microcontroller can perform the work of multiple older microcontrollers without increased system cost or power consumption.

CONTROLLING OS-BASED 32-BIT MICROCONTROLLER APPLICATIONS

GRAEME CLARK, PRODUCT MARKETING SPECIALIST IN THE INDUSTRIAL BUSINESS GROUP AT RENESAS ELECTRONICS EUROPE, LOOKS AT A NEW REAL-TIME STRUCTURE FOR MICROCONTROLLERS THAT CAN PROVIDE FASTER REAL-TIME PERFORMANCE IN A SYSTEM WITHOUT INCREASING THE CODE SIZE OR COMPLEXITY

One of the biggest challenges facing developers today is to find ways of increasing real-time system performance without increasing clock speed (due to EMC and power constraints), software complexity and costs.

For developers of many microcontroller-based applications, this has not been a big problem until recently. However, this is becoming more of an issue as applications begin to encompass operating systems (OS) based designs and their necessity for software abstraction outside the microcontroller. This is true even for the developers of traditional 8- and 16-bit microcontroller applications, but more so for those using 32-bt microcontrollers.

This demand is being driven by a number of conflicting requirements:

- The increasing trend towards using operating systems or schedulers in a growing number of applications – using an OS can add greatly to the system overhead in handling real-time events.
- The addition of increasingly complex functions in even the simplest of applications, especially in the areas of user interface and communications. This has resulted in less processor time being available for the real-time tasks, even though they may be at the core of the application.
- The requirement for simplified testing and, hence, the need to increase the reliability of the system.

With this in mind, at lot of research is going into finding suitable solutions. As no one has yet found a perfect one, we are now looking at the fundamental way of how microcontrollers handle real-time events, and seeing if this can be improved.

Figure 1: The RX210 32-bit microcontroller feature set

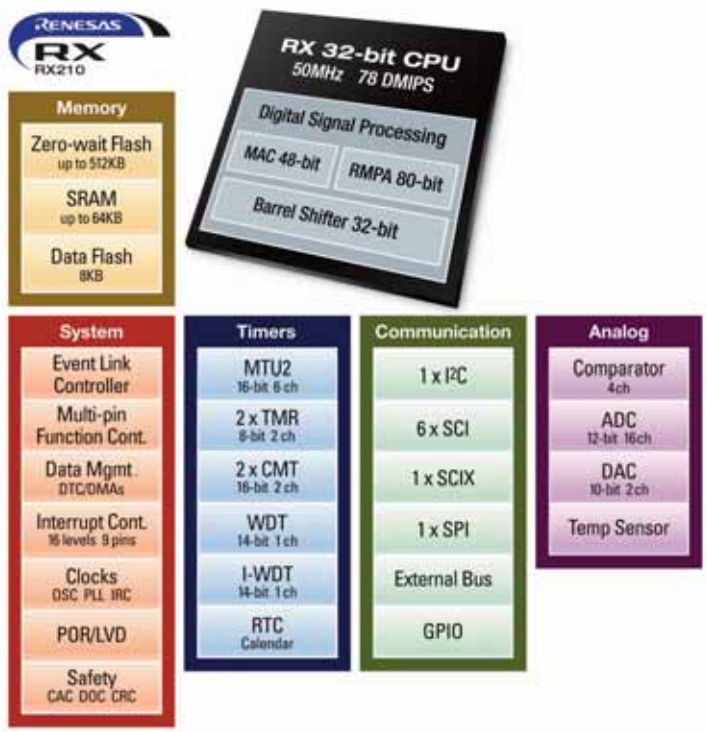

Reinventing the Basic Real-Time Structure

One solution is to reinvent the basic real-time structure of the microcontroller, and to provide it with additional hardware to manage real-time events, until the operating system can "catch up". By using this technique, the fundamental real-time response issues can be moved from the OS and the application software (where they currently reside in a classical microcontroller design) to a more hardware-based solution. This can result in faster response to real-time events and a large reduction in software size and complexity. Such a new structure allows the hardware of the chip itself to handle many of the real-time events that in the past would have required CPU intervention.

Here we will present the implementation of this new structure, on the new RX210 32-bit microcontroller from Renesas, seeing the many benefits it can bring, including significantly faster real-time response as well as a reduction in code complexity.

The RX210 is the latest in a long line of powerful 32-bit microcontrollers based on the popular RX CPU core from Renesas. The features of this device can be seen in Figure 1.

This new structure for managing real-time events allows the creation of up to three types of "events" inside the

Figure 2: The "interrupt " controller on the Rx210

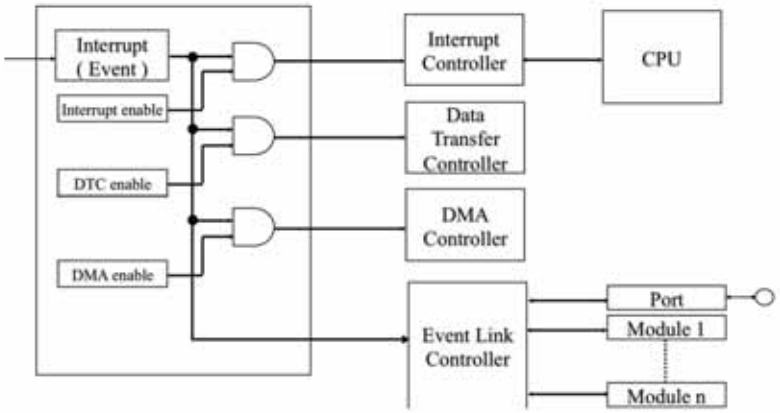

microcontroller; and each may be generated from any external or internal interrupt source. These can include a wide variety of real-time stimuli, ranging from a simple occurrence such as a changing state of an input pin, to more complex events such as a timer timing out or a byte arriving in a serial interface.

Figure 2 shows the block diagram of the RX210 event controller and its associated logic, which can be used to handle real-time events by a number of different mechanisms, depending on the users' requirements, each offering a number of advantages and disadvantages.

RX210 with the New Structure

The RX210 implements standard CPU interrupts, using a 16 priority level interrupt controller, however, as can be seen in the diagram, you also have a choice of generating an automatic data transfer using either the high speed but inflexible Direct Memory Access (DMA) controller or the slightly slower but much more flexible Data Transfer Controller (DTC) instead of an interrupt. The final choice is to generate an "event" using the Event Link Controller (ELC). You can choose to generate any one of these or almost any combination of them in response to any real-time event that occurs inside or outside the microcontroller. However, the real benefit of this device is in handling such events using either the DTC or ELC, as these are then handled without any CPU time, and so are being executed faster and without any software overhead.

The DMA controller is designed to provide a high speed method of transferring data from a peripheral to memory or memory to a peripheral, without interfering with CPU execution. It provides an extremely fast method of transferring data, but is very inflexible as each channel of the DMA controller is fixed to generate a particular transfer.

The DTC provides mostly the same functions as the DMA controller, as it is designed to allow the transfer of one or more bytes between memory and peripherals and vice versa, with the transfer being triggered by an event on the chip. However, it is designed to be a lot cheaper to implement than a DMA controller, and more flexible.

The DTC Controller

The DTC controller provides a high level of flexibility in its operation by using a simple programmable controller to make the transfers rather than a large block of dedicated hardware. This makes it extremely programmable; it holds its set-up information in a small block of on-chip SRAM (rather than in a fixed set of on-chip registers), and as such it can be used to create 10 or 20 channels of data transfer and not just one or two. A small block of SRAM is dedicated to holding the configuration information for one transfer channel. The disadvantage is that for each transfer a few cycles are taken to read the configuration held in the SRAM before each transfer takes place.

The DTC can transfer 1 byte or more between a peripheral and memory or memory and peripheral, up to 256 times. The address in memory can be the same or it can be incremented or decremented, as such creating a buffer structure.

At the end of the transfer, the DTC can generate an interrupt, to tell the CPU the data is ready, or it can trigger a second DTC transfer; it can be used to daisy-chain a number of transfers, if required. Such a chain mode can be especially useful if multiple blocks of data need to be moved between peripherals. And, one interrupt can generate a complex sequence of different transfers. Some of the different operating modes of the DTC are shown in Figure 3. The DTC can also be placed into repeat mode, where it will continue to repeat the transfer as long as specified. Automated transfers between any peripheral and memory can be created almost without limit, and most importantly without any CPU intervention.

The Event Link Controller (ELC)

The final method of handling real-time events is to use the new ELC, which is designed to allow events on one peripheral to be automatically linked in hardware to another peripheral. For instance, by using the ELC, an event, such as an I/O pin

Figure 3: Data transfer controller operation modes

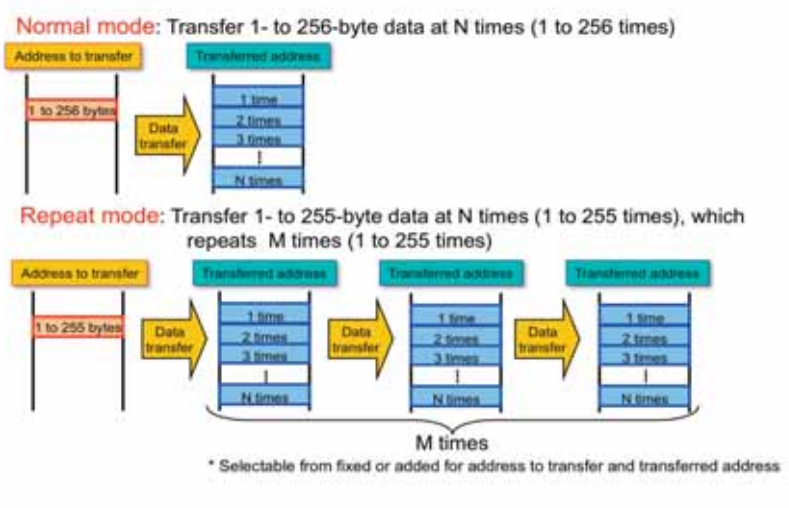

Figure 4: Block diagram of the event link controller

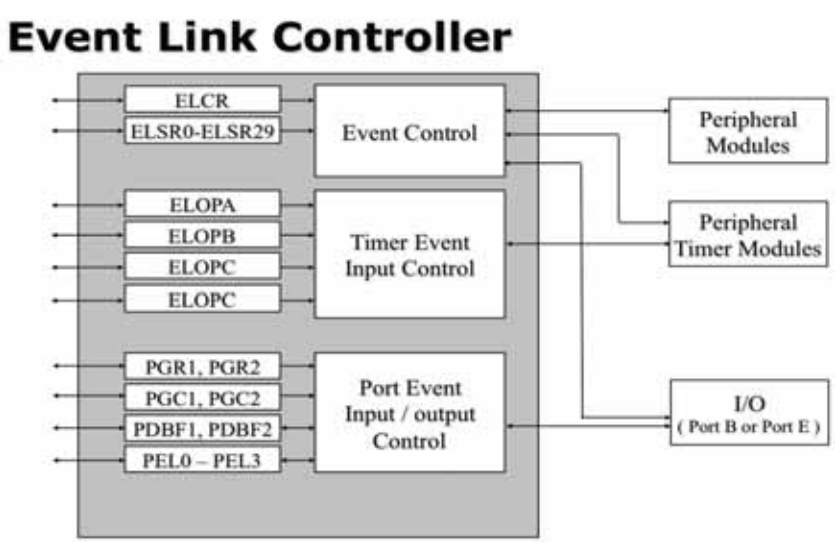

changing state, or a timer timing out, can trigger another action, such as starting another timer or kicking off an ADC conversion, or even causing an I/O pin to toggle – completely automatically.

The main functional blocks of the ELC are shown in Figure 4. It comprises four main modules: the Peripheral Event Controller for the peripherals, the Timer Event Input controller, the Port Event Controller and the Event Generation Timer.

The ELC allows the user to choose which peripheral is controlled by which event. For instance, the overflow from a 16-bit timer can start the ADC. The Timer Event Input controller selects what each timer will do when it is triggered by an event. A list of all the possible actions that can be generated in response to an event is shown in Figure 5.

One of the most powerful features of the ELC is the ability to generate events from the I/O ports. The port pins can be used individually or in groups of 4 pins. Each pin or group of pins can be used to generate an event, which then can control other functions on the chip. However, they can also be controlled themselves by the ELC.

When an event in generated, you can choose, for instance, to toggle a pin or group of pins or drive them to a specific logic level. Behind each group of pins, there is also a port buffer, which can then use an event to drive the contents of that buffer onto the pins, and if required, rotate the buffer ready for the next event. This type of feature is ideal for generating output waveforms of various types, such as for driving stepper motors and LCD displays. The buffer can also be used to latch the input values on the pin when an event occurs.

These features allow the developer to automate a number of real-time events in a system; an event will take 3 clock cycles from the time it is recognised before the resulting

Figure 5: Rx210 event link controller operations

Module Operation when Event is Input

Module	Operations when Event is Input		
MTU2 CMT TMR	Each timer operates differently depending on the ELOPA to ELOPD registers as below. · Starts counting when an event signal is input. · Restarts counting when an event signal is input. · Counts the input events (CMT, TMR). · Performs input-capture operation when an event is input (MTU2).		
A/D converter	Starts A/D conversion when an event signal is input.		
D/A converter	Starts D/A conversion when an event signal is input		
Output ports	The value of PODR (output port register) changes when an event signal is input. (The value of the signal to be output from the relevant external pin changes.)	Port group	The port group operates differently depending on the settings as below. · Changes the PODR value to the specified value. · Transfers the PDBFn value to PODR. · Rotates out the bit value.
		Single-port	Changes the PODR value to the specified value.
Input ports	When the signal value of the input pin changes	Port group	Generates an event.
		Single-port	
	When an event is input	Port group	Transfers the signal value of the external pin to PDBFn.
		Single-port	Event connection is not possible.
Clock oscillator	Switches the clock source to the low-speed on-chip oscillator operation		
Interrupt controller	Issues an interrupt request to the CPU, and the DTC starts to transfer data.		
POE2	Drives the MTU2 complementary PWM output pin and MTU0 output pin to the high-impedance state.		

action starts, for instance an I/O pin toggling or a timer making a capture.

One Example

The use of the ELC alone enables the users to handle many real-time events in hardware on the chip automatically, without the overhead of CPU intervention.

Figure 6 shows a simple example of the use of an ELC. An external interrupt causes a timer to start, after the timer overflows this then causes the ADC to begin a conversion. When this system is implemented on a typical microcontroller, it requires 3 interrupts: one for the external interrupt, one for the timer overflow interrupt and one for the ADC end of conversion, with the resulting CPU overhead and the software required for each Interrupt Service Routine.

The system will also experience some jitter, depending on when the respective interrupts can be serviced and the length of the respective ISRs. This is especially true if the system is using an operating system, which may be handling a higher priority task when one or more of these interrupts occur, adding to the time delay before the event can be handled.

After the ELC is initialised, the whole process can be handled automatically, without any CPU involvement, until the complete process is finished. Figure 7 shows the comparison between the process flow when using interrupts and associated Interrupt Service Routines (ISRs), and the new Event Link Controller.

In a typical microcontroller, after each interrupt is generated, it has to be 'serviced', resulting in a delay and, depending on exactly when the interrupt can be serviced, some jitter in the actual execution time. In a microcontroller with an ELC, each event triggers the next peripheral automatically. In this example, the CPU is only interrupted after the ADC conversion has finished. Therefore there's no jitter, as each peripheral starts exactly 3 clock cycles after the generating event and is not effected by the condition of the CPU, whether it's handling a priority task or not. There's the additional saving too, as, apart from the initialisation, there is no software required for this process and thus the code space required for the application is smaller.

Figure 6: Using the Rx200 event link controller

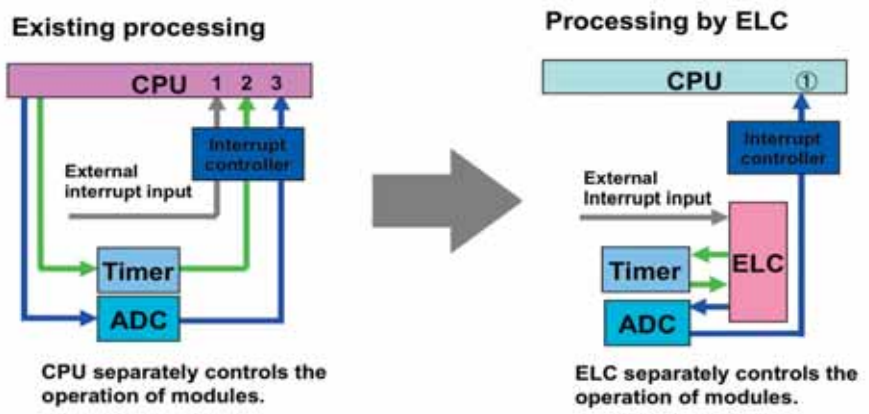

We could also go even further to automate this example and reduce the CPU load by using the DTC. Data could be automatically transferred from the ADC into a buffer in an SRAM. The CPU can be interrupted perhaps only after every 100 or 200 readings.

The combination of the DTC and ELC means that many low-level tasks, especially those based around the I/O and timers, can be handled automatically without CPU intervention. This in turn reduces the load on the CPU and greatly reduces system reaction time to external events.

Increasing Demands on the MCU

In many of today's systems we are asked to add more and more functions while still maintaining and even improving the real-time operation of the system.

The RX210 family from Renesas is one of the first devices to combine features like the Event Link Controller and the Data Transfer Controller, which together can be used to remove much of the real-time response requirements of the system and can automate many of these low level and repetitive tasks.

These devices are aimed at a wide range of communications and control applications such as motor control, measurement applications, low power modems and many others, and come with a wide range of memory size and package options.

Figure 7: Processing time comparison

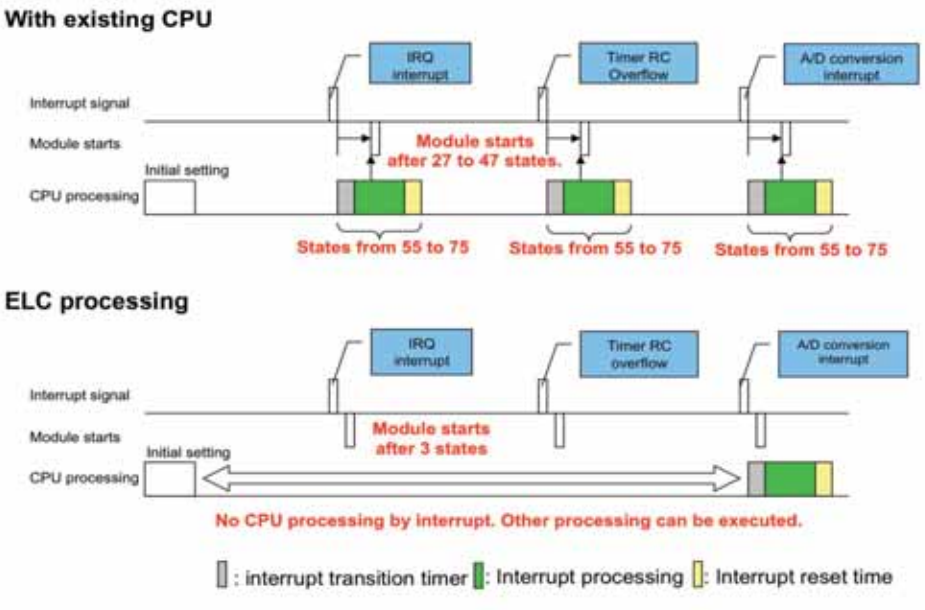

DESIGN OF A COST-EFFECTIVE BARCODE SCANNER USING A 32-BIT MIXED-SIGNAL MCU

PARKER DORRIS, SENIOR APPLICATIONS ENGINEER FOR MICROCONTROLLER PRODUCTS AT SILICON LABS, PRESENTS THE DESIGNING A COST-EFFECTIVE AND RELIABLE BARCODE SCANNER

The design challenge for any developer entering a mature market is to create a 'disruptive' product, i.e. something that will redefine the market, and barcode scanners are a prime example.

Barcode scanners are a ubiquitous technology that touches our lives on almost a daily basis. Despite the prevalence of barcode scanners in commercial and industrial settings, it's easy to take the complexity of their underlying designs for granted. They require multiple discrete integrated circuits (ICs) and an array of passive and active circuitry to provide the functionality and reliability that users have come to expect.

Many barcode scanners generally use an optical sensor, such as a charge-coupled device (CCD), which outputs an analog representation of what is visible to the sensor to an analog-to-digital converter (ADC) controlled by a microcontroller (MCU). The MCU interprets the ADC's output as a sequence of thick and thin black and white bars and processes this sequence further to derive a string of characters from the pattern. The character sequence may be encoded in any one of a number of well-defined barcode protocols, such as Code 39.

Among the additional features found in barcode scanners are regulators, USB connectivity, audio output – usually driven by a discrete digital-to-analog converter (DAC) or simple pulse-width modulation (PWM) – and glue logic for level-shifting between components running from 3V such as the MCU, and components running from 5V such as the CCD.

CCD Overview

The barcode scanner's CCD optically scans a narrow area and converts optical information into a sequence of analog pulses. Each pulse's voltage level represents the state of one pixel within the CCD's scan range. Many commonly used CCDs in handheld barcode scanners produce output at a resolution of a few thousand pixels.

A CCD relies on a controlling MCU to provide a clock source. Clock pulses fed into the CCD's shift register initiate a scan, and the CCD synchronously outputs a voltage proportional to the light. The voltage level describing the pixel state usually contains a direct current (dc) bias, and the voltage swing between optical black and optical white can vary depending on the CCD used, as well as other design factors.

Advantages of an Integrated Solution

In traditional barcode scanner designs, each system feature requires a collection of discrete components. So, the barcode sensing subsystem requires a CCD surrounded by supporting passive components, a discrete ADC and a controlling MCU. The audio subsystem often requires a discrete DAC, a filtering circuit and a small speaker. The USB subsystem requires an external crystal to meet timing requirements imposed by the USB protocol, and so on.

Developers who find ways to eliminate some of these components can minimize the total cost of the board, shrink its size, and reduce the system's average current draw. Combining more features into fewer components also streamlines the development process and speeds time-to-market. A highly integrated solution (see Figure 1) will improve system reliability, as a design with fewer components reduces the number of potential points of failure in the system.

On-Chip ADC Sampling

When choosing a 32-bit MCU with an on-chip ADC, developers must beware misleading specifications that can narrowly define ADC performance as an operating region in which the device happens to perform well. For instance, some MCU data

Figure 1: Barcode scanning system based on a highly integrated 32-bit MCU

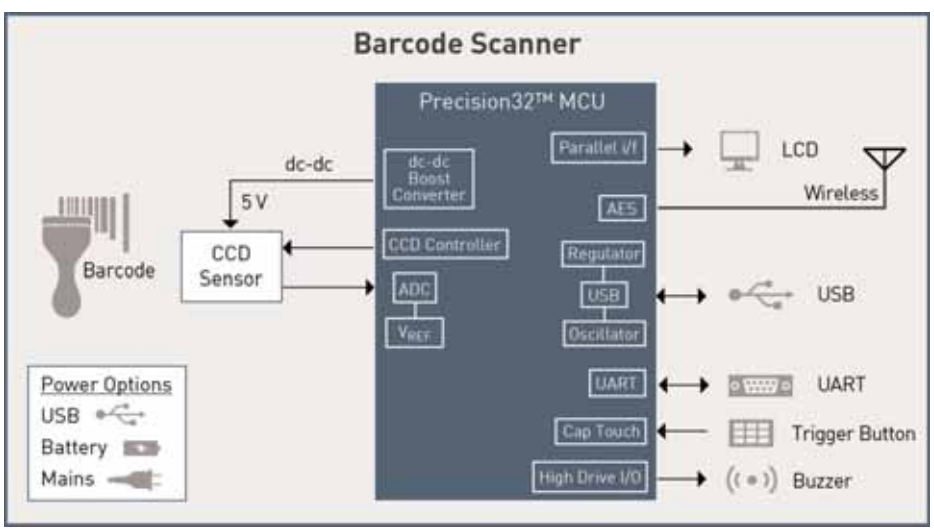

Figure 2: Example schematic showing ADC and port interface

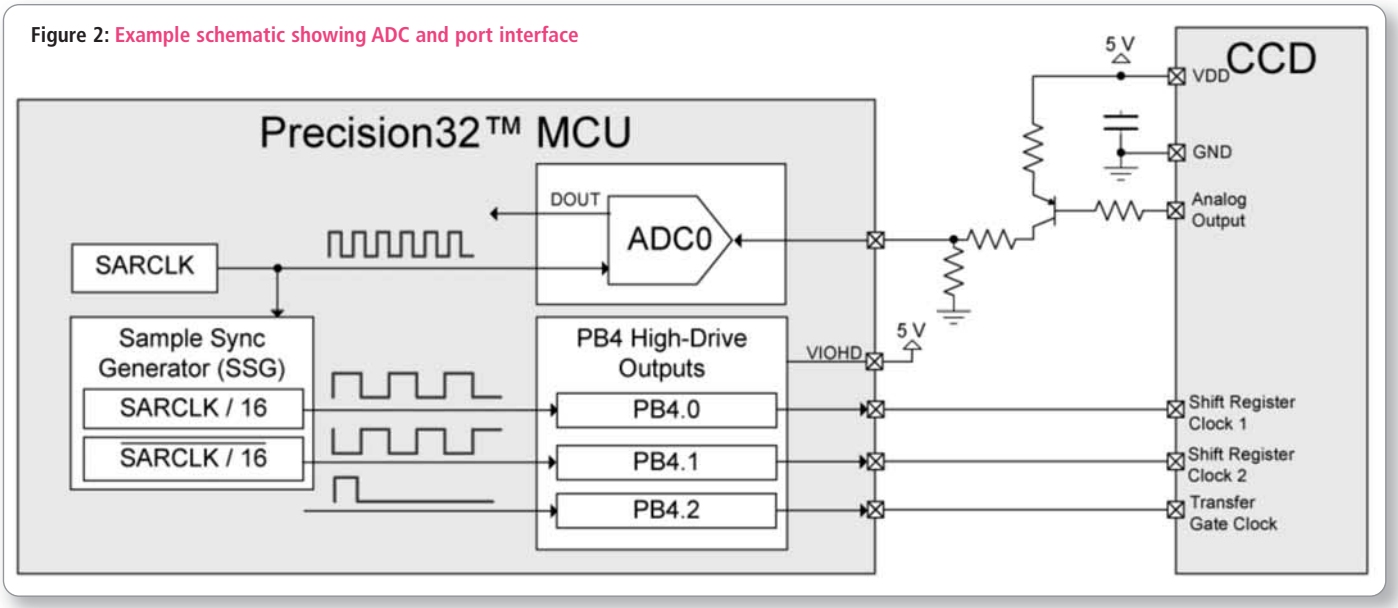

sheets will not provide a specification that spans the MCU's entire operating temperature and voltage range. Choosing an MCU whose performance is so narrowly defined leaves a design vulnerable to performance degradation when operating outside the confines of the ADC's specification. As an example, MCU's ADC might begin to see non-monotonic behavior when operating at high temperatures or when operating at the low end of the operating voltage range.

Therefore, to ensure that the barcode scanner products will function reliably across a wide spectrum of end-customer usage cases, a developer should select an MCU that has been exhaustively characterized, and with a performance specified to cover the broadest possible operating range.

Designers will also find useful an on-chip

ADC subsystem that can provide additional functionality which will make highly integrated mixed-signal MCUs even more compelling for barcode scanning applications. For example, an MCU's sample sync generator (SSG) can be configured to output a clock source derived from the ADC's sample clock. The ADC can be configured to autonomously execute multiple captures and store the resulting data in memory using a direct memory access (DMA) module. This combination of a fast, accurate ADC that can scan the entire range of the CCD's pixel resolution autonomously means that the developer's firmware only needs to initiate the first scan and then wait for the final DMA transfer to complete.

Figure 2 shows an example circuit connecting the ADC and port pins of a 32-bit MCU to a CCD.

Why Not Use PWMs?

An alternative solution to using an on-chip SSG block for synchronized clocking and sampling involves using one or more PWM channels, which are readily available on most 32-bit MCUs. While a PWM-controlled output offers the essential clocking features necessary for CCD operation, the SSG block provides some significant advantages. When a CCD receives a clock pulse, the device initiates the output of a voltage level to be captured by an ADC. However, the output of the anticipated voltage never appears immediately after the clock edge. Instead, most CCDs begin by outputting a signal representing optical black before the signal settles to the voltage describing the pixel state. The ADC's start-of-conversion process must be delayed until the CCD's analog output has settled.

Silicon Labs's Precision32 MCUs, for example, enable hardware configuration to control that timing offset before the conversion begins. When using a PWM, firmware is responsible for delaying the start of conversion, potentially by using a timer module or another PCA channel. Timing jitter introduced by firmware intervention in such a system can result in sample inaccuracies and lead to incorrect or unreliable barcode decoding.

Many barcode scanning designs require the clocking signal pins to hold state under certain circumstances. In most systems with PWM output, the holding state requires removal of the PWM from the pin followed by configuration of the pin for general-purpose use. During the brief period of time that PWM control is swapped out for general-purpose control, the pin state can be indeterminate, resulting in the output of runt pulses or other unintended signals that can cause the MCU to lose synchronization with the CCD.

Managing Multiple Supplies

A system in which devices are powered by both 5V logic and 3V logic traditionally requires that level shifting or some other type of glue logic be placed on interfacing signal traces. Additionally, these systems can require separate, discrete voltage regulators to supply 3V or 5V where appropriate. Choosing an MCU with an integrated voltage regulator and adjustable port drive levels eliminates the need for this external circuitry. In some MCUs, the high-drive port pins used by the SSG to drive the CCD's clock register can be configured to use the 5V supply as their logic high supply. The high drive port pins may be configured for current-limiting to directly drive an array of LEDs. An on-chip boost converter can be used to generate a 5V supply in systems in which a USB-based 5V supply is not available, such as battery-powered handheld barcode scanners.

Audible Feedback

Many barcode scanners provide some form of audible feedback. For instance, the device might beep when the user presses a button to initiate a scan, or the device may audibly indicate that a scan has completed successfully.

Traditionally, implementing this feature has required the MCU to connect to a DAC or use a simple PWM to activate a buzzer. Integrating a flexible, high-precision PWM capable of interfacing with Class D amplifier circuits helps, as they are more power- and cost-efficient and have smaller PCB footprint.

Great Integration

Achieving maximum system functionality with the fewest possible components can yield a number of design advantages for barcode scanning products. For example, integrating the sensing, USB, power, and audio output features of a traditional barcode scanner design can eliminate multiple discrete components, thus reducing BOM cost and board size.

Integration also provides additional benefits, such as more precise and controlled power management, simplified code development and potential improvements in system reliability. ●

SILICON LABS'S PRECISION32 FAMILY OF 32-BIT MIXED-SIGNAL MCUs

THE HIGHLY INTEGRATED 32-BIT MIXED-SIGNAL MCUS FROM SILICON LABS, CALLED PRECISION32, HANDLE MANY OF THE SYSTEM-LEVEL TASKS TRADITIONALLY DISTRIBUTED AMONG MULTIPLE DISCRETE COMPONENTS.

The family provides a highly integrated, peripheral-rich embedded control solution that enables developers to eliminate many discrete components in a barcode scanner design while enhancing system reliability and performance. Such a high level of integration, enabled by a peripheral-rich MCU solution, results in differentiated barcode scanner designs.

The Precision32 family's two on-chip 12-bit ADCs provide a maximum throughput of 2MSPS with linearity of less than a single code across the voltage and temperature operating ranges. The on-chip ADC enables the developer to take advantage of the MCU's ultra-low-power features by placing the MCU in low-power states during autonomous scanning and enabling code to easily toggle power to the ADC between scans.

The MCU's sample sync generator (SSG) can be configured to output a clock source derived from the ADC's sample clock. The SSG's clock output toggles at a rate of 1/16 of the ADC's clock, enabling the system to output an external clock that is synchronized to the ADC's start of conversion. This output clock, when routed to the CCD's shift register input, along with another MCU-provided start-of-conversion signal, effectively synchronizes the ADC's scan rate to the output rate of the CCD's pixel information.

Most barcode scanners transmit decoded barcode information to a host system across a USB port. Unfortunately, demanding USB timing requirements force many designs to use an external high-precision crystal as a timebase, which adds to the cost, increases board size and causes reliability issues. Precision32 family's USB peripheral eliminates this need by providing an on-chip clock recovery mechanism that generates an internal USB clock compliant with full-speed USB timing specifications without an external crystal. Systems that communicate with a non-USB enabled host can pass information across one of the many other serial interfaces provided by the Precision32 MCU, including UART, USART, SPI and I2C.

PROGRAMMABLE VOLTAGES
FOR EVALUATING MULTIPLE POWER SUPPLY SYSTEMS

STEVEN XIE, ADC APPLICATIONS ENGINEER WITH CHINA DESIGN CENTER IN ADI BEIJING, **KARL WEI**, PRODUCT ENGINEER WITH THE MICROCONVERTER GROUP, AND **CLAIRE CROKE**, APPLICATIONS SUPPORT FOR PRECISION ADCS AT ANALOG DEVICES PROPOSE AN EASY AND COST-EFFECTIVE METHOD FOR DETERMINING THE BEHAVIOR OF A SYSTEM WHEN SUBJECTED TO SUPPLY TRANSIENTS, INTERRUPTIONS OR SEQUENCE VARIATIONS

High-voltage switches, bipolar ADCs and other devices with multiple power supplies often require that supply voltages be applied or removed in a particular sequence. An example of a device using multiple supplies is the AD7656-1 (Figure 1), a 16-bit, 250 kSPS, 6-channel, simultaneous-sampling, bipolar-input ADC. The ADuC7026 precision analog microcontroller's four 12-bit DACs provide the DUT's programmable supply voltages. Using the AD7656-1 evaluation board and the ADuC7026 evaluation board, prototyping can be performed with a minimum of hardware and software development.

Table 1 shows the typical voltage and maximum current for each of the ADC's power supplies. The programmable sequence-controllable voltage waveforms generated by the four DACs on the ADuC7026 are scaled by the ultralow noise-and-distortion AD797 op-amps on the AD7656-1 evaluation board to provide the specified supply voltages and currents. The microcontroller's speed and programmability facilitate control of voltage level, period, pulse width and ramp time of the power supply voltages.

For example, using external power supplies, the AD797 amplifiers on the AD7656-1 evaluation board, configured for a gain of 5, can generate a voltage range of 0V to 12.5V to drive the ADC's VDD supply rail. The high output drive capability of the AD797 allows up to 50mA to be provided to each supply rail. Figure 1 shows the connections to the ADC.

The ADuC7026 DAC data register can be updated at 7MHz with a 41.78MHz core clock, which maximizes the voltage update rate. The following sections describe the development process and provide measurement results obtained using the evaluation boards.

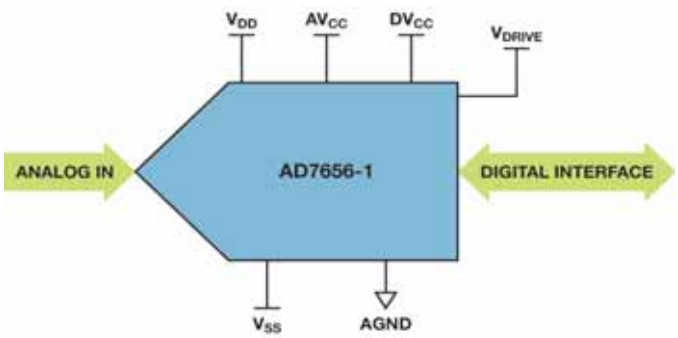

Figure 1: AD7656-1 connection diagram

Supply	AVCC, DVCC	VDRIVE	VDD	VSS
Voltage (V)	5	3.3	10	–10
Current (mA)	30	10	0.25	0.25

Table 1: AD7656-1 typical supply voltages and maximum supply currents

DAC Channel	DAC0	DAC1	DAC2	DAC3
Output Range	0V to 1.250V	0V to 0.825V	0V to 2.500V	0V to 2.500V
AD797 Gain	4	4	5	–5
AD797 Output Swing	0V to 5.00V	0V to 3.30V	5.00V to 12.50V	–12.50V to –5.00V
Nominal Voltage	5.00V	3.30V	10.00V	–10.00V
AD7656-1 Power Supply	AVCC, DVCC	VDRIVE	VDD	VSS

Table 2: Power supplies for AD7656-1

Hardware Development and Setup

The hardware connection and test setup are shown in Figure 2. Four DAC output pins and AGND on the ADuC7026 evaluation board are connected separately to the four AD797 inputs and AGND on the AD7656-1 evaluation board. An Agilent E3631A external power module provides ±15V for the AD797. A computer connected via USB to the ADuC7026 evaluation board provides the 5V power supply and serial communications.

Schematic Design

The only hardware changes required on the AD7656-1 evaluation board relate to the AD797. R1 and R2 can be selected for

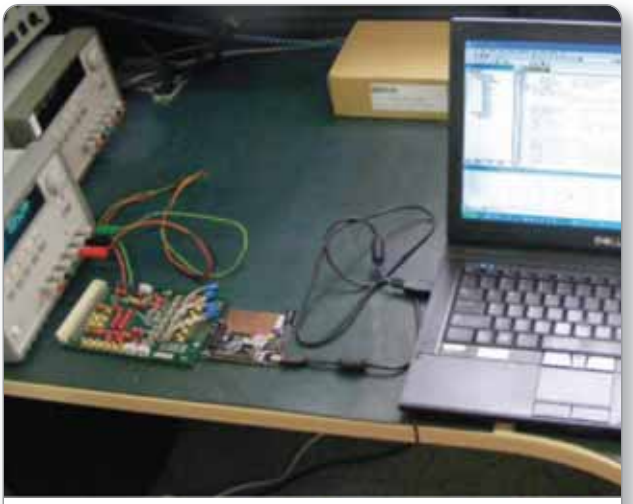

Figure 2: Hardware connections and test bench

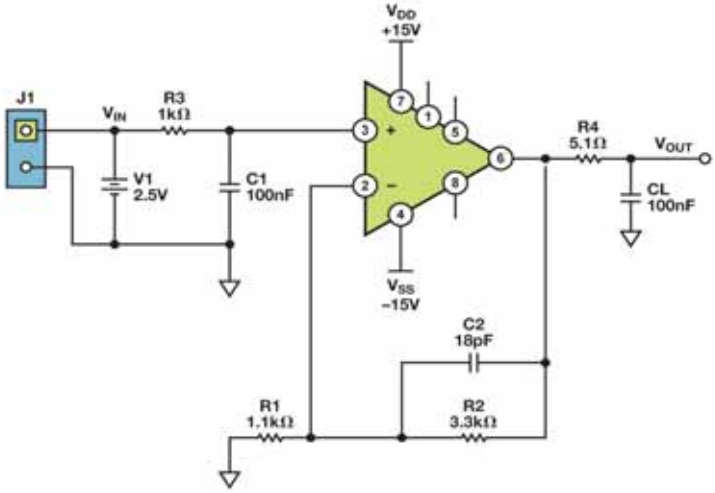

Figure 3: AD797 schematic design with gain of 4

different gain and bandwidth requirements. Figure 3 shows the AD797 set for gain of 4 to provide a 0V to 10V output from the 0V to 2.5V output of the ADuC7026 DAC. R3 and C1 form a low-pass filter to reduce high frequency noise. C_L is used as a load capacitor on the power rail.

Figure 4 shows the frequency response of the AD797 with gain of 4, from an NI Multisim simulation. The 1.0MHz bandwidth and 73° phase margin provide fast transient response and stable operation.

AD797 Design Notes

The AD797 ultralow-distortion, ultralow-noise op-amp features 80µV maximum offset voltage, excellent DC precision, 800ns settling time to 16 bits, 50mA output current and ±13V output swing with ±15V power supplies, making it well-suited to driving power-supply rails.

It is not internally compensated for substantial capacitive loads though, so external compensation techniques must be used to optimize this application. Figure 5 shows oscillation on the AD797 output caused by driving capacitive loads.

For stable drive with capacitive loads on the power rail, Resistor R4 is placed between the output and the load. This resistor isolates the op-amp output and feedback network from the capacitive load and introduces a zero in the transfer function of the feedback network, reducing the phase shift at higher frequencies. The feedback capacitor, C2, compensates for the capacitive loading, including C1, at the input of the op-amp.

Applying the DACs

The ADuC7026 precision analog microcontroller features four 12-bit voltage-output DACs with rail-to-rail output buffers, three selectable ranges and 10µs settling time.

Each DAC has three selectable ranges: 0V to VREF (internal band gap 2.5V reference), 0V to DACREF (0V to AVDD), and 0V to AVDD. The range is set using the control register DACxCON. The DAC accepts an external reference with a range of 0V to AVDD. When using the internal reference, a 0.47µF capacitor must be

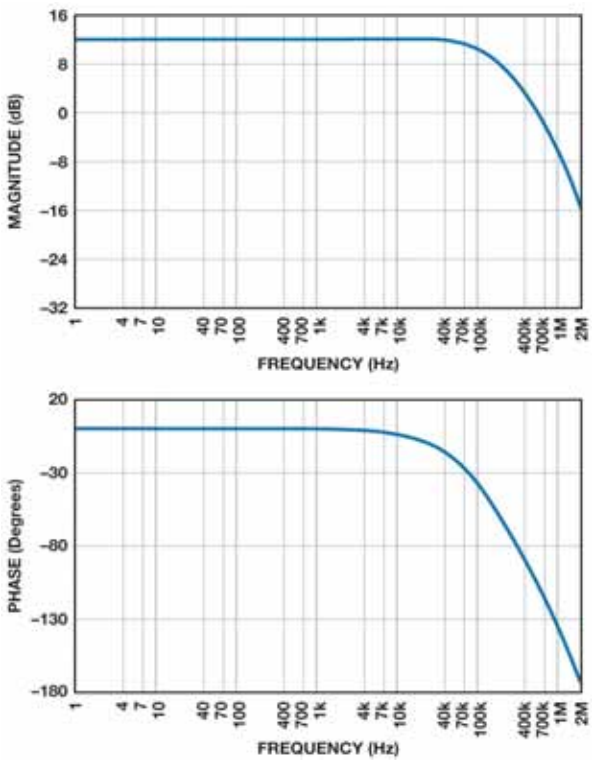

Figure 4: Frequency response of the AD797 with gain of 4

Capacitive Load	10 nF (V/µs)	0.1 µF (V/µs)	1 µF (V/µs)	10 µF (V/µs)
Rising Edge	6.90	0.97	0.07	0.01
Falling Edge	5.71	0.93	0.06	0.01

Table 3: Ramp time with capacitive load

Power Supply	AVCC, DVCC (5.00V)		VDRIVE (3.30V)	VDD (10.00V)	VSS (–10.00V)
200MHz (mV)	20.8	28.0	25.6	30.4	
20MHz (mV)	12.8	24.8	15.2	18.4	

Table 4: The ripple of each power supply

connected from the VREF pin to AGND to ensure stability.

Each of the four DACs is independently configurable through control register DACxCON and data register DACxDAT. Once the DAC is configured through the DACxCON register, data can be written to DACxDAT for the required output voltage level.

The four DAC outputs are easy to control using C or assembly language. This C code example shows how to choose the internal 2.5V reference and set the DAC0 output to 2.5V, as follows:

```
//connect internal 2.5V reference to
VREF pin
REFCON = 0x01;
//enable DAC0 operation
DAC0CON = 0x12;
//update DAC0DAT register with
data 0xFFF
DAC0DAT = 0x0FFF0000;
Using assembly language,
DAC0CON[5] is cleared to update
DAC0 using core clock (41.78MHz)
for fast update rate;
DAC0CON[1:0] is set to '10' to use 0V
to VREF (2.5V) output range
'DAC0DAT = 0x0FFF0000' can be
compiled to assembly code with two
instructions:
MOV R0, #0x0FFF0000
STR R0, [R1, #0x0604]
```

These two instructions take a total of six clock cycles to execute, corresponding to a 7MHz update rate with a 41.78MHz core-clock frequency. Thus, the time delay between voltage rails can be as accurate as 144ns.

Measurement Results

The four DACs in the ADuC7026 provide four power supplies to the AD7656-1 to test its behavior with power-supply transients or sequence variations. Table 2 shows the ADC's power supplies and voltage levels.

The waveforms from the four DAC outputs, as described in Table 2, were captured using a scope and are shown in

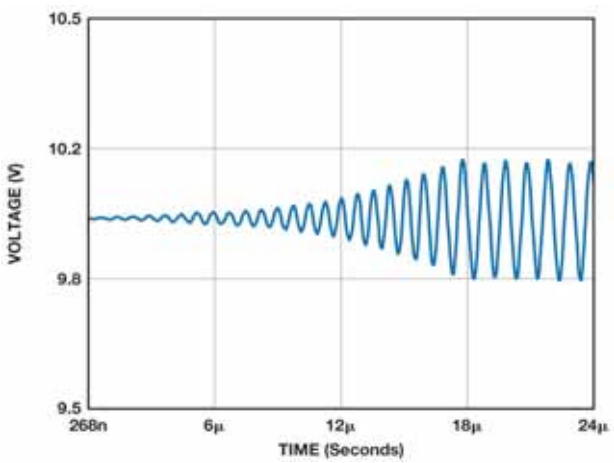

Figure 5: Oscillations without compensation

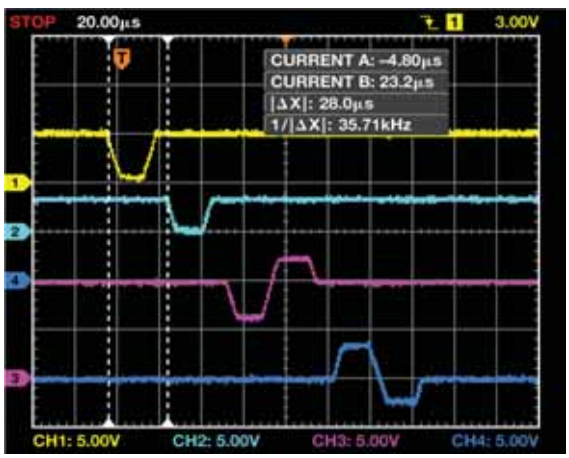

Figure 6: Four-channel voltage waveform

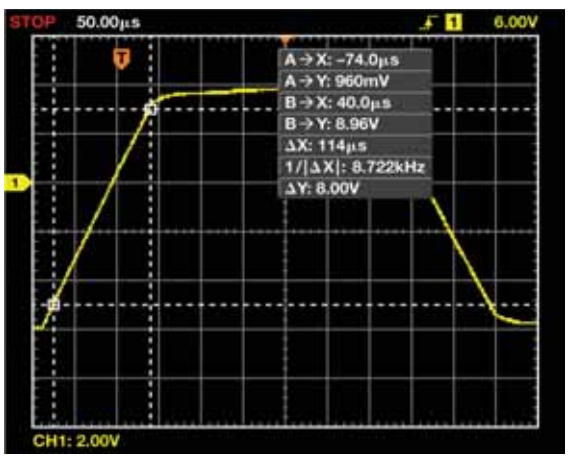

Figure 7: Rise time with 1µF capacitive load

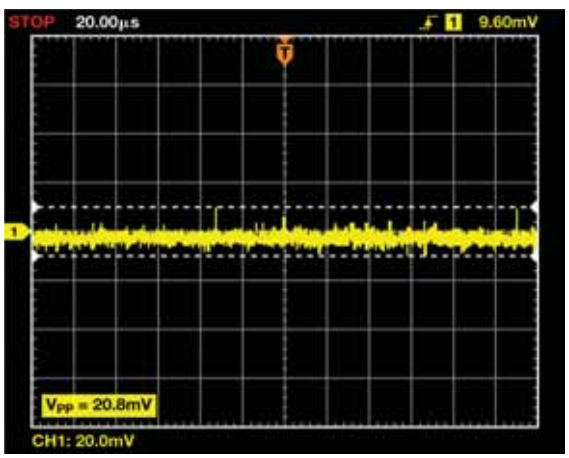

Figure 8: Ripple of 5V supply on AVCC and DVCC

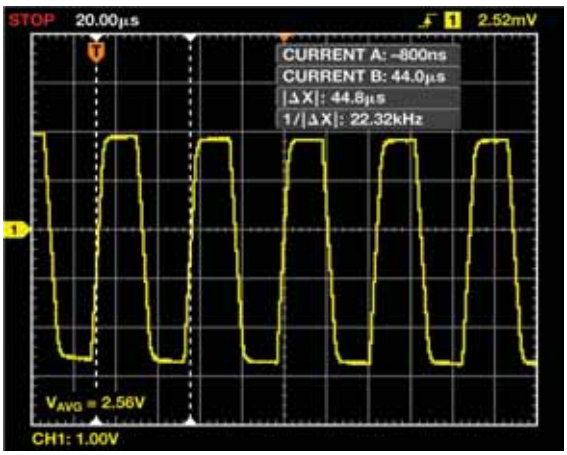

Figure 9: 22.32kHz square waveform

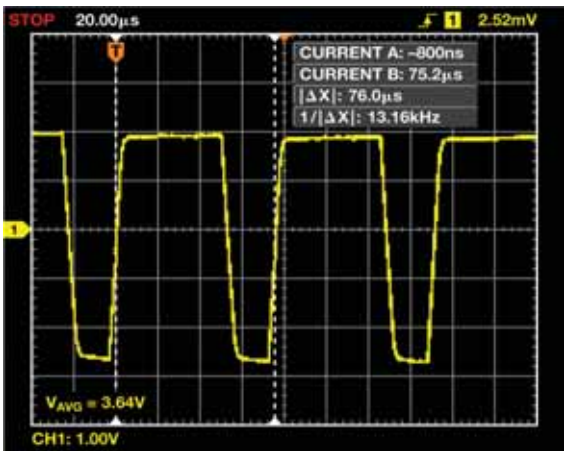

Figure 10: 13.16kHz pulse waveform

Figure 6. The voltage level, period, pulse width and ramp time of each channel are programmable and easy to control.

To achieve an accurate voltage level for each power supply, an adjustable resistor can be used for R1 in Figure 3. The voltage level was calibrated by adjusting R1 with an Agilent 34401A digital multimeter.

Rising and falling ramp times are measured to determine the maximum frequency of the voltage waveforms. The ramp time is related to the value of Resistor R4 and the capacitive load, C_L. For slower ramp times, larger resistor and capacitor values can be used for R4 and C_L. The rising and falling ramp time of AVCC and DVCC were tested with different load capacitors, with the results shown in Table 3. The rising waveform with a 1µF capacitor is shown in Figure 7. The ramp time is measured between 10% and 90% of 10V.

Power-Supply Ripple

The DC precision of the AD797 makes it easy to provide accurate nominal voltage levels for the AD7656-1 by adjusting the feedback resistor, R1. The peak-to-peak ripple of the power supplies was measured at nominal voltage levels with 200MHz and 20MHz bandwidths, a 0.1µF capacitive load and a DS1204B scope. Table 4 shows that the ripple is less than 1% of the nominal voltage, so the four supplies are qualified.

Generating Waveforms

With simple modifications to the ADuC7026 source code, many different sequences of voltage waveforms can be generated for a variety of different applications that require evaluation of device operation under different supply conditions. Typical waveforms that can be generated are shown in Figure 9 and Figure 10.

The LabVIEW GUI shown in Figure 11 can be used to generate the power supply waveforms. The voltage level, ramp time, period and sequence delay time of the four channels are easy to configure. The serial port is used for communication between GUI and the ADuC7026.

Evaluate the Effects of Supply Sequencing

An easy, cost-effective way to evaluate the effects of supply sequencing was developed and verified using the AD7656-1 and ADuC7026 evaluation boards. The ADuC7026 evaluation board generates a controllable programmable sequence for four voltage supplies to evaluate the operation of the ADC under different supply sequence/ramp conditions. The 3-phase, 16-bit PWM generator in the microcontroller can provide a total of seven voltage channels.

With a standard ±15V DC power module, this portable power supply evaluation systems allows designers to evaluate ADCs, especially for those with larger number of supplies. ●

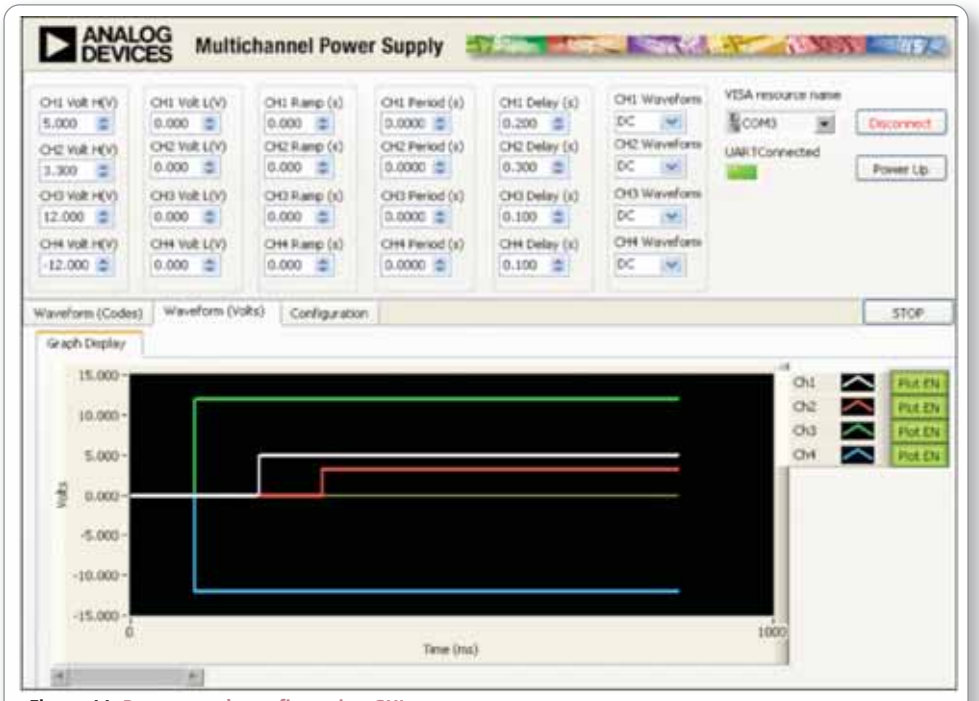

Figure 11: Power supply configuration GUI

ACCELERATED PROCESSING UNITS: BIG PERFORMANCE FOR SFF APPLICATIONS

CAMERON SWEN, DIVISIONAL MARKETING MANAGER AT THE AMD EMBEDDED SOLUTIONS DIVISION, EXPLAINS HOW THEIR EMBEDDED R-SERIES APUS MAY FILL THE GAP BETWEEN HIGH PERFORMANCE AND LOW POWER FOR EMBEDDED SMALL FORM-FACTOR (SFF) STANDARDS

As embedded system design marches steadily forward with ever shrinking form-factor standards, a gap has formed between performance and low power. New and emerging graphics-intensive and visually immersive applications need higher performance yet low power is needed to put these applications into small and highly mobile systems. Innovative processor architectures, such as the new Accelerated Processing Units (APUs), are able to provide the required performance/watt levels combined with the integration and packaging efficiency that really make it feasible to design high-end processing and discrete level graphics into mobile and handheld systems.

Over the past decade, advances in semiconductor technology have continued to follow Moore's law by roughly doubling the number of transistors available in a given area of silicon every two years. With such increasing transistor budgets, architects of traditional x86 CPUs have focused on techniques such as increasing clock rates, expanding the size and number of on-chip caches and adding additional processor cores. As such, performance gains have been tremendous, enabling PCs to become much more efficient at multitasking.

However, as fast as these modern PC processors are, they alone still cannot deliver the image, video and digital signal processing horsepower that many of today's emerging interactive multimedia embedded applications demand, and at the low power required by small form factors. Unlike traditional PC applications primarily built on scalar data structures and serial algorithms, emerging embedded applications, such as medical imaging and intelligent cameras, require processors that can handle vast amounts of data consisting of hundreds, if not thousands, of individual threads that must be manipulated and processed in parallel.

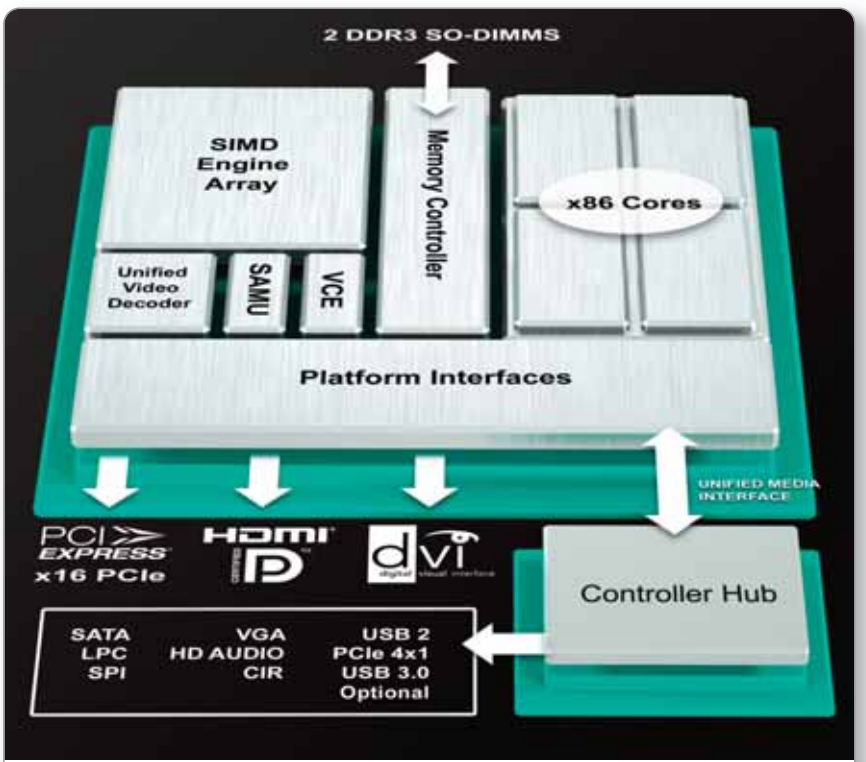

Figure 1: The architecture of the AMD Embedded R-Series platform integrates x86 cores, GPU vector (SIMD) engines and I/Os

> Accelerated Processing Units (APU) combine both x86 processing cores and discrete-level graphics processing units on a single die

Heterogeneous Approach

Graphics processing units (GPUs) were originally intended to enhance and accelerate the rendering of 3D images, but have evolved into powerful, programmable vector processors that can accelerate a wide variety of data-intensive algorithms and applications (commonly referred to as "stream processing").

THE EVOLUTION OF THE SMALL FORM FACTORS

SMALL FORM FACTOR STANDARDS SUCH AS PC/104 AND EBX WERE DESIGNED SPECIFICALLY FOR EMBEDDED APPLICATIONS AND HAVE BEEN AROUND FOR A LONG TIME.

But developers trying to take advantage of the standardization and economies of scale developed around standardized PC motherboards have worked to bridge the gap between standardized PC motherboards and small form factor embedded motherboards.

Comparing boards from just 10 years ago to today's boards, the smallest, high-volume motherboard for embedded systems in 2001 was a derivative of small PC motherboards and was known as Mini ITX, at 170 x 170mm. Mini ITX was then replaced by Nano ITX as the smallest standard embedded motherboard size and by 2007 it had evolved into Pico ITX, which at only 100 x 72mm, represents a size reduction of 75% from Mini ITX. New, smaller platforms continue to be developed today like Mobile ITX, which, at 60 x 60mm, is driving a new generation of handheld devices.

Computer-on-module (COM) solutions, consisting of individual CPU and I/O-based modules, have typically been smaller than the standardized embedded motherboards due to their modular/stacking nature and their size, fitting onto small boards to implement I/O.

At 90 x 96mm, PC/104 has been an industry standard size for embedded system modules since 1992. Both PC/104 and COM standards have varied in size and connections to accommodate various CPU architectures. Figure 2 below shows a comparison of some of the common small form-factor embedded boards.

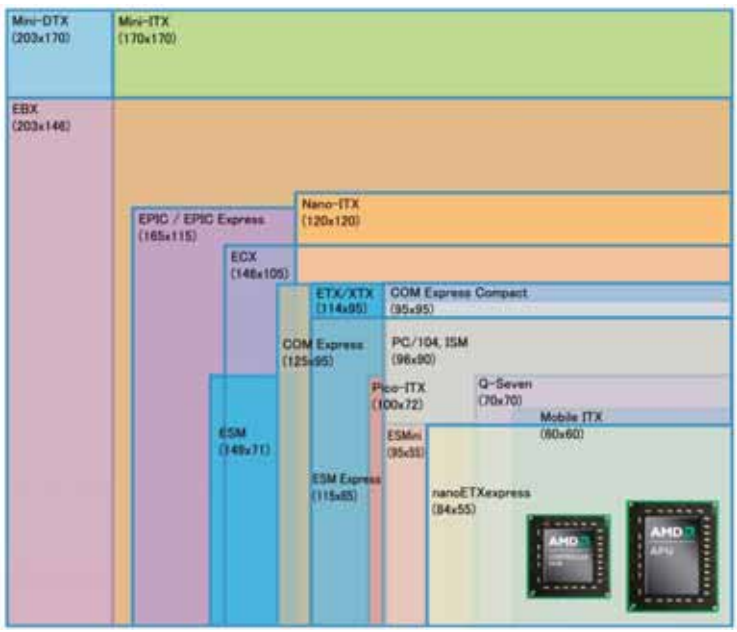

Figure 2: Size comparisons of the latest and most popular small form-factor boards and modules for embedded systems

GPUs implemented on PC add-in cards can pack Teraflops of floating-point compute horsepower onto a single PCI Express graphics card. With each generation, new features are added and limitations removed: from vertex processing stream operations to flexible branching and array manipulation and, even, to append buffers and atomic operations. Today's hundreds of processing cores in modern GPUs have deep pipelines and are nearly identical to each other, making them highly scalable and allowing GPGPU (General-Purpose GPU) computations of highly parallel workloads. As opposed to the conventional sequential-processing CPUs, which have been enhanced with only modest parallelism in the form of multi-threading and multiple cores, modern GPUs are optimized for massive parallel computing – for graphics or otherwise.

Smaller die geometries and new silicon-design innovations have enabled processor vendors to create the first family of single die CPU+GPU solutions. With hundreds of computing cores, these heterogeneous multi-core processors, or APUs, can help reduce the size and power of embedded systems dramatically, whilst increasing performance.

Software development tools which maximize the benefits of the APU architecture have also come a long way. Limited, proprietary tools have given way to advanced, open and portable standards. As such developers have access to a broad range of APIs and other software development tools, including DirectCompute (as part of DirectX 11) and OpenCL. Processor vendors such as AMD also provide suitable Software Development Kits (SDKs) for the task of Accelerated Parallel Processing (APP). These tools are gaining momentum, enabling developers to create standards-based applications that leverage the combined power of CPU and GPU cores, and that can run on a wide variety of hardware platforms.

The APU Architecture in Detail

APU stands for Accelerated Processing Unit, which combine both x86 processing cores and discrete-level graphics processing units on a single die. With an APU, the GPU becomes fully programmable, offering acceleration to the

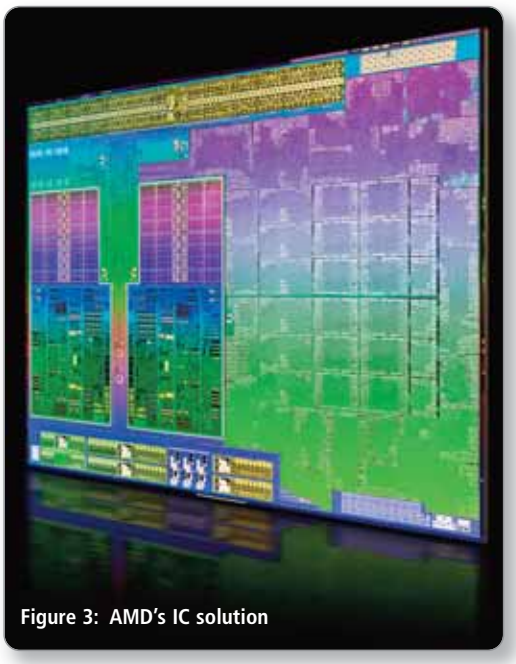

Figure 3: AMD's IC solution

processor for compute-intensive tasks.

Until very recently, the limited number of transistors meant having a two-chip solution. This in turn created performance constraints, as the external busses add latency when accessing memory. When integrated onto the same die, the APU's x86 CPU cores and SIMD GPU engines share a common and much faster speed path to system memory that helps to overcome these constraints. Furthermore, APU implementations divide this shared memory into multiple regions: those managed by the operating system running on the x86 cores and others managed by software running on the GPU engines.

The APU architecture provides high-speed block transfer engines that move data directly between the x86 and GPU memory partitions, without needing the additional bus transactions that occur when the frame buffer and system memory are separate. By structuring the code properly, it is also possible to overlap – or interleave – memory transactions between the CPU and GPU memory partitions to gain even higher performance.

Scalar processors operate on arrays of data one element at a time. Vector processors, like those used in advanced GPUs, have dozens or even hundreds of calculating units that operate simultaneously. The x86 technology (which efficiently cuts through scalar and vector workloads), when enhanced with GPU technology improves the total system performance, and compared to competing platforms it offers either greater performance within the same power envelope or much lower power consumption within the same performance envelope. In fact, the latest APUs can now deliver higher graphics performance than many of today's standard desktop and notebook PCs at lower power.

APU Model	Core Clock Speed PR/boost	Max TDP	Average Power 3DMark 06 (APU/Hub)	Average Power Video (APU/Hub)
R-464L	2.3/3.2 GHz	35w	13.105w / 0.694w	5.240w / 0.762w

Table 1: AMD Embedded R-Series platform TDP vs average power

AMD Embedded R-Series APUs

A new processor platform that is suited to bring the benefits of APUs to the mid-range and high-performance SFF designs is the AMD Embedded R-Series APU. With just over 1400mm² of combined chip real estate and under 2400 pins, the AMD Embedded R-Series platform is highly integrated, with a reduced pin-count, low power consumption and rich I/O support. Using this solution will eliminate many external signals and buses found in traditional solutions, making it much easier to route within a small footprint and without the costly board layers.

The architecture features up to 384 parallel compute units resulting in up to 563GFLOPS of single precision performance, whereas the CPU integrates up to four x86 cores for sophisticated multi-purpose x86 performance.

Depending on the task, the application code can be executed on the most power efficient core in the system, resulting in overall low-power consumption. This is the strength of the heterogeneous multi-core architecture of an APU. So, for example, on the AMD R-464L APU for graphics or parallel processing applications, the power allocation can be shifted to the GPU for up to a 38% boost in graphics speed, while staying within the defined power limits. For CPU-intensive applications, the power allocation can be shifted to the CPU for up to a 39% boost in CPU speed. Therefore, even in demanding applications the overall power consumption can stay well below the specified TDP (thermal design power) of the processor. For example, the top-of-the-line AMD R-464L APU, rated at 35W TDP, consumes just over 13W while running 3DMark06. This represents significantly lower power consumption than what TDP alone would indicate.

For applications requiring high performance multimedia capabilities, the new AMD Embedded R-Series APU also offers a variety of new power-saving features. These include the Video Compression Engine that provides a dedicated hardware video encoder, for quick and efficient video encoding. The Secure Asset Management unit enables GPU-assisted encryption and decryption of content. Enhancements to the Unified Video Decoder extend the capabilities of the AMD Embedded R-Series platform to include dual, high-definition video decode. All of these updates help to minimize CPU utilization when dealing with video and reduce overall APU power consumption. ●

FOR MORE INFORMATION

ABOUT AMD'S APU ARCHITECTURE AND THE AMD EMBEDDED R-SERIES PLATFORM, INCLUDING PRODUCT SPECIFICATIONS, REFERENCE DESIGNS, PERFORMANCE COMPARISONS AND APPLICATION INFORMATION, go to AMD's website at http://www.amd.com/R-Series.

To learn more about applications and demos created by customers and technology partners using AMD Accelerated Parallel Processing technology, including compilers, libraries and a multitude of multimedia applications, go to AMD's Accelerated Parallel Processing (APP) Developer Showcase at: http://developer.amd.com/sdks/AMDAPPSDK/samples/showcase/Pages/default.aspx

electronica

13 – 16 November 2012, Messe München

Opening hours: Tue – Thur: 9am – 6pm
Fri: 9am – 5pm

electronica is the world's leading trade fair for electronic components, systems and applications.

As in the past, the 2012 fair will feature well-known major manufacturers, as well as interesting specialists and newcomers, in a range of exhibits that cover technologies, products and solutions in the entire electronics industry.

The program consists of many related events, which include three international conferences, four forums, a CEO roundtable and more than 300 lectures.

About electronica

electronica has been held every two years in Munich since 1964 and presents innovations from the entire range of products and services in the electronics industry. electronica is the most important international meeting-point for the electronics industry and, in addition to the exhibition areas, it features a supporting program containing first-class conferences – electronica Automotive Conference and Wireless Congress – and discussions.

In 2010 electronica was attended by some 72,000 visitors and nearly 2,600 exhibitors. The total exhibition space was 142,500 square meters.

hybridica, a trade fair of innovative solutions made of plastics, metal, ceramics and composites, has been staged concurrently with electronica since 2008 and produces numerous synergy effects.

In addition, and for the first time ever, there will be an embedded platforms conference on 14 and 15 November. Together with the exhibition and the forum, it is part of the show's three-pronged approach to embedded electronics.

When developing a successful product, selecting the right embedded module – i.e. the processor platform, operating system, drivers and networking technology – is not the only thing that plays an important role. It is also essential that the individual components interact correctly and that they are adapted to suit the specific application.

"The conference is our way of addressing this challenge for embedded system engineers, who can come here to gather specific information about selecting a platform that is right for them," said Nicole Schmitt, electronica Exhibition Director. "Processor and controller manufacturers and their software partners can also attend seminars and lectures to find out what solutions they can use to support developers in the areas of component selection and system design. Sponsors that have already indicated interest include Infineon and TQ Components."

Embedded technology is very significant throughout the entire electronics industry. At electronica 2010, that fact was documented by 230 German and international companies in this sector and by the embedded forum. As of 2012, the embedded platforms conference not only gives exhibitors the opportunity to present their developments and solutions, but it allows system architects and developers to gather information about pioneering technologies and their potential applications.

More information about electronica 2012 is available online at www.electronica.de.

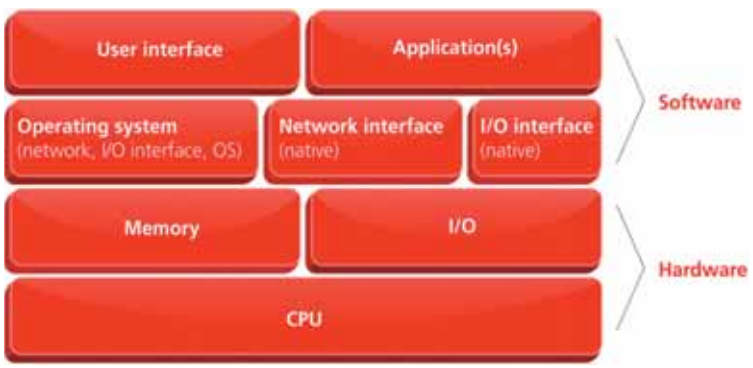

Figure 1: Modules of the electronica embedded platform conference

- Semiconductors
- Embedded systems
- Displays
- Micronano-systems
- Sensor technology
- Test and measurement
- Elektronic Design (ED/EDA)
- Passive components
- Electromechanics/System peripherals
- Power supplies
- PCBs, other circuit carriers and EMS
- Assemblies and subsystems
- Automotive
- Wireless
- Information gathering and services

Table 1: Exhibition areas

Pickering Electronics's Preview at electronica 2012

Pickering Electronics will highlight the following products and other products at electronica 2012 in Munich – the international summit for the electronics industry:

Pickering Series 117 SIL Reed Relays, Pickering Series 113 SIL Changeover Reed Relays, Series 104, Series 60/65 fit on a board area of 0.65 x 2.3 inches among others.

In addition, on the stand Pickering will be offering a free copy of its '*Reed RelayMate Book – Expertise at first hand*'.

Pickering Electronics's 117 Series Reed Relays are ideal for very high density applications such as ATE switching matrices or multiplexers. Its Series 113 changeover single-in-line reed relay is its smallest changeover. They're magnetically screened and require a board area of only 0.15 inches (3.8mm) by 0.5 inches (12.7mm).

www.pickeringrelay.com
Hall A1, Booth 530

Bulgin Push-Pull Coupling Buccaneer Connectors Launched

Bulgin, an Elektron Technology connectivity brand, has unveiled the latest addition to its rugged Buccaneer range of connectors, giving engineers and product design professionals even greater flexibility in systems design.

The all new Buccaneer 6000 Series of waterproof power, signal and data connectors features a unique and easy to use, patent-pending push-pull locking system that connects up to 10 times faster than a traditional screw thread mechanism. This latest addition to the popular Buccaneer range is designed to withstand the harshest environments and meets the IP66, IP68 and IP69K standards.

www.bulgin.co.uk
Hall B3, Booth 271

AMD Embedded R-Series Platform in COM Express Type 6 Module

congatec AG has set another milestone by offering support for the AMD Embedded R-Series platform with the new conga-TFS COM Express module based on Type 6 pin-out. This expands the COM Express standard with a new module that combines power efficient high-performance x86 processing, unprecedented integrated graphics performance and high performance parallel processing support.

The conga-TFS enables multimedia advancements and an exceptional visual experience across applications.

"We expect AMD Embedded R-Series APUs to be a great addition to our computer-on-module family which will be ideally suited for demanding graphics applications that require lower power consumption," explained Gerhard Edi, CEO at congatec AG.

www.congatec.com
Hall A6, Booth 306

Geotest Introduces Performance PXI DMM

Geotest has recently introduced the GX2065, a 6.5 digit PXI DMM for high performance measurement applications. The GX2065 is full-featured, multifunction DMM and digitizer, offering all of the capabilities associated with standard bench-top DMMs including DCV, ACV, 2 and 4-wire resistance measurements, and current measurements. Additionally, the GX2065 features a 3MS/s, 16-bit isolated input digitizer which allows users to acquire and analyze waveforms.

Featuring 6½ digit resolution, 0.005% basic DCV reading accuracy and up to 3,500 reading per second (rps), the GX2065 provides measurements that are accurate, fast and repeatable. All measurement functions including digitizing functions are isolated from the PXI bus, providing the ability to make true differential, floating measurements. An on-board controller performs all necessary DMM and digitizer calculations, minimizing PXI control bus overhead.

www.geotestinc.com
Hall A1, Booth 261 (Schneider and Koch Booth)

Ericsson Sets a Standard for Board Power in Advanced Bus Converters

Ericsson's Advanced Bus Converter (ABC) footprint is fast becoming a new standard in digitally-controlled power converters. Introduced in 2008, the Ericsson ABC footprint was designed to provide additional functionality to system architects developing advanced power architectures to reduce energy consumption and to increase flexibility in their operation and supply chain.

The Ericsson ABC footprint is used in more than 90% of the installed base of digitally controlled PMBus-based DC/DC products worldwide, making it the de facto industry standard for Advanced Bus Converters. Additionally, a growing number of companies are now developing Ericsson ABC footprint compatible products.

www.ericsson.com/powermodules
Hall A4, Booth 260

IQD Introduces Ground-Breaking Quartz Crystal Range

IQD has launched two new quartz crystal models housed in unique innovative packages that deliver significant space and cost saving, whilst still utilising tried and tested technology.

The new surface mount IQXC-74 measures just 7 x 4mm with a height of 2.3mm – this compares with the similar current industry standard HC49/4HSMX design which measures 11.4 x 4.9 x 4.3mm. Whilst IQD's latest ceramic package designs do offer smaller dimensions than this, they are significantly more expensive due to the artificially high cost of the ceramic packaging. The IQXC-74 offers an attractive compromise at significantly lower cost.

IQD is also launching a thru-hole version, the IQXC-75 which is even smaller at 6 x 4 x 1.8mm with similar benefits. Both models are designed to provide a low-cost, high reliability timing solution in a wide range of electronics applications.

www.iqdfrequencyproducts.com
Hall B5, Booth 314

A subscription to Electronics World costs just £51 and provides you with:

- ✔ 12 print and digital issues
- ✔ Weekly e-newsletters
- ✔ Regular topical supplements

Follow us on Twitter @electrowo

OR CALL OUR HOTLINE:

+44 (0)1635 879361

Plug-in Data Acquisition System and Signal Conditioning

IN THIS SERIES, **MAURIZIO DI PAOLO EMILIO**, TELECOMMUNICATIONS ENGINEER, SOFTWARE DEVELOPER AND DESIGNER OF ELECTRONIC SYSTEMS, PRESENTS A TUTORIAL ON DATA ACQUISITION SYSTEM DESIGN

Many scientific applications use plug-in boards to acquire data and transfer it directly to computer memory. Others use DAQ hardware remotely from the PC via a parallel or serial port, GPIB bus or network.

Typically, DAQ plug-in boards are general-purpose data acquisition instruments well suited to measuring voltage signals. Most PC-based DAQ systems (Figure 1) include some form of signal conditioning in addition to the plug-in DAQ board.

Applications requiring real-time processing of high frequency signals need a high-speed, 32-bit processor with a coprocessor, or a dedicated plug-in processor such as a digital signal processor (DSP). If the application only acquires a reading once or twice per second, however, a low-end PC can be sufficient.

Analog input (A/D) boards convert analog voltages (Figures 2 and 3) from external signal sources into digital format, which can be interpreted by the host computer. The functional diagram of a typical A/D board consists of the following main components:

- Input multiplexer;
- Input signal amplifier;
- Sample and hold circuit;
- A/D converter;
- Memory (DMA);
- Timing system and filtering;
- Bus interface;
- Digital signal processing;
- Microprocessor and Field-Programmable Gate Array (FPGA).

The data transfer capabilities of the board used can significantly affect the performance of the DAQ system.

All PCs are capable of programmed I/O and interrupt transfers. However, DMA transfers (which may not be available on some computers) increase the system throughput by using dedicated hardware to transfer data directly into the system memory. Using this method, the processor is not burdened with moving data and is free to engage in more complex processing tasks.

Without DMA, when the CPU is using programmed input/output, it is typically fully occupied for the entire duration of the read or write

Figure 1: Functional diagram of a basic data acquisition system

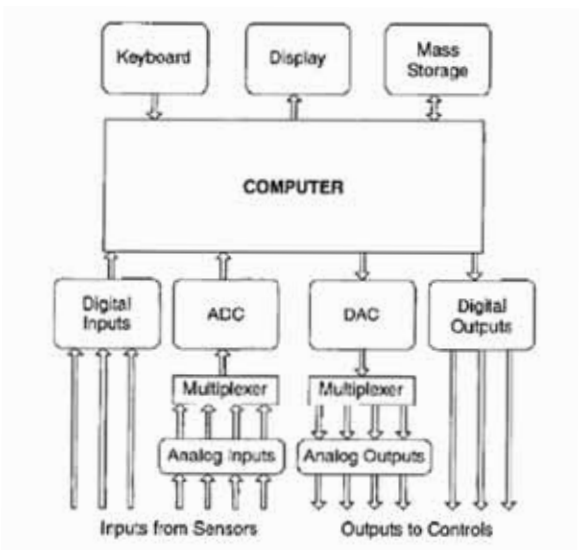

Figure 2: Functional diagram of a generic A/D board

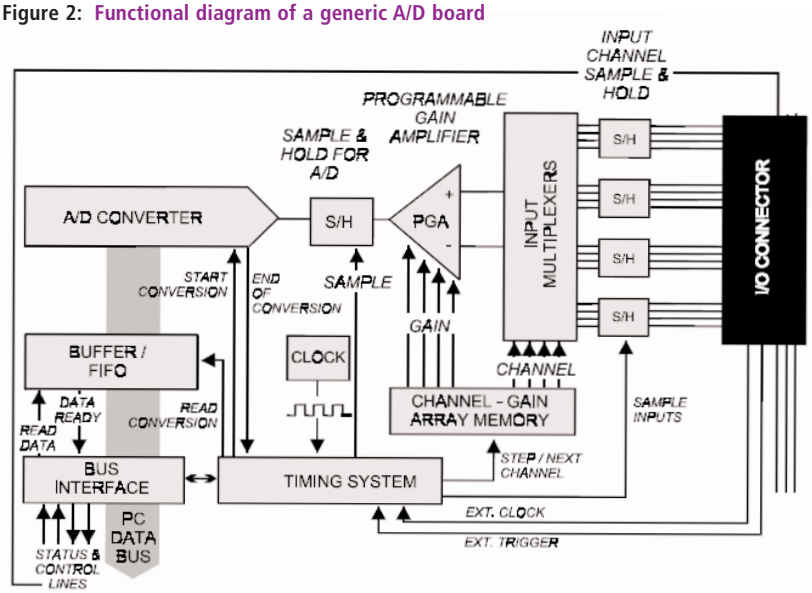

> *Signal conditioning is useful for all types of functions, including amplification, filtering and isolation*

operation, and is thus unavailable to perform other work. With DMA, the CPU initiates the transfer, does other operations while the transfer is in progress, and receives an interrupt from the DMA controller when the operation is done. This feature is useful any time the CPU cannot keep up with the rate of data transfer, or where the CPU needs to perform useful work while waiting for a relatively slow I/O data transfer. Many hardware systems use DMA, including disk drive controllers, graphics cards, network cards and sound cards.

In order to get information from a sensor into a computer, the signal from the sensor must first be sent to an interface device of some form and from there to the computer. However, in order to be useful to the interface device, the signal from the sensor must often undergo some form of conditioning. Almost all interface devices linking computers to sensors are designed to accept a voltage signal in the range of 0 to 5 volts and digitize it. The aim of the signal conditioning circuit is to take the sensor's output, whether voltage or resistance, and convert it to a 0 to 5 volt signal. This process generally involves a combination of one or more simple steps: converting a resistance to a voltage, dividing a voltage, amplifying a voltage and shifting a voltage.

Regardless of the types of sensors or transducers used, proper signal conditioning equipment can improve the quality and performance of the system. Signal conditioning functions are useful for all types of signals, including amplification, filtering and isolation.

Signal conditioning circuitry with amplification, which applies gain outside the PC chassis and near the signal source, can increase the measurement's resolution and reduce the effects of noise (Figure 4). An amplifier, whether located directly on the DAQ board or in the external signal conditioners, can apply gain to the small signal before the ADC converts the signal to a digital value. Boosting the input signal uses as much of the ADC input range as possible.

There are types of sensors that produce a voltage output that is symmetrical around 0V; in the form of -X to +X volts. As analog-to-digital converters generally require a positive voltage, it is necessary to shift these voltages upwards so that they operate from 0 to 2X volts.

The circuit for shifting a voltage upwards is more complex than the previous circuit. It uses a dual op-amp to add a fixed voltage to the input signal (Figure 5). Filters may be used to reject unwanted noise within a certain frequency range. For example, applying low-pass filters in the signal conditioning circuit can eliminate unwanted high-frequency components.

Improper grounding of the DAQ system is the most common cause of measurement problems and damaged DAQ boards. Isolated signal conditioners can prevent most of these problems by passing the signal from its source to the measurement device without a galvanic or physical connection. Isolation breaks ground loops, rejects high common-mode voltages and protects expensive DAQ instrumentation. ●

Figure 3: Functional diagram of a commercial A/D board

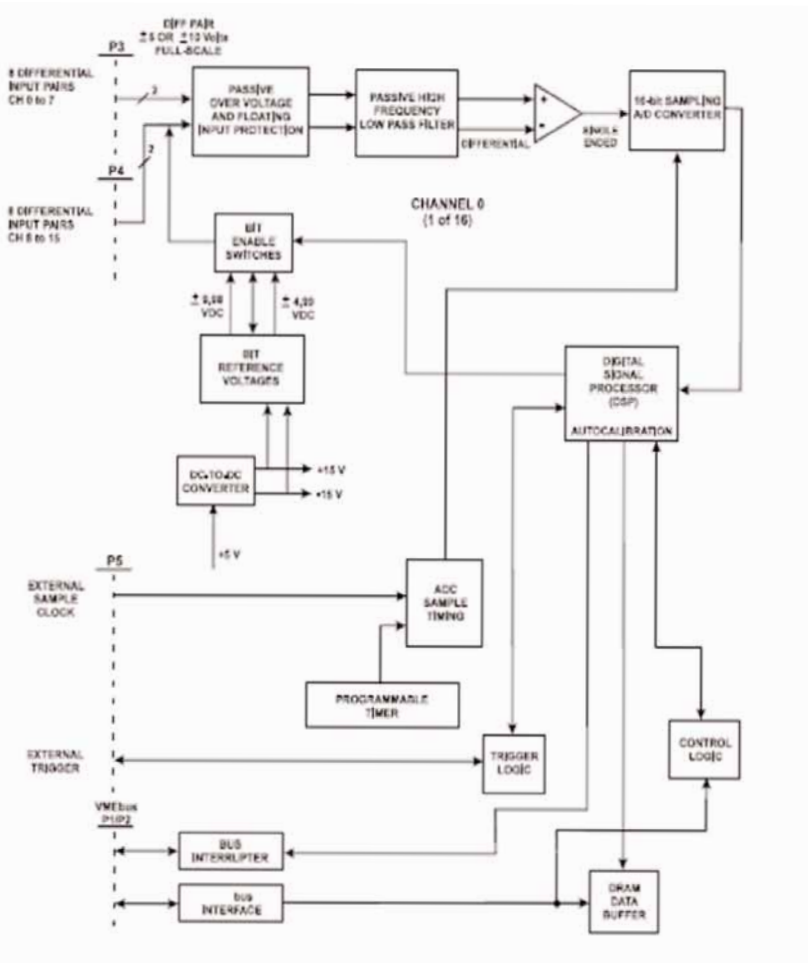

Figure 4: Signal conditioning – amplification

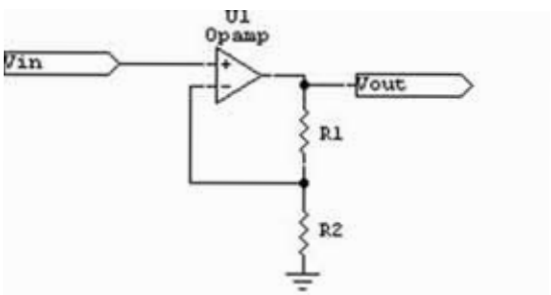

Figure 5: Signal conditioning – shifting voltage

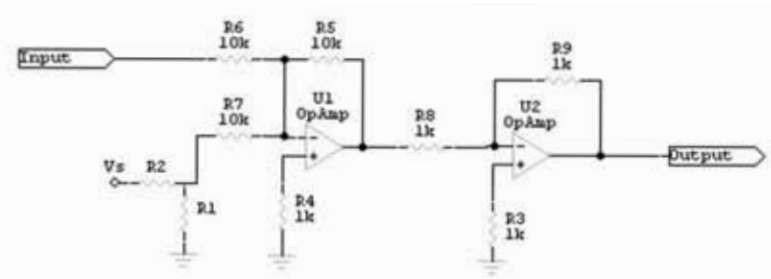

Agilent Technologies Introduces World's Fastest PXI Vector Signal Generator

Agilent Technologies introduced the world's fastest vector signal generator in a PXI form-factor. The Agilent M9381A is a 1MHz to 3GHz or 6GHz VSG that combines fast switching and excellent RF parametric performance: high output power, linearity and superior level accuracy, outstanding adjacent channel power ratio performance (for output levels up to +10dBm or more) and wide modulation bandwidth (up to 160MHz) for testing RF devices.

Agilent designed the M9381A for applications such as testing and validating the design of wireless power amplifiers and transceivers, public safety and military radios, and cellular base stations, primarily picocell and femtocell. The new PXIe VSG does more tests in less time, which reduces the cost of testing. Agilent's exclusive baseband-tuning technology innovation enables frequency and amplitude switching speed as fast as 10µs in list mode and 250µs from its programming interface. The switching speed is further enhanced by excellent linearity and repeatability associated with Agilent's signal generators.

www.agilent.com

NEW 1.5GHZ SPECTRUM ANALYSER FROM RIGOL, FOR LESS THAN £900

Designed to speed RF measurement tasks, Rigol's DA800 series of wide-screen, yet compact, instruments is helping engineers to accomplish a range of testing by combining many functions into one easy-to-use instrument. Model DSA815 comes with pre-amplifier as standard, and there's an optional built-in 1.5GHz tracking generator in model DSA815-TG.

Extensive connectivity including LAN, USB host and USB device makes it easy to integrate the instrument into complete T&M solutions with software control and data capture for recordings and results. There's no compromise on low-noise, with 135dBm (DANL f/1MHz typ.) and a total amplitude uncertainty of less than 1.5dB, this spectrum analyser has more measuring power as standard including AM/FM-demod and built-in pre-amplifier. Plus there are additional options for built-in EMI filter and a quasi-peak detector kit.

The DSA800 spectrum analysers are fast, easy-to-use and reliable test instruments, and at prices starting from £899, they are very affordable too.

www.rigol-uk.co.uk

APACER PCIE SSD ADVANCES TOWARD THE ERA OF OUTSTANDING PERFORMANCE

Apacer launched its latest mPDM (mini PCIe Disk Module) modular SSD. It is based on the "One lane PCIe 2.0" interface to meet the application needs of high-speed transmission. Compared to PCIe 1.0, the broadband speed is boosted to 5Gb/s and the product capacity reaches 128GB, which will offer higher transmission speed and larger storage capacity for high-end application such as network storage or enterprise servers.

The mPDM complies with the AHCI advanced motherboard controller interface and the enterprise standard NVMHCI non-volatile memory host controller interface. Its 'hot plug' and 'plug-and-play' functions ensure the compatibility and competitive edge of the ultra-small (50.8 x 29.8 mm) SSD.

The newly released mPDM modular SSD has standard Mini PCIe interface, in tune with the PCI SIG 2.0 standard, and is backward compatible with the PCIe 1.1 motherboard. Based on the single-level cell flash memory, it delivers sequential read/write speed of up to 300/200MB/sec and IOPS up to 18,000.

www.apacer.com

BELDEN EXTENDS HIRSCHMANN BRAND WITH NEW OCTOPUS RANGE IP54 SWITCH

Belden has launched a new managed Octopus switch from its Hirschmann brand in the EMEA region. This switch is the first of its kind to feature two IP54-rated SFP slots that can optionally be fitted with Gigabit Ethernet or Fast Ethernet transceivers for single-mode or multimode fibers. The switch can be used even in rugged conditions, such as in railroad or mechanical engineering applications, to flexibly implement fast data networks based on various transmission technologies and expand them at low cost when necessary.

Eight Fast Ethernet ports (10/100 BASE-TX) are available for connecting devices. These feature vibration-proof, D-coded M12 connection technology and support PoE. They make it possible to power IP cameras, voice-over-IP phones or WLAN access points without the need for additional cables.

A version with another eight Fast Ethernet ports is also available. The switch can also be integrated in Profinet and Ethernet/IP networks.

www.belden.com

HIGH-SPEED DIFFERENTIAL PROBE FOR HIGH-VOLTAGE SCOPE MEASUREMENTS

The Yokogawa Model 701927 is a new high-speed differential probe for high-voltage floating signal measurements using the company's DLM2000 Series of mixed-signal oscilloscopes and other instruments equipped with the Yokogawa probe interface.

The 701927 has a 3dB bandwidth of DC to 150MHz, and can measure differential or common-mode voltages up to ±1400V (DC plus AC peak) or 1000V RMS. For ease of use, the compact unit features automatic attenuation detection and obtains its power supply via the Yokogawa probe interface. Input attenuation ratio is switchable between 1/50 and 1/500.

A one-metre long extension cable allowing the probe to be used in thermostatic chambers or for measurements on large or hot devices is included as standard. Even when this extension is used, a high bandwidth of 100MHz is maintained.

The new probe is ideally suited to the requirements of the mechatronics, automotive and power electronics sectors for testing devices such as motors, power supplies and invertors among others.

www.tmi.yokogawa.com

New Family of 600V Discrete IGBTs with Low Saturation Voltage

Advanced Power Electronics Corp (USA), a Taiwanese manufacturer of MOS power semiconductors for DC-DC power conversion applications, has recently released a new family of high speed 600V discrete IGBTs for use in AC and DC motor control, home appliances (such as air conditioning, refrigerators, microwave ovens), UPS, solar inverters and induction cookers.

All of the AP20GT60 family's IGBTs feature a collector/emitter voltage rating (VCE) of 600V, a peak gate emitter voltage rating (VGE) of ±20V, and a very low saturation voltage.

AP20GT60P-HF-3 benefits from a typical saturation voltage, VCE(sat), of 1.8V at Ic = 20A, a maximum power dissipation at 25degC of 104W, and comes in a fully RoHS-compliant and halogen-free TO-220 package. AP20GT60ASP-HF-3 (with internal diode) has a typical saturation voltage of 1.7V at Ic = 19A, a maximum power dissipation at 25degC of 78W, a diode forward current rating of 8A, and a diode pulse forward current rating of 40A.

www.a-powerusa.com

NEW 12.1" SUPER WIDE VIEWING ANGLE LCD WITH LED DRIVER

Japanese technology corporation Kyocera has added a new 12.1-inch Super Wide Viewing Angle SVGA LCD module with integrated LED driver circuit.

Using the latest display technology, Kyocera's new SVGA 12.1-inch LCD comes with a super wide viewing angle of 170 degrees. The long-life LED backlight provides lower power consumption with up to 70,000 hours of operating time.

For simple usage and installation, this module is already equipped with an integrated LED driver circuit, so no further components are required to drive the backlight. The backlight itself is especially designed to provide lower power consumption by using the latest technology for LED chips and light guides.

The Kyocera Group is actively engaged in the development of environmentally friendly products and continually strives to reduce the use of environmentally harmful materials and processes in its manufacturing. The company complies with the EU's RoHS (Restriction of Hazardous Substances) directive.

www.kyocera.eu

AVX Expands Inductor Portfolio With Multilayer Chip Inductors

AVX Corporation has expanded its portfolio of thin film, MLO, power and conical inductors with a new series of multilayer chip inductors designed for impedance matching, RF chokes and high frequency filters. The new LCMC Series consists of multilayer chip inductors that provide excellent Q, high SRF, a wide range of stable and accurate inductances and tight tolerances, down to 0.3nH.

AVX's LCMC Series multilayer chip inductors are available in standard miniature 0201, 0402 and 0603 chip sizes, feature tight tolerances of ±5%, ±10%, or ±0.3nH, and are rated for use between -40°C and +85°C. The series is ideal for a variety of high frequency applications, including mobile communication devices, such as cell phones, notebooks, netbooks and tablet computers, WLAN and PHS devices, and EMI counter measure in high frequency circuits.

The LCMC Series is currently in production, and samples, literature and pricing are available upon request.

www.avx.com

NARROW PROFILE HI-REL DATAMATE FROM HARWIN FITS INTO SMALL SPACES

Harwin has expanded its range of Datamate high reliability connectors with the new L-Tek Single-in-line (SIL) male crimp contact connector. The SIL design provides a narrow profile which enables the connector to fit into small spaces. The low weight of the product makes it ideal for use in airborne and portable applications.

Datamate 2mm pitch connectors provide military-level performance at industrial price levels. Datamate L-Tek SIL male crimp contact connectors feature gold plated contacts capable of carrying 3A per contact. Devices are available in unlatched, latched or friction latch versions – the latched versions provide additional mating security.

Available with 2 to 7 contacts for use with either 24-28 AWG or 22 AWG wires, Datamate L-Tek SIL male crimp contact connectors complement the L-Tek product family and enable SIL cable-to-cable and cable-to-board configurations.
New narrow SIL Datamate connectors can withstand high vibration, shock and extremes temperatures.

www.harwin.co.uk

FMICRO NTC THERMISTOR SENSORS FOR MEDICAL APPLICATIONS

ATC Semitec's newest product is a micro thermistor sensor designed primarily for use in medical applications.

Based on ATC Semitec's latest FT thin-film technology and combined with laser-trimming techniques, the Fmicro thermistor sensor is only 0.5mm diameter by 2.3mm long. It is based on one of Semitec's smallest FT thermistors encapsulated in a polyamide tube and fitted with 38AWG insulated leads.

Ideally suited to catheter applications, the Fmicro is accurate to ±0.2K at 37°C. Its operating temperature range is -10/+70°C.

The FT thermistors are highly reliable, characterized by fast response time, which was made possible by the miniturisation of the thermistor dimensions. FT thermistors are the best in class of today's chip thermistors.

www.atcsemitec.co.uk

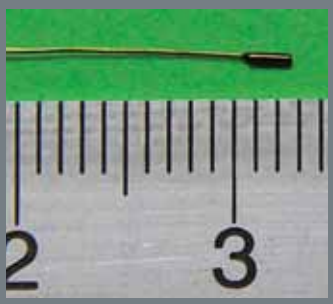

World's Smallest 5 and 10 Watt AC-DC Power Supplies

XP Power announced the launch of what it claims are the world's smallest 5 & 10 Watt single output AC-DC power supplies. Measuring just 25.4 x 25.4 x 15.24mm (1 x 1 x 0.6in) the miniature 5W ECE05 series of encapsulated board mount power supplies are ideal for product designs where the available board space is at a premium and high power density supply is required.

Complementing the ECE05 series, the ECE10 series provides 10W output from an equally compact encapsulated package measuring just 38.1 x 25.4 x 15.24mm (1.5 x 1 x 0.6in).

Both series comprise seven models covering the popular nominal output voltages from +3.3V to +48VDC, all from a wide universal input of 85-264VAC with Class II construction, which means that no earth or ground connection is required. All models meet the internationally recognized no-load power consumption limit of less than 0.3W and average efficiency limits ensuring end-products can comply with world-wide energy efficiency standards.

www.xppower.com

KEITHLEY EXPANDS SOURCEMETER SMU LINE WITH THREE NEW MODELS

Keithley Instruments has added three new source measurement unit (SMU) instruments suited for benchtop and R&D applications. Its Series 2600B line now includes new capabilities that enhance productivity and ease of use. As such, the test and measurement industry's most powerful, fastest and highest resolution SMU instruments now provide $6^1/_2$-digit resolution, software emulation of Keithley's widely installed Model 2400 SourceMeter instrument and USB 2.0 connectivity.

The Series 2600B's models 2604B, 2614B and 2634B are optimized for benchtop research, product development, student labs and other settings where tightly integrated operation and high accuracy are important, but leading-edge, system-level test automation is not. They are also priced to fit as easily into a researcher's budget as they do onto a benchtop.

Series 2600B offers both single- and dual-channel models combining a precision power supply, true current source, $6^1/_2$-digit multimeter, arbitrary waveform generator, pulse generator and electronic load.

www.keithley.com

Constant-Current 3A Buck Regulator IC For Driving High-Power LEDs

The new A6211 from Allegro MicroSystems Europe is a switching buck regulator IC that provides a constant-current output of 3A for driving high-power LEDs.

The new device is designed for use in consumer, industrial and architectural lighting applications, including MR16 bulb replacements and light bars for scanners and multifunction printers. It has significant fault protection advantages including design techniques to address pin-pin shorts, pin-ground shorts and open or shorted external components.

With a wide input supply voltage from 6 to 48V, the new device integrates a high-side N-channel MOSFET switch for DC-to-DC step-down conversion. A true average current is produced using a cycle-by-cycle controlled "on" time method.

The LED current is user-configurable with an external current sense resistor, allowing the output voltage to be set to optimise constant-current supply for a wide variety of LED configurations including multiple LEDs in series, in a single string or in parallel.

www.allegromicro.com

CYNTECH NOW OFFERS ITS KIT FOR THE RASPBERRY PI COMPUTER

Cyntech Components, a specialist supplier of electro-mechanical components and power supplies, now has stock of its custom-designed enclosure for the Raspberry Pi computer. Cyntech is also offering the enclosure with a full kit of cables and power supply unit (PSU). Both the enclosure and the full kit with cables are in stock for immediate delivery with higher volumes subject to a lead-time of just two weeks.

The compact, low-cost case is tailored to fit the Raspberry Pi board and provides openings for all the connectors. This includes an exit slot in the side of the case allowing a ribbon cable to be fitted to the board's 40-pin GPIO header. Another feature is the light pipe window for viewing the board's LED indicators. The case, which is supplied in the distinctive Raspberry colour and printed with the Raspberry Pi logo, provides an attractive solution for protecting the printed circuit board.

www.cyntechcomponents.com

MAJOR INVESTMENT IN IQD FOQ GMBH MADE BY HONG KONG X'TALS LTD

Asian frequency products manufacturer Hong Kong X'tals Ltd, part of the Kolinker Group, has acquired a 40% shareholding in IQD's specialist OCXO manufacturing division IQD FOQ GmbH for an undisclosed sum.

IQD FOQ (previously FOQ Piezo Technik) has over 30 years' experience in developing high precision crystals and oscillators including specialist Clock Oscillators, VCXOs, TCXOs and OCXOs. Its modern 1500m^2 facility based in Bad Rappenau, Germany, has a high degree of automation allowing for competitive manufacturing costs and flexible delivery lead-times. Processing of both AT and SC-cut crystals is carried out on-site supported by a class 1000 clean room, in-house PCB design & assembly and automated test facilities.

The Kolinker Group is a maker or frequency products test and measurement equipment. Established in 1989, Hong Kong X'tals owns four product factories in China and Hong Kong with a total annual production capacity of over 500 million units.

www.iqdfrequencyproducts.com

RS COMPONENTS CELEBRATES 75 YEARS OF BUSINESS SUCCESS

RS Components (RS), the trading brand of Electrocomponents plc, a leading distributor of electronics and maintenance products, is celebrating 75 years in business. From humble beginnings in a small London garage in 1937, RS has grown into a successful multinational organisation by retaining its focus on high service-level customer support while evolving into a multi-channel business with e-commerce at its heart.

Over the last 75 years RS has transformed itself from a UK company into a global business. With direct operations in 30 countries and distributors in a further 37 markets, RS is the number one high service distributor across Europe and the Asia Pacific region. The company's 16 distribution centres located worldwide enable it to provide high service levels to over one million customers.

The company's success is a direct result of its ability to embrace change, adapting to the needs of the market and its customers.

www.rs-components.com

High-Voltage Transmission Industrial Connector Module

A new high-voltage single module has been added to the Harting Han-Modular industrial connector series, allowing the transmission of extremely high voltages with a minimum space requirement.

The new Han HV single module incorporates Teflon material to provide high-voltage insulation, allowing a Han 24 B connector, for example, to be designed with up to 12 power contacts rated at 16A and 2500V. The module can also be combined with signal and pneumatic modules.

The new connector module features a high working voltage of 2500V, and is based on a space-saving design with two power contacts per module. It uses crimp terminal technology, and can handle conductor cross sections of between 0.14 and 4 mm^2.

Harting's Han-Modular series, as an open connector system, can provide flexible configuration to the users' individual requirements. It gives the user the freedom to assemble a customised connector from a range of more than 40 modules for electrical, optical and gaseous signals.

www.harting.com

NEW FLAME RETARDANT EMC SHIELDING GASKET FROM KEMTRON

Kemtron, the British manufacturer of RFI/EMI shielding gaskets, materials and components has launched a flame retardant, low smoke, low toxicity EMC shielding gasket. It's tested and approved to the international standard UL94V-0 by Underwriters Laboratories for flame retardancy, file number E344902. Also tested for smoke density to BS 6853:1999: Annex D.8.3and oxygen index to BS EN ISO 4589-2:1999, confirming the material meets the requirements for minor internal use on vehicles category 1a, such as gaskets for electronic enclosures, is making it highly suitable for applications in underground transportation, trains and other safety critical applications.

The material is nickel-coated graphite loaded into silicone elastomers, product code SNG-FR. This allows the gasket to provide a highly electrically-conductive path between mating flanges of an electronics equipment enclosure giving a high level of RFI/EMI shielding. The material can be supplied as an extruded strip in various profiles, "O" rings or flat die cut gaskets.

www.kemtron.co.uk

UK HIGH-SPEED BROADBAND TIMETABLE SLIPS

Despite the ambitions of the UK government and local service providers to build a superfast national high-speed broadband network the most advanced in Europe by 2015, its planned rollout will be missed this year.

Access to high-speed Internet networks is seen of great economic and social importance, in particular for those working and living in more remote areas, but the timetable has been set back by three months at least. This is partly due to the majority of the local authorities in the UK having not yet begun procurement. Another parameter causing delays are "the lengthy negotiations with the EU over the state aid required to subsidise the project" according to the government spokespeople. Under the Broadband Delivery UK programme, £530m has been allocated to help fund infrastructure in far-flung places of the UK, with matched sums from the private sector.

PROFESSOR DR DOGAN IBRAHIM, Near East University in Nicosia, Cyprus: High speed broadband will become increasingly important as the number of users and the volume of data increases rapidly. However, speed is only one aspect of having Britain well connected. There are other factors which are overlooked by the government, such as coverage, reliability, competition and free Internet. The fact that the project is to be delayed is not surprising as this is a massive undertaking and it is taking place whilst other communications related topics are being discussed and debated. These include proposals to monitor Internet exchanges between individuals, and also the provision of the 4G services in the UK. It is preferable to have the project delayed, than to undertake a project without the adequate understanding, planning and procurement associated with it.

BARRY MCKEOWN, RF and Microwave Engineer in the Defence Industry, and Director of Datod Ltd, UK: As Corporal Jones of "Dad's Army" would say: "Don't Panic!". Instead look to the rollout in Cornwall (rural) with the fastest, and Belfast (urban) which has the highest broadband penetration with 97% coverage.

Next, exclude all marketing people from the room so the grownups can take a strategic overview. As previously outlined in EW, we are enduring three technology (both fixed and mobile) eras. The all-analogue with interference era (a) is now past. An analogue and digital with interference era (b) is the present. An all-digital without interference era (c) is the future. It is for scientists and engineers to deliver this outcome, not politicians.

The speed aspect is being directed by the ITU/WRC and other International Standard bodies by agreement; accordingly the industry focus on speed is a global strategic imperative, whereas rural vs urban rollout penetration is dictated by national governments locally. However, once this (primarily fibre) broadband infrastructure is in place we then can move to era (c) when the real fun shall begin. This is where DSA and Spectrum Sharing technologies shall come into their own.

The strategic DSA technical specifications are still to be agreed at ITU/WRC level in 2015. The UK is well placed for era (b) but not for era (c), largely because of a failure of political leadership by both the previous and the current government and an over-reliance on Ofcom. In strategic, national security terms, era (b) is about economic wellbeing whereas era (c) is about war-fighting capabilities, and the US with its policy of "full spectrum dominance" has absolutely no intention of voluntary relinquishing its global dominance here.

Also note that it is not clear, at least to me, that politicians have grasped the differences between fixed and mobile broadband and the applications that flow from each. Naturally, the marketing people have mudded the waters here.

> What is more important is the House of Lords' critical comment on the focus on very high speed rather than on 100% coverage

JAN DIDDEN, Audio Expert and Publisher of Linear Audio, UK: Projects like these always have a mixture of political, economic, social and – yes! – prestige drivers. There's no question about the importance of (super)fast Internet for the future economic and social health of any nation, and the various public and private partners in this undertaking have recognised this and are acting on it. Yet, it has to be questioned whether a completion delay of three months spells the end of the world, as clearly it doesn't!

Public procurement processes have always been somewhat convoluted, driven by the need to get the best bang for the pound. Blaming these processes for the 'disaster' of a three months' delay seems more politically motivated than out of concern for those lacking superfast connections.

Having been involved in public procurement (NATO) in another life, I am aware that there is always a tension between the desire of industry for swift decision making, preferably with a single counterpart, and the public sector's desire for transparency and following local regulations and budget rules. That will never change. Using an insignificant delay to try to influence the procurement process is always part and parcel of projects of this magnitude and, as always, will not change anything. Business as usual!

What is more important here is the House of Lords' critical comment on the focus on very high speed rather than on 100% coverage. Undisputable, the bandwidth requirements of major cities are not identical to far-flung rural communities so trying to roll out a one-size-fits-all superfast network nationwide doesn't seem to be most efficient.

I noticed that current 'broadband' connectivity is defined as up to 2Mbps while the new superfast network will offer up to 1Gbps capacity. That's a large difference and while 2Mbps is insufficient for today's increased use, 1Gbps is certainly not required for all UK users in all locations.

A two-tiered approach could be to strive for 100% coverage for 'fast' access, loosely defined as, say, minimum 50Mbps, with future growth built-in, while rolling out the superfast network to those places where it makes economic sense. Perhaps that can even be achieved within the original timetable.

If you are interested in becoming a member of our panel and comment on new developments and technologies within the electronics sector please register your interest with the Editor by writing to Svetlana.josifovska@stjohnpatrick.com

ROUTE FASTER !

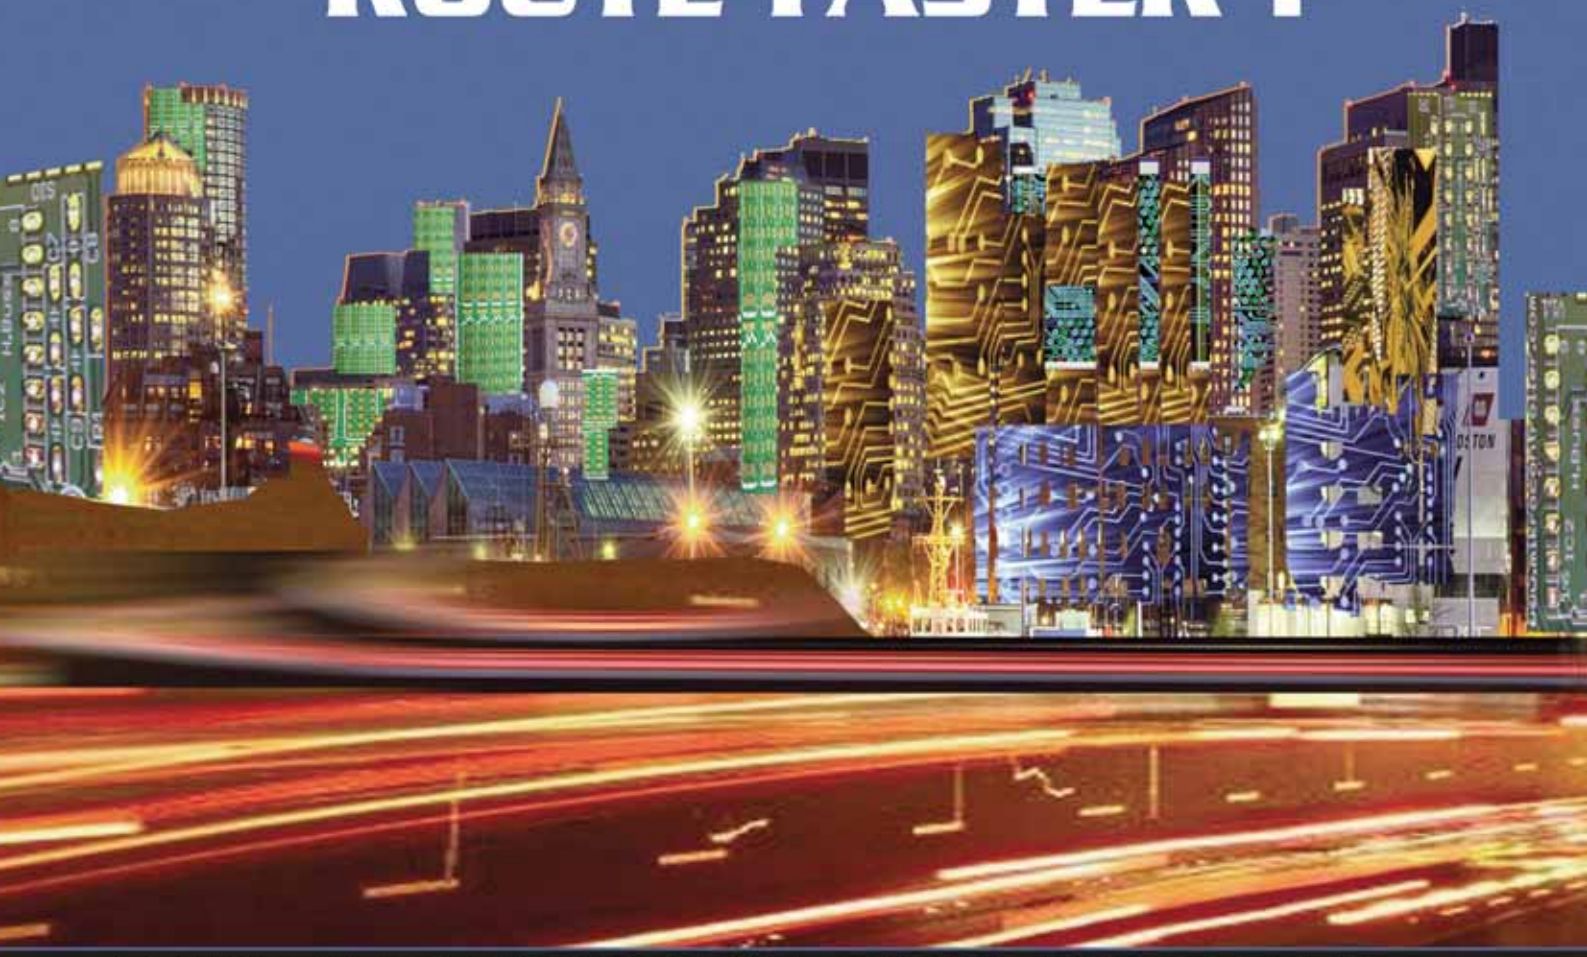

WITH PROTEUS PCB DESIGN

Our completely new manual router makes placing tracks quick and intuitive. During track placement the route will follow the mouse wherever possible and will intelligently move around obstacles while obeying the design rules.

All versions of Proteus also include an integrated world class shape based auto-router as standard.

PROTEUS DESIGN SUITE Features:

- Hardware Accelerated Performance.
- Unique Thru-View™ Board Transparency.
- Over 35k Schematic & PCB library parts.
- Integrated Shape Based Auto-router.
- Flexible Design Rule Management.
- Polygonal and Split Power Plane Support.
- Board Autoplacement & Gateswap Optimiser.
- Direct CADCAM, ODB++, IDF & PDF Output.
- Integrated 3D Viewer with 3DS and DXF export.
- Mixed Mode SPICE Simulation Engine.
- Co-Simulation of PIC, AVR, 8051 and ARM7.
- Direct Technical Support at no additional cost.

Prices start from just £150 exc. VAT & delivery

Labcenter Electronics www.labcenter.com

Labcenter Electronics Ltd. 53-55 Main Street, Grassington, North Yorks. BD23 5AA.
Registered in England 4692454 Tel: +44 (0)1756 753440, Email: info@labcenter.com

Visit our website or phone 01756 753440 for more details

NEW!

Primary & Rechargeable Lithium Batteries

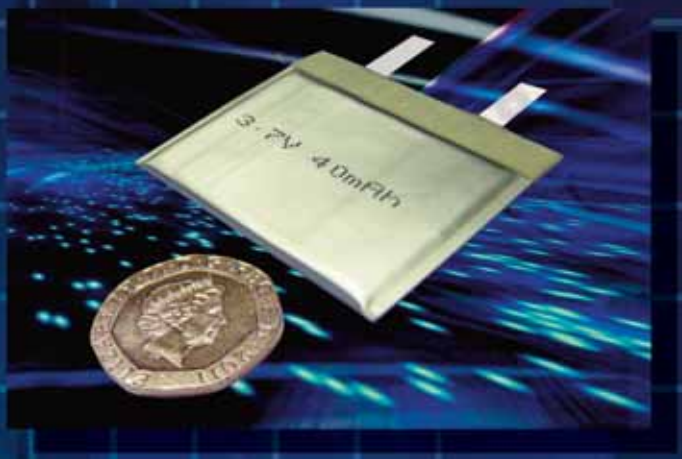

CP Series Lithium Polymer Primary Thin Cells

- 3V DC nominal output
- Standard products from 16mAh up to 3000mAh
- Very low profile can be to made to <0.4mm thick
- Very light weight
- Batteries are safe and explosion proof
- Very low discharge rate of around 1% per annum
- Laminated aluminium thin foil construction
- Wide operating temperature range of -20°C to +60°C
- Can be manufactured in custom sizes and ratings to suit customer specific requirements

PP3 Size Lithium Primary & Rechargeable Cells

- **CP9V** 9V Lithium Manganese Dioxide Primary Cells
- High capacity (1200mAh) in compact lightweight case
- Low discharge rate, <2% per annum.
- Wide operating temperature range -20°C to +60°C
- Restricted for UL, UN38.3 and ROHS
- **G600** 9V Lithium Polymer Rechargeable Cells
- 600mAh capacity similar with primary alkaline cells
- Fitted with protected circuit module (PCM) to protect against over charging, over discharging and over current
- Fully safety tested for high impact, crushing, forced discharge, short circuit and piercing with no fire or explosion caused

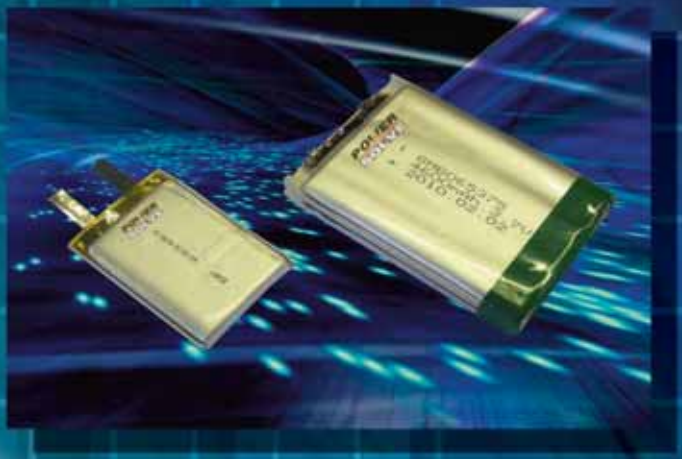

GMB Series Lithium Polymer Rechargeable Cells

- 3.7V per cell with capacities up to 10Ah (standard models)
- 150 standard sized cells down to <2.0mm thickness
- Can be supplied as single cells or packaged in shrink sleeving for increased voltage or current
- Custom molded versions are available subject to quantity
- All versions supplied with protected circuit modules (PCM) to protect against over charging, over discharging and over current.
- Standard charge time 8 hours, fast charge time only 2.8 hours
- Operating Temperature -10°C to +45°C
- Fully safety tested against fire and explosions

POWER SOLVE

TEL: 44-(0)1635-521858 Email: sales@powersolve.co.uk

www.powersolve.co.uk