July 2012
Volume 118
Issue 1915
£5.10

www.electronicsworld.co.uk

Electronics WORLD

THE ESSENTIAL ELECTRONICS ENGINEERING MAGAZINE

Technology
Fast-acting electronic nose uses QTC material

Special report
Embedded design

Products
Various PCB hardware products

Any Questions?

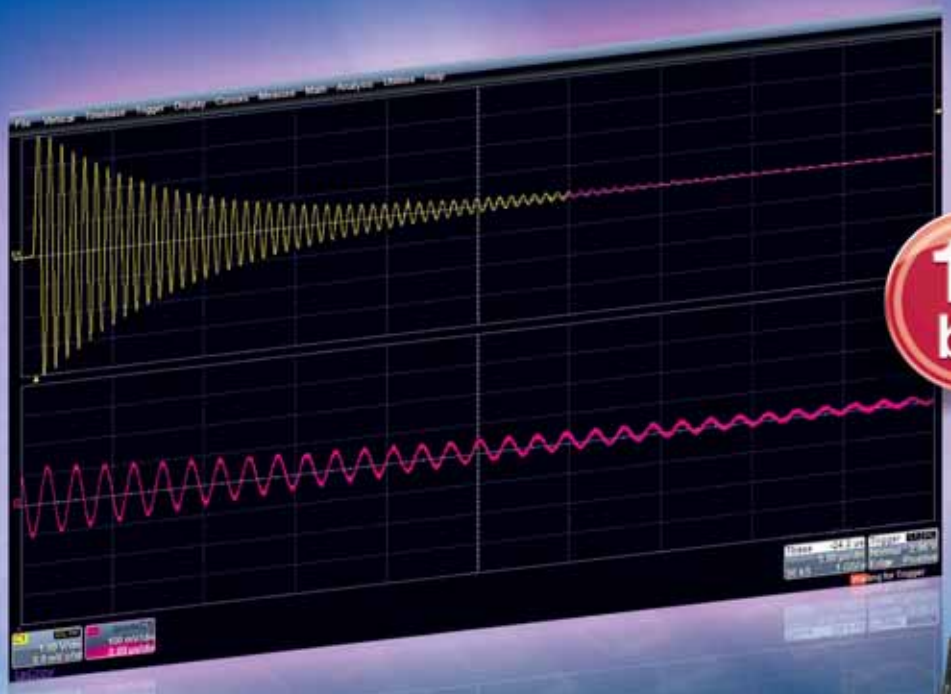

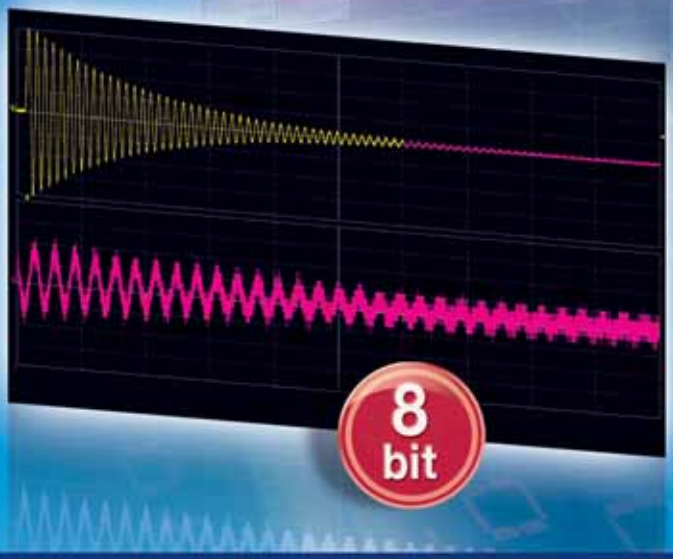

Unmatched Signal Fidelity and Excellent Measurement Precision

WaveRunner High-Resolution Oscilloscopes

- 12-bit ADC Resolution, 15-bit with ERES
- ± 0.5 % DC Accuracy
- Complete Analysis Capabilities incl. Spectrum Analysis Package
- 1 mV/Div at full Bandwidth
- 1.95 µV Smallest Voltage Step
- 400 V maximum Offset
- 36 Ch. Mixed Signal Solution

Phone: 01753 725371
www.lecroy.co.uk

LeCroy

Cover supplied by MOUSER ELECTRONICS
More on pages 8-9

REGULARS

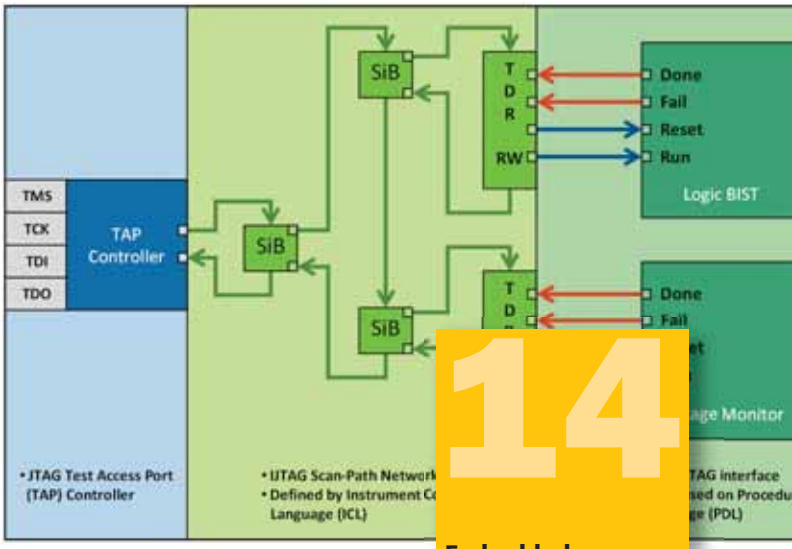

14 Embedded instrumentation

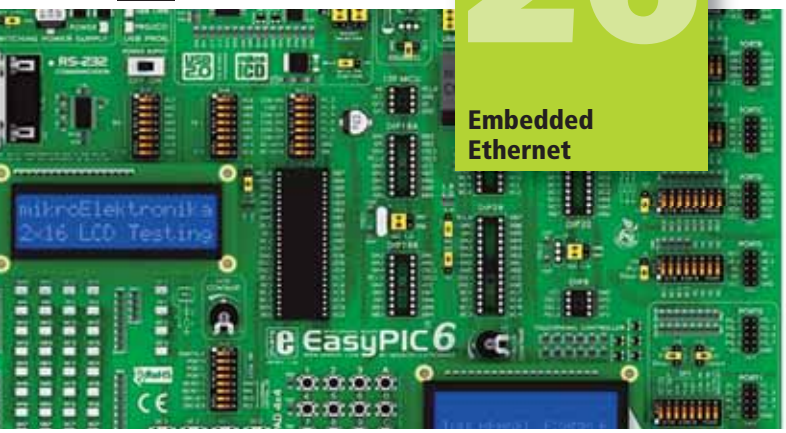

26 Embedded Ethernet

FEATURES

COMPETITION WINNER
Mark Abbott
ST-Ericsson BU Connectivity & Broadcast, UK

Mark wins a Rigol DS1052E oscilloscope
Congratulations from the EW team!

One platform for 8-, 16- and 32-bit development - with Microchip's MPLAB® X IDE

MPLAB® X IDE is the free, integrated toolset for all of Microchip's 900+ 8-, 16- and 32-bit PIC® Microcontrollers, dsPIC® Digital Signal Controllers, and memory devices. Based on the open-source NetBeans platform, MPLAB X runs on Windows® OS,MAC® OS and Linux, supports many third-party tools, and is compatible with many NetBeans plug-ins.

MPLAB XC compilers help increase code speed of any PIC® Microcontroller or dsPIC® digital signal controller by 30%, whilst also cutting code size by 35%. These new compilers give designers the choice of Free, Standard or Pro code optimisation levels for 8-bit, 16- or 32-bit development, or a single C compiler suite to support all Microchip Microcontrollers and digital signal controllers.

Microchip's tool chain of compatible compilers and debugger/programmers operate seamlessly within the universal, cross platform and open-source MPLAB® X integrated development environment, reducing both learning curves and tool investments.

START DEVELOPING TODAY

Download a free copy of MPLAB X and choose from a choice of new C compilers:

- MPLAB XC8 for 8-bit MCUs
- MPLAB XC16 for 16-bit MCUs and DSCs
- MPLAB XC32 for 32-bit MCUs
- MPLAB XC Suite for all 900+ PIC MCUs and dsPIC DSCs.

Evaluate MPLAB X today! www.microchip.com/get/eumplabx

MICROCHIP

Microcontrollers • Digital Signal Controllers • Analog • Memory • Wireless

DESIGNING THE NETWORK FOR BIG DATA

Migrating to Ethernet as an underlying transport will benefit all aspects of the transmission of Big Data. However, Ethernet on its own is not sufficient to handle the torrent of Big Data

The tidal wave of Big Data is washing over today's businesses. Big Data is not only measured in the quantity of data traversing the network, but according to Forrester Research, it is defined as the techniques and technologies that make handling data at extreme scale economical.

Global data volume is increasing 40% annually, much of this generated by the growth in mobile traffic. The value to public and private organisations is enormous; global management consulting firm McKinsey estimates that in the European public sector alone, Big Data will contribute €250bn.

How can today's networks cost-effectively handle the high volume of interactive, multimedia traffic? Just adding more bandwidth will not solve the problem. The value chain of network providers spanning mobile backhaul, through carrier transport, to the data centre must reconsider networking to make it more cost-effective and efficient to accommodate Big Data.

Big Data Requires Big Changes in the Network

How much data are we actually talking about with Big Data? The size of the datasets varies by sector, but they range from a few dozen terabytes to multiple petabytes. Moreover, Big Data introduces new technologies such as Hadoop and MapReduce. Unlike previous data transfer technologies that moved gigabytes of data in a single job, MapReduce can move multi-petabytes of data.

Big Data requires that the networking industry accelerates change, moving from legacy technologies in data transport such as SONET/SDH to high-value-per-bit and lower-cost-per-bit technologies such as Carrier Ethernet. In the data centre, networking switching is moving rapidly to 10-40Gbps and Ethernet is rapidly replacing older Storage Area Networking (SAN) technologies.

Big Data Requires Flatter Data Centre Networks

Migrating to Ethernet as an underlying transport will benefit all aspects of the transmission of Big Data. However, Ethernet on its own is not sufficient to handle the torrent of Big Data.

Network architectures also need to change to become flatter and more flexible. For example, in today's typical data centre, the network architecture consists of one or more L3 core routers, multiple L3 access routers, L2 aggregation switches, load balancers, as well as top-of-the-rack switches. A hierarchical network forces data centre operators to oversubscribe network resources up to 200:1. Taking advantage of new 10Gbps and 40Gbps Ethernet and higher performance silicon will help drive the price performance up. Taking advantages of new innovations in network virtualisation enable flatter and better performing networks.

The Future of Big Data Networking is Software

The key to delivering higher performance, more optimised networks is a software-defined networking architecture using a centralised control plane and fast forwarding data planes based on merchant silicon. The Open Networking Foundation (ONF) is championing the standardisation of this approach, along with OpenFlow, a protocol to enable distributed control plane implementations.

Key software networking vendors have been offering modular, portable routing/switching software for more than a decade to leading network equipment manufacturers. By leveraging innovations and the ONF standards, network equipment providers can optimise the price-performance of Ethernet networks.

Software Defined Networking is in its infancy, but the growing demands of Big Data will continue to drive new innovations in networking. Today, service providers and data centre operators can exploit Ethernet to offer more economically-attractive mobile and Big Data transport services. In the near future, the network equipment industry will accelerate the adoption of SDN to support Big Data. The successful companies will ride the wave of Big Data, transform their networks and capitalise on new business opportunities.

Sandy Orlando is Vice President for Worldwide Marketing at IP Infusion

EDITOR: Svetlana Josifovska
+44 (0)845 4790343
Email: svetlanaj@stjohnpatrick.com
DISPLAY SALES: John Steward
Tel: +44 (0) 20 7933 8974
Email: johns@stjohnpatrick.com
DESIGN: Tania King
Email: taniak@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

PERATECH CREATES FAST-ACTING ELECTRONIC NOSE USING QTC TECHNOLOGY

Granular Quantum Tunnelling Composite (QTC) material has been used in Peratech's innovative electronic nose

Touch technology innovator Peratech is developing an 'electronic nose' using its award-winning Quantum Tunnelling Composite (QTC) material. This new sensor technology detects the presence of Volatile Organic Compounds (VOCs) very rapidly and can recover in a matter of seconds – beating all other techniques currently available.

QTC materials change their resistance when a force is applied and, in this case, the polymer content of the composite swells when exposed to VOCs. One form of Peratech's sensor uses a granular type of QTC material that provides a high surface area for absorption, enabling it to detect levels of VOCs in the region of 10-100ppm. The sensor rapidly recovers once the VOCs have gone from the surrounding atmosphere – and it is the speed of sensing and recovery that marks the difference between QTC sensors and those using other sensing technologies. An additional feature of the QTC technology is that it has very low power requirements.

The conductive particles used in the QTC electronic nose have nano-sized features and are distributed in a non-conductive polymer. When a force is applied or swelling occurs, the particles move close enough for the electron flow between the particles to alter due to an effect called quantum tunnelling. The polymer used is selected for its response to the particular VOCs to be monitored. The company believes licensing the technology will lead to a variety of products.

"The electronic nose application was developed in conjunction with the Quantum Tunnelling Composite research group at the University of Durham," said David Lussey, CTO of Peratech. "We are now looking for companies who are interested in licensing the technology from us to develop products."

DuPont Novel Formulation Brings Greater Efficiency to Solar Cells

DuPont Solamet PV17x photovoltaic metallization paste raises efficiency in standard solar cell constructions and its unique properties enable Lightly Doped Emitter (LDE) cell designs, says the firm.

LDE is a strong differentiator for solar cell producers as it can boost efficiency by up to 0.4%. "This product has a two-stage advantage, since on its own it can raise efficiency by up to 0.2%, and when used to enable LDE, it can raise efficiency by up to 0.4%," said Peter Brenner, photovoltaics global marketing manager at DuPont Microcircuit Materials. "Solamet PV17x series continues to advance technology to help meet the industry's goal of 20% efficiency by 2012."

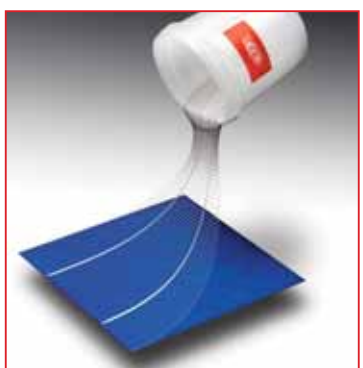

DuPont Solamet photovoltaic metallizations boost the efficiency of photovoltaic solar cells

DuPont says that Solamet PV17x is the most advanced composition in the market, allowing contact to be made to the most lightly doped junctions. Doping diffusion optimization is a key area of experimental study in the photovoltaic industry for the design of high efficiency cells. Diffusion optimization has been significantly limited by the inability of traditional frontside photovoltaic silver pastes to contact lightly surface-doped emitters. Prior to Solamet PV17x being released, the industry had no real commercially available option for making a screen printed frontside metallization that could economically and practically enable an LDE.

Extensive testing is underway within DuPont and in collaboration with several research organizations as well as in customer trials to fully characterize and continue to advance this technology. RWTH-Aachen University recently published a comparative study involving Solamet PV17x and four competing metallization pastes. Solamet PV17x outperformed four competing products, demonstrating its ability to contact 100 Ohm/sq emitters on multicrystalline cells – the first time this had been achieved – with lightly doped phosphorous surface concentration. This enabled an efficiency increase of one full percent versus the homogenous emitter base line and 0.4% higher efficiency was confirmed versus laser doped selective emitter technologies.

"We presented a characterization of POCl3 parameters influencing the electrically active phosphorus concentration profiles by electrochemical capacitance voltage measurements," said Ali Safiei, PhD researcher at the Institute of Semiconductor Electronics at RWTH Aachen University. "For the first time we could demonstrate a successful direct contacting of an optimized high sheet resistance emitter at 100Ω/sq by increasing the n++ layer and at the same time reducing the dead layer. Multicrystalline silicon solar cells were fabricated using five different silver pastes resulting in an absolute efficiency gain of $\Delta\eta$ = 1% in comparison to a standard 55Ω/sq emitter. Based on these investigations we evaluated a 160Ω/sq emitter and could successfully demonstrate by laser doping that a n++ layer of up to 25nm depth (a Lightly Doped Emitter) leads to high FF and an absolute efficiency gain of $\Delta\eta$ > 0.6%."

The breakthrough formulation of Solamet PV17x also enables cell makers to use up to 15% less material, in line with the company's intent to accelerate product developments that help the photovoltaic industry reduce its dependence on silver metals and offset some of the impact that rising silver prices have on the cost of producing solar cells and modules.

Research Projects Combine to "End the Segregation Between Control, System and Software Engineers"

The research efforts "Project "P" and "Hi-MoCo" (High-Integrity Model Compiler) have combined to provide an open-source, tunable and qualifiable code generation framework for domain-specific modeling languages.

The key idea is to allow control engineers, system engineers and software engineers to easily collaborate for system-level model integration, verification and final optimized code generation targeting the Ada 2012, C/C++ and VHDL languages. Typically, control engineers use Simulink, Stateflow and Scicos/XCos; system engineers use SysML/MARTE and AADL; and software engineers use UML.

"The ultimate goal of these projects is to end the segregation between the control, system and software engineers," said Franco Gasperoni, Managing Director of AdaCore, one of the project's participating companies and technical coordinator.

AdaCore is working closely with the IB Krates team and members from IRIT (Institut de Recherche en Informatique de Toulouse), the principal architects of the ITEA GeneAuto project and technology on which "Project P" and Hi-MoCo are based. The three companies will play a major role in the cross-domain qualification effort spanning the avionics, space and automotive sectors.

"A major bottleneck in the model-driven development of software for avionics, space and automotive systems is the integration of heterogeneous models and the lack of comprehensive verification and code generation technologies," said Gasperoni.

The current state-of-the-art is to perform integration on generated sources. The trio are proposing to do this at the model level to verify integration issues well before models are mature enough for code generation.

"Project "P" and Hi-MoCo are open-source research efforts, which are supported and partly funded by the French and Estonian national governments and the European EUREKA agency.

Project P

Project P is a three-year research project financially supported within the French FUI 2011 funding framework. Headed by Continental Automotive France, it involves the collaboration of 19 partners including major industrial users from the avionics, automotive and space sectors that count Airbus, Astrium, Continental Automotive, Rockwell Collins, Safran, Thales Alenia Space and Thales Avionics among its ranks, as well as technology providers such as AdaCore, Altair, STInformatique and Scilab Enterprise; service companies ACG Solutions, Aboard Engineering and Atos Origins; and research centres ENPC, IRIT-INPT/CNRS, INRIA, ONERA, Lab-STICC/Université de Bretagne Sud. To find out more go to http://www.open-do.org/projects/p.

Project Hi-MoCo

Hi-MoCo P is a two-year research project financially supported within the Eurostar 2011 funding framework. It supports the collaboration of IB Krates (Estonia), IRIT and AdaCore (France).

MOUSER ELECTRONICS: DISTRIBUTING THE MOST ADVANCED EMBEDDED PRODUCTS FROM THE WORLD'S TOP MANUFACTURERS

The creation of every new smartphone, portable medical device, or automotive driver assist feature serves as a constant reminder of how technology never stands still. "Quickly innovate or risk being left behind" is the industry maxim heard everywhere. And nowhere is this cry more prevalent than among design engineers faced with complex task of delivering the next technological wonder or "it" item.

As the convergence of communications and computing functions within components continues to advance, embedded systems are transforming into complex systems. Major components of embedded systems are microcontrollers (MCUs), digital signal processors (DSPs), analog, real-time operating systems (RTOS), industry-specific protocols and interfaces. Embedded systems also span digital and analog, sensors and actuators, software, mechanical items and other components.

Kevin Hess, Mouser VP Technical Marketing

Ever present are new opportunities to develop and build more powerful, energy-efficient solutions around minimization of size and low cost. Important design characteristics at work are low power, real-time responsiveness, low thermal dissipation, small physical form factor, memory, regulatory concerns, ruggedness in design and other factors. All of this creates a major design challenge for engineers when it comes to the selection of processors, operating systems, etc., as demands for functionality increase with rising time-to-market pressures.

"In today's world, convenience means portable, which in turn means low power systems running on batteries for sleep mode," explains Kevin Hess, Mouser Vice President Technical Marketing. "Although low power has always been a goal, the trend towards microcontrollers with ultra-low-power consumption has accelerated in recent years."

Mouser Electronics stands at the forefront in embedded processors and development tools, stocking and quickly shipping the latest, leading-edge solutions that are smarter, greener and more convenient. With its eight Customer Support Centers across Europe's largest technology hubs, it's not surprising to learn that Mouser has rapidly grown into a global engineering resource and leading worldwide distributor for advanced semiconductors and electronic components among design engineers across Europe.

Mouser UK Locally Supports Design Engineers by offering the Newest Products for Embedded Designs

Being green doesn't necessarily mean low power, Hess adds. To reduce the carbon footprint, engineers need to make the most efficient use of power consumed. Embedded system engineers also face the need to update their programming, debugging and testing skills to draw the best out of designs. Thus, many semiconductor companies are aligning with third-party development tool manufacturers for their expertise. Hess says open source technology is greatly influencing the design of low-cost, highly functional and adaptable solutions.

Mouser has established itself as the go-to source for the latest MCU and embedded products from the world's leading manufacturers, including Texas Instruments, Microchip Technologies and Freescale Semiconductor, to name a few. Mouser also stocks the newest embedded products from Laird Technologies, Olimex, Digi Int'l, Lantronix, Netburner and Roving Networks, carrying a full suite of development tools, including starter kits, evaluation boards, programmers, debuggers and emulators for rapid evaluation and testing of designs. Product numbers are updated daily on www.mouser.com and labeled as new technology or flagged NRND (Not Recommended for New Design) and/or EOL (End of Life) and made available via online, catalogue or mobile.

"With the requirements of both the end consumer and electronics industry need for embedded design increasing so rapidly, every link in the chain – from semiconductor manufacturer to development tool maker to distributor – must constantly be on the cutting edge," explains Hess. "Our goal is to help design engineers around the world get their products to market faster. We do that by supplying the most advanced products from the top suppliers. Therefore, customers can count on Mouser to deliver What's Next in authorized embedded processors and development tools."

Mouser Electronics UK
Suite C, First Floor, Artisan Building
Hillbottom Road, High Wycombe, Bucks. HP12 4HJ,
United Kingdom
Tel: 01494 467490
Email: uk@mouser.com

The Design Engineer's Perspective on the Supply Chain

WHERE DO ENGINEERS GO FOR MORE INFORMATION ABOUT THEIR PRODUCTS AND PROJECTS? **CHRIS SULLIVAN**, HEAD OF GLOBAL SOLUTIONS MARKETING AT PREMIER FARNELL FOR ELEMENT14 HAS SOME ANSWERS FOLLOWING A SURVEY

The electronics industry's key focus in the past decade has been on the supply chain: globalizing it, reducing its costs and risk, optimizing it and auditing it for ethical conduct. As important as these issues are, this intense focus on suppliers has overshadowed perhaps the most important determinant of a successful supply chain: the design phase.

Like captains of huge ships, electronics design engineers make hourly and daily decisions that steer their companies' supply chains toward new technologies, divergent regions, lead-time factors, environmental sustainability and competitiveness. But it's easier said than done.

The four stages of design – concept, designing, prototyping and pre-production – all have their challenges for electronics engineers. Today, electronics design engineers are burdened with lengthy and disparate processes comprising dozens of technical documents, design tools, product searches, evaluations, regulatory requirements and promises of availability.

Like captains of huge ships, electronics design engineers make hourly and daily decisions that steer their companies' supply chains toward new technologies, divergent regions, lead-time factors, environmental sustainability and competitiveness

In a recent survey, element14, Premier Farnell's online community for electronic design engineers, teamed up with Technology Forecasters Inc (TFI) to ask 328 design engineers around the world how they think about their design process. An overwhelming majority of respondents voiced similar top three concerns in the design process:

1. **Increased time pressures:** Initial design stages – before prototype assembly and testing – typically require the most time and effort to gather all the necessary information. There is never enough time to properly utilise every relevant source.
2. **Incomplete or inaccurate information from relevant sources:** Information overload is not as much an issue, as is it's quality and completeness. This prolongs the design process further.
3. **Difficulty comparing options and alternatives:** While top-level information about component characteristics is easily available, the more detailed, sophisticated and time-specific information and tools can be very difficult to locate.

Naturally, the Internet has become a critical tool in the engineering design process.

Design stage	Average percentage of time spent finding and aggregating information, data, and tools	Average degree of difficulty in finding and aggregating the information, data, and tools (1 = extremely easy, 5 = extremely difficult)
Concept Development	41%	3.0
Detailed Design	45%	3.0
Prototype Testing and Evaluation	36%	2.7
Pre-Production	25%	2.5

Figure 1: Time spent and the degree of difficulty by stage in the design process

Figure 2: Challenges in the design process

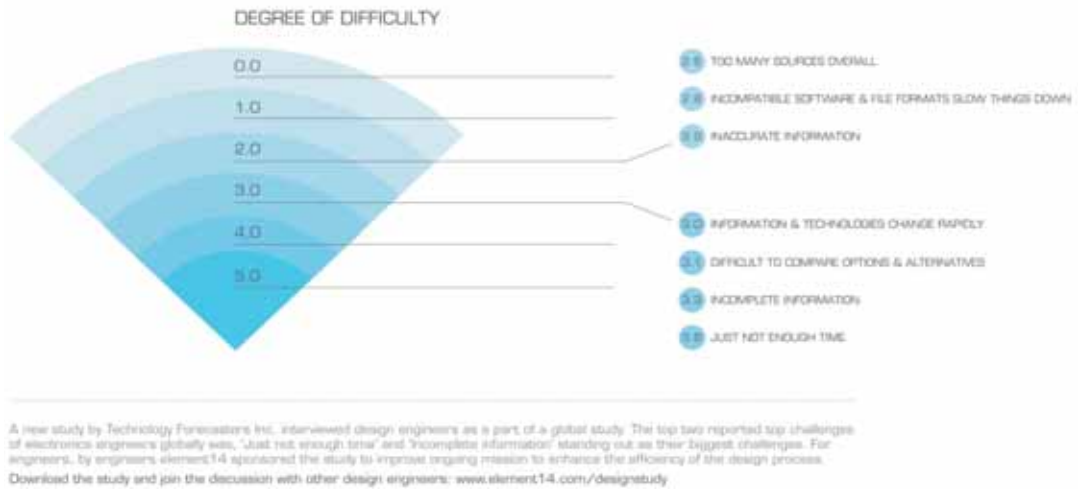

Engineers rely most heavily on four primary sources of information: search engines, manufacturers' websites, distributors' websites and increasingly on professional social media tools.

Making Sense of the Information Overload

An Internet search is typically the first step in seeking out information. However, contrary to the general perception, search engines do not always present unbiased results. It's important to be aware of the bias that search engine optimisation can create. Nonetheless, to overcome the design challenges, engineers spend 50% of their time looking for information and tools online, coupled with the remaining time spent talking with vendors, customers and using internal tools.

A majority of respondents cited that the earlier stages of design are the most challenging, with an average of 41% of design time spent on concept development. For the earlier stages of the design process, designers may use up to a few dozen sites, while the final two stages usually require fewer than 10. The range of online sources consulted is highly variable, based upon the requirements of the design as well as the habits of an individual engineer.

ELEMENT14

LAUNCHED IN 2009, ELEMENT 14 IS THE FIRST COLLABORATIVE COMMUNITY AND ELECTRONICS STORE FOR DESIGN ENGINEERS AND ELECTRONICS ENTHUSIASTS. It distributes technology products and solutions for electronic system design, maintenance and repair and in the process brings together the latest products, services and development software. In addition, it is an innovative online engineering community where purchasers and engineers can access peers and experts, a wide range of independent technical information and helpful hints and tools.

Created by engineers for engineers, the Knode on element14 is an intelligent online search and knowledge tool that helps to select and buy the right solutions for a design, right from the start. Whether researching a new technology, designing an electronic product, or looking for parts to repair an existing system, element14 promises to find the answers and parts needed to keep a project on the fast track from the beginning.

The element14 community and the Knode address many of the top challenges facing engineers globally, helping save hundreds of hours in the design process.

Furthermore, specialised information as well as performance failure rates and component lifecycle data are particularly difficult to collect. The four toughest types of information to get are:
(1) Reference designs;
(2) Application notes/technical papers;
(3) Simulation models; and
(4) Component pricing and availability.

The challenge of finding regulatory information is also becoming more and more critical. Regulations concerning hazardous materials, energy use and other environmental considerations are increasing in number and differ greatly by country. Increasingly design engineers are expected to have some competence in sorting through the complicated set of applicable regulations.

The Promise of Social Media

With so much information at hand, credibility of the information source assumes significant importance in the design process. In the survey, more than 70% of design engineers said they relied heavily on online forums, blogs and engineering communities to collaborate with peers and share insight on components and design processes. Technical forums, blogs and engineering communities offer the potential to provide a useful new dimension to how engineers find, aggregate and utilise critical information for their designs. They provide a direct link between engineers, including those working at different stages of the design cycle and across different industries. They especially appeal to the younger generation of engineers who prefer social media over other forms of communication in general.

It is evident that the Internet provides a vast quantity of information and tools that are now essential for the electronics design process, but engineers are challenged to integrate all this data effectively within the short time-frame they have to complete projects. A lack of consolidated online tools and databases hinder their ability to make accurate comparisons. Some 80% of design engineers said they would prefer online resources that: (1) cover multiple stages of the design cycle and (2) provide consolidated information across vendors, distributors, publications and other players in the supply chain.

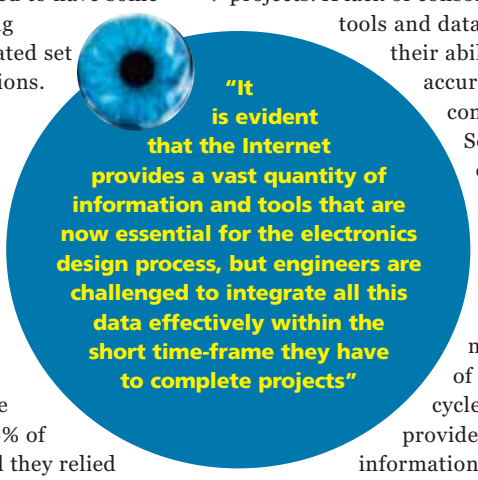

The Silver Lining for Design Engineers

The study underscores design engineers' unrestrained call-to-action for better ways to access the information they need. The design process can be enhanced by giving engineers consolidated design resources, increased collaboration with peers worldwide who are working on similar stages of the product design lifecycle and access to reliable search sources.

"I'd like to search by function of the component, or to a lesser degree by application," said one engineer in the survey. "I don't like application searches, because they assume that only particular functional devices pertain to a given application, and it's too restrictive of my ability to use my imagination and creativity. They might say 'only this memory is appropriate for automotive'. Instead, I'd like to see the parameters: size and type of memory, package type, cost, DRAM, SRAM, etc., with check boxes to select the parameters that I'm interested in. By searching for parameters, you can zero in on what you want."

Intelligent Consolidation of Information

It may not be soon that a single resource aggregates all component information, design tools, regulatory data, parts availability, pricing and other relevant data that the engineers report they need to ensure a successful start to the supply chain. In fact, from studying engineers' processes around the world across a variety of company types and sizes, it is clear that engineers would not want to start and end their search with only one resource; some engineers explicitly stated that they would not trust a single "go to" entity. But by consolidating necessary and even "nice to have" information in a few objective, well-run search communities, engineers' processes will be more efficient and decisions more exact. The end result will be better products going to market, faster cycle times, fewer design changes and fewer surprises about cost, availability and regulatory compliance. ●

Pulling Up the Ladder

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

The loss of analogue TV broadcasting in the UK is no longer news. It is an established fact and, with the availability of low-cost receiving equipment for the digital terrestrial and subscription-free satellite services, this loss is little mourned.

The issue I wish to highlight here relates to the analogous switch-off of the analogue radio broadcasting services, as DAB and internet radio services become more widely adopted. This cut-off has been mooted several times in recent years, with service end dates ranging from 2015 to "after 2019", depending on the particular political mouthpiece or industry pundit speaking.

In my opinion the current digital audio broadcast service is a valuable addition to the existing spectrum of radio services we currently enjoy in the UK, but the idea that it could replace analogue broadcasting fills me with horror. My reason for this is not related to the already widely discussed issues of coverage, cost of receiving equipment or the fundamental quality of the received audio signal; my concerns are related to education.

Let me explain: I am an electronics engineer, specialising in radio frequency circuitry. I design radios and have done so for a very long time. The skills necessary to follow my career cannot be learned purely from the theory, in a classroom. There is a huge amount of practical, on the bench tinkering, which can only be learned by years of practical experimentation.

To follow this path, it is not enough to make a decision in the final year of, usually mostly theoretical, electronics degree "to be an RF engineer". You already need to be an enthusiast, an "amateur", an experimenter, and that is something that starts in childhood.

Like most engineers of my type, I started early with "crystal sets" and simple discrete transistor circuits, slowly increasing the sophistication of my receivers until (long before the GCSE examinations loomed) I had a good practical grounding in the practicalities of "radio circuits". My childhood knowledge was by no means cutting edge, and much of the time I had spent was applied to re-inventing things which would have looked old-fashioned in the 1940s, but it was fascinating and it gave me a fundamental interest in, and grasp of, the subject. There is a unique satisfaction to be enjoyed when listening to a radio station on a receiver that you have constructed yourself, from basic components.

Without analogue broadcasting (medium wave and later band II FM) such simple juvenile experimentation is quite simply impossible. An AM crystal set can be built with less than ten (mostly home made or improvised) parts, while a simple FM receiver kit might include less than thirty simple components. The complex digital coded OFDM transmissions of the DAB system cannot be received by such elementary, even "primitive", circuitry at all. There is no "entry level" to this sort of receiver design: even the simplest digital receivers are built around highly complex custom made integrated circuits, not available to the general public. Even with a complete, fully trained, engineer's knowledge the task of putting together a bench-prototype of a DAB radio is daunting. Without access to very specialised parts, I would need at least the processing power of a desktop PC to even decode the signal into audio. Even then, the specific "codec" algorithms used in the decoders are not public domain – they are "owned" by specific corporations.

By taking away analogue broadcasting, you take away the mechanism by which future radio engineers (and military signallers, and amateur radio operators) first encounter the technologies we need them to master, if the future generations of engineers our society needs are to thrive.

To find an analogy: it is as if we were to prohibit children from riding bicycles, while still expecting those same children to grow up into professional motorcycle-racers.

By taking away analogue broadcasting, you take away the mechanism by which future radio engineers first encounter the technologies we need them to master, if the future generations of engineers our society needs are to thrive

EMBEDDED INSTRUMENTATION – CLOSING THE T&M GAP FOR BREAKTHROUGH TECHNOLOGIES

EMBEDDED INSTRUMENTATION IS FILLING A GAP IN TEST AND MEASUREMENT THE ELECTRONICS INDUSTRY HAS HAD FOR SOME TIME. BY **REG WALLER**, EUROPEAN REGIONAL DIRECTOR AT ASSET INTERTECH

Since the dawn of the electronics industry there's been a gap between technological advancement and the test and measurement (T&M) equipment needed to test it. This gap is finally closing, thanks to embedded instrumentation.

Typically, technology advances by becoming faster, more complex, more integrated and capable of performing a wider range of tasks. Processors have become smaller and denser, chip-to-chip buses faster, and circuit board design and fabrication techniques much more complex.

The recent rise of embedded instrumentation has shown that certain segments of the industry are realizing that T&M requirements should be considered an integral part of any breakthrough in technology. Doing so accelerates the adoption of new technology by reducing the risks and speeding up return on investment associated with it.

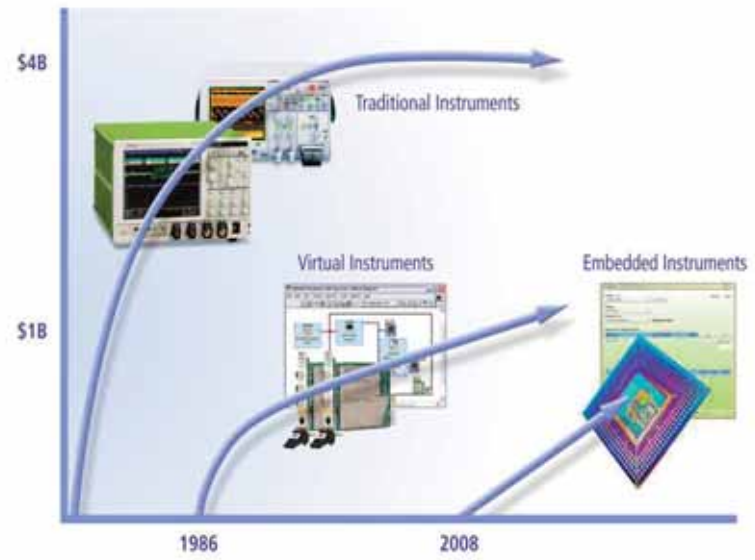

Figure 1: T&M is moving away from external probe-based instruments toward non-intrusive instruments embedded in chips and on circuit boards

Embedding T&M

Embedded instrumentation can be defined as the logic designed into semiconductors, used for design validation, test, debug, yield analysis and other T&M activities.

By its very nature, embedded instrumentation is a non-intrusive technology. Instead of physically placing a probe on a PCB test-pad or a chip pin, embedded instruments are able to launch test or validation routines from within a chip. When test results are reported, the embedded instrumentation typically accesses an external validation/test platform through a bus. In most cases today, this access technology is usually the boundary-scan interface as defined by the IEEE 1149.1 Boundary-Scan Standard, also known as JTAG.

Until recently much of this embedded instrumentation was intended for a narrow task or purpose. For example, a chip vendor might decide to embed an instrument to debug the device during manufacturing and then render the instrument inaccessible to the board designer or manufacturing test engineer. Or an instrument designed into an ASIC for board test purposes might be abandoned once the board was deployed in the field where support personnel might be interested in employing it in trouble-shooting applications.

Such a limited view of embedded instrumentation is changing today as the industry is realizing that it is much more effective and cost-efficient to migrate validation, test and debug routines from chips to circuit boards and eventually into systems in the field, than it is to develop new tests at each phase in the life-cycle.

The benefits of embedded instrumentations are cumulative. That is, they accrue over the entire life-cycle of a system, beginning with the initial stages of chip design and culminating in the field.

Why is Embedded Instrumentation Needed?

Embedded instrumentation has become critically important in the industry as intrusive external or modular testers which rely on probing a device or a circuit board are losing their effectiveness. They simply can't provide the results they once did.

Software-based embedded instrumentation which tests from the inside out rather than from the outside in is much

more cost-effective, efficient, agile and better suited to a wide range of today's emerging computer and communications technologies. A few examples best illustrate the increasing inadequacies of traditional external testers.

The rapidly accelerating speeds of PCI Express, Fibre Channel, 100Gb Ethernet, InfiniBand, Intel's Quick Path Interconnect (QPI) and other high-speed serial buses can wreak havoc on the accuracy of results obtained by an oscilloscope. When a scope's probe is placed on a test pad on one of these high-speed buses, the probe itself introduces capacitance anomalies onto the bus. The validation or test engineer can't tell the difference between an instrument-induced anomaly and a fault in the design or manufactured assembly. To further exacerbate the problem, semiconductor vendors are designing in signaling enhancements such as pre-conditioning, pre-emphasis and equalization to help move signals at higher frequencies. Unfortunately, these techniques make accurate measurements even more difficult for traditional instruments such as oscilloscopes and logic analyzers.

Embedded instrumentation is able to provide validation, test and diagnostic data that traditional intrusive instruments cannot. For example, Intel has developed embedded instrumentation methodology called Interconnect Built In Self Test (IBIST). Tools that make use of IBIST are able to stress and, therefore, test high-speed I/O buses well beyond the capabilities of traditional OS-based testing. In addition, at any one time, traditional measurement techniques typically only measure signal integrity margins on one or a few of the high-speed serial lanes on chip-to-chip interconnect buses on a PCB.

Because it is embedded into chips and onto circuit boards with those chips, IBIST can test and measure all lanes on all buses concurrently. This makes the test more robust and complete, and it drastically reduces the amount of time required to validate a system design. Validating the signal integrity of the high-speed buses on a circuit board could involve several tools, such as pattern generation and checking, bit error rate testing (BERT), and/or margining tests (see Figure 2).

In addition, sub-100 nanometer chips have demonstrated problematic parametric performance, including thermal conditions, timing issues, clock propagation, power distribution and others. Traditional testing techniques are coming up short at identifying these problems. External instruments at the corners of a chip cannot identify the rampant variations across the chip. Only on-chip instruments can effectively monitor parametric characteristics.

> Embedded instrumentation is a non-intrusive technology: instead of physically placing a probe on a PCB test-pad or a chip pin, embedded instruments are able to launch test or validation routines from within a chip

Using Embedded Instrumentation

The concept of a tools platform or a unified tool environment is most appropriate for embedded instrumentation since different tools may be needed at any given point in the life-cycle of a chip, circuit board or system. As mentioned, Intel IBIST might validate the performance of a design's high-speed buses on prototype boards prior to volume manufacturing. These IBIST tests could then be redeployed into manufacturing along with other types of tests, such as boundary-scan (JTAG) test and processor-controlled test (PCT) in a manufacturing test strategy.

In board test applications, PCT takes advantage of the most powerful instrument on most circuit boards, the processor. PCT tests are executed when control of the processor is temporarily given over to the PCT tools. The board's CPU then reads and writes to memory and input/output (I/O) registers on the board's addressable devices. In this way, PCT exercises the

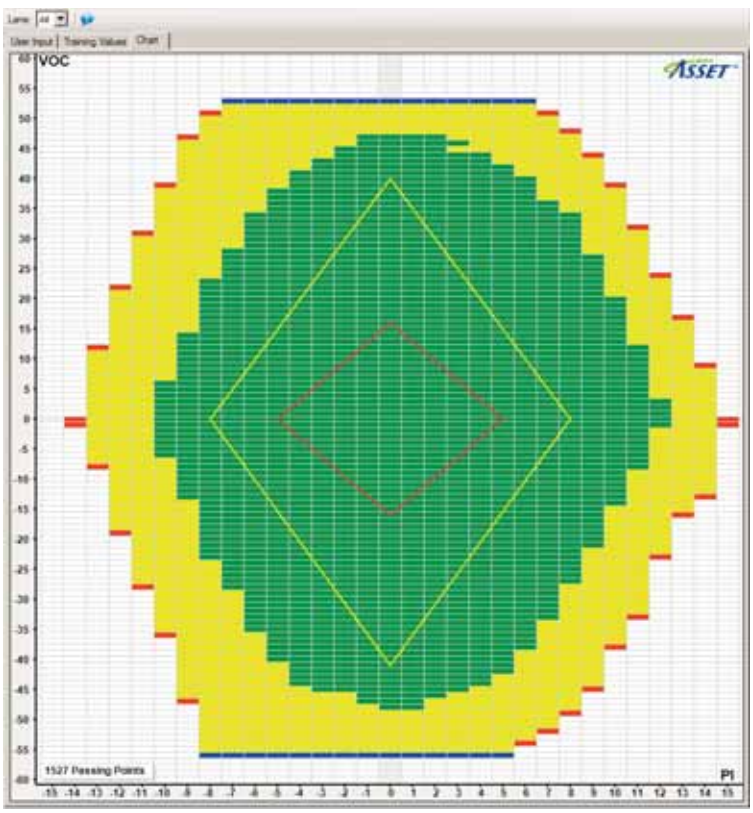

Figure 2: A margin diagram generated by embedded instrumentation attests to the signal integrity of a high-speed I/O or memory bus

Figure 3: Bi-directional TMSC combines TMS and TDI/TDO

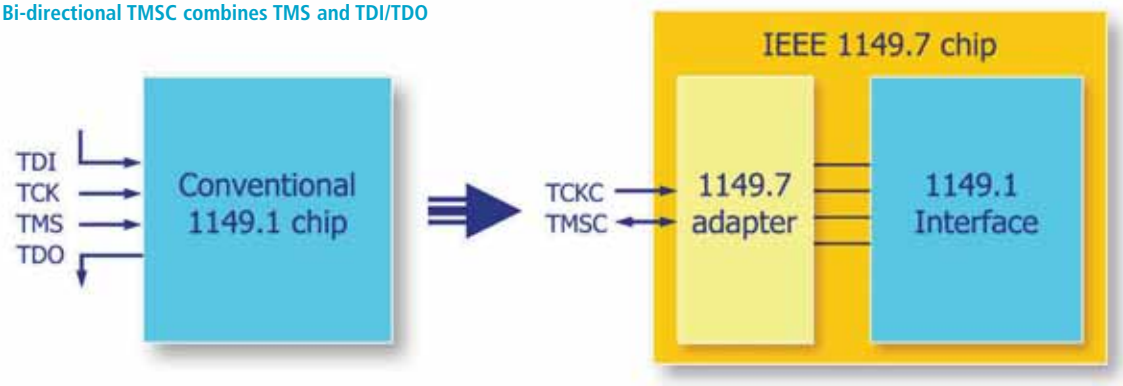

functionality of the board and, as a by-product, detects and diagnoses structural faults like opens and shorts on the traces.

PCT is an 'at speed' test technology, meaning it will detect faults which are manifested when the board is running at operational speeds. In contrast, boundary scan is a static test technology. Taken together, these two embedded instrumentation technologies provide test coverage for a broad spectrum of structural and electrical faults.

Embed Your Own Instruments

Much of the infrastructure surrounding embedded instrumentation is based on industry standards. As a result, tools vendors are not limited solely to the instruments embedded by chip vendors at manufacture. Suppliers like ASSET InterTech are adding considerable value to the entire T&M lifecycle by developing innovative applications of embedded instrumentation.

One of the foundations for embedded instrumentation is the IEEE 1149.x boundary-scan family of standards. In addition to its ability to detect, isolate and diagnose structural shorts-and-opens faults at the board level, the original IEEE 1149.1 Boundary-Scan Standard (also known as JTAG) provides non-intrusive access via a four-wire interface to instrumentation embedded in chips and on circuit boards. Developed in the 1990s and widely adopted by the industry, boundary scan serves as a robust and well understood structural test and infrastructure technology.

The boundary-scan family of standards has been expanded several times with new specifications that enhance and extend the original document. IEEE 1149.6 is a methodology for testing high-speed AC-coupled chip-to-chip interconnects on circuit boards. Another addition to the family, IEEE 1149.7, builds on the original specification but offers the option to reduce the footprint from four wires to two. It also adds functional and architectural features for testing system-on-a-chip (SoC) devices and other types of chips like multi-die 3D components with vertical through-silicon vias (TSV).

> Processor-controlled test (PCT) exercises the functionality of the board and, as a by-product, detects and diagnoses structural faults like opens and shorts on the traces

The provisions for multi-core and multi-die architectures in IEEE 1149.7 complement quite effectively the capabilities being defined in a preliminary standard under development, the IEEE P1687 Internal JTAG (IJTAG) standard. IJTAG is being developed specifically to support embedded instruments. P1687 IJTAG standardizes the interface to embedded instruments, defines an architecture for on-chip instruments and specifies many of the mechanisms for managing them.

Since embedded instrument IP (intellectual property) can come from a number of sources, such as chip suppliers, third-party providers, EDA tools or in-house design groups, P1687 is intended as a standard way of connecting, analyzing, describing and using embedded instrumentation regardless of where the instrument has come from.

In chip test-and-debug applications, IJTAG technology (see Figure 4) embedded

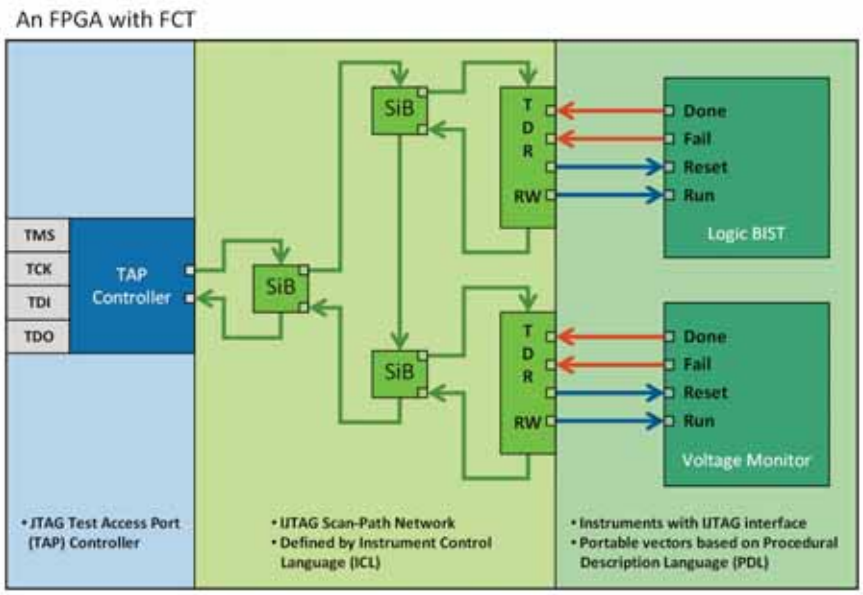

Figure 4: The IEEE P1687 IJTAG standard defines the interface to embedded instruments as well as an on-chip architecture for accessing and automating their application

in various types of chips, either logic chips or field programmable gate arrays (FPGA), could automate and schedule the parallel operation of multiple embedded instruments. An embedded logic BIST engine for chip test might be operated simultaneously with a voltage monitor intended for yield analysis. The resulting simultaneous operation of these two embedded instruments could determine whether failures identified on chip-level automatic test equipment (ATE) or board-level tests correlate with voltage starvation.

Agile T&M

Deploying an IJTAG architecture and multiple embedded instruments either temporarily or permanently in an FPGA demonstrates the versatility and agility of embedded instrumentation. Such a strategy can be employed in chip test and debug, or board test applications. In fact, an entire board tester might be temporarily inserted into an FPGA to extend the test coverage of software-driven non-intrusive board test (NBT) and validation methodologies.

Designers would find this approach particularly effective for bringing up prototype circuit boards prior to a new design's transition into volume manufacturing. First prototypes often are delivered before the board's support package or FPGA's firmware have been completed. At this point, the structural, functional and performance capabilities of the prototype hardware must be tested to validate the design in preparation for software integration. Previously, functionally validating the hardware was severely limited without the board's firmware or OS in place. As a result, development schedules might be delayed while software was being completed. Now, an FPGA-based deployment of embedded instruments offers an alternative.

A network of IJTAG embedded instruments in the FPGA could function as a tester to validate, test and debug prototypes. In addition, this embedded tester might be needed later in the lifecycle of the circuit board. If so, then some or all of the embedded instruments might remain in the FPGA to perform manufacturing test, ongoing fault analysis, remote diagnostics or troubleshooting by field service.

Filling Gaps

Embedded instrumentation is indeed filling a gap that the electronics industry has had for some time. Specifically, attitudes toward validation, test and debug – the mainstays of T&M – are changing. Addressing test requirements sooner during the development of new technological advancements is a huge breakthrough, as are the vastly improved levels of validation, test and debug coverage made possible by non-intrusive software-based embedded instrumentation. ●

DESIGNING LOW-POWER EMBEDDED SYSTEMS

WOJCIECH GELMUDA AND **ANDRZEJ KOS** FROM THE AGH UNIVERSITY OF SCIENCE AND TECHNOLOGY IN POLAND LOOK AT HOW ENERGY DEBUGGING TECHNIQUES AND TOOLS CAN HELP A DEVELOPER CREATE MORE POWER-EFFICIENT BATTERY-RUN DEVICES

The consumer electronics market is flooded with different types of mobile devices and all of them have at least one thing in common – they run on batteries. But for several decades we have not seen a major breakthrough in batteries that might revolutionize the way mobile devices run, for example when they can operate for months or years without the need to recharge. Therefore, limited battery capacity-to-weight ratio is the mobile devices' bottleneck.

To overcome the insufficient battery capacity issue, new lines of low-power electronic components are being presented by manufacturers worldwide. These products have low-voltage power supplies and need less current to run. ARM, for example, offers the Cortex-M series of cores for energy-efficient microcontrollers to be used in embedded devices. Still, such new electronic devices, with their low leakage currents and sophisticated sleep modes, do not entirely solve the insufficient battery capacity issue.

Most mobile devices nowadays, and in particular smartphones and tablets, employ multicore microcontrollers and many sensors. Even though most of them are energy-efficient devices, the sum of their power consumptions is usually not so low.

A low-power embedded system developer has to have extensive knowledge of existing electronic components to choose the best ones for the design to meet the proposed functionality. The battery used in the project has limits in terms of its capacity and maximum output current. With the given components and battery types there is one more thing the developer can do to prolong system working time on one battery charge and that is to write the code for the MCU that drives the whole system in a way that all the energy-saving functionality of the used components is efficiently utilized – in other words one has to use energy debugging.

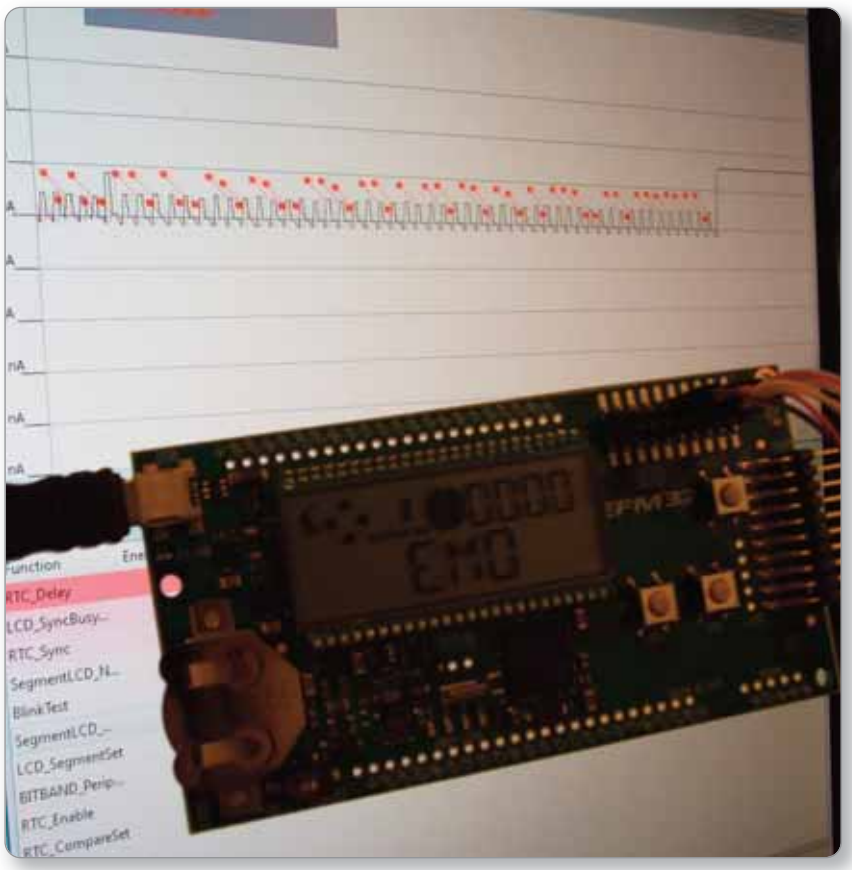

Figure 1: One of Energy Micro's Starter Kits with the energyAware Profiler software running in the background

Energy Debugging

The easy way to determine the power consumption of an electronic device is to measure the voltage drop across a small resistor in series with the main DC power supply. With the known value of this resistor, the supply current can be calculated.

There are some power monitor chips available, for instance, LTC2990 from Linear Technology. Using this type of monitor is the proper approach when the system being designed does not have many different work modes – in other words, the power consumption does not change often in time – and the embedded software engineer is able to determine when a part of the system was turned on and consumed this extra current, which was measured by the power monitor.

Another example where the mentioned approach can be useful is when one wants to determine how much power the system consumes over a long period of time (for instance a day or a week) to select the right

battery capacity for it. Some reduction in overall power consumption can be accomplished by employing this kind of power monitor. Nevertheless, when the system being developed implements a fast microcontroller and the current is being drawn from the battery supply in a non-deterministic way (due to the synchronous and asynchronous signals and interrupts), another approach has to be used in order to eliminate short- and long-term unnecessary power consumption. Writing and debugging a code for the microcontroller that drives the whole system has to be linked to the real-time current consumption of the system. Thus, the energy consumption has to be debugged.

Developing and Analyzing the Code

Years of working with low-power embedded systems have taught us that developing a code optimized for a specific microcontroller's current consumption is crucial. In some of our projects we used chips and tools from Energy Micro. Microcontrollers based on the ARM Cortex-M cores are suitable for our mobile applications. All of Energy Micro development boards have the J-Link debugger and Advanced Energy Monitoring (AEM) on board. The AEM tool makes it possible to gather information about current consumption from the microcontroller VMCU power domain.

Current sense amplifiers measure the voltage changes across a small resistor between the LDO output and the VMCU domain. Depending on the current consumption range, the measurements are 1µA and 0.1mA for ranges below and above 250µA respectively. For small current consumptions up to 10µA the monitor can detect changes of even 100nA.

> At the early stages of prototyping, algorithms can be optimized and tailored for the specific system architecture to enhance their operation and, at the same time, made to consume as little power as possible

The data gathered by the AEM can be presented in real-time on a TFT display available in some of the Energy Micro kits or on a computer screen via the energyAware Profiler application provided free as part of the Simplicity Studio software.

The most helpful AEM functionality is its cooperation with the debugger. The developer can test and debug the microcontroller code and then compare the power consumption over time against the code. If a sudden power consumption peak occurs on the graph, it can be pinpointed to the specific location in the code. As such, the developer can detect what caused this increased current consumption and decide if it is normal or not, and then decide what next.

The developer can analyze the code line-by-line and disable any subsystems that are not needed at certain times to preserve power. There are also labels attached to some data points in the plot. They indicate when specific, important events happened, for instance interrupts, to energy debug the system faster and more efficiently.

The energy debugging technique is important at all levels of system development. At the early stages of prototyping, when the code is written for the microcontroller on the development board, algorithms can be optimized and tailored for the specific system architecture to enhance their operation and, at the same time, made to consume as little power as possible. Next, when the code is being developed for the final system board, it is possible to energy debug the system with the use of a debug interface and by connecting the VMCU domain power supply as the main power supply of the final system board. However, the maximum current consumption of the VMCU domain is limited and it can be less than the board needs. Nevertheless, in most low-power embedded systems it is sufficient. Another approach to overcome this issue is to split the main power supply on the final board into several power domains.

Example

At the AGH University of Science and Technology in Poland one project deals with navigation assistants for blind people (see Figures 3 and 4). Since holes and imperfections in pavements and roads can be dangerous while navigating blind people,

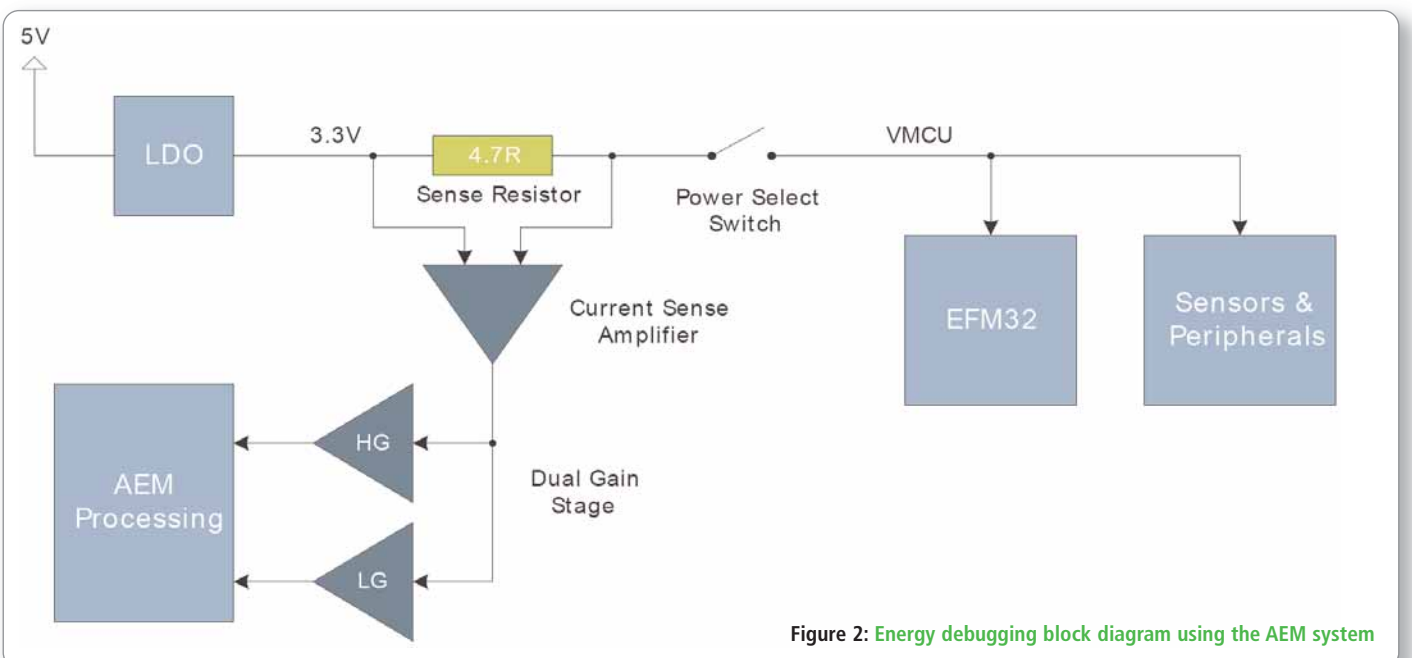

Figure 2: Energy debugging block diagram using the AEM system

Figure 3: Block diagram of the navigation assistant being developed at the AGH University of Science and Technology in Poland

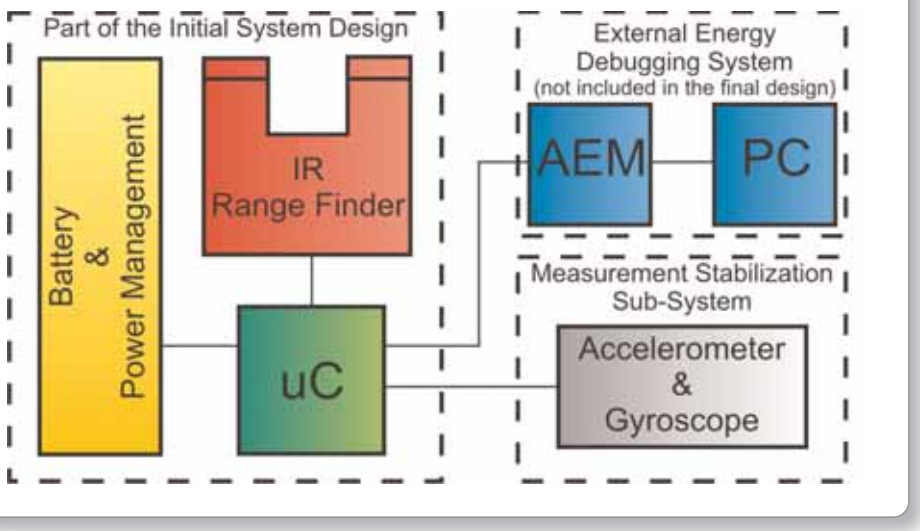

one task was to design a device for detecting and informing the blind pedestrians of road imperfections.

The device was built based on a low-energy microcontroller from Energy Micro with a Cortex-M3 core and an infrared range finder for scanning the area in front of the user. Since the device is mobile, the code for the microcontroller was developed to use all possible power-saving functionalities, so the battery used in the device is light and ensures a day's operation.

The preliminary tests showed the range measurements had a large error coefficient due to the user's movements. These errors have been so great that it was impossible to detect the dangerous 5-10cm deep holes with high reliability. Thus, a stabilization subsystem had to be implemented. For that, an accelerometer and gyroscope were used.

Using adaptive filtering and integration the velocity are calculated. The analysis of the pedestrian movement scheme showed that, based on detecting points in every step where the velocity equals zero, the algorithm for the infrared measurements stabilization could be created. Unfortunately, adding extra components to the existing system with a previously-selected battery, DC/DC converters and LDOs, caused problems like unexpected microcontroller resets, incorrect data being read from the sensors and a noticeable reduction of working time between battery recharging. After elimination the possibility that these could be caused by faulty components, insufficient maximum output current was the obvious cause.

Figure 4: Prototype of the navigation assistant for obstacles and road imperfections avoidance

Power Awareness

Working closely with the energy debugging software and tools for the Cortex-M3 microcontroller and going over the lines of the code made detection of the sections where the output current was maxed out was both relatively easy and accurate. This also confirmed the insufficient output current assumptions.

> The power consumption debugging makes it possible to eliminate both software and hardware errors at the early stages of the design

Making changes in the code and adapting the stabilization algorithm to the system architecture made it possible to eliminate any unexpected resets and faulty data readings. The energy debugging technique also helped with better optimization of the current consumption, thus enhancing the energy-efficient functionality of the components used, which led to increasing the working time to an acceptable level.

The optimized stabilization system reduces the errors caused by the user's movement such that even 5cm holes or bumps on the road – as well as stairs – can be detected with great reliability.

Nowadays, where time-to-market is so important to all companies, electronic systems developers should not only follow news of low-energy chips but also news of energy debugging tools and software. Power consumption debugging makes it possible to eliminate both software and hardware errors at the early stages of the design.

Furthermore, when a device is already on the market and customers notice disappointingly high power consumption, maybe some of the problems can be detected with energy debugging and quickly eliminated by releasing new firmware, which is usually much cheaper than even the smallest hardware change. That is why understanding energy debugging techniques and being aware of the available tools is a must for every electronics engineer who is developing mobile and power-efficient devices. ●

Now available from RS Components

New family of digital Multimeters from Keithley, offering a wide range of price and performance options – the perfect balance of high precision and affordable cost.

FIND IT.
DESIGN IT.
BUY IT.

Find it online at:
www.rs-components.com/keithley

RS

EMBEDDED PROCESSORS GET SMARTER AS ENERGY EFFICIENCY TAKES CENTRE STAGE

JO STEIN HAUGEN, APPLICATION ENGINEER AT ENERGY MICRO, GIVES A PRACTICAL GUIDE TO TACKLING ENERGY EFFICIENCY IN DESIGN THROUGH SELECTING THE RIGHT MCU

Figure 1: Smart meters need to operate for many years on a single battery

Energy efficiency is increasingly being recognised as a key performance attribute of many electronic systems and in particular of embedded designs. In products ranging from smart meters to wireless sensor nodes and from mobile health monitors to intruder alarms, designers need to take every step possible to limit energy consumption.

The detailed operational and design requirements of such applications may vary, but they share a common feature: they need to operate reliably over a period of years or even decades, using the limited energy resources available from batteries or an energy harvesting system. At the heart of designers' efforts to achieve this are a new generation of microcontrollers (MCUs) specifically designed with energy efficiency in mind.

It is worth considering in more detail the scale of the design challenge involved. Many of these microcontrollers must deliver real-time (or close to real-time) performance throughout a service life that is spent predominantly in a low-power sleep mode that inherently reduces their speed of response to external events. A device such as a smart meter will often be buried or otherwise difficult to access, requiring 10 to 20 years of operation on a single battery.

As a result, energy budgets in such applications are aggressively trimmed. Given the current state of battery technology, it is not unusual to see an energy requirement for the lifetime average current draw to be between 5µA and 20µA. If a system designer budgets 20% of the available battery capacity for data collection and 'housekeeping' functions, 20% for receiving commands and 30% for transmitting data, then the remaining 30% must support the system's idle mode. This may sound generous, but the meter will likely spend in excess of 99% of its life in some kind of sleep state, whose efficiency becomes absolutely crucial in reaching the design targets.

This is why sleep modes are one of the most important features of a low-power MCU.

The Factors that Make a 'Good' Energy-Saving MCU

In the past there has been a somewhat casual assumption that because of their lower active power consumption, 8- or 16-bit solutions must be the best choice for low-energy applications. But the fact is that a modern 32-bit CPU core like the ARM Cortex-M3 delivers a processing power of 1.25DMIPS/MHz; that is, it can perform a given task in around one quarter of the time required by older 16-bit CPU cores.

Since energy is 'power times time', using the 32-bit Cortex-M3 in preference to a less powerful alternative can prove more energy efficient, particularly if the designer opts for an MCU such as Energy Micro's Tiny Gecko, that even in active mode consumes only 150µA/MHz.

So 'low power' does not necessarily mean 'low energy'.

As observed, an energy efficient processor stands and falls on its sleep modes and the designer's ability to choose the right mode of operation for the job. In a power-constrained application, subsystems such as displays and RF and PHY-layer transceivers will be designed to spend only as much time as necessary in their active state, MCU functions will therefore represent the only "low-hanging fruit" for power optimisation.

Most modern MCUs have more than one low-power mode (these will typically be something like "Sleep/Standby", "Deep-Sleep" and "Off"). The specific CPU, memory and I/O functionality available will vary from microcontroller to microcontroller, but the general principles involved remain the same.

Although it may appear to make sense to use the deepest possible sleep mode in every case, there is in fact a trade-off involved, the deeper the sleep, the longer the wake-up overhead. The overhead consists of hysteresis and propagation delays in the wake-up circuitry, start-up and stabilisation of the power supply and re-starting the CPU clock. The device then needs to restore its registers and other internal configurations, before doing the same for system memory. Only when these processes are completed can the MCU service the event that triggered the wake-up process.

This trade-off has two implications: first, if the MCU is in an overly deep-sleep mode, it may not be able to wake up in time to deal with an external event; second, during the wake-up phase the device doesn't do any useful work. So waking too frequently from a deep sleep mode in order to perform a small amount of processing could conceivably burn more power than staying an a "higher" activity state with less time wasted on the wake-up process.

A smart meter will likely spend in excess of 99% of its life in some kind of sleep state, whose efficiency becomes absolutely crucial in reaching the design targets; this is why sleep modes are one of the most important features of a low-power MCU

Dealing with Events

If there is a need to deal with unpredictable events it is much more difficult to reduce power consumption than for scheduled events such as real-time clock interrupts, or the time-slotted communications processes common in ZigBee and other mesh protocols, that require the processor to check for messages from other network nodes at a certain time.

Unfortunately, most applications involve events of the less predictable type and require the MCU to respond in real time or almost real-time. Examples include threshold/alarm conditions, power interruptions and wireless protocols that can transmit and receive at random intervals. In such situations, failure to wake within an appropriate timeframe may mean losing critical data or failing to respond to a user command. In the worst case, there may be a risk of system damage if an alarm condition is not dealt with in a timely fashion.

The task for the designer, therefore, is to find the lowest-power sleep mode that provides adequate response to anticipated (if unscheduled) events. Energy Micro's EFM32 Gecko family of energy-optimized 32-bit MCUs provides a good example of the options available. Although the Gecko architecture is optimized to perform beyond the basic requirements of each mode, most MCUs used for these types of applications have similar operating states and characteristics (Figure 3).

At the lowest level, most MCUs have a mode designated as "off" or shut-down mode, that preserves the minimum functionality needed to trigger wake-up from an external stimulus: in our example of the Gecko series this is designated EM4. In this mode, the entire device is powered down, other than the interrupt monitoring circuitry on the reset pin and GPIO pin wake up. The EFM32 draws around 20nA of current in this mode, although a typical 32-bit MCU would require nearer 1.5µA.

Restarting from such an "off" mode is essentially a device reset, a process that for the EFM32 takes around 160µs. Main memory contents will have been lost and

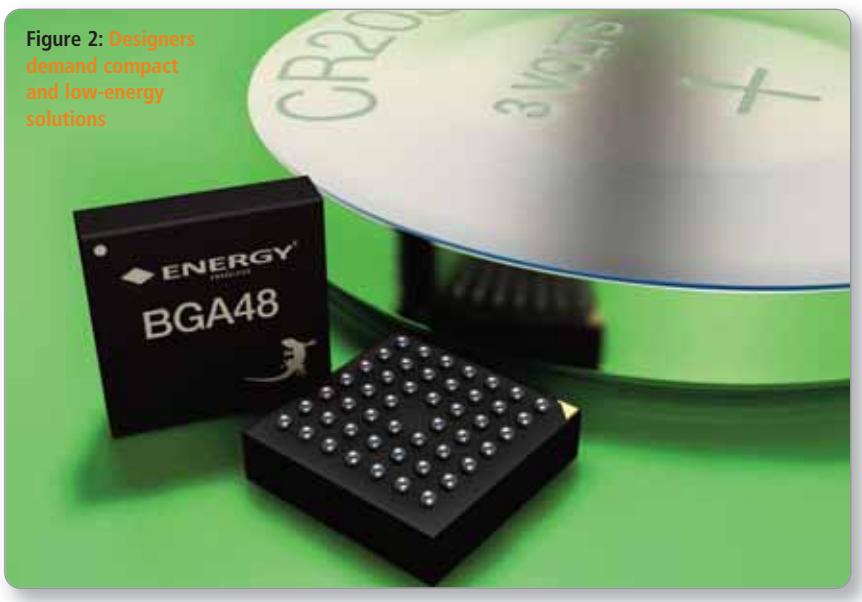

Figure 2: Designers demand compact and low-energy solutions

must be re-loaded. Some processors include a small (512 byte in the case of the EFM32) block of memory whose content is preserved for use on start-up.

A kind of 'mezzanine' state above EM4 preserves a few more critical functions, in particular the real-time clock and 512 bytes of backup memory. Drawing only 400nA, this state would consume the capacity of a pair of AA cells in roughly seven years. Although intended primarily for use in the event of power supply failure, it can be an excellent alternative when RTC and a faster wake-up time are desirable.

With more functionality still, 'stop' mode (EM3 in Figure 3) enables a limited degree of autonomous peripheral activity and faster wake-up. Here, the high- and low-frequency oscillators are disabled, but the MCU's full RAM and configuration register states are preserved. In addition to the elements active in EM4, the power-on reset and brown-out detector are active, and the CPU can be woken up by an asynchronous external interrupt or via a number of internal sources, such as the device's analogue comparators (ACMP) and pulse counter (PCNT).

For a designer, the key specifications to note in this mode are the time to wake up – which should be of the order of a few microseconds – and the absolute power consumption (10 to 30µA for a typical MCU, 590nA for a best-in-class device such as the EFM32). Perhaps the most important point for the designer to remember is that some MCUs do not include full internal memory retention in this mode.

> Waking up too frequently from a deep sleep mode in order to perform a small amount of processing could conceivably burn more power than staying an a "higher" activity state with less time wasted on the wake-up process

Deep sleep mode (EM2 in Figure 3) leaves all of the MCU's critical elements active, while disabling high frequency system clocks and other non-essential loads. In addition to the EM3 functionality, the 32kHz oscillator used to clock on-chip peripherals remains active. This allows selected low-energy functions including the RTC, watchdog timer and some external interfaces to remain active.

As in EM3, the designer needs to take care, since not every MCU series provides full register and RAM retention. Choosing a device with this capability allows the device to return to active state and resume program execution quickly. Best-in-class figures for current draw in this mode can be as low as 900nA (with RTC running from a precise clock source), with time to return to active mode of as little as 2µs.

Finally, in sleep/standby mode (EM1), the MCU's pointer and configuration register states are fully preserved, eliminating the need to save them on power-down and restore them on power-up. This typically saves hundreds to thousands of instruction cycles for each wake-up. The high-frequency oscillator remains active with the CPU clock tree disconnected, allowing the device to resume execution on the next clock cycle after a wake-up event. High-frequency peripherals – for instance direct memory access (DMA), analogue/digital and digital/analogue converters (ADC/DAC) and hardware encryption – remain active.

Sleep mode current draw ranges from 45µA/MHz to 200µA/MHz, depending on the choice of MCU.

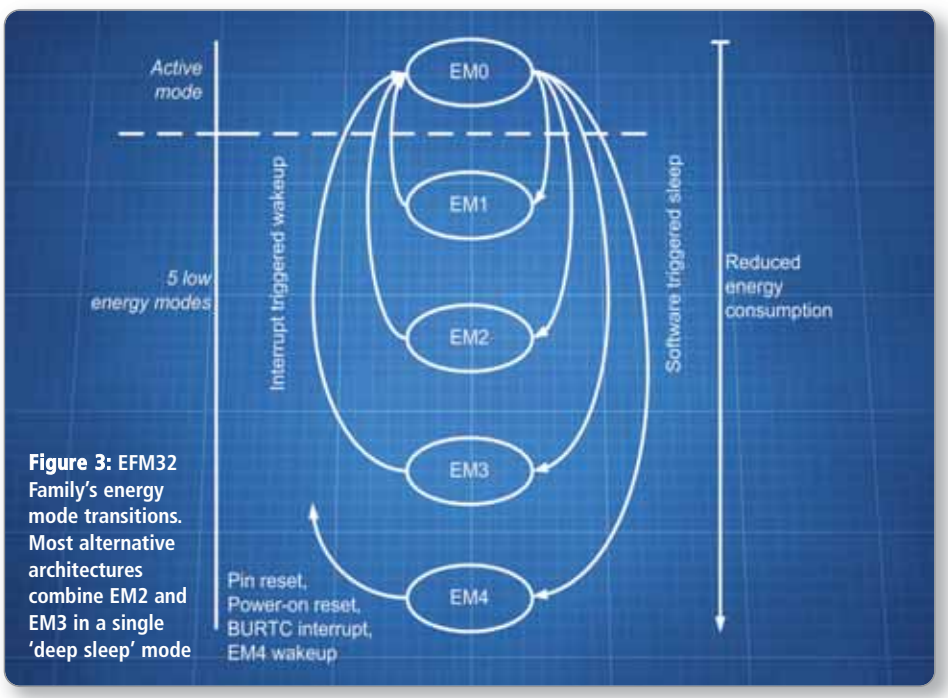

Figure 3: EFM32 Family's energy mode transitions. Most alternative architectures combine EM2 and EM3 in a single 'deep sleep' mode

Selecting an MCU

The process of selecting an MCU and the correct mode of operation for the application is relatively straightforward once these typical sleep mode specifications are understood. The designer first needs to list all of the alert, alarm and interrupt conditions within the application, and identify the one that requires the fastest response. For instance, in a logistics application, a device may perform both GPS-based location reporting and shock and vibration monitoring. The first has a wake-up tolerance that is measured in seconds; but when the MCU's accelerometer inputs exceed pre-set threshold levels, the device will need to wake up and start measuring and recording within a millisecond or less.

The second step is to quantify the exact wake-up time from each sleep mode in the application that is required. This includes not just the time taken to return to a ready state, but also the time needed to execute the code that services the event in question. Executing an interrupt service routine, for instance, will take longer than jumping straight to a block of code dedicated to the specific event.

Comparing the maximum tolerable delays with the minimum attainable service time figures will identify the appropriate sleep mode.

Finally, it is necessary to verify that the sleep mode selected has the appropriate resources: for instance RTC or fast ADC/DAC functionality.

But what if the selected sleep mode still does not provide the required battery life?

Figure 4: Using a faster clock speed to reduce an MCU's execution times can reduce its overall power consumption

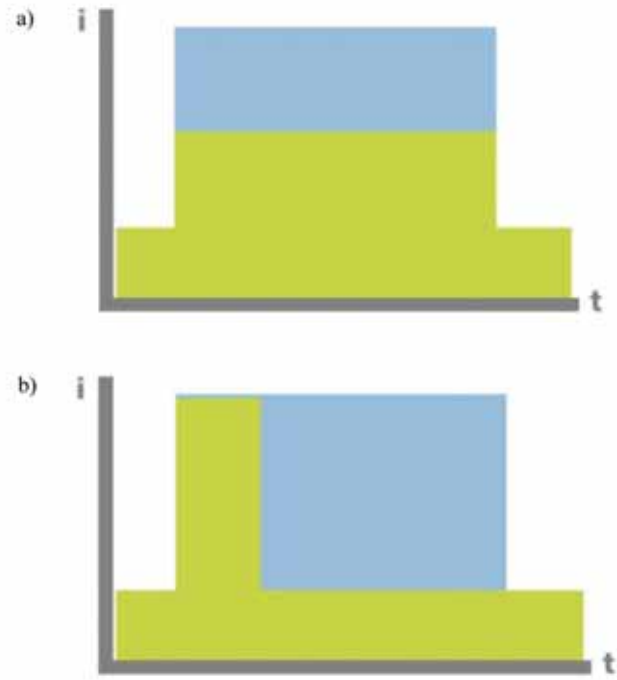

The Gecko architecture takes this approach one step further with a 'peripheral reflex' system. This allows peripherals to be interconnected without CPU intervention, allowing the designer to create quite complex state-machine type structures that can monitor outside events while the CPU is inactive. EFM32 processors also include a sensor interface called LESENSE, which allows autonomous monitoring of up to 16 external analogue sensors, including capacitive and resistive touch pads, and the inductive rotation sensors commonly used in metering products, all while the CPU remains in deep sleep mode.

Finally, it is worth remembering once again that energy is 'power times time'. In many cases, the processor speed that results in the lowest overall energy use may not be the slowest clock speed. Since total energy is the sum of both static and active energy consumption, it may be that increasing the clock frequency, thereby reducing the amount of time the processor spends in active mode, is the most energy-efficient strategy (Figure 4). ●

The emergence of low-energy MCUs has provided a number of tools that address this question.

Modern MCUs often include on-chip peripherals which can function without CPU intervention. In particular, monitoring analogue and digital alarm thresholds without CPU intervention enables the designer to wake the system without the overhead of interrupt-driven techniques.

DEEP SLEEP MODE

STAND-BY CURRENT FOR THE EFM32 IN DEEP SLEEP (EM2) MODE, WAS MEASURED UNDER THE FOLLOWING CONDITIONS: POWER SUPPLY = 3V, AMBIENT TEMPERATURE = 25 DEGREES C.

Typical operating current was measured for the complete microcontroller and any external oscillators, including the RTC (running on the internal precision oscillator), with brown-out detection, full RAM and register retention active.
Current in Deep Sleep (EM2) RTC mode was measured with the RTC running from the internal 32kHz low-frequency oscillator.

Figure 5: This annotated block diagram of an EFM32 MCU illustrates how many of its peripherals remain available in each of its low-power modes

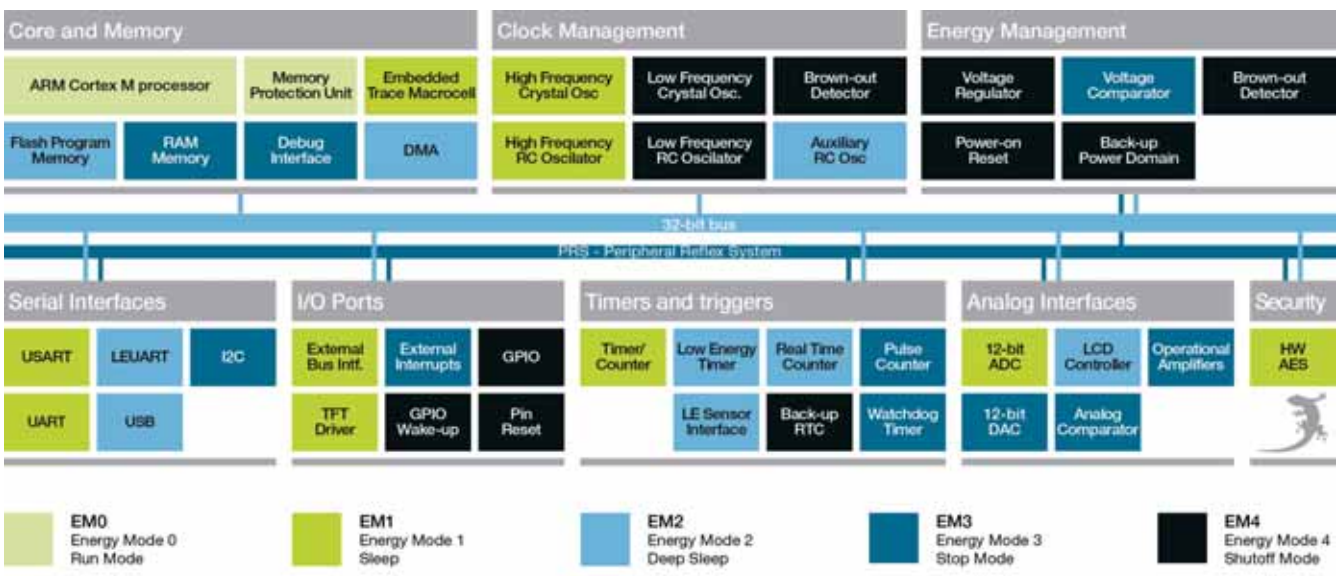

EMBEDDED ETHERNET DEVELOPMENT

PROFESSOR DOGAN IBRAHIM OF THE NEAR EAST UNIVERSITY IN CYPRUS DESCRIBES THE BENEFITS OF EMBEDDED ETHERNET AND PRESENTS AS AN EXAMPLE EMBEDDED ETHERNET-BASED MICROCONTROLLER APPLICATION

Today's embedded systems are small, fast and very powerful. They are commonly used in many diverse industrial and consumer applications in communications, automation/control, entertainment and many more.

Figure 1: Ethernet packet format

Embedded systems are based on microcontrollers. With so many embedded devices being part of systems we use and encounter every day, their connectivity and networking are important issues. For example, we may want to connect a printer to a mobile phone to print pictures; or send those pictures to another mobile phone; send e-mails via wireless routers using mobile phones; and so on.

Several technologies are used to achieve embedded connectivity. Some commonly used domestic solutions are based on the infrared (IR), Bluetooth, ZigBee and Ethernet technologies. This article focuses on Ethernet connectivity in embedded applications.

Ethernet – the Traditional Way?

Ethernet has traditionally been implemented in PCs and laptops and has been widely used in the home and office too, to access the Internet and intranet networks. Today, the Internet can also be accessed using smaller, handheld, devices such as smart mobile phones, PDAs, IPADs and so on, most of which use microcontrollers and single-chip Ethernet controller devices for connectivity. The advantages of embedded Ethernet connectivity using single-chip Ethernet controllers are:

- **Low-cost.** Single-chip Ethernet devices can be integrated into cost-sensitive applications. A typical design will require an MCU (microcontroller unit) and an Ethernet controller chip.
- **Remote monitoring and control.** Once an embedded system is on a network, it can be accessed from anywhere. The user can be located at any location, however remote, and yet control the embedded system with ease.
- **PC software is not required.** In many applications the interface to the embedded system can be achieved using Web tools such as a Web-browser, HyperTerminal and e-mail and, thus, there's no need to develop PC software.
- **Uses the existing infrastructure.** Ethernet is the most widely implemented network and is found nearly in all homes, offices and industry. It is low cost and convenient for attaching embedded systems to existing infrastructure – only a suitable cable is required to make the connectivity.

Ethernet Connectivity

Ethernet was originally invented by Xerox in 1972 and then developed jointly by Xerox, DEC and Intel. It is a frame-based networking technology, using the IEEE 802.3 standard. The physical medium of a typical Ethernet-based LAN network can be coaxial cable, twisted pair wires, fibre optic, or it can be in the form of Wireless LAN. Currently, the most common form of Ethernet is called 100Base-T and it provides transmission speeds of up to 100Mbps. Slower Ethernet or 10Base-T is also commonly used in lower speed control and monitoring projects.

Devices on the Ethernet are all connected together and communication is based on the Carrier Sense Multiple Access with Collision Detection (CSMA/CD) protocol. In comparison with other network topologies, such as the Token Ring or Token Bus, all nodes on an Ethernet network can see each other at all times. Only one node transmits its data while the others listen to avoid collision. In case of a possible collision, transmitting nodes wait for a random time and attempt to re-transmit, hoping to avoid the collision.

The maximum length of an Ethernet cable depends on the speed of transmission and the type of cable used. For standard twisted cables operating at 10Mbps, the maximum cable length is specified as 100m. If using fibre-optic cables, this length can be extended to several kilometers.

As shown in Figure 1, an Ethernet packet consists of:

- 6 byte destination address;
- 6 byte source address;
- 2 byte data type;
- 45 to 1500 byte data;
- 4 byte Cyclic Redundancy Check (CRC).

In addition, when transmitted on the Ethernet medium, a 7-byte pre-ample field and Start-of-Frame delimiter byte are appended to the beginning of the Ethernet packet.

Various network communication protocols are embedded inside Ethernet packets. For example, DECnet, IP, ARP, etc – all use the Ethernet as the communications protocol. In this article we will be using a Web browser command to establish communication between the PC and the microcontroller system. Web browser is based on the TCP and uses port 80.

TCP is an advanced protocol requiring connection and providing guaranteed packet delivery with re-transmission if an error occurs. TCP packets are acknowledged to confirm safe packet-delivery.

Embedded Ethernet Controller Chips

There are several Ethernet controller chips on the market. Although such chips can be purchased as components, in most applications it is easier and usually cheaper to use boards with incorporated Ethernet controller chips and network connection sockets (e.g. RJ45).

The ENC28J60 (www.microchip.com) is a standalone 28-pin Ethernet controller chip that meets the IEEE 802.3 specifications and is controlled using the SPI interface. This chip is used in the project given later in this article. This chip has the following basic features:

- Compatible with 10Base-T networks;
- Supports both half-duplex and full-duplex operation;
- Supports automatic polarity detection and correction;
- Automatic re-transmit on collision;
- 8K byte transmit/receive buffer;
- Supports unicast, multicast and broadcast addresses;
- Link and Activity LED interface;
- Differential signal interface to RJ45 connector.

Figure 2 shows the block diagram and connection of the ENC28J60 Ethernet controller chip to a microcontroller. Basically, the interface requires the SPI signals SI, SO and SCK to be connected to the microcontroller. In addition the CS pin can also be connected to a microcontroller I/O pin.

Another popular Ethernet controller chip is the ENC624J600 with 10/100Base-T interface and 24K byte transmit/receive buffer. Some high-end PIC microcontrollers incorporate Ethernet controllers. For example, the PIC18F97J60 is a microcontroller that includes a 10Base-T Ethernet controller with 8K byte transmit/receive buffers. Similarly, PIC32MX6xx is a 32-bit PIC microcontroller with built-in 10/100Base-T Ethernet controller and 512K byte flash memory. The advantage of using a microcontroller based chip is that it provides analog and digital I/O ports and many other microcontroller features.

In general there are four access methods that can be used to establish connectivity between a PC and an embedded Ethernet controller (Application Note AN292, Silicon Labs):

- Using a Web browser on the PC
- Using a HyperTerminal on the PC
- Having the embedded system to send e-mail
- Using a custom application based on developing software on both the PC and the embedded system.

Using a Web Browser on the PC

This is perhaps the easiest and the most reliable method of establishing connectivity with no software development on the PC. This method is based on HTTP which has been in use since 1990s as the most widely used protocol to transfer data on the Internet. The aim of HTTP protocol is to allow the transfer of HTML files between a browser (usually a PC) and a Web server where the data item is located. In this method, the PC is termed the Client and the microcontroller system is termed the Server. The client sends a request by entering the url of the server. Assuming that the server url is 192.168.10.15, then entering the following command on the PC will establish a link to the microcontroller system:

http://192.168.10.15

The microcontroller system, for example, can then send an HTML page as a response

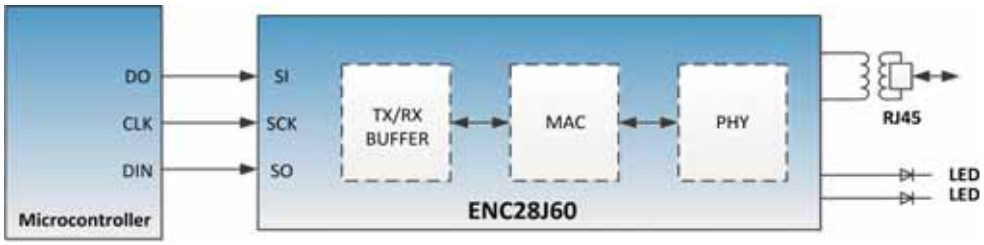

Figure 2: Connecting the ENC28J60 Ethernet controller to a microcontroller

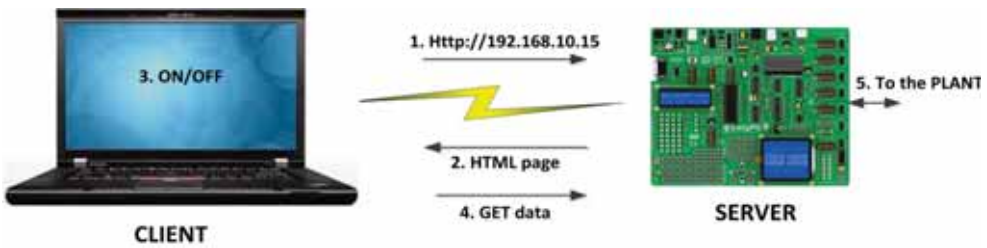

Figure 3: Web Browser connectivity

Figure 4: Connectivity using the E-mail method

Feature	TCP	UDP
Speed	Slow. Packets acknowledged, lost packets re-transmitted	Fast. No acknowledgement, no re-transmission
Complexity	High	Low
Connectivity	Connection required between two devices	Connection is not required
Packet delivery	Guaranteed	Not guaranteed. Lost packets are not re-transmitted
Packet overhead	Large	Small

Table 1: Comparing the TCP and UDP protocols

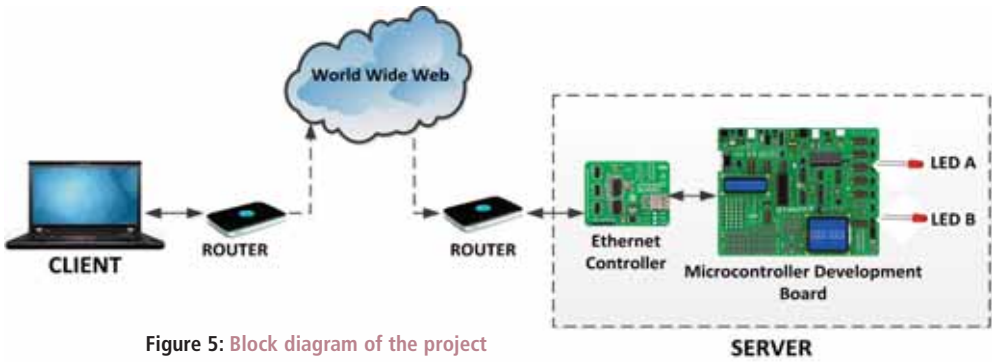

Figure 5: Block diagram of the project

to the client to display a menu with buttons. By clicking a button on the menu, a command (e.g. GET) will be sent to the server with the appropriate command tail. The server can decode this command tail and take the appropriate actions, in this case, for example, toggle the required LEDs.

Figure 3 shows the connectivity using a Web browser interface.

The HyperTerminal interface is also known as the Telnet interface. Here, the user connects to the microcontroller system by issuing Telnet commands and specifying its IP address. This kind of interface is usually an interactive interface and requires the connectivity to be initiated and terminated by the user on the PC.

Embedded System Sending E-mail

In this method the microcontroller system sends its data using the e-mail protocols. The outgoing e-mail is handled by SMTP, while the incoming mail is handled by mail servers such as POP, IMAP, or HTTP. Using this method has disadvantages that an incoming mail may stay in the input buffers for a long time until it is read by the user.

Figure 4 shows the connectivity using the e-mail method.

The development of custom applications for network connectivity provides highly flexible interface. This method, however, has the disadvantage that the network software should be developed both, on the PC and on the microcontroller system.

Two protocols are usually used for custom development: UDP and TCP. The TCP protocol is used in applications where a guaranteed packet delivery is required with the delivery of each packet being acknowledged. Lost packets are re-transmitted. The UDP protocol on the other hand is used where high transmission speed is more important than the safe delivery of packets. There is no acknowledgement and no re-transmission, should a packet is lost.

Table 1 compares the TCP and the UDP protocols.

Example of an Ethernet-Based Embedded Control Project

This section describes the design of a simple microcontroller-based automation system using the Ethernet as the communication medium. Figure 5 shows the block diagram of the example project. The project hardware is in two parts, connected

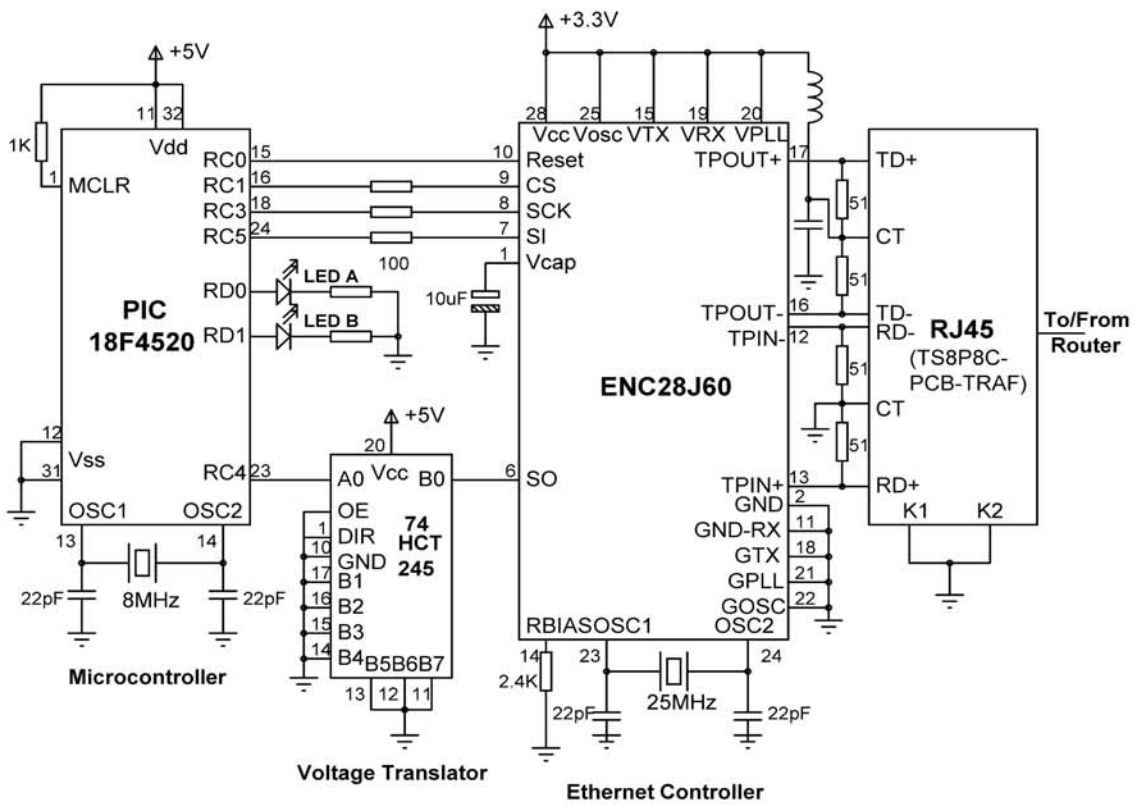

Figure 6: Circuit diagram of the project

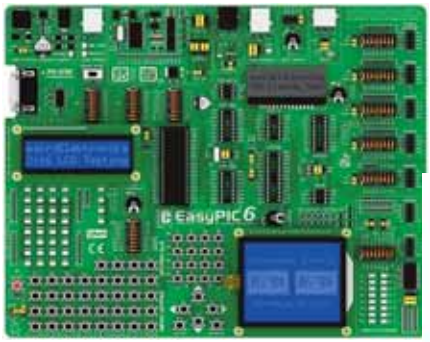

Figure 7: EasyPIC6 microcontroller development board

Figure 8: The Serial Ethernet Board

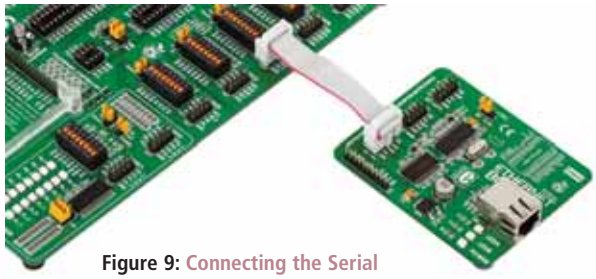

Figure 9: Connecting the Serial Ethernet Board to EasyPIC6

using a network hub or a switch: the Ethernet controller (or the server) and the PC (or the client).

The Ethernet controller consists of a microcontroller and an ENC28J60 Ethernet controller chip. Two LEDs (LED A and LED B) are connected to the microcontroller to simulate two lamps. These lamps are toggled under the control of a standard Web browser command initiated on the PC. There is no software development on the PC.

Figure 6 shows the circuit diagram of the system. The microcontroller system is designed around a PIC18F452 microcontroller chip, operating at 8MHz. The Ethernet controller is based on the ENC28J60 chip, operating at 25MHz. The interface between the microcontroller and the Ethernet chip is based on the SPI bus protocol, where SI, SO and SCK pins of the Ethernet chip are connected to SPI pins (PORT C) of the microcontroller.

The Ethernet controller chip operates at 3.3V and thus its output pin SO cannot drive the microcontroller input pin without a voltage translator. In Figure 6, a 74HCT245 buffer is used to boost the output level of pin SO. Other lower cost chips such as 74HCT08 (AND gate), 74ACT125 (quad 3-state buffer) could also have been used.

The internal analogue circuitry of the ENC28J60 chip requires that an external resistor is connected from RBIA to ground. Some of the device's digital logic operates at 2.5V and an external filter capacitor should be connected from Vcap to ground. Transmit output pins of the Ethernet chip (TPOUT+ and TPOUT-) and the receive inputs (TPIN+ and TPIN-) are connected to an RJ45 network socket with an integrated Ethernet transformer (T58P8C-PCB-TRAF).

Two LEDs on the Ethernet controller board provide visual indication of the Link and Activity on the line (the RJ45 socket has a pair of built-in internal LEDs, but are not used in this project). A 5V to 3.3V power supply regulator chip (e.g. MC33269DT-3.3) is used to provide power to the Ethernet chip. If the PC and the Ethernet controller are on the same network and close to each other, then the two can be connected together using a crossed network cable, otherwise, a hub or a switch may be required. If the PC and the Ethernet controller are located on different networks and are not close to each other, then routers may be required to establish connectivity between the two.

Two LEDs – LED A and LED B – are connected to port D pins RD0 and RD1 of the microcontroller respectively. These LEDs are toggled under the control of a Web browser command issued from the PC.

The Construction

The project was constructed using the EasyPIC6 Microcontroller Development Board (see Figure 7) and the Serial Ethernet Board (www.mikroe.com). EasyPIC6 is a fully integrated development board with the following basic features:

- Sockets for 8, 14, 18, 20, 28 and 40-pin microcontrollers;
- 36 LEDs and 36 push-button switches;
- LCD and GLCD sockets;
- 2x16 COG display;
- USB programmer;
- Port expander module;
- In-circuit debugger (mikroICD);
- 4x keypad and menu keypad;
- PS/2 connector.

The developed software can be easily compiled and then loaded to the program memory of the target microcontroller located on the development board. An in-circuit

```
MAIN program
BEGIN
        Configure I/O ports
        Initialise SPI bus
        Initialise Serial Ethernet Library
        DO FOREVER
                Check for packets
        ENDDO
END

User TCP Code
BEGIN
        IF a GET packet is received THEN
                IF two characters starting at index 6 are "TA" THEN
                        Toggle LED A
                ELSE IF two characters starting at index 6 are "TB" THEN
                        Toggle LED B
                ENDIF
        ENDIF
        Send HTML response page to the Client
END
```

Figure 10: Operation of the software (PDL)

Figure 11: Program listing of the project

```
const char HTTPheader[] = "HTTP/1.1 200 OK\nContent-type:";
const char HTTPMimeTypeHTML[] = "text/html\n\n";
const char HTTPMimeTypeScript[] = "text/plain\n\n";

char StrtPage[] =
"<html><body><h1 align=center>Electronics World Embedded Ethernet</h1>\
<form name=\"input\" method=\"get\"><table align=center width=500 \
bgcolor=Red border=4><tr><td align=center colspan=2><font size=7 \
color=white face=\"verdana\"><b>LED CONTROL</b></font></td></tr>\
<tr><td align=center bgcolor=Blue><input name=\"TA\" type=\"submit\" \
value=\"TOGGLE LED A\"></td><td align=center bgcolor=Green> \
<input name=\"TB\" type=\"submit\" value=\"TOGGLE LED B\"></td></tr>\
</table></form></body></html>";

// Ethernet NIC interface
sfr sbit SPI_Ethernet_Rst at RC0_bit;
sfr sbit SPI_Ethernet_CS  at RC1_bit;
sfr sbit SPI_Ethernet_Rst_Direction at TRISC0_bit;
sfr sbit SPI_Ethernet_CS_Direction  at TRISC1_bit;

unsigned char MACAddr[6] = {0x00, 0x14, 0xA5, 0x76, 0x19, 0x3f} ;
unsigned char IPAddr[4] = {192,168,1,15};
unsigned char getRequest[10];

typedef struct
{
  unsigned canCloseTCP:1;
  unsigned isBroadcast:1;
}TethPktFlags;

unsigned int SPI_Ethernet_UserTCP(unsigned char *remoteHost,
unsigned int remotePort, unsigned int localPort,
unsigned int reqLength, TEthPktFlags *flags)
{
  unsigned int Len;
  for(len=0; len<10; len++)getRequest[len]=SPI_Ethernet_getByte();
  getRequest[len]=0;
  if(memcmp(getRequest,"GET /",5))return(0);

  if(!memcmp(getRequest+6,"TA",2))RD0_bit = ~ RD0_bit;
  else if(!memcmp(getRequest+6,"TB",2))RD1_bit = ~ RD1_bit;

  if(localPort != 80)return(0);
  Len = SPI_Ethernet_putConstString(HTTPheader);
  Len += SPI_Ethernet_putConstString(HTTPMimeTypeHTML);
  Len += SPI_Ethernet_putString(StrtPage);
  return Len;
}

unsigned int SPI_Ethernet_UserUDP(unsigned char *remoteHost,
unsigned int remotePort, unsigned int destPort,
unsigned int reqLength, TEthPktFlags *flags)
{
  return(0);
}



void main()
{
    TRISD = 0;
    PORTD = 0;
    SPI1_Init();
    SPI_Ethernet_Init(MACAddr, IPAddr, 0x01);

    while(1)
    {
      SPI_Ethernet_doPacket();
    }
}
```

debugger incorporated on the board helps to debug software during the development cycle.

The Serial Ethernet Board (see Figure 8) is a small board that plugs-in directly to the EasyPIC6 development board via a 10-way IDC plug (see Figure 9), simplifying the development of embedded Ethernet projects. The board is equipped with an EC28J60 Ethernet controller chip, a 74HCT245 voltage translation chip, three LEDs, a 5V to 3.3V voltage regulator and an RJ45 socket with integrated transformer. Thus, considering Figure 6, all of the components except the microcontroller and the two LEDs are located on the Serial Ethernet Board.

The Software

The software of the example is developed using the mikroC PRO for PIC language. This language has been developed for the PIC microcontrollers and it provides a large number of built-in functions to support peripheral devices such as RS232, RS485, CAN bus, SD card, Compact flash card, USB, LCD, GLCD, touch screen, Ethernet and many more. Figure 10 shows operation of the software as a PDL.

The operation of the Ethernet controller software is very simple and the complete program listing is given in Figure 11. At the beginning of the main program, the I/O ports are configured and the SPI bus is initialized by calling the built-in library function SPI_Init. Then the serial Ethernet library is initialised by calling the built-in function SPI_Ethernet_Init and specifying the MAC address, the IP address and the mode of operation as full-duplex.

The MAC address of the Serial Ethernet Board is set at factory to “0x00, 0x14, 0xA5, 0x76, 0x19, 0x3F”. The program sets the IP address of the board to “192.168.10.15”. The main program then enters an infinite loop, where built-in library function SPI_Ethernet_doPacket is called to check for the arrival of packets and also to send any outstanding packets to their destinations. The Ethernet library requires both the UDP and TCP functions to be present in the program, even though they may not be used. Only the TCP is used in this example as the Web browser communication is based on TCP.

Inside the TCP function any received packets are checked and the function continues if the packets are of type “GET /”. Then, the command tail is checked and the LEDs are toggled as required. The transmit buffer is loaded with the HTML response

and the length of the buffer is returned from the function, which then sends the buffer to the client.

The connectivity of the system can be checked by using PING to send packets from the PC to the Ethernet controller. If everything is working as expected, then PING replies should be displayed on the PC screen.

- Upon receipt of this request the Ethernet controller sends the HTML code shown in Figure 12 to the Web browser, together with the HTTP header. The Content-Type field is used by the browser to tell which format the document it receives is in. HTML is identified with "text/html" and ordinary text is identified with "text/plain".
- The Web browser then displays the form shown in Figure 13.
- The user can toggle LED A or LED B by clicking on the appropriate buttons. Assuming button LED A is clicked, the Web browser sends the following command to the Ethernet controller:

GET /? TA=TOGGLE_LED+A

Similarly, if button B is pressed, the Web browser sends the following command to the Ethernet controller:

GET /? TB=TOGGLE_LED+B

- The Ethernet controller checks the received command (inside function TCP) and toggles LED A or LED B as required. ●

Figure 12: HTML code sent to the Web Browser

```
<html>
<body>
<h1 align=center>Electronics World Embedded Ethernet</h1>
<form name="input" method="get">
  <table align=center width=500 bgcolor=Red border=4>
   <tr>
     <td align=center colspan=2><font size=7 color=white face="verdana"><b>LED CONTROL</b></font></td>
   </tr>
   <tr>
     <td align=center bgcolor=Blue><input name="TA" type="submit" value="TOGGLE LED A"></td>
     <td align=center bgcolor=Green><input name="TB" type="submit" value="TOGGLE LED B"></td>
   </tr>
  </table>
 </form>
</body>
</html>
```

Figure 13: Form displayed by the Web Browser on the PC

EMBEDDED SYSTEM PERFORMANCE TRADEOFFS

TODAY'S DIVERSE COMPETING SYSTEM REQUIREMENTS ARE PUSHING DESIGNERS TO UTILIZE THE LATEST TECHNOLOGIES TO CREATE DESIRED SYSTEM SOLUTIONS. **KEN PERDUE**, SPANSION FIELD APPLICATION ENGINEER, EXPLORES THE TECHNOLOGIES THAT ENABLE HIGH PERFORMANCE, RELIABLE AND COST-EFFECTIVE DISPLAY SOLUTIONS

Many of today's innovations in embedded design are the result of a growing demand by customers for state-of-the-art services. Such services may be vehicle infotainment, medical uses, in home control systems, and other market applications. These systems continue to utilize greater amounts of graphic-intensive content, resulting in more sophisticated and informative interfaces having additional levels of functionality and increasing system complexity.

Today's newest embedded systems are enabling a more user-friendly and informative environment, promoting overall improved user experience. These system requirements are creating new challenges for designers to deliver real-time high-performance levels similar to PC-based consumer products, while many embedded application still require fast system boot, along with long-term operational reliability and product availability.

To meet this diverse set of requirements many embedded designs employ methodical planning, design, verification and approval processes to ensure full functional operability while mitigating reliability issues. Many applications require the embedded designer to select electronics suppliers whose products are designed and qualified for high reliability and long-term availability. Today's diverse set of competing system requirements such as instant boot, fast real-time performance, long-term support, high reliability and reduced costs are pushing designers to utilize the latest technology innovations to create the desired system solutions. We will explore how today's technology is enabling high performance, reliable and cost-effective solutions.

> The system architecture and design implementation can have a major effect on an embedded system's performance

Figure 1: Cluster block diagram

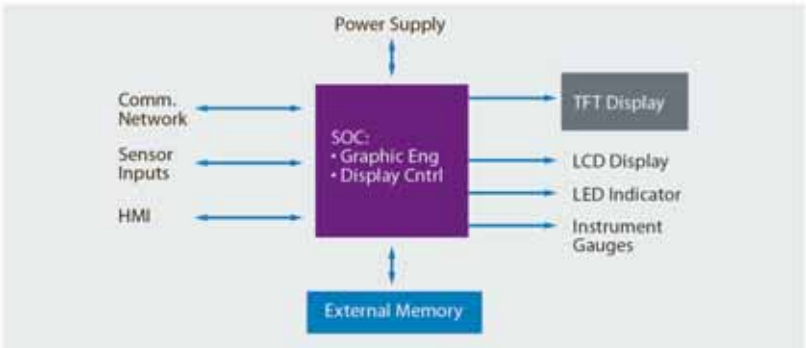

Digital Displays

A typical digital display that you can find in a multitude of embedded systems uses a system-on-a-chip (SoC), which substantially derives data for driving a TFT display through its communication network or internal/external memory (Figure 1). These digital displays employ architectures with some similarities to high-performance PC-based display systems.

Most markets such as automotive, medical, gaming, or in-home display platforms now require high-density memory to support an increasing amount of digital content. For instance, the density requirement can be more than 4Gbits for high-resolution large displays with 3D graphics content. This content typically consists of large character sets, multiple fonts, graphic images and extended support for multiple languages to create the informative display system.

These embedded systems employ a digital frame buffer which is generated and displayed via industry-standard techniques. The SoC display controller renders one frame on the system display; in parallel, the SoC/graphics engine accesses internal or external memory to obtain, process and store the next set of data in the frame buffer for subsequent display. Today's SoC require high bandwidth access to reliable code and data residing in external memory to enable fast system start-up and high-speed real-time processing.

The system architecture and design implementation can have a major effect on an embedded system's performance. To understand its impact, it's best to review a couple of memory architectures that illustrate certain performance/cost tradeoffs when using today's technology.

Figure 2: Code shadowing architecture

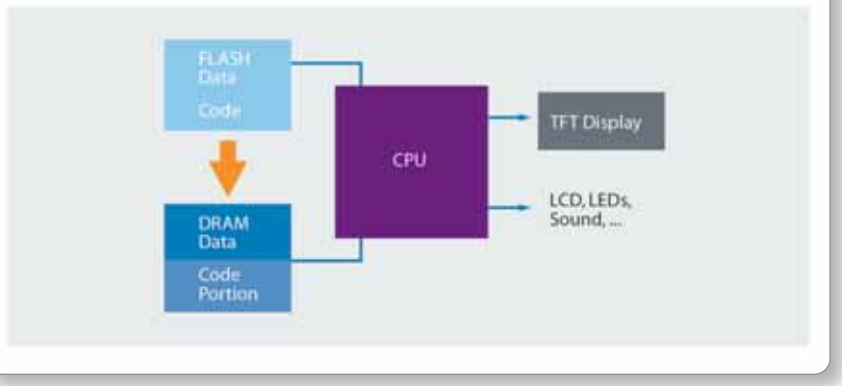

Figure 3: Hybrid memory architecture

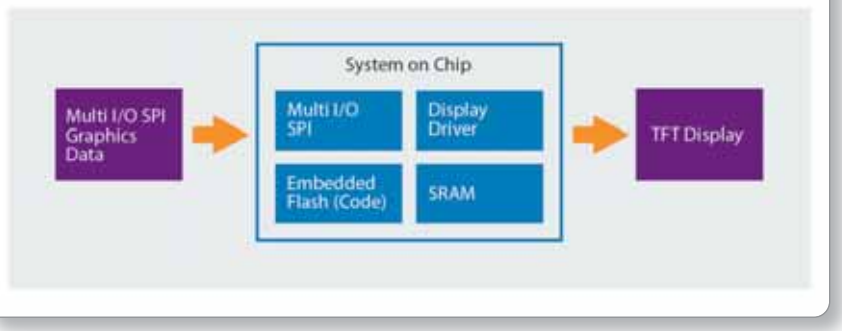

Different Memory Architectures

First, let's look at standard code-shadowing memory architecture, which is very typical for a high-performance embedded-display system (Figure 2).

The system controller or general processing unit (GPU) has a high level of integration, including graphics engines, display controllers, options for limited embedded RAM and flash, while continuing to offer external memory interfaces to support a digital display's high-performance and high-density memory requirements.

The external memory falls into two standard offerings: DRAM and flash. The embedded system start-up can be highlighted by a three-step process: code/data shadowing from flash to DRAM, initialization of the processor, DRAM and other key components, and then application execution. Once code execution begins, useful information can be displayed.

There are a couple of key takeaways with this architecture. First the SoC-DRAM high-speed access capabilities facilitate the high-performance real-time capabilities. Second, the system's initial start-up time is defined by several factors like code shadowing and initialization times. The flash access bandwidths and the data densities transferred from flash to DRAM substantially characterize the shadowing time.

Many high-end digital displays in automotive, medical and other applications must provide an almost instant-on experience at power-up. The US Food and Drug Administration is mandating fast start-up times for some Emergency Room medical equipment. Automotive digital clusters require instant-on during vehicle "Key ON" => the drive train position (i.e. P R D 1 2 3) must be displayed in less than one second. This requires flash memory that supports high-speed access, high densities and long-term data integrity to meet basic design functionality and quality requirements.

Today, flash suppliers offer multiple non-volatile-memory (NVM) technology. Two of the most widely used technologies are NOR and NAND flash memory. Each provides different features, such as access capabilities, reliability levels, product life cycles and cost per bit. Today's SoC typically support multiple interface configurations to access external parallel and serial NOR flash, while NAND flash interfaces continue to garner increased support.

NOR and NAND flash memory offer key tradeoffs and options for the design and sourcing teams to consider.

One very dynamic variable for the SoC NAND interface involves defining the required level of ECC support. Today's NAND has very fast lithography migrations which can provide lower initial cost per bit but can introduce long-term sourcing issues, depending on the product life cycle.

Continued Innovations

Continued innovations by both chipset and flash memory makers are enabling new, cost-efficient, hybrid, digital display architectures (Figure 3).

Hybrid architectures and flash use can require hundreds of millions of read cycles that affect the viable use of technologies like NAND. (NAND devices can exhibit additional bit disturb errors when used in applications that have mid- to high-read cycling requirements.)

But there are areas where embedded designers have innovated using today's latest technologies to provide state-of-the-art solutions for their customers (see 'Example of a Hybrid System' below). ●

EXAMPLE OF A HYBRID SYSTEM

A HIGH-END DIGITAL CLUSTER DESIGN LIKE THE VOLT'S E-FLEX CLUSTER USES FREESCALE'S QORIVVA POWER ARCHITECTURE PROCESSOR AND THE SPANSION GL PARALLEL FLASH MEMORY.

The system controller (Freescale Power Architecture Family) and flash (Spansion FL Serial flash with DDR capability) make it possible to optimize both the TFT display performance and enable a more cost-effective architecture and systems solution. These system controllers can execute code from embedded flash, while loading graphics data from external flash via the high-bandwidth Multi I/O serial peripheral interface (SPI). Graphics data is processed and stored in the internal frame buffer, which is directly displayed on a TFT.

The Spansion FL serial flash memory communications protocol is used for data transfers between the system controller and external SPI flash. Accesses to Multi I/O SPI can be configured to support one to four data connections and utilizes single, dual or quad data rates with optional double data rate (DDR) capability.

This hybrid architecture meets system performance and reliability requirements, while reducing BOM cost since there's no need for external DRAM support.

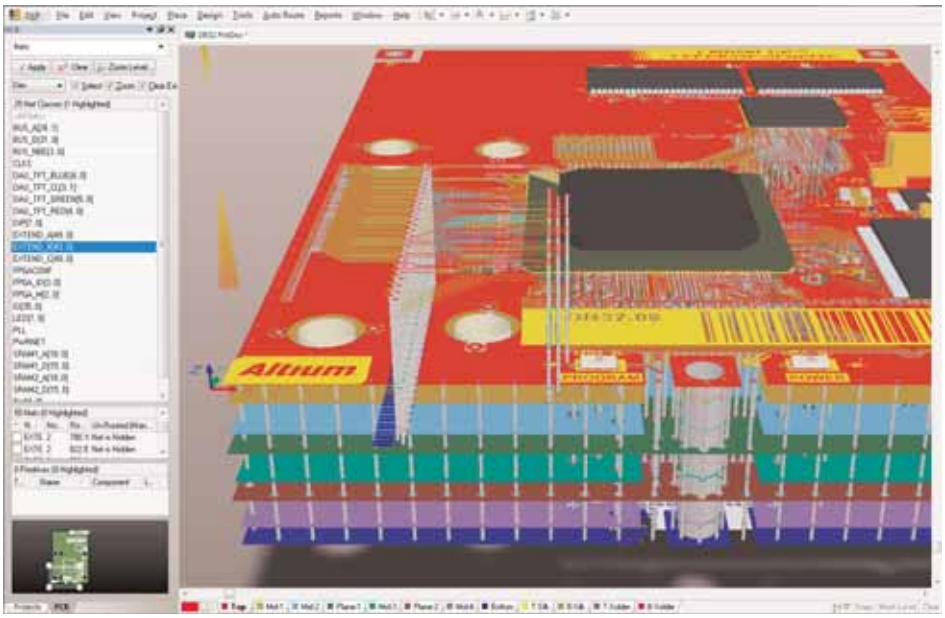

Figure 1: Designing modern electronic products demands greater visibility and manipulation of multiple components and, increasingly, multiple PCBs

E-CAD MEETS M-CAD

FRANK KRÄMER, TECHNICAL MARKETING DIRECTOR FOR EMEA AT ALTIUM EUROPE, LOOKS AT WAYS ENGINEERS CAN INCREASE PRODUCTIVITY THROUGH SPECIFIC TOOLS

It is often felt that because manufacturing processes naturally migrate to the region with the lowest cost, developed economies suffer. However, another viewpoint is that by outsourcing repetitive tasks, the resources previously occupied with low value activities are presented with greater opportunities. The same is true with design; by employing design automation tools, engineers are able to increase their productivity. It would be unthinkable today for an engineer to use a non-computerised approach to designing a PCB for instance, even though it is entirely possible to do so.

Equally, just as economies must adapt to changing global conditions, engineering teams are encouraged to make use of every productivity tool at their disposal. For electronic design engineers the use of EDA tools has vastly improved the design process, from component to end product. As well as accelerating the design process, it is now becoming possible to simulate or emulate every part of an electronic design before a single component is even purchased.

> It is now commonplace for mechanical design engineers to have access to three-dimensional representations of their designs, rendered in real-time

In the field of integrated circuit development for example, the very last step in a long and expensive process is to commit – or 'tape out' – a design to silicon. This is largely due to the enormous costs involved with manufacturing integrated devices, a cost that is only borne by the opportunity of selling the final device in large volumes. The same is not yet true for all electronic products developed today; most incur much lower non-recurring engineering costs, but for mechanical design there are significant cost implications of re-designing tooling to, for example, a clearance issue with the PCB inside the product.

Relying on Tools

The advances made in EDA tools means it is becoming easier to simulate a product's form, fit and function before it is manufactured. Even so, with the obvious exception of IC design, the progress of design automation tools in the electronics sector has been relatively focused on niche applications and vertical markets. PCB design is an example; there are any number of low-cost PCB design tools that are adequate for creating simple single or double-sided layouts, but fewer that are capable of tackling multi-layer PCBs with very high speed signals and mixed-signal domains, and fewer still that incorporate effective analysis tools that ensure signal integrity isn't compromised.

For those designs that need these features, the tools are invaluable. They offer the only realistic solution to developing the kind of end products that we now take for granted to deliver our digital lifestyle. For example, mobile telecommunications would not be possible without sophisticated EDA tools; they enable talented engineers to develop the complex mixed-signal devices and systems needed to make 3G networks and smart phones a reality.

The examples are numerous but the underlying trend is that the more complex the design, the more

sophisticated the tools. However, there is one aspect of design that is applicable to practically every single product developed, irrespective of its functional complexity or end market value.

Crossing Design Domains

The integration of electronic and mechanical design is inexorable; with few exceptions PCB design is not only influenced by the components it carries but by the space it can occupy. Many products today feature a single PCB and in such cases the size and shape of the PCB is dictated less by its functionality than by its environment. In fact, in some cases – particularly in consumer devices – the shape and size of the end product will actually define the available space for the PCB and its components. In these cases mechanical design dictates both domains, yet there remains limited interaction between CAD tools targeting each specific one.

While the focus of electronic CAD tool vendors has, understandably, been on addressing the complexity of electronic design, their counterparts have been industrious in improving M-CAD tools to make full use of the processing and graphical capabilities of the latest PCs and desktop computers. It is now commonplace for mechanical design engineers to have access to three-dimensional representations of their designs, rendered in real time. As a productivity enhancement it is hard to deny the value of seeing the product of an engineer's efforts in a 3D environment that supports real-time manipulation to vary the viewing angle.

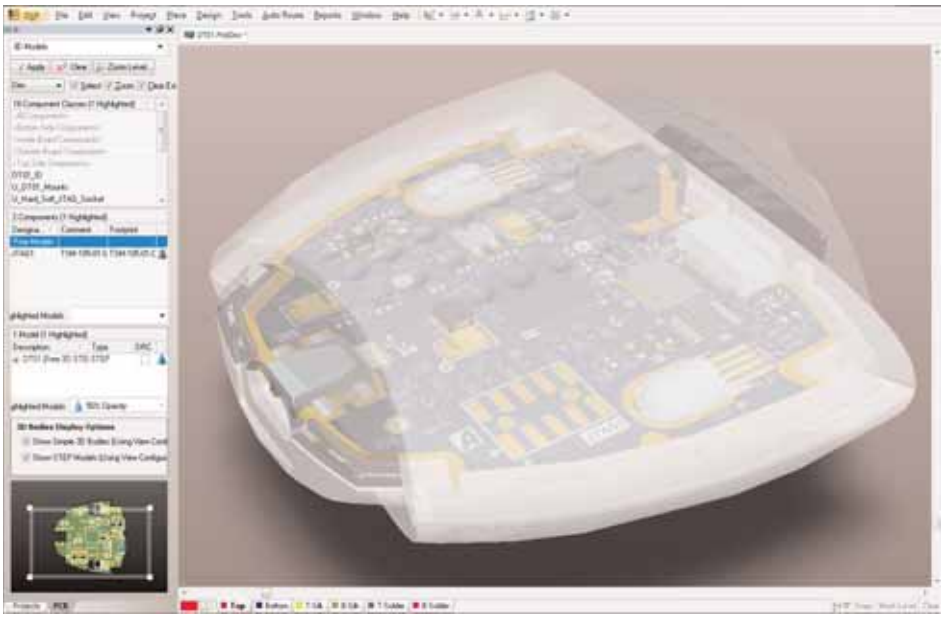

Figure 2: The interaction between the mechanical and electronic design domains is now critical at every stage of product development, as form, fit and function become inexorably linked

It's also valid to mention that while ICs continue to shrink in size, their supporting components are less likely or able to do so. Specifically, fundamental principles limit the physical size of passive components such as transformers, resistors, capacitors and inductors, and while there is less need for numerous connectors in modern equipment, those that remain have physical restrictions on how small they can be and where they must be placed on a board. This has its benefits, however, as there now exist numerous 3D models for standard components such as passives and connectors that can be used and manipulated in a growing number of CAD packages.

The widespread creation of these 3D models indicates renewed commitment from vendors towards E-CAD/M-CAD integration, a trend that many in the industry believe will continue, bringing significant productivity gains to engineers across both domains.

Perhaps the most significant development along this path to full integration was the introduction of a design interchange protocol that both E-CAD and M-CAD tool vendors could adopt with confidence.

While there have been many attempts in the past to integrate the two domains, they have been impeded by a lack of cooperation between vendors, which resulted in adding complexity rather than removing it. But, with the introduction of STEP (Standard for the Exchange of Product model data), specifically version AP214 which defines 3D models, the interchange of design data has become much simpler. The M-CAD sector has been quick to

Figure 3: Only EDA tools that fully implement the STEP file format effectively can provide the level of mechanical and electronic design domain interaction necessary for the development of the latest products

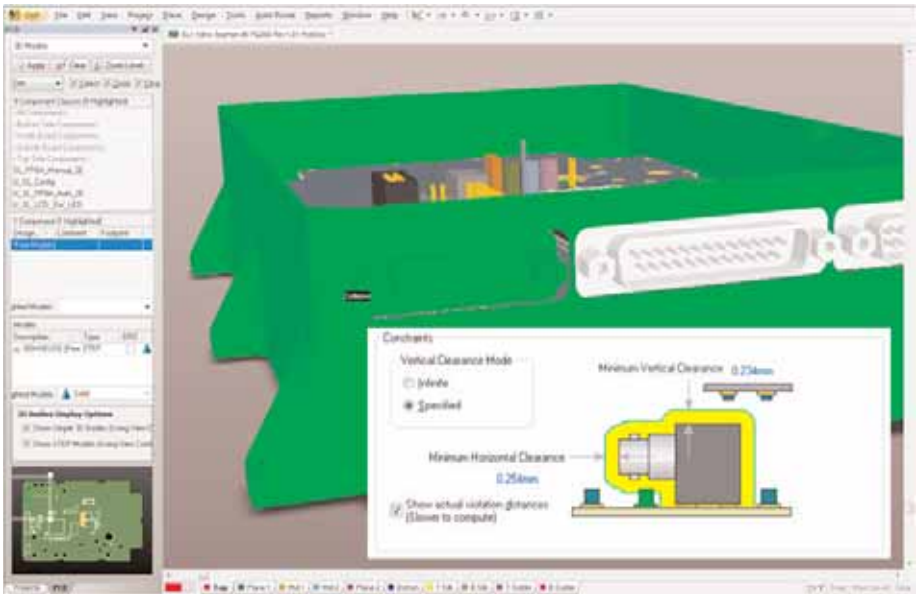

Figure 4: STEP brings a level of visibility and interactivity at the modelling level that is both unprecedented and invaluable to today's cross-domain design teams

integrate STEP AP214 compliance into its value chain, but the same isn't true for the E-CAD sector. However, Altium Designer, the unified design environ-ment from Altium, does support STEP file import/export and manipulation, and coupled with its comprehensive PCB design functionality it brings a new level of productivity to all electronic engineers.

3D Capabilities in the PCB Space

Many M-CAD tools now support 3D models of populated PCBs created by a third-party tool, but while this offers valuable visualisation of how the PCB and the casing will ultimately interact, it doesn't provide the ability for the mechanical environment to feed back to the PCB designer critical dimensions, clearances or other spatial compliance issues. In addition, mechanical design engineers are less equipped to appreciate the need for certain component positioning, particularly when high speed, mixed signal or high voltage signals are present.

The adoption of the STEP format within Altium Designer overcomes this restriction. It enables engineers not only to use a 3D model of an enclosure to visualise the end product, but actively adopt a three-dimensional approach to their design. Embedded within the AP214 format is enough data to allow an imported model of an enclosure to actually be used to define the dimensions of a PCB, overcoming the issue of accurately and manually transferring critical measurements from one domain to another. This capability takes a significant step towards enabling electronic design engineers to design for manufacture, by closely linking the mechanical domain to the electronic design process.

> With the introduction of the Standard for the Exchange of Product model data (STEP), specifically version AP214 which defines 3D models, the interchange of design data has become much simpler

Furthermore, the ability to define clearance requirements within a 3D format means engineers in both domains can instantly see the impact of design changes. By combining the models of the enclosure and the PCB, within Altium Designer an engineer can manipulate the resulting 3D representation to actively measure clearances. This unprecedented feature means the electronic engineer will have complete confidence in the compliance of the final PCB long before manufacturing.

To make the process even more productive, models can be linked so that any changes made in one domain are reliably reflected in the other. This means that any changes made to the enclosure will be seen by the electronic engineer and any alterations to the PCB or its components will be automatically relayed to the mechanical engineer.

Key to this functionality is the ability not only to manipulate a single 3D model, but coordinate multiple models in a virtual 3D space, using reference points. By precisely aligning models for the parts of an enclosure and a populated PCB, design engineers can verify critical clearances, as well as the way the PCB will fit within the enclosure or the need for support bosses and fixings, while preserving the product's overall market objectives.

Another benefit of working in a virtual world is that engineers can experiment without cost. For example, if a component is aligned using three reference points, it is possible to place one component such that it passes through the second. Imagine a PCB positioned so that it protrudes through an enclosure; it may seem unconventional but it could just solve a design dilemma. Achieving this using real models would require hours of effort but in a virtual domain it is as simple as changing a single reference point. This close interaction between the electronic and mechanical domains is only possible today through the use of the STEP format. Embodying the STEP format in to a PCB design environment marks a significant move towards creating a unified approach to product development. ●

Europe's Premier Microwave, RF, Wireless and Radar Event

EUROPEAN MICROWAVE WEEK
RAI Amsterdam
28 October – 2 November 2012
www.eumweek.com

Space for Microwaves

EuMW2012 returns to this city for what promises to be an important and unforgettable event. Bringing industry, academia and commerce together, European Microwave Week 2012 is a SIX day event, including THREE cutting edge conferences and ONE exciting trade and technology exhibition featuring leading players from across the globe.

THE EXHIBITION (29th - 31st October 2012)

- 7500 sqm of gross exhibition space
- 5000 key visitors from around the globe
- 1700 - 2000 conference delegates
- In excess of 250 exhibitors

THE CONFERENCES

Spanning the length of the week, choose from three separate but complementary conferences:

- European Microwave Integrated Circuits Conference (EuMIC) 29th - 30th October 2012
- European Microwave Conference (EuMC) 29th October - 1st November 2012
- European Radar Conference (EuRAD) 31st October - 2nd November 2012
- Plus, Workshops and Short Courses (From 28th October 2012)

In addition, for the third year, the free, full-day Defence and Security Executive Forum.

EuMA European Microwave Association

Official Publication: Microwave Journal

Organised by: horizon house

Supported by: IET The Institution of Engineering and Technology

Co-sponsored by: MTT-S

Co-sponsored by: IEEE

EuRAD 2012 The 9th European Radar Conference

42nd EUROPEAN MICROWAVE CONFERENCE 2012 The 42nd European Microwave Conference

EuMIC 2012 The 7th European Microwave Integrated Circuits Conference

Co-sponsored by: ELECTRON DEVICES SOCIETY, GAAS

Interested in exhibiting?
Call +44(0) 20 7596 8742 or visit www.eumweek.com

ARM TECHNOLOGY ON X86 STANDARD FORM FACTORS

Taking the ARM strategy to the next level

With the availability of the PC-like ARM processor technology, there is a huge chance to improve time to market and to reduce R&D costs by facilitating processor implementations via x86 form factors. One example is the availability of the NVIDIA® Tegra™ 2 processor on Pico-ITX™ boards. What are the benefits?

The current performance level and low power consumption of ARM technology, which is applied in standard tablet applications, make it particularly interesting for embedded SFF applications. Consequently x86 embedded platform vendors now also enter the ARM technology market, as an addition to x86 technology. One strategy of these vendors is to level out the technological barriers between ARM and x86 using scalable building blocks. This strategy is very appealing to a large number of OEMs - as this means they can obtain very scalable platforms with completed Board Support Packages for all popular operating systems. And thanks to the corresponding hardware-specific software the underlying hardware can be abstracted so that more and more homogenous, application ready platforms come into existence. For OEMs, swapping from one board, module or system to another becomes a relatively easy procedure. But this goal can only be achieved if vendors will ensure this with the right standardizations on board and software level and by offering extensive software services. Depending on the programming of the application and which operation system is employed, no software changes have to be made, or only very few. Further to this, ARM solutions are always available as full-custom designs on board and system level, so that OEMs can fully concentrate on their application development, without having to consider any individual differences.

STANDARDS ALLEVIATE IMPLEMENTATION

The hardware design on board level goes to demonstrate how easy the selection of the right CPU building blocks for the application can be. For example, the interface feature set of a Pico-ITX™ board with Tegra™ 2 from NVIDIA® (see box 1) is hardly any different from the Intel® Atom™ or AMD Embedded G-Series designs which are already available. The major difference is to be found in the processor and consequently in its performance class (see table 1).

A major difference is that ARM processors are more dedicated and thus offer less generic interfaces like SATA or PCI Express which in x86 designs are often used for connecting individual extension options. Having said this, many of the ARM SOCs have several UARTs, I2C and SPI interfaces. So, theoretically speaking, the generic interfaces could be levelled out with additional components and some development effort. That would, however, also level out the valuable energy-saving advantages which make ARM designs so attractive: The need for cooling is reduced, fanless designs are possible which makes them more failsafe and results in a better MTBF. Developing and manufacturing the systems becomes easier - and the systems lighter in weight - due to the lack of heat pipes, cooling elements or fans.

But bringing back these generic interfaces is superfluous as, especially in SFF designs, the trend is towards less not more generic interfaces. Subsequently, the difference between the feature set of the Pico-ITX™ board is of little relevance. As the Pico-ITX™ format is standardized, the application-specific choice of the right x86 or ARM designs can be carried out within one single ecosystem - and no technological barriers have to be considered. The mechanical compatibility to the whole existing product portfolio is a major advantage and goes to simplify system design. Even more simplified desings can be achieved if besides the mechanical design advantage, ARM based boards additionally offer a range of software benefits. For example, the

Figure 1: SFF Singleboard Computer with NVIDIA® Tegra 2 Dual Core processor

The Small Form Factor board in Pico-ITX™ format (100 mm x 72 mm), which is currently being developed, is equipped with a 1 GHz NVIDIA® Tegra™ 2 Dual Core processor and integrates a completely passive cooling concept while boasting very low power consumption of 3 watts and an attractive feature set: in addition to 10/100Mbit Ethernet, five USB 2.0 ports and up to 24 configurable GPIOs the ARM Cortex A9 architecture-based mini-board has a slot for Micro SD cards and 512MB or 1GB 32bit DDR-2 memory. The audio-visual experience is also worth looking into too: the integrated ultra low-power (ULP) NVIDIA® GeForce® GPU delivers graphic performance for mobile devices in high-quality gaming console quality and can simultaneously stream two HD videos (1080 p). Displays can be connected via DVI-I for analog and digital signal transmission and via a 24 bit LVDS converter. Backlight support is provided by either an internal 5V intern or an external 12 V. Following audio support is available: via SPDIF and stereo line-in and line-out as well as MIC. A whole range of hardware accelerations for flash, video and audio codecs ensures fluent and brilliant playback of multimedia and web content.

extensive support for all operating systems which are currently available for these processors. With such application-ready platforms the time-to-market can be reduced significantly and with it the costs for development.

Table 1:

	KTT20/pITX	KTA55/pITX	pITX-SP
			
Processor platform	NVIDIA® Tegra™ 2 1 GHz Dualcore	AMD Embedded G-Series, 1 GHz Dualcore	Intel® Atom™ Z5x0 1.1 / 1.6 GHz Singlecore
Ethernet	1x 10/100 Mbit	1x 10/100/1000 Mbit	1x 10/100/1000 Mbit
USB	3x USB 2.0	6x USB 2.0	6x USB 2.0
Graphic output	DVI-I, 24 bit single channel LVDS, Display Serial Interface	DVI-I, 24 bit dual channel LVDS	DVI-D, 24 bit single channel LVDS
RAM	512 / 1024 MB DDR2	Max 4 GB DDR3	Max 2 GB DDR2
Storage interfaces	MicroSD Card Slot, Onboard NAND flash	2x SATA, MicroSD Card Slot	2x SATA, 1x PATA
Audio	2ch. In/out, Mic, SPDIF	HD Audio	HD Audio analog / SPDIF
Other I/Os	16x GPIO, 2x RS232, CSI, I2C, SPI, JTAG	8x GPIO	8 Bit GPI/O, SDIO

SCALABILITY ACROSS ALL PROCESSOR PLATFORMS

Are ARM processors now integrated into the realm of x86 form factors? Has this brought an end to the whole technology argument? If it was up to the embedded standard form factor vendors to decide, then, yes. Because at the end of the day, it is simply these vendors goal to deliver standard platforms to its embedded customers and to embrace new applications, which to date were not possible with all other existing processor implementations. A high level of scalability of the fitting standard form factors across all processor platforms makes a lot of sense, as then OEMs can port their applications more easily between RISC and CISC architectures. If additional hardware-specific software services are provided to realize the code modifications which are sometimes necessary, then the underlying processor architecture is less and less a fundamental criterion for decision-making.

Two other factors will play principle roles: energy consumption and performance per watt. One could also say that the market has now entered an era – after the enormous successes which the x86 technology inspired – in which thanks to the extensive software support, the borders of the processor technology are disappearing as the software ecosystem can be extended to further technology platforms. Consequently, the standard form factors on board level have to be extended to accommodate these new processor platforms.

WIDE RANGE OF SERVICES FOR DIRECT ENTRY

In order to enable customers to make an immediate entry into the ARM technology, embedded platform vendors need to offer their ARM-based building blocks in a bundle with extensive custom design services, so that OEM customers can get integrated 'application-ready platforms' on board and system level either as standard or as customer-specific versions. Besides this individual hardware development service on board and system level, Kontron for example also focuses on providing extensive services for software development, which ranges from driver development and OS code adjustments to a broad range of application porting and validation services as well as hw/sw bundles including quantity licensing. Application developers profit from efficient migration, fast time-to-market and, parallel to this, can reduce development risks and cost, as application-ready platforms are available which are already certified, so that the customer can completely concentrate on his core competence: application development.

Needless to say, that embedded platform vendors need to support all the ARM-relevant operating systems. Apart from Windows CE 6 and Windows Embedded Compact 7 (WEC7) in particular Linux-based operating systems are supported on ARM products. VxWorks support is planned for TI processors. These operating systems are especially interesting for applications, which demand the highest availability and best real time behaviour. Furthermore, support of an ARM-native version of Windows 8 is in the pipeline. The Android operating system which is at home in the smartphone and tablet market is a must to open the door to the vast market of networked multimedia-oriented applications based on ARM technology which this relatively young operating system caters for. The BSPs are validated right up to system level which enables OEMs to focus on the application without having to train accordingly and this minimizes time-to-market and TCO.

When comparing the three Pico-ITX™ feature sets it becomes apparent, that in terms of the most important interfaces like USB, Ethernet, graphics and storage for SFF devices, there are hardly any differences, so that basically, with the extension into ARM technology, scalability is further increased.

Author: Daniel Pieper is Product Marketing Manager at Kontron

WHAT THE READERS SAY

BEST RESPONSE TO THE CATT QUESTION

The response to the Catt Question by Raymond Boute [Letters, June] is one of the best I have seen, which both deals with the objections raised by Catt and presents an understandable model of what is really going on, without recourse to complex mathematics or arguments. I always felt there was a fundamental error by Catt, but could not quantify where it was.

The reality of the "electron gas" moving at high speeds within the copper conductor, and the influence electrons have on each other at great distances with respect to the effective electron size, was always going to mean that reality was somewhat different from the idealised model presented by Catt. Consider the effect of relatively slow average electron movement in the coils of a magnetic deflection system. These have been used for years without really understanding the process. A magnetic field is surely nothing more than the influence of one set of moving charges on another set. In permanent magnets it is the aligned spin of unpaired orbital electrons that creates a magnetic field; in conductors carrying current it is the relative drift speed of the electron gas.

Much as scientists dislike the concept of force at a distance, the influence of one set of charged particles on another is profound and happens over distances several orders of magnitude greater than the size of the particles.

I wonder therefore whether the propagation of an EM wave is nothing more that the force at a distance of one charged particle on another. Although we consider the vacuum of space to be empty, there are in fact many particles per cubic metre, many of which are charged, all of which can be affected by a movement of other charged particles. We are not able to detect EM waves except by using instruments composed of atoms and molecules, so how can we know they exist in isolation?!

I don't know whether anyone has done the experiment to time the propagation of magnetic influence, or whether it is even practically possible or not. If one could switch on the current in an electomagnet instantly, how long would it take a beam of electrons one metre away to be deflected?

The logical answer would be that it is no less than the speed of light, in that that is the speed at which the sphere of influence appears to propagate. That being the case there would be no way to determine whether a magnetic field really exists in its own right, or is simply the artefact of moving charges influencing one another. In a whole universe concept it is impossible to determine cause and effect if these movements of charges each spread their own sphere of influence at the speed of light. As soon as one intercepts or changes something in order to measure it, the sphere of influence of the measurement spreads at the speed of light changing the surrounding environment.

Scientists and engineers have over the years worked out usable models that enable us to make use of the phenomena of EM radiation and magnetic and electric charges. Raymond Boute's explanation of the transmission line problem is a good usable model. In a universe where everything has an influence on everything, there will always be room for other interpretations and we will never fully understand, but to my mind a simple model is sufficient to enable us to use the phenomena to our advantage.

Ray Lee

MIXING THEORY WITH PRACTICE

[In response to Ivor Catt's letter 'Where are They' in the Letters section of the May issue] Engineers are not trained or have as an objective the upsetting of physics, so we have lived in the dual world of "theory says" versus how things "really work". The real world and the theory world are compartmentalized. They are not allowed to influence each other. I think this is Catt's objection.

It is correct that the theory world – that is the world of platonic idealism – has always been the refuge of the elite. There was a brief period from, say, 1600s to maybe 1900s wherein the real world was supposed to be important. This may be something of a myth. The actual fact may be more along the lines that the theory world, in looking for a reason to reject the old theory, looked to the empirical world as an aid in doing away with the old theory. But once that was accomplished, the theory world ceased to dirty its hands in the facts.

> *Engineers are not trained or have as an objective the upsetting of physics, so we have lived in the dual world of "theory says" versus how things "really work"*

Today we have a theory-driven empirical science. That is something of a contradiction, in that we are told that theories are tested. But unless there is a motivation to replace a theory the actual testing doesn't seem to be effective in changing the thinking of the theory world.

Harry Ricker

ELECTRONS AND PHOTONS

The letter by Raymond Boute in the June 2012 issue of Electronics World is very plausible. It recognizes the facts that a current pulse will propagate along a transmission line at near-light velocity and that the bulk movement of electrons along a conductor progresses at a snail's pace. However, he writes: "It is the collective (average) movement along the wire that constitutes current". He is confusing cause and effect.

Current must be caused by the movement of a physical entity, and the only entities capable of moving at the velocity of light are sub-atomic particles. We have given these particular particles the name 'photons'.

The simplest configuration which can be used to illustrate this effect is the dipole antenna. Consider a setup where a length of wire is cut at the middle and the two halves connected together via a battery and switch, as illustrated by Figure 1. When the switch closes, a step voltage is applied between the two conductors. This creates a current step which propagates to the right at velocity *v* along the right-hand conductor, whilst another step propagates to the left along the left-hand conductor at the same velocity. Figure 2 illustrates the distribution of current *Ia* after a time period of 1ns has elapsed. It is assumed that *v* = 300Mm/s.

But that is not the whole picture. Over the same time-period, an electromagnetic wave propagates outwards in all directions, forming the spherical wavefront shown in Figure 3. Current *Ie* is departing from the conductor into the environment. Antenna theory analyses this effect and introduces the concept of the 'radiation resistance', *Rrad*. That is, the effect of the environment can be represented by the circuit model of Figure 4.

Not all of the current departs into the environment. Most of it propagates to the end of each wire. *Ia* and *Ie* can be described as 'partial currents'.

At any instant, the wavefront intersects with a cross-section of each conductor. At the right-hand section, electrons are moving radially into the body of the conductor to leave a net positive charge on the surface. At the left-hand section, electrons move radially from the body of the conductor to the surface, to leave a net negative charge. The charged surface propagates at the same velocity as the wavefront.

There is a superabundance of electrons in the conducting material of the wires. So the electrons need not move very far to compensate for the action of the photons, which are the real charge carriers. The terms 'field propagation', 'current' and 'photon movement' have essentially the same meaning.

If a second conductor is laid alongside the first and the experiment repeated, the same thing happens. The wavefront propagates outwards. But this time it propagates along both wires in parallel. This constitutes antenna-mode current which flows along the outer surfaces of the conductors.

The second conductor also acts as a receiving antenna, where the induced current flows back towards the center. Most of the current flows out along the transmitting conductor and back along the receiver. This is the differential-mode current.

The conductor pair is now acting as a transmission line with a differential-mode current pulse following behind the antenna-mode wavefront. Most of the electrical energy is concentrated in the region between the two conductors. Since the two wires are both insulated and are close together, there is a significant amount of dielectric material in the path of the differential-mode wavefront. This is in contrast with the antenna-mode wavefront where the dielectric material is mostly air. Solid dielectric has a higher relative permittivity than air.

Since the velocity of propagation is inversely proportional to the square root of the relative permittivity, the antenna-mode current propagates at a higher velocity than the differential-mode current. This explains the phenomenon observed by Ivor Catt and reported in the January 2011 issue of Electronics World. As a short, sharp, current pulse propagates along a transmission line, it separates out into two components; one travelling faster than the other.

Ian Darney

Figure 1: Long wire configured as a dipole antenna

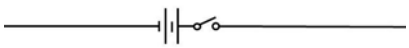

Figure 2: Current in conductors 1ns after switch closure

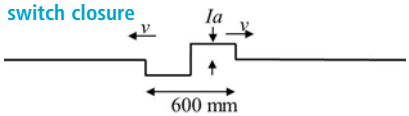

Figure 3: Wavefront 1ns after switch closure

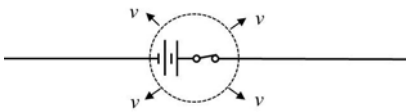

Figure 4: Circuit model

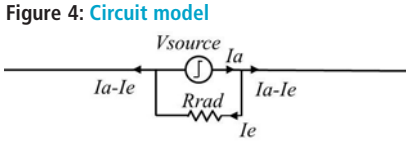

IF YOU WOULD LIKE TO COMMENT

on this subject or any other that you have read in Electronics World, please write to the Editor at Svetlana.josifovska@stjohnpatrick.com

The publisher reserves the right to edit and shorten letters due to space constraints

PLEASE EMAIL YOUR LETTERS TO: SVETLANA.JOSIFOVSKA@STJOHNPATRICK.COM

ECOC Exhibition

17th-19th September 2012, Amsterdam RAI

www.ecocexhibition.com

The ECOC Exhibition is one of the largest optical communications events in the world. It is the meeting place for all those involved in the fibreoptic communications industry, so if you are a manufacturer, supplier or service provider in this field, you can't miss ECOC 2012. Held from the 17th-19th September at the Amsterdam RAI, the event offers many new and exciting features and events that you could be part of.

Optical Connections, Market Focus

Walking the show floor among the 350 exhibitors, one area not to be missed is the Optical Connections Market Focus, sponsored by Opnext. Now in its seventh year, the Market Focus sessions boast an even broader selection of informative and ontrend optical communications topics planned for discussion, with the 2012 schedule set to be the largest to date. Past presenters include Google, JDSU, Openxt, Ericsson, Nokia Siemens Networks, Huawei, Deutsche Telekom and many more.

The content for the sessions is overseen by a panel of highly-respected industry experts that are the decision-making force behind the presentations. Market issues to be discussed this year include: service provider optical transmission, future technology applications, optical network agility, optical integration and digital photonics, data centre and optical interconnect technologies, packet-optical transport, test and measurement and massmarket broadband fibre access.

FTTx Centre

Returning to the show floor for 2012 is the FTTx centre. The centre will have 11 specialist zones all featuring the latest cutting edge technologies and various live interactive engineering demonstrations including fibre optic network delivery methods and OSP and ISP for FTTx vendor independent products. Run in conjunction with CTTS training with zone sponsors including Anritsu, FTTH Council Europe, ITW Chemtronics, Prysmian Group and Sticklers by MicroCare Corp it is certainly worth a visit, CTTS technical experts will be on hand to give tours of the centre and explanations of the products.

ECOC TV

Once again the camera crews of ECOC TV will tour the exhibition floor and market focus areas inviting views, news and opinions from the exhibitors and visitors alike – so don't be scared to get in front of the camera and promote your company and industry news and views!

This year will also see the introduction of a giant Twitter wall at the show's entrance, where exhibitors and visitors will be able to keep their followers up to date on the latest event news, gossip and highlights.

The ECOC Exhibition and all feature areas are free to attend and registration is now open.

To see who is already exhibiting, as well as plan your visit, go to: **www.ecocexhibition.com**

Amsterdam
ecoc 2012
Exhibition 17 – 19 September
Conference 16 – 20 September

Flash Analog to DIGITAL CONVERTER (ADC)

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

There are several types of ADCs available, depending on the application. They are usually classified into three main categories, depending on their speed of operation. The three types of ADCs are low-speed/serial ADC, medium speed ADC and high speed ADC. Typically, the serial ADCs have very high resolution, which means they support high accuracy, whereas high-speed ADCs operate at very high frequencies but have relatively low resolution.

Many applications require high-speed ADCs with a conversion speed of one clock-cycle. Of many analog-to-digital converters, flash ADC, also known as parallel ADC, holds its importance because of its high speed operation. With a conversion speed of only one clock-cycle flash ADC is the fastest architecture available, limited only by comparator and gate propagation delays.

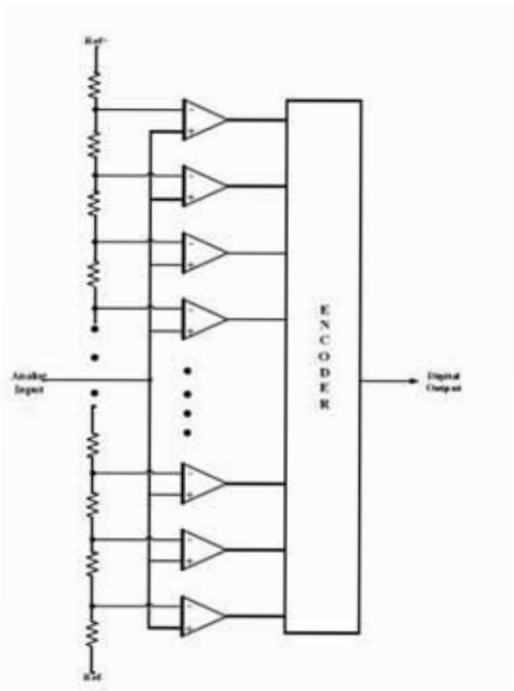

Figure 1: Flash ADC

The concept of a flash ADC (Figure 1) is straight forward. It basically compares the analog input to a set of reference voltages and determines the threshold to which the input is the closest. In general, an N-bit flash ADC consists of a resistor string, a set of comparators and a digital encoding network. The resistor string is composed of 2N resistors which are connected between Ref+ and Ref- to produce a unique reference voltage for each of the comparators as shown in Figure 1. The difference between these reference voltages is equal to the least significant bit (LSB) voltage.

> *Serial ADCs have very high resolution, which means they support high accuracy, whereas high speed ADCs operate at very high frequencies but have relatively low resolution*

The 2N-1 comparators produce the thermometer code (TC). It is called 'thermometer code' because, as the amplitude of the analog input increases, the number of ones in the output increases linearly similar to the mercury rise in a thermometer, and the digital encoding network converts 2N-1 inputs to N-bit binary code (BC). The digital encoding network comprises 1-out-of-N code generator circuit for intermediate conversion, which is usually implemented with XOR gates.

For example, a four-bit flash ADC consists of sixteen resistors generating fifteen different reference voltages for the comparators. The comparators generate a fifteen-bit thermometer code, which is encoded to four bits digital output using an encoder.

The main disadvantage of the high speed architectures is that they compromise speed with area and so does the flash

ADC. Unfortunately, it is the most component-intensive ADC architecture for any given number of output bits. With each additional output bit, the number of required comparators doubles. The increased transistor count increases power dissipation. This also results in significant capacitive loading and large die size, which directly affects cost.

Flash ADCs have been implemented in many technologies, varying from silicon-based bipolar (BJT) and complementary metal oxide FETs (CMOS) technologies to rarely used III-V technologies. Often this type of ADC is used as a first medium-sized analog circuit verification.

An additional advantage of the flash converter, often overlooked, is its ability to produce a non-linear output. With equal-value resistors in the reference voltage divider network, each successive binary count represents the same amount of analog signal increase, providing a proportional response. For special applications, however, the resistor values in the divider network may be made non-equal. This gives the ADC a custom, nonlinear response to the analog input signal. No other ADC design offers this signal-conditioning behavior with just a few component value changes.

Designs with power-saving capacitive reference ladders have been demonstrated. In addition to clocking the comparator(s), these systems also sample the reference value on the input stage. As the sampling is done at a very high rate, the leakage of the capacitors is negligible.

Recently, offset calibration has been introduced in the flash ADC designs. Instead of properly designing the analog circuit (which actually means increasing the component sizes to suppress variation), the offset is removed during use. A test signal is applied and each offset of each comparator is calibrated to below the LSB size of the ADC. Due to the heavy calibration effort the designs have – up to now – always been limited to 4-bits.

Another recent improvement in many flash ADCs is the inclusion of error correction. When the ADC is used in harsh environments or fabricated in very small integrated circuit processes, there is a heightened risk of a comparator randomly outputting a wrong code. Bubble error correction is a digital correction mechanism that will prevent a comparator that has tripped high from outputting a high code if it is surrounded by comparators that have not tripped high.

A pipelined ADC employs a parallel structure in which each stage works on a few bits of successive samples concurrently. This design improves speed at the expense of power and latency, but each pipelined stage is much slower than a flash section.

The pipelined ADC requires accurate amplification in the DACs and inter-stage amplifiers, and these stages have to settle at the desired linearity level. By contrast in a flash ADC the comparator only needs to be low offset and to resolve its inputs at a digital level; there is no linear settling time involved. Some flash converters require pre-amplifiers to drive the comparators. Here, gain linearity needs to be specified carefully. ●

Harwin Adds New Size Coin Cell Holders To Its EZ-Boardware Range

Harwin has expanded its EZ-BoardWare range of PCB hardware products with four new coin cell holders. As with all the EZ-BoardWare products, the EZ Coin Cell Holders, which cover battery sizes BR1225/CR1225 and BR2032/CR2032 (12.5 and 20.0mm diameter), are designed to reduce manufacturing costs by eliminating time-consuming secondary assembly operations.

The new parts include a single-piece SMD holder designed to allow 20.0mm diameter by 3.25mm thick coin cells to be easily and securely connected to the PCB. Manufactured in phosphor bronze and tin plated, these products are available on tape and reel packaging, ideally suiting them to automatic placement systems. Also newly launched are three contact/insulator style products, two of which are low profile, surface mounted products available in tape and reel, while the other is a departure from the rest of the EZ-Boardware range, being a space-saving, through-board (vertical) mount device rather than an SMD.

www.harwin.com

AVX's New 2-Piece Vertical Plug Connector Provides Highest Temperature Range Available

AVX Corporation has developed a low pin-count, vertical 2-piece connector for commercial and industrial applications.

The 9159 Series of robust 2-piece board-to-board (BTB) connectors were originally developed for the harsh environments of linear solid state lighting (SSL) with coplanar PCB mating. These connectors have been tested to 1000 hours at 125ºC.

Providing a 5A current rating with a pin count of 2 through 6 positions, the low profile 9159 Series connector features a gold-plated BeCu contact system that delivers higher signal and mating integrity than tin contact systems.

The vertical 9159 Series provides perpendicular PCB mating, which satisfies a more traditional PCB interconnect scheme typically used in standard commercial and industrial applications.

In addition to solid state lighting fixtures, the 2-piece vertical plug 9159 Series connector is well-suited for portable applications, including handheld terminals, scanner, instrumentation meters, patient monitoring devices and industrial electronic modules/controls.

www.avx.com

3.3V HIGH-CURRENT/HIGH-TEMPERATURE LINEAR HALL-EFFECT CURRENT SENSOR ICS

The new ACS759 from Allegro MicroSystems is a high-current/high-temperature linear Hall-effect current sensor IC that provides an economical and precise solution for AC or DC current sensing in 3.3V supply applications.

The device consists of a precision, low-offset linear Hall sensor circuit with a copper conduction path located near the die. Applied current flowing through this copper conduction path generates a magnetic field which is sensed by the integrated Hall IC and converted into a proportional voltage.

Device accuracy is optimised through the close proximity of the magnetic signal to the Hall transducer. A precise, proportional voltage is provided by the low-offset, chopper-stabilised BiCMOS Hall IC, which is programmed for accuracy at the factory.

High-level immunity to current conductor dV/dt and stray electric fields, offered by Allegro proprietary integrated shield technology, guarantees low ripple at the output and low-offset drift in high-side high-voltage applications.

www.allegromicro.com

HAN-FAST LOCK PROVIDES SIMPLE INTERFACE BETWEEN PCBS AND HIGH-CURRENT CONNECTORS

Han-Fast Lock from HARTING is a new circuit-board connection technology that allows the company's Han range of high-current connectors to be connected directly to printed circuit boards (PCBs) with ease and flexibility and without the need for additional components.

The Han-Fast Lock circuit-board connector is a compact, easy-to-handle interface that can handle currents up to 60A and requires only one contact point to make the connection. Conductor lines are through-plated, and are each provided with a drill hole and supporting point.

Han-Fast Lock is supported by a range of pre-assembled system solutions. For example, a Han Q 4/2 connector with pre-connected litz multi-stranded wires and circuit board contact can be supplied as a pre-tested system for direct connection to the circuit board.

Litz wire is a type of cable that is designed to reduce the skin-effect and proximity-effect losses in conductors at frequencies up to about 1MHz.

www.harting.com

Lightest And Smallest Industrial SSD Is Available Now

Apacer markets one of the lightest and smallest industrial SSD (32.5mm x 29.4mm), weighing only 6 grams – the SDM4 7P/180D-LPH. The unique hook design of the product, along with the patented innovative concept of ridding of power cables, fully demonstrates the advantages of the shockproof feature and built-in power supply, maintaining signal integrity and greatly increasing product stability. Additionally, Apacer also provides the general power-cable solution for the traditional designed PCB.

For the exterior of the product, this SATA interfaced SSD is designed to be only 12% the size of a 2.5-inch hard disk. Different from the monotonous design of typical ones, it is finely proportioned at 2:1 ratio with smile-like curve in the front side. The U-shaped troughs at both sides give better grip when users are removing/plugging the device, while adding more options to the sophisticated design. The slim size enables a compact design space for client's system.

www.apacer.com

High Voltage System SourceMeter for High Power Semiconductor Test

Keithley Instruments introduced Model 2657A High Power System SourceMeter instrument. This model adds high voltage to the company's Series 2600A System SourceMeter family of high speed, precision source measurement units. Together, these instruments allow Keithley's customers to characterize an even broader range of power semiconductor devices and materials.

There is a built-in 3,000V, 180W source and a high speed 6-1/2-digit measurement engine that enables 1fA (femtoamp) current measurement resolution to support the low-leakage requirements of next-generation power semiconductor devices.

The Model 2657A is optimized for high voltage applications such as testing power semiconductor devices, including diodes, FETs and IGBTs, as well as characterizing newer materials such as gallium nitride (GaN), silicon carbide (SiC) and other compound semiconductor materials and devices. It is also useful for characterizing high-speed transients and performing breakdown and leakage tests on a variety of electronic devices at up to 3,000V.

www.keithley.com

CISSOID LAUNCHES VOLGA, A HIGH-TEMPERATURE, HIGH-SPEED, RAIL-TO-RAIL COMPARATOR

CISSOID has launched VOLGA, the first high-speed comparator guaranteed for operation from -55°C up to +225°C.

VOLGA is a comparator available in a tiny TDFP (thin dual flat pack) surface mount, hermetically-sealed ceramic package, with footprint of only 5 x 5.5mm. It requires a 5V (±10%) power supply and offers rail-to-rail input/output capability, as well as a shutdown mode that places the device in a low power consumption mode when the comparator function is not needed. The comparator features internal hysteresis for improved noise immunity. Its output implements a push-pull CMOS stage, that can sink or source up to 16mA. Typical propagation delay at 20mV overdrive is less than 30ns and has a low variation with temperature up to +225°C.

Applications include window comparators, threshold detectors and discriminators, zero-crossing detectors, clock regeneration, sensing and control. Compared to operational amplifiers implementation, the comparator VOLGA decreases cost and PCB area, with fewer passive components, higher speed and improved stability.

www.cissoid.com

Maxim Introduces Third-Generation TINI Power SoC Chipset

Maxim Integrated Products announced that its newest Power SoC chipset in the Galaxy S III operates with the Exynos 4412 quad-core applications processor to provide a smaller, much thinner and significantly more efficient smartphone.

Maxim's newest Power SoC chipset covers all power management, charging and USB multiplexing needs. It provides an optimized balance between size and flexibility to power the Samsung applications processor and baseband processor. It maximizes battery usage and enhances USB connectivity.

Managing the power for upwards of 60 channels, the chipset offers up to 20% more conversion efficiency compared to the previous generation. Maxim's unique Green Mode regulators and subregulation architecture, along with the company's proprietary low-power, submicron geometry process, extend standby and active battery life. The chipset also enables fastest battery charging with minimal heat generation. High integration and advanced design techniques reduce the size, thickness and number of external components so smartphones become even thinner.

www.maxim-ic.com

MULTI-FUNCTION JTAG MODULE ADDRESSES ANALOG TESTING

JTAG Technologies has introduced the JT 2149/DAF, a compact, mixed-signal (digital/analog/frequency) measurement module. The JT 2149/DAF is the first unit of its type to offer both digital and analog test access to PCBs via JTAG Technologies's widely-used QuadPod signal conditioning interface.

The JT 2149/DAF module has been designed to slot into JTAG Technologies's regular QuadPod transceiver system as used by the renowned DataBlaster series of boundary-scan/JTAG controller hardware. When connected to a circuit board via edge connector or test fixture/jig test pins the module enhances standard digital boundary-scan tests by enabling a series of analog and frequency measurements to be made.

Capabilities of the JT 2149/DAF module include 16 dual-purpose digital pins capable of digital I/O stimulus and response at voltages of 1.0 to 3.6V, plus frequency measurements of up to 128MHz on any pin. Twelve additional analog measurement channels can capture values from 0 to 33V with better than 10mV resolution.

www.jtag.co.uk

ONE MBIT VERSATILE SPI/8-BIT PARALLEL BUS SRAM

The VLSI Solution VS23S010 is an easy to use static RAM device. With one megabit capacity it is the largest capacity SPI RAM on the market. The memory can be accessed via a standard SPI compatible serial bus (1, 2 and 4-bit modes), or via an 8080-type 8-bit parallel bus similar to NAND flash devices. The main applications of the device are microcontroller RAM extension and audio buffer for internet and personal streaming.

One of the main applications for the VS23S010 is microcontroller RAM extension. Internet streaming applications, which are becoming increasingly popular, will benefit from the large buffer memory provided by VS23S010.

In addition to being used purely as a memory storage device, the VS23S010 also contains pattern generation logic. When VS23S010's RAM is accessed through the SPI bus, it can simultaneously output a repeating memory data pattern through four of its parallel bus pins.

www.vlsi.fi/samples/vs23s010

Fujitsu Expands Line of New 8FX 8-bit Microcontrollers for DC Motor Control

Fujitsu Semiconductor Europe (FSEU) expanded its "New 8FX" family of high-performance 8-bit microcontrollers. These now include nine new products from the 48-pin MB95690 series, featuring brushless DC motor control functionality, and six products from the 64-pin MB95810 series featuring analog comparator technology. As low-cost and low-power microcontrollers, products from both series offer 'right-sized' solutions for a wide range of applications, including home appliances, home electronics, personal care products, power tools and office automation equipment.

Brushless DC motors are superior in terms of their noiselessness and long life expectancy. Since such motors account for a significant percentage of a product's overall cost, there is a significant demand for lower cost microcontrollers to control brushless DC motors. In addition, there is also considerable demand for microcontrollers that can support relatively high-end products requiring fine-grained system control, such as temperature sensors, light sensors and other applications that monitor a range of environmental parameters for a device.

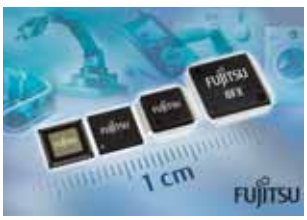

www.fujitsu.com

USING ESSEMTEC'S SMD TOWER FOR IGBT STORAGE IN A CLEAN ROOM

The HiPak from ABB Semiconductors is a high-performance semiconductor switching module for high current and voltage, and is known for its short switching time and reliability. As such, the product is used where high power must be switched fast, for example, on a train or in a wind power plant.

A HiPak module consists of several IGBT modules which during production need to be held in so-called washing baskets. These are very specific in nature, having to hold the required IGBTs, and ABB Semiconductors long searched for a suitable production storage system for its HiPak IGBT modules production. During a visit to Essemtec AG, the Swiss manufacturer learned about the SMD Tower, which is an automatic and scaleable storage system for electronic components in the form of reels and trays.

The Tower features a small footprint of only one square meter, and several Towers can be grouped for larger scale storage, which was ideal for HiPak production at ABB Semiconductors.

www.essemtec.com

New RF Building Block from CML Microcircuits

CML Microcircuits has released a new integrated quadrature IF/RF demodulator and quadrature modulator.
The CMX973 quadrature modulator/demodulator is the latest RF building block following on from the CMX970 and CMX971. The RF building block range is designed to provide flexible, high-performance ICs required for HF/VHF/UHF, professional radios, wireless data terminals, wireless microphones, marine and avionics radio systems.

The CMX973 is the basis of a high-performance RF transceiver system combining a flexible IF/RF quadrature modulator and quadrature demodulator in a single chip. Wide signal bandwidth at RF, 0dBm output, low noise and serial bus control delivers maximum flexibility and high-performance, all at an attractive low operating power.

The demodulator is suitable for superheterodyne architectures with IF frequencies of 20MHz to 300MHz and it may be used in low IF systems or those converting down to baseband. The modulator converts directly from baseband to the desired transmit frequency 20MHz to 1GHz and features quadrature phase correction to minimise unwanted spectral components.

www.cmlmicro.com

ALTERA'S QUARTUS II SOFTWARE DELIVERS UP TO 4X FASTER COMPILE TIMES

Altera released the latest version of its industry-proven Quartus II development software for FPGA design. Quartus II software version 12.0 provides customers additional productivity and performance advantages, such as 4X faster compile times for high-performance 28nm designs. Additional upgrades include broadened 28nm device support, enhanced Qsys system integration and DSP Builder tools, and improved intellectual property (IP) core offerings.

Stratix V FPGA users will achieve on average 35% faster compile times, while Cyclone V and Arria V FPGAs users will see on average 25% faster compile times with this software version, as compared to the company's previous release of the software.

Also with this release, Altera is adding support for the ARM AMBA AXI-3 interface in its Qsys system integration tool, giving users the flexibility to connect or mix IP cores and IP sub-systems based on different standard interfaces. Qsys is the FPGA industry's first system integration tool based on network-on-a-chip (NoC) technology.

www.altera.com

PANASONIC OFFERS EV 200A – ELECTRIC VEHICLE RELAYS

Panasonic Electric Works will soon release a 200A EV relay for electric drive trains, closing the gap between the 120A and 300A types. The 200A relay can carry 200A continuously or 300A for 15 minutes. Higher peak currents, due to acceleration of vehicle for example, for a few minutes or seconds are possible much beyond 200A. This means many sport cars and light commercial vehicles can be realized with this main contactor.

The 200A EV is built on the proven design and continuously developed EV relay family. Its contacts are integrated in a hermetically-sealed ceramic chamber. This contact chamber is constructed of high quality ceramic that separates the arc from other components, electrically and thermally. The chambers superior mechanical stability increases the system safety even in the event of a malfunction.

Another special feature is that the contact chamber, including the actuating mechanism, is encapsulated and permanently filled with hydrogen gas.

www.panasonic-electric-works.co.uk

ELEMENT14 DEVELOPS NEXT GENERATION COMPONENT CATALOGUES WITH CUSTOMERS

Farnell element14 is pioneering a new generation of component catalogue designed to make the buying process easier. It has developed a new online catalogue, named MROCAT, in BETA, to engage customers in finalising the design of what essentially will be their very own dynamic catalogue.

Farnell element14 is launching MROCAT in BETA version catalogue to a wide range of its customers who are focused on maintenance and repair products. Accessible from PC, Mac, iPad and Android tablets, MROCAT provides a fully illustrated catalogue page style browsing experience with real-time stock and price information and encourages users to build shopping lists, add personal notes and bookmarks.

MROCAT Beta contains over 500 pages and covers the connector, electromechanical, automation, power, test and measurement, tools and maintenance sections of the Farnell element14 range. This first edition lists 17,000 products from 300 brands and includes live daily stock and pricing.

http://farnell.com/mrocat/

SemiSouth Garners 30th US Patent

SemiSouth Laboratories has announced its 30th US patent granted by the US Patent and Trademark Office. SemiSouth designs and manufactures silicon-carbide (SiC) power semiconductor transistors and diodes, which are rapidly gaining market share in the solar, UPS, traction, wind, automotive and aerospace industries for their superior performance in high-efficiency, harsh-environment power applications.

US Patent 8,169,022 was issued on May 1, 2012, and is entitled "Vertical Junction Field Effect Transistors and Diodes Having Graded Doped Regions and Methods of Making". It was co-invented by Dr. Michael Mazzola, a co-founder of SemiSouth in 2000 when the company spun off of Mississippi State University.

"The underlying technology in this patent allows SemiSouth to fine tune its already performance-leading vertical channel junction field effect transistors and diodes to get ever closer to the unipolar theoretical limit. Customers can expect even better value from the products based on this patent," said Mazzola.

www.semisouth.com

£2.5M INVESTMENT TO GROW UK SPACE TECHNOLOGY INDUSTRY

The UK government is awarding £2.5m to 22 British companies to support them in developing commercial products and services using space technology or space-derived data. The involved businesses will match this funding, which will bring the total value of the R&D to nearly £5m.

The grant funding – from the UK Space Agency, the Technology Strategy Board and the South East England Development Agency – will support 28 fast-track research and development projects, each lasting between six and nine months. Projects will cover a broad range of growth opportunities: from novel propulsion for cubesats, through technology to exploit the Galileo navigation satellite system, to techniques for crop monitoring from space.

The investment forms part of the UK Space Agency's National Space Technology Programme (NSTP), which will see government investment of £10m to help UK industry exploit growth opportunities in the space sector and improve the UK's space technology capabilities.

"The UK space industry is one of the fastest growing sectors in the country, contributing £7.5bn annually to the UK economy," said Dr David Williams, Chief Executive of the UK Space Agency. "The National Space Technology Programme will help us further this success by providing an opportunity for promising UK space technologies and applications to be developed to meet their full commercial potential and for businesses to explore collaborations with other sectors to establish services in new markets."

PROFESSOR DR DOGAN IBRAHIM, Near East University in Nicosia, Cyprus: I think it is about time that the UK research and development wakes up! It is music to the ears to hear that 22 British companies will participate in developing space-derived technologies. In addition to creating new professional jobs, this initiative will once again show that the UK is capable of developing state-of-the-art space-related projects and will help to expand existing capabilities of the UK space technologies.

BARRY MCKEOWN, RF and Microwave Engineer in the Defence Industry, and Director of Datod Ltd, UK: Space evokes strong emotions and imaginations, so in this case a dose of reality is needed to bring matters back to earth. The UK Space Agency represents just £313m in public funding, set against a global budget of $65bn (approx £42bn). Off the back of 15% growth in this sector a recent IoD (Institute of Directors) report has also got in the act, calling for a UK spaceport, funded via private initiatives once the proper regulation is addressed.

I note that Avanti and Inmarsat are not among the usual suspects gaining funding. I would also reference the difficulties arising with Light Squared satellite broadband here.

According to the IoD report, space tourism is highlighted, as there are many people willing to pay $200,000 for a few hours in space, which puts this funding program into its correct perspective.

I am not really against space initiatives, it is just that the level of funding and particularly the matched-funding/publicly-funded seed models adopted are wholly inadequate for entry into this global industry – it is all just window dressing.

Note that at the recent WRC-12 spectrum was allocated for supporting lunar missions, earth observation, and climate monitoring and digital maritime communications and that there were developments for improving the registering of satellite networks. The next WRC-15 shall examine allocating spectrum for command and control of unmanned aircraft from space. So this field shall eventually all come down to China and a few billionaires with the time and imagination to match, who really don't care about the cost just as they back football clubs without any economic reality applying.

> The level of funding and particularly the matched-funding/publicly-funded seed models adopted are wholly inadequate for entry into this global industry – it is all just window dressing

MAURIZIO DI PAOLO EMILIO, Telecommunications Engineer, INFN – Laboratori Nazionali del Gran Sasso, Italy: Space exploration plays a big part in our day-to-day life whether we know it or not.

Space exploration has aided the military and many modern-day technologies were discovered by NASA through its space studies, so without space exploration and research, things like GPS and breast cancer detection wouldn't be possible.

There are dangers and disadvantages to research in this field, but if we continue to explore it, test its technologies and take some risks, I am sure that most of the problems will be and can be fixed. Of course, nothing about space exploration will ever be perfect, but nothing is ever completely risk-free. In the past, space exploration has proven to be extremely beneficial and will continue to do so in the future, just so long as we continue to work hard to make it better.

Even though there are a few disadvantages to space travel, there are also many benefits including: the jobs it provides, the medical discoveries, all of the new technologies and innovations, and, perhaps most importantly, a greater understanding of Earth as well as everything beyond it.

HAFIDH MECHERGUI, Associate Professor in Electrical Engineering and Instrumentation, University of Tunisia: The decision taken by the UK government to invest, support and encourage the space industry comes across as an intelligent vision as this initiative will make possible to develop this field on a large scale.

It is notable how many companies have been invited to take part and will be financed by this program. But to have successful co-operation between all of them, it is necessary to prepare a clear coordination structure of the activities and programs to be developed. Generally, research into space applications remains very expensive, so this will benefit from financial support.

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 !

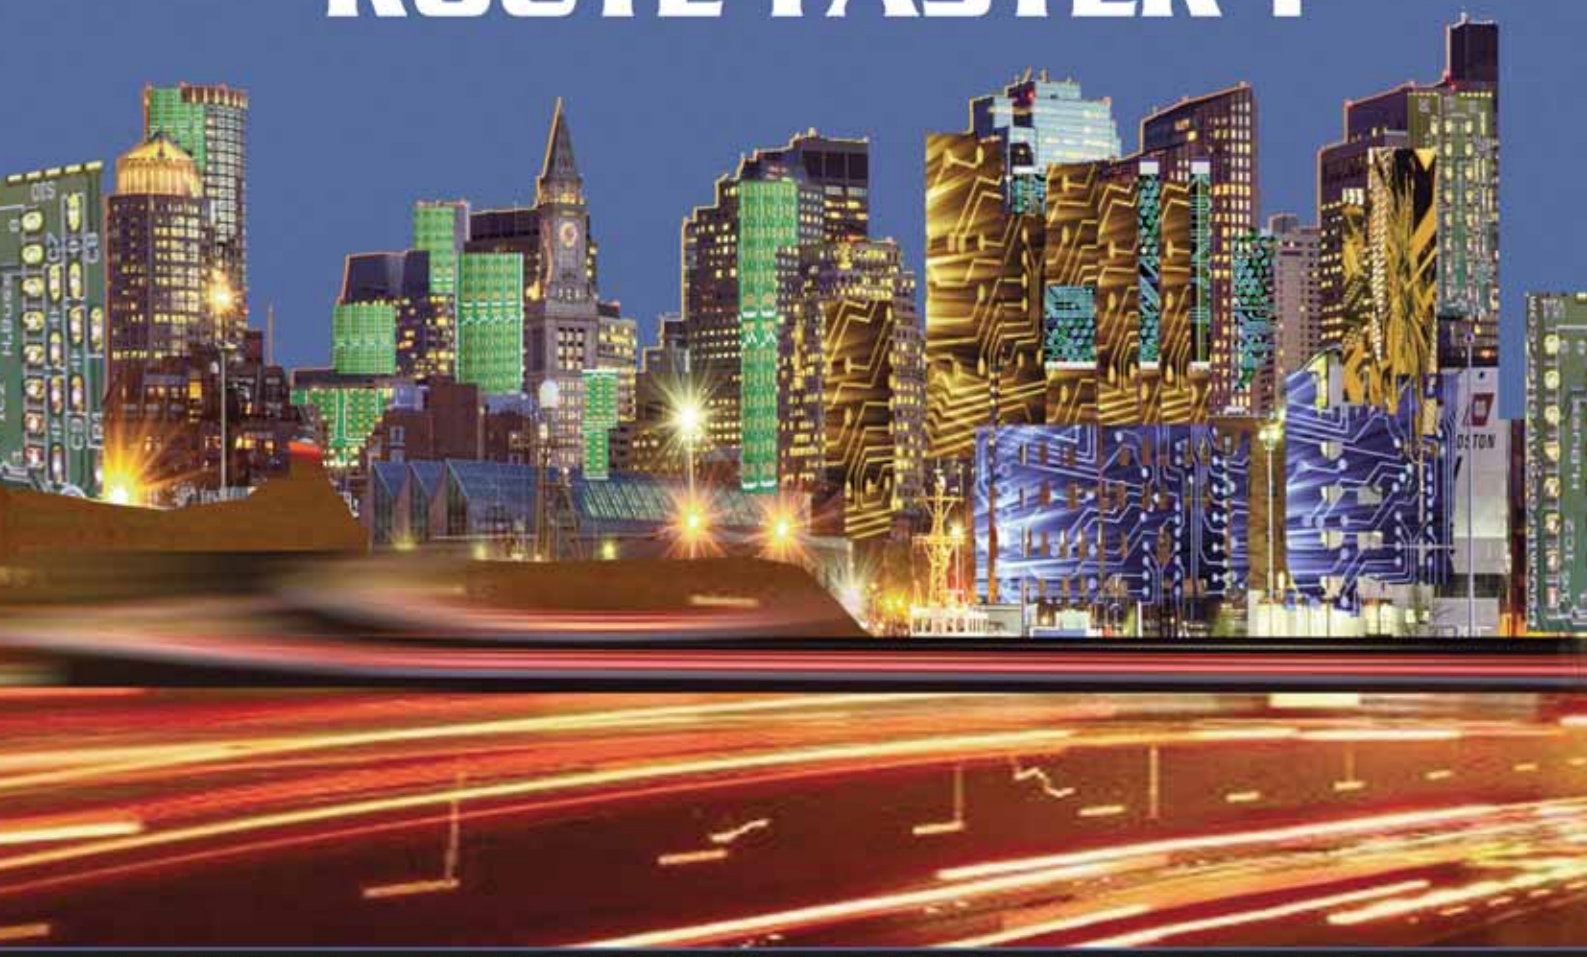

WITH PROTEUS PCB DESIGN

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

Generates complex conditions. And extreme confidence.

Now available from Microlease, the Agilent MXG X-Series signal generator lets you simulate complex signal technologies and extreme conditions with unmatched performance. Like the cost-effective EXG, it features industry-leading ACPR, EVM and output power. So you can be sure that the products you're developing are perfectly ready for real-world performance.

MXG X-Series Signal Generators
Analog and vector; 9 kHz up to 6 GHz
Lowest phase noise in the industry: -146 dBc @ 1 GHz
160 MHz bandwidth; ± 0.2 dB flatness
Real-time 4G, video, GNSS with Signal Studio software

microlease

Agilent Authorised Technology Partner

Contact us:
Tel: +44 (0) 208 4200205
www.microlease.com/agilent

Trade in your current signal generator today and receive up to 50% off your upgrade
www.agilent.com/find/MXG_wrls

Anticipate ___Accelerate ___Achieve

Agilent Technologies