Volume 116 • Issue 1886
February 2010 • £4.60

ELECTRONICS WORLD

www.electronicsworld.co.uk

THE ESSENTIAL ELECTRONICS ENGINEERING MAGAZINE

SOFTWARE REPORT:

- COMPLIANCE
- CLEANLINESS
- VIRTUALIZATION
- TOOLS

RTOS:

- SCHEDULING AND TIMING ANALYSIS TECHNIQUES

SIMPLE AND EFFICIENT DEBUG

4-CHANNEL 40MHZ-300MHZ OSCILLOSCOPES

TECHNOLOGY
POWER EFFICIENCY BREAKTHROUGH FOR MMI

RF COLUMN
AN ENGINEER'S NEW YEAR RESOLUTIONS

FOCUS
BATTERIES REGULATIONS: A ROUND-UP OF RESPONSIBILITES

Find your distributor:

http://www.mikroe.com/en/distributors/

DEVELOPMENT TOOL JUST THE WAY YOU LIKE IT

Everything you've always wanted from a development tool
Experience the ease of creating your own electronic devices!

Complete PIC development solution

Speed up your prototype development with the EasyPIC6 Development System. The EasyPIC6 comes packed with examples that makes your PIC development fast and easy.

Save time & money!

Get extra value for your money with this first-class tool for PIC development solutions. Get into the PIC world faster and easier than ever before with the EasyPIC6 Development System.

Designed to suit your needs

Your development time can be considerably reduced, resulting in an early prototype design and fast time-to-market for your end product.

Thanks to many new features, you can start creating your own devices immediately. **EasyPIC6** supports **8-, 14-, 18-, 20-, 28- and 40- pin PIC** microcontrollers. The **mikroICD** (Hardware In-circuit Debugger) enables very efficient step by step debugging. Examples in **C**, **BASIC** and **Pascal** are provided with the board.

Hardware In-Circuit Debugger for step by step debugging at hardware level

On-Board 2x16 Serial LCD Display

Port Expander provides easy I/O expansion (2 additional ports) using serial interface

Keypad enables easy and fast data entry

Full-featured and user-friendly development board for PIC microcontrollers

High-Performance USB 2.0 On-Board Programmer

lead-free

RoHS 2002/95/EC

We deliver our products across the globe and our satisfied customers are the best proof of our first-rate service. The company is an official consultant on PIC microcontrollers and a Third Party Partner of Microchip Technology. We are also an official consultant and Third Party Partner of Cypress Semiconductors since 2002 and official consultants to Philips Electronics. All our products are RoHS compliant.

mikroElektronika

SOFTWARE AND HARDWARE SOLUTIONS FOR EMBEDDED WORLD

www.mikroe.com

16

FEATURES

REGULARS

Electronics World is published monthly by
Saint John Patrick Publishers Ltd,
6 Laurence Pountney Hill, London, EC4R 0BL.

Quasar Electronics Limited
PO Box 6935, Bishops Stortford
CM23 4WP, United Kingdom
Tel: 01279 467799
Fax: 01279 267799
E-mail: sales@quasarelectronics.com
Web: www.quasarelectronics.com

Postage & Packing Options (Up to 0.5Kg gross weight): **UK Standard 3-7 Day Delivery** - £4.95; **UK Mainland Next Day Delivery** - £9.95; **Europe (EU)** - £8.95; **Rest of World** - £12.95 (up to 0.5Kg)
!Order online for reduced price UK Postage!
We accept all major credit/debit cards. Make cheques/PO's payable to Quasar Electronics. Prices include 15.0% VAT.
Please visit our online shop now for details of over 500 kits, projects, modules and publications. Discounts for bulk quantities.

QUASAR electronics
The Electronic Kit Specialists Since 1993

01279 Credit Card Sales 467799

Motor Drivers/Controllers

Here are just a few of our controller and driver modules for AC, DC, Unipolar/Bipolar stepper motors and servo motors. See website for full range and details.

Computer Controlled / Standalone Unipolar Stepper Motor Driver

Drives any 5-35Vdc 5, 6 or 8-lead unipolar stepper motor rated up to 6 Amps. Provides speed and direction control. Operates in stand-alone or PC-controlled mode for CNC use. Connect up to six 3179 driver boards to a single parallel port. Board supply: 9Vdc. PCB: 80x50mm.
Kit Order Code: 3179KT - **£15.95**
Assembled Order Code: AS3179 - **£22.95**

Computer Controlled Bi-Polar Stepper Motor Driver

Drive any 5-50Vdc, 5 Amp bi-polar stepper motor using externally supplied 5V levels for STEP and DIRECTION control. Opto-isolated inputs make it ideal for CNC applications using a PC running suitable software. Board supply: 8-30Vdc. PCB: 75x85mm.
Kit Order Code: 3158KT - **£23.95**
Assembled Order Code: AS3158 - **£33.95**

Bi-Directional DC Motor Controller (v2)

Controls the speed of most common DC motors (rated up to 32Vdc, 10A) in both the forward and reverse direction. The range of control is from fully OFF to fully ON in both directions. The direction and speed are controlled using a single potentiometer. Screw terminal block for connections.
Kit Order Code: 3166v2KT - **£22.95**
Assembled Order Code: AS3166v2 - **£32.95**

DC Motor Speed Controller (100V/7.5A)

Control the speed of almost any common DC motor rated up to 100V/7.5A. Pulse width modulation output for maximum motor torque at all speeds. Supply: 5-15Vdc. Box supplied. Dimensions (mm): 60Wx100Lx60H.
Kit Order Code: 3067KT - **£17.95**
Assembled Order Code: AS3067 - **£24.95**

Most items are available in kit form (KT suffix) or assembled and ready for use (AS prefix).

Controllers & Loggers

Here are just a few of the controller and data acquisition and control units we have. See website for full details. Suitable PSU for all units: Order Code PSU445 £7.95

8-Ch Serial Isolated I/O Relay Module

Computer controlled 8-channel relay board. 5A mains rated relay outputs. 4 isolated digital inputs. Useful in a variety of control and sensing applications. Controlled via serial port for programming (using our new Windows interface, terminal emulator or batch files). Includes plastic case 130x100x30mm. Power Supply: 12Vdc/500mA.
Kit Order Code: 3108KT - **£64.95**
Assembled Order Code: AS3108 - **£79.95**

Computer Temperature Data Logger

4-channel temperature logger for serial port. °C or °F. Continuously logs up to 4 separate sensors located 200m+ from board. Wide range of free software applications for storing/using data. PCB just 45x45mm. Powered by PC. Includes one DS1820 sensor.
Kit Order Code: 3145KT - **£19.95**
Assembled Order Code: AS3145 - **£26.95**
Additional DS1820 Sensors - **£3.95 each**

Rolling Code 4-Channel UHF Remote

State-of-the-Art. High security. 4 channels. Momentary or latching relay output. Range up to 40m. Up to 15 Tx's can be learnt by one Rx (kit includes one Tx but more available separately). 4 indicator LED 's. Rx: PCB 77x85mm, 12Vdc/6mA (standby). *Two and Ten channel versions also available.*
Kit Order Code: 3180KT - **£49.95**
Assembled Order Code: AS3180 - **£59.95**

DTMF Telephone Relay Switcher

Call your phone number using a DTMF phone from anywhere in the world and remotely turn on/off any of the 4 relays as desired. User settable Security Password, Anti-Tamper, Rings to Answer, Auto Hang-up and Lockout. Includes plastic case. Not BT approved. 130x110x30mm. Power: 12Vdc.
Kit Order Code: 3140KT - **£74.95**
Assembled Order Code: AS3140 - **£89.95**

Infrared RC Relay Board

Individually control 12 on-board relays with included infrared remote control unit. Toggle or momentary. 15m+ range. 112x122mm. Supply: 12Vdc/0.5A
Kit Order Code: 3142KT - **£59.95**
Assembled Order Code: AS3142 - **£69.95**

New! 4-Channel Serial Port Temperature Monitor & Controller Relay Board

4 channel computer serial port temperature monitor and relay controller with four inputs for Dallas DS18S20 or DS18B20 digital thermometer sensors (£3.95 each). Four 5A rated relay channels provide output control. Relays are independent of sensor channels, allowing flexibility to setup the linkage in any way you choose. Commands for reading temperature and relay control sent via the RS232 interface using simple text strings. Control using a simple terminal / comms program (Windows HyperTerminal) or our free Windows application software.
Kit Order Code: 3190KT - **£69.95**

PIC & ATMEL Programmers

We have a wide range of low cost PIC and ATMEL Programmers. Complete range and documentation available from our web site.

Programmer Accessories:
40-pin Wide ZIF socket (ZIF40W) £14.95
18Vdc Power supply (PSU120) £19.95
Leads: Serial (LDC441) £3.95 / USB (LDC644) £2.95

USB & Serial Port PIC Programmer

USB/Serial connection. Header cable for ICSP. Free Windows XP software. Wide range of supported PICs - see website for complete listing. ZIF Socket/USB lead not included. Supply: 16-18Vdc.
Kit Order Code: 3149EKT - **£49.95**
Assembled Order Code: AS3149E - **£59.95**

USB 'All-Flash' PIC Programmer

USB PIC programmer for all 'Flash' devices. No external power supply making it truly portable. Supplied with box and Windows Software. ZIF Socket and USB lead not included.
Assembled Order Code: AS3128 - **£49.95**

See website for full range of PIC & ATMEL Programmers and development tools.

No.1 FOR KITS

www.quasarelectronics.com

Secure Online Ordering Facilities • Full Product Listing, Descriptions & Photos • Kit Documentation & Software Downloads

OPTICAL COMMS: TOP FIVE TRENDS AND AREAS TO WATCH IN 2010

What this piece states is that there are several issues to watch out for in optical communications in 2010

1. Continued Consolidation and Economic Recovery

BY SINCLAIR VASS

In 2010, it will be widely accepted that economic recovery is well underway across the world. And while there is significantly more life in the optical communications marketplace, the intense cost-reduction focus that occurred during the recession has resulted in even lower inventory levels and shorter lead-times requirements, making it even more challenging for NEMs and component suppliers to meet the volatile end carrier demand.

Optical vendors will, therefore, continue to optimise lean manufacturing practices and operations to support unpredictable fluctuations in market requirements. There will be industry consolidation as the large number of smaller players look to partner in their core segments, acquire or divest.

2. The Asia Effect

As large Asian NEMs rapidly gain traction in the worldwide market with low-cost system offerings, European and American NEMs will continue to feel the pressure. Some Asian NEMs have a business model that focuses on buying components at lower levels of integration and building systems in-house with cheap labour; often with support of local government subsidy. For other NEMS the opposite approach may work best, outsourcing the integration of optical hardware solutions at the circuit pack level, whilst focusing internally on core strengths. It is hard to predict whether both models can co-exist, but the winners will obviously be the companies with the best product portfolio and cost structure.

3. Evolution in the Supply Chain

In the wireless industry, there have been significant changes to the operating model, with service providers outsourcing network management to the NEMs. This trend may also reach the wireline business over the next five years. In this scenario, service providers would focus on creating new applications that drive more revenue over their networks. NEMs would become less component-focused and take on the responsibility of managing and optimising the network. This in turn would allow optical vendors to focus on building and testing critical solutions. Many players believe that the converged network will be a future reality and the boundaries between wireless and wireline networks will disappear.

4. 40G – 100G

The industry continues to deal with fragmentation in the 40G space and cost challenges for the development of both 40G and 100G products. Both markets will likely ramp over the next 2-4 years, but these current challenges need to be addressed in order to move forward. Standards bodies need to quicken the pace on ratifying new standards to curb the implementation of multiple proprietary products that support several different modulation schemes. 40G will in the future likely be supported by three dominant modulation formats superseding the current ODB scheme. These new modulation schemes will co-exist in the network but over time the most powerful coherent receiver intradyne solution will dominate.

During 2010, the advanced development functions within the larger NEMs will turn their focus towards even faster transmission rates to 400Gb/s and beyond.

5. Increased Use of WSS Solutions and Tunable Transceivers

WSS modules will see significant uptake over the next five years. Solutions will include growth in a broad range of WSS products.

Particularly over the next two years, solutions for directionless and colourless, meshed network designs will be vital. This will drive a range of new reconfigurable, colourless and non-blocking modules that will provide the necessary functionality to support this architecture. Particularly new WSS switching modules, flexible amplifiers and integrated tunable XFP transceivers solutions will emerge.

Sinclair Vass is Senior Director of Communication Sales for EMEA at JDSU

> **"MANY PLAYERS BELIEVE THAT THE CONVERGED NETWORK WILL BE A FUTURE REALITY AND THE BOUNDARIES BETWEEN WIRELESS AND WIRELINE NETWORKS WILL DISAPPEAR"**

EDITOR: Svetlana Josifovska
Email: svetlanaj@stjohnpatrick.com
PRODUCTION MANAGER: Tania King
Email: taniak@stjohnpatrick.com
DISPLAY SALES: Neil Coshan
Tel: +44 (0) 20 7933 8977
Email: neilc@stjohnpatrick.com
PUBLISHING DIRECTOR: Chris Cooke

PUBLISHER: John Owen
SUBSCRIPTIONS:
Saint John Patrick Publishers
PO Box 6009, Thatcham,
Berkshire, RG19 4QB
Tel: 01635 879361
Fax: 01635 868594
Email: electronicsworld@circdata.com

SUBSCRIPTION RATES:
1 year: £46 (UK); £67.50 (worldwide)
MISSING ISSUES:
Email: electronicsworld@circdata.com
NEWSTRADE:
Distributed by Seymour Distribution Ltd,
2 East Poultry Avenue, London, EC1A 9PT
Tel: +44 (0) 20 7429 4000

PRINTER: William Gibbons Ltd
ISSN: 1365-4675

John Patrick Publishers

Power Efficiency Breakthrough for Mobile Memory Solutions

Memory architectures licencing company Rambus announced it has achieved a new breakthrough level of power efficiency with its latest silicon test vehicle developed through its Mobile Memory Initiative (MMI).

The silicon-validated results demonstrate that through the use of MMI innovations, a high-bandwidth mobile memory controller can achieve a power efficiency of 2.2mW/Gbps. This is nearly a one third improvement over the initial MMI silicon and significantly better than the estimated 10mW/Gbps of an LPDDR2 400 memory controller.

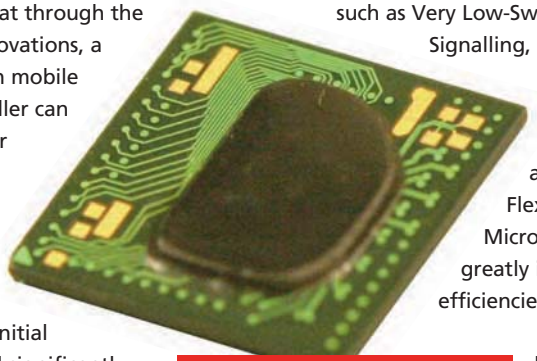

Rambus's MMI test silicon chip

"The performance demands of next-generation mobile devices are vastly outstripping the pace of battery technology improvements," said Martin Scott, senior vice president of Research and Technology Development at Rambus. "With the innovations developed through our Mobile Memory Initiative we can deliver advanced applications and maintain long battery life through our breakthroughs in both, bandwidth performance and power efficiency."

Rambus's MMI encompasses innovations such as Very Low-Swing Differential Signalling, FlexClocking Architecture and Advanced Power State Management. In addition, Rambus's FlexPhase and Microthreading technologies greatly improve the power efficiencies of mobile platforms.

The MMI was launched in February last year, with focus on achieving high bandwidth at extremely low power to enable advanced applications in next-generation smartphones, netbooks, portable gaming and portable media products. Operating at 4.3Gbps, a memory system using MMI innovations can deliver over 17GB/s of memory bandwidth from a single mobile DRAM device.

New Alchimer Seedless Wet Deposition Technology Eliminates Entire Step From TSV Film Stack Process

Massy, France, based Alchimer, a provider of technology for the deposition of nanometric films used in semiconductor interconnects and 3D through-silicon vias (TSV), announced a groundbreaking advance in TSV formation that eliminates one of the traditional metallization steps.

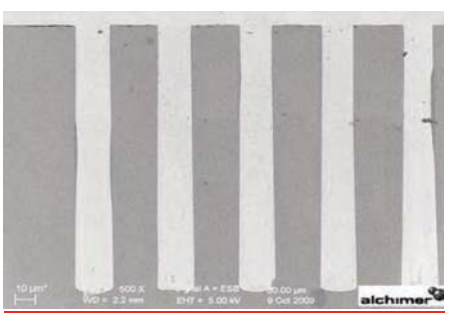

AquiVia XS is now available for demonstration

Alchimer already offers AquiVia wet deposition processes for isolation, barrier and seed layers in TSV metallization that uses electrografting, a nanotechnology solution based on surface chemistry formulations and processes to grow highly conformal and uniform layers in TSVs with aspect ratios up to and beyond 20:1.

But in its newest solution, AquiVia XS, by eliminating the seed-layer step from the standard isolation-barrier-seed process flow, Alchimer allows bulk fill to take place directly after application of the barrier layer. This is also the first deposition solution to support both nickel and copper metallization.

"The AquiVia XS approach provides a previously impossible combination of film quality, broad applicability and cost benefits," said Alchimer's CEO Steve Lerner.

Like AquiVia, AquiVia XS enables the use of existing plating equipment for layer deposition and completely eliminates all dry processing techniques from TSV metallization. As a result, cost of ownership of the via stack metallization process is reduced by up to 80%. Both products also result in highly conformal and uniform layers for TSVs with aspect ratios of 20:1 and beyond, even on the highly scalloped etch profiles produced by the DRIE/Bosch process.

AquiVia XS is now available for immediate demonstration and licensing.

■ Companies are slow to take up the innovation challenge that plastic electronics presents, says research from the Advanced Institute of Management Research (AIM Research).
The Plastic Electronics industry is predicted to be worth $335bn within 20 years, and although plastic electronics products are already being manufactured on a commercial scale, what is lacking, says the research, is involvement by product designers and market-led end users. In addition, companies in the retail, healthcare, transport, electronics and packaging industries have yet to understand the revolution that is taking place.
With Plastic Electronics electronic materials can be formed into circuits using cheap core materials. Products can be printed onto flexible surfaces like paper, film or fabric allowing the manufacture of thin, lightweight devices. The costs are low enough that wallpaper, for example, can be designed with integral lighting, and clothing can be printed with wearable electronics and solar cells.

■ Cranfield University has developed a new technology that promises to significantly reduce the manufacturing cost of complex micro-mechanical and micro-optical devices.
"Standard micro-fabrication techniques are often incompatible with high quality transducer materials such as shape memory alloys and functional ceramics," said Stephen Wilson, Senior Research Fellow at Cranfield University. "This is one of the major bottlenecks for the development of novel micro-scale systems. The new technology enables multi-material devices to be made that do not conform to the usual silicon MEMS stereotype."
The new methods can be used in the manufacture of many components and systems ranging in size from a few millimetres to a few 100s of nanometres. Applications include biomedical devices that can diagnose disease and electronically administer drugs, electronic noses to sniff out explosives or dangerous chemicals and environmental control systems for personal healthcare.
The project is part of a €3.2m research consortium entitled Q2M (Quality to Micro).

TSMC ANNOUNCES PROCESS TECHNOLOGIES FOR INTEGRATED LED DRIVERS

Taiwan Semiconductor Manufacturing Company unveiled modular BCD (Bipolar, CMOS DMOS) process technologies targeting high voltage integrated LED driver devices.

The new BCD technologies feature a voltage spectrum running from 12 to 60V to support multiple LED applications, including LCD flat panel display backlighting, LED displays, general lighting and automotive lighting. It covers process nodes from 0.6-micron to 0.18-micron with a number of digital core modular options for varying digital control circuit gate densities. The CyberShuttle prototyping service supports the 0.25-micron and 0.18-micron processes for preliminary function verification.

12-inch wafer fab

"The new BCD technologies for LED drivers are very leading edge in driving device integration. The associated PDKs feature highly accurate SPICE models that really enhance the potential for easy single chip design," said George Liu, Director, Industrial Business Development. "In addition, mismatching models help optimise current mismatching performance in multi-channel LED driver designs."

The DMOS process supports 72mohm per mm^2 at BV > 80V for a specific 60V NLDMOS Rds(on) performance and its high current driving capability optimises device sizes that enhance power efficiency. A robust safe operating area (SOA) makes it ideal for both power switch and driver design. Fine detailed characterisation also provides a useful reference to optimise the design budget for optimum chip size.

On the CMOS side, a 5V capability supports analogue Pulse Width Modulation (PWM) controller design elements, and the 2.5V and 1.8V logic cores are optional modules for higher-level digital integration. In addition, logic compatible one-time programmable (OTP) and multi-time programmable (MTP) memory options are available for enhanced digital programming design.

The integrated component options include high voltage bipolar transistors, high voltage, high precision capacitors, high resistance poly and Zener diodes to reduce external passive component count and circuit board area.

THE GREAT FAB-LITE DELUSION

Malcolm Penn is CEO and chairman of market analyst firm Future Horizons, based in the UK

OVER THE PAST few years, squeezed by declining ASPs (average selling prices) and a zero-growth market (in value terms), the so-called fab-lite business model was born and unashamedly embraced. Encouraged by the financial community, seized upon by struggling IDMs, driven by the fabless firms' success, edged on by OEM disinterest, firm after firm has signed up to the concept, lauding the benefits and turning a blind eye to the flawed logic reality… if something is seemingly too good to be true, it usually is.

Not so, goes the argument. Wafer fabrication is a service operation, a simple make-buy decision best left to outsourcing. Foundries are fundamentally more efficient than IDMs, meaning they can make wafers much cheaper than in-house production. In any event, the chip industry has been outsourcing back-end manufacturing for decades without any problem and fabless companies have constantly out-performed IDMs' growth with no competitive disadvantage to not having a fab. Contrast that with the IDM corporate liability of owning and operating a fab, tying up cash and management resources; fab-lite elegantly solves these inherent problems, levelling the playing field with the fabless competition.

If only it were that simple. Not only is the justification logic muddled, superficial, flawed and confused, there is no industry consensus on what exactly fab-lite means; witness the fact the term is littered with a variety of alternative colourful euphemisms, such as asset-lite and asset-smart, implying these are something similar to fab-lite but subtly different and implicitly better.

The fact is, unlike fabless or IDM, fab-lite is a chameleon meaning different things to different people, but said with such brash and reassuring gusto that no one questions the strategic reality that fab-lite is nothing more than an illusion … Emperor's new clothes.

Euphemisms aside, there are two fundamental fab-lit varieties. Option 1: Maintain a small in-house wafer fab to prove out each process node but then outsource to a foundry the bulk of production. This is essentially the current STMicroelectronics approach and has the advantage of keeping up with technology, provided the facility is constantly upgraded for future node transitions.

On the face of it Option 1 seems an elegant solution were it not for the fact a small pilot fab will never be cost-effective versus a foundry. Proponents of this route will thus face perpetual hostility from investors and shareholders: "Why are you wasting money, tying up capital in expensive assets and depleting shareholder value when your outsource supplier is clearly much cheaper than you?" It would take a strong CEO and board to fend off this criticism, more likely than not they would all be fired and replaced by a more 'investor-compliant' team. In the long-term, Option 1 will likely default to Option 2.

Option 2: Stop building fabs completely at a certain process node and then use a foundry for new wafer production. This is the route NXP and Infineon have taken. This is clearly a bipolar structure; IDM up to a certain node and then fabless thereafter. Unlike Option 1, this strategy is process-terminal, once you exit a fab node it will be virtually impossible to re-enter the wafer manufacturing business. The more generations missed, the greater the impossibility.

"NOT ONLY IS THE JUSTIFICATION LOGIC MUDDLED, SUPERFICIAL, FLAWED AND CONFUSED, THERE IS NO INDUSTRY CONSENSUS ON WHAT EXACTLY FAB-LITE MEANS"

Option 2 thus combines the worst of both worlds. For the legacy fabs it ignores the fundamental reality that today's leading edge is tomorrow's commodity, meaning these fabs will slowly become more and more obsolete and harder to fill. Firms will thus be subjected to a constant closure and restructuring effort, damaging employee morale and affecting costs and productivity. Finding and keeping good operations personnel will be difficult given this business strategy is an operational dead end.

For the new fabless future, Option 2 sidesteps the fact that simply not having a fab does not make you a fabless firm, just a

firm without fabs; not the same thing at all. The legacy overhead infrastructure and costs and inefficiency will be much higher than with a truly fabless company and overall competitiveness will continue to erode. In short, restructuring from an IDM to a fabless business model will be both 'operationally challenging' and unlikely to make the organisation more structurally competitive. To the contrary, it will more likely have the totally opposite effect and competitivity will continue to erode. Option 2 thus represents death by slow strangulation.

Aside from the definition and implementation issues, there is also the fundamental underlying problem that fab-lite is based on the shaky assumption that foundry wafers will always be freely available at ever-cheaper prices. This is the chip industry equivalent of the 'debt is freely available and cheap' corporate business model that came to such an abrupt and catastrophic halt in the recent financial crisis. Just as with cheap debt, ever-reducing prices (and profits) whilst simultaneously investing in new process and production technology cannot be sustained forever; they result would be bankruptcy. Structurally prices must eventually increase.

At the same time the chip world is now staring into the eyes of a wafer fab famine, triggered by two or more years of rampant underinvestment. With no hope of fixing this problem near-term – 2010's capacity is already cast in stone, determined by 2009's (lack of) Cap Ex spend – foundry wafer prices will be hit with supply-and-demand price increases. Worse still, they will also be on allocation. Paying a higher price is one thing; not getting the wafers is another matter entirely. No wafers equals no sales – that is a 100% correlation and a major competitive threat.

Structurally deceitful, operationally faulty, 2010 will see the fab-lite 'model' fall to pieces. The so-called fab-lite option is thus nothing short of deception. Fab-smart therefore remains the only true solution, i.e. continuing to build in-house fabs but outsource a modest amount (say 10-15% maximum) to foundries to both smooth the supply and demand peaks and built external fab demand, high enough to justify equipping the next modular in-house expansion. In this way any expansion in capacity enters production 'fully loaded' from the beginning, whilst simultaneously improving response time to near-term demand fluctuation. The foundries do not like this option, of course, but this is the only real competitive reality.

The final part of the delusion? IDM fabs do not have to be wholly owned; there is no reason at all why a jointly owned fab could not work. ■

Malcolm Penn can be contacted at mail@futurehorizons.com. Future Horizon's website can be found at www.futurehorizons.com

FASTER TIME TO MARKET THROUGH RAPID PROTOTYPING

THE TERM 'time to market' is often used when referring to the pressure OEMs face when developing new products, but it can be interpreted in many ways. Some may feel that it is a general expression of the overriding need to reduce development costs, while others may say it is in reference to meeting a window of opportunity; crucial to a product's commercial success. While both of these statements are undoubtedly true, the process of reducing time-to-market is more complex and has greater impact than either of these explanations suggest.

The design phase of any product must mix the freedom to develop new ideas quickly, with the discipline of documenting that process. Once in production, the cost effectiveness of a product comes down to two things; how much it costs to produce and how long it can remain in production. Therefore, it is crucial to reach maximum production volume as quickly as possible and for as long as possible. Meeting the former at the cost of the latter can spell disaster for any company; the impact of a product recall is massive, in any vertical sector. Cutting corners at the development stage, therefore, isn't an option.

The process of turning a proof-of-concept into a production-ready device demands stringent adherence to best engineering practice and quality assurance; design validation isn't something that can be rushed or avoided. However, there are ways to cut time to market at the research stage, without sacrificing quality in development.

While the time needed to carry out stringent design validation is difficult to compress, there are effective methods for reducing the initial phases of a design cycle, particularly in reaching proof-of-concept, and rapid prototyping is one approach that we can look at.

Today's embedded products heavily rely on software to define their functionality; software that, wherever possible, is reused or repurposed. Normally, it will also include software developed by a third party. For new products, it is likely that any or all of these software components will need to be revised for their new target. This requires easy access to software modules, and allowing them to be altered or modified without threatening the 'golden' code base. Increasingly, and in response to the trend towards software defined products, software configuration management (SCM) tools are being used to create this so called 'agile' design environment.

Inherently, people are more creative when they aren't constrained to working within rigid parameters. However engineering, as with all disciplines, imposes certain constraints if efforts are to return results.

By definition, rapid prototyping and proof-of-concept are a means to an end; something that has been developed to justify further investment is rarely intended to form the foundation of a production model without going through rigorous redesign.

It follows, therefore, that reaching the proof-of-concept stage rapidly should not impose undue restrictions on creativity. SCM tools designed to enable this process include features that allow engineers to create branches in the design line of a code base, to evaluate ideas quickly that could, later, simply be removed from the code repository without repercussion, or equally form part of a product variant. For this reason, SCM must also facilitate simple merging of new code back into the code base.

"THE PROCESS OF TURNING A PROOF-OF-CONCEPT INTO A PRODUCTION-READY DEVICE DEMANDS STRINGENT ADHERENCE TO BEST ENGINEERING PRACTICE AND QUALITY ASSURANCE"

This 'branch and merge' process isn't unique to SCM tools, but is significantly enabled and eased by them, particularly in those tools that can record the differences between sets of files at different stages of the development cycle. A proven code base can comprise millions of lines of code, while a variant may only represent a small change, by only recording the 'delta' at the

Market dynamics demand faster response times, but not at the expense of good engineering practices, says **Dave Robertson**, director of European operations at Perforce Software

branch and merge stages, the storage space and search time needed to control both the golden code base and any branches are kept to a minimum. The same isn't true with a 'copy and change' philosophy.

The use of SCM tools to control assets described by software IP isn't restricted to just software components; today, most design elements are created and stored digitally, and SCM tools are equally able to manage any asset described in software.

This asset management process can enable the creation of standard platforms, configured when needed, which may be branched to create completely different families of products. These product variants, where relatively small changes at a low level can propagate into significant differences at a higher level, could be created – and recreated – much more quickly and efficiently using SCM.

It is a practice already being used throughout the industry. As it takes hold, the concept of rapid prototyping will change; no longer will it be used as a differentiator employed by some, but the norm expected by many.

The definition of time-to-market could once again come under scrutiny. Currently, the ability to create a proof-of-concept rapidly offers a competitive advantage. However, customers may come to expect this faster reaction using SCM, because it is both feasible and, perhaps more significantly, makes good commercial and engineering sense. ■

A NEW Year


Myk Dormer

IT IS TRADITIONAL around this time of year to optimistically make a list of New Year's resolutions. Maybe we engineers should follow suit?

Here are mine:

- I will always design-in a good, solid ground plane. If it isn't needed it does no harm, but it's really hard to add one later.
- I will think about the aerial at the start of the project. I will not just cram in a roughly quarter wavelength long piece of bell-wire somewhere, as an afterthought.
- I will allow for sufficient shielding at the start. If radiated emissions aren't as high or the circuitry is less sensitive than I thought, then I can always leave the can off in the next version.
- Decoupling is easier to allow for and not fit, than try to squeeze in later or bodge onto early production runs by hand.
- A maximum spec figure is a maximum. It isn't a "Well, I can go a bit over that if I'm careful" figure. It's likewise for minimums.
- Average specifications don't mean very much really. See previous resolution.
- The heat-sink *will* be big enough; under all conditions of airflow and ambient temperature. Just using the ground-plane is not always enough.
- Everyone makes errors. I will allow for one more design/PCB iteration in the project plan than I think I need. (In fact, I think I'll allow for two!)
- Quoted lead times are not there for fun. I will order components well ahead of when they are needed.
- Factory yields are never 100%, so production batches will always have one or two extra units added. (Any excess can always be used as sample stock).
- I will not use single sourced parts unless I am absolutely certain of the supplier, or there is really and truly no alternative. Even then, I will lay in a plan of action in the event of the critical part going obsolete/unavailable/extended lead-time; even if that plan requires a re-design.
- I will only design in components that I know I can source. If the MOQ is three thousand and I'll only ever need fifty, I'll find another part; not just blame the purchasing guy.
- I will not design in the over-hyped, easy-to-use component. I will do the work and design in the right parts, even if it takes more effort.
- I will not mistake my simulator results for real test data.
- I shall hold grudges. I will remember which component suppliers have withdrawn good parts in the past or who have unreliable supply policies, or simply bad quality. I will blacklist them.
- I will not keep a bad design in production. Factory production yield figures matter. I will keep track of build failure rates and if they are unacceptably high, I will do the work and find out why. And then, if necessary, re-design.
- The technical author is not a mind-reader or a psychic. I will write copious notes on which they can base the data sheet, and if they need more information, they'll get it.
- When I write specifications, I will not lie. Not even about maximum radio link range. I will not "just double it, to make it look good" nor will I use free-field calculations for terrestrial links.
- My code will include comments, version control numbers, and will be documented and backed up. Old versions will be archived (just in case my 'upgrade' isn't all I hoped).

"I WILL THINK ABOUT THE AERIAL AT THE START OF THE PROJECT. I WILL NOT JUST CRAM IN A ROUGHLY QUARTER WAVELENGTH LONG PIECE OF BELL-WIRE SOMEWHERE, AS AN AFTERTHOUGHT"

Like all New Year resolutions, I imagine most of these will barely last into February. The pressures of real production and design schedules will begin to bite, the number of working hours in the day will move back into two figures and the weekend becomes just two more workdays again … but there's one resolution that's got to hold:

Test everything! Only real measurements on real hardware are worth anything.

Happy 2010!

Myk Dormer is Senior RF Design Engineer at Radiometrix Ltd
www.radiometrix.com

The World's Lowest Sleep Current MCUs: PIC® MCUs with nanoWatt XLP Technology

Microchip's PIC® Microcontrollers with nanoWatt XLP Technology offer the industry's lowest currents for sleep, where extreme low power applications spend up to 99% of their time.

- **Extend Battery Life**
  - Sleep current down to 20 nA
  - Brown-out Reset down to 45 nA
  - Watchdog Timer down to 400 nA
  - Real time clock down to 500 nA
- **Extreme Flexibility**
  - 5 different low power modes to improve power & performance in your application
  - Many low-power supervisors, alarms, and wake-up sources
- **Expanded Peripheral Set**
  - Integrated USB, LCD, RTCC & touch sensing
  - Eliminates costly external components

GET THE MOST FROM YOUR BATTERY IN YOUR NEXT DESIGN!

1. View the Low Power Comparison demo
2. View free Webinars and Application Notes
3. Download the Low Power Tips 'n Tricks
4. Order samples and development tools www.microchip.com/XLP

Intelligent Electronics start with Microchip

microchip DIRECT
www.microchipdirect.com

www.microchip.com/XLP

MICROCHIP

WaveAce Oscilloscopes

LeCroy Expands Entry-level Oscilloscope Series to Include 4 Channel and 40 MHz Models from just £550 for Simple and Efficient Debug now from 40 MHz to 300 MHz

LeCroy® Corporation today announced that it is expanding the popular WaveAceTM oscilloscope line to include 4 channel models from 60 MHz to 300 MHz and adding a new 2 channel, 40 MHz model. The 4 channel models provide 10 kpts/ch memory and up to 2 GS/s sample rate; the 40 MHz model provides 4 kpts/ch and a sample rate of up to 500 MS/s. All models offer long memory, large colour displays, extensive measurement capabilities, and advanced triggering to improve troubleshooting and shorten debug time. With USB host and device ports, plus a LAN connection, the WaveAce oscilloscopes easily connect to a memory stick, PC or printer for saving data or remote control. Combined with the streamlined, time-saving user interface, these features make the WaveAce oscilloscopes the ideal tools for affordable design, debug and troubleshooting from 40 MHz to 300 MHz.

Deep Feature Set for Better Debug

Available in bandwidths of 60 MHz, 100 MHz, 200 MHz and 300 MHz, the new 4 channel models provide a maximum sample rate of 2 GS/s and up to 10 kpts/ch memory or 20 kpts when interleaved. The long memory allows users to capture full sample rate acquisitions that are two to three times longer than the competition. With 32 built-in automated parameters, including advanced timing parameters for skew, phase and edge-to-edge measurements between channels, the WaveAce oscilloscope broadens the ways a user can understand and analyze waveforms. Additional features such as Pass/Fail testing, user definable digital filters and a waveform sequence recorder simplify and shorten debug time.

Remote Control via USB and LAN Ports

A new remote control command set has been implemented in the 4 channel models to make the WaveAce even more useful and versatile than ever before. This modern remote control command set provides access to all of the WaveAce controls, functions and measurements from a remote PC via the rear panel USB or LAN ports. These connections also allow the user to remotely view the waveform display and access a virtual front panel.

Saving and Documenting Results

Documenting results and saving screenshots, waveforms and setups are easy with the WaveAce. Internal storage can hold up to 20 waveforms and 20 setups. Mass storage can be done by connecting a USB memory device directly to the front panel of the oscilloscope. The rear panel USB port allows for direct printer connection to quickly generate hard copies of the screen image.

Easy to Use for Faster Debug

The high performance and large feature set of the WaveAce is controlled by an intuitive user interface with 11 different languages and streamlined front panel. All important controls and menus are accessed from the front panel with a single button press. All position and offsets can be reset by simply pressing the knob, pressing the V/Div knob will switch between fixed and variable gain and pressing the T/div knob will toggle between zoom modes. Buttons on the front panel that open and close menus or switch modes are backlit to let the user know exactly what mode the WaveAce is operating in.

Pricing

The 40 MHz WaveAce 101 price is £550 while prices for the new 4 channel models range from £1,160 to £ 2,210.

LeCroy

WaveAce Oscilloscopes

40 MHz – 300 MHz

FROM £550*

The New Affordable Oscilloscopes from the Manufacturer of the Fastest Oscilloscopes in the World

- 40 MHz – 300 MHz Bandwidth
- 2- and 4-Channel Models
- Up to 2 GS/s Sample Rate
- Up to 20 kpts Memory
- 5.7" Colour Display in all models
- Class Leading Math and Measurements
- Advanced Trigger
- LAN/USB Connectivity
- 3 Years Warranty

www.lecroy.co.uk

Phone: 01235 - 536973

*Launch Special Offer valid until 31 March 2010

BATTERIES REGULATIONS: A ROUND-UP OF RESPONSIBILITIES

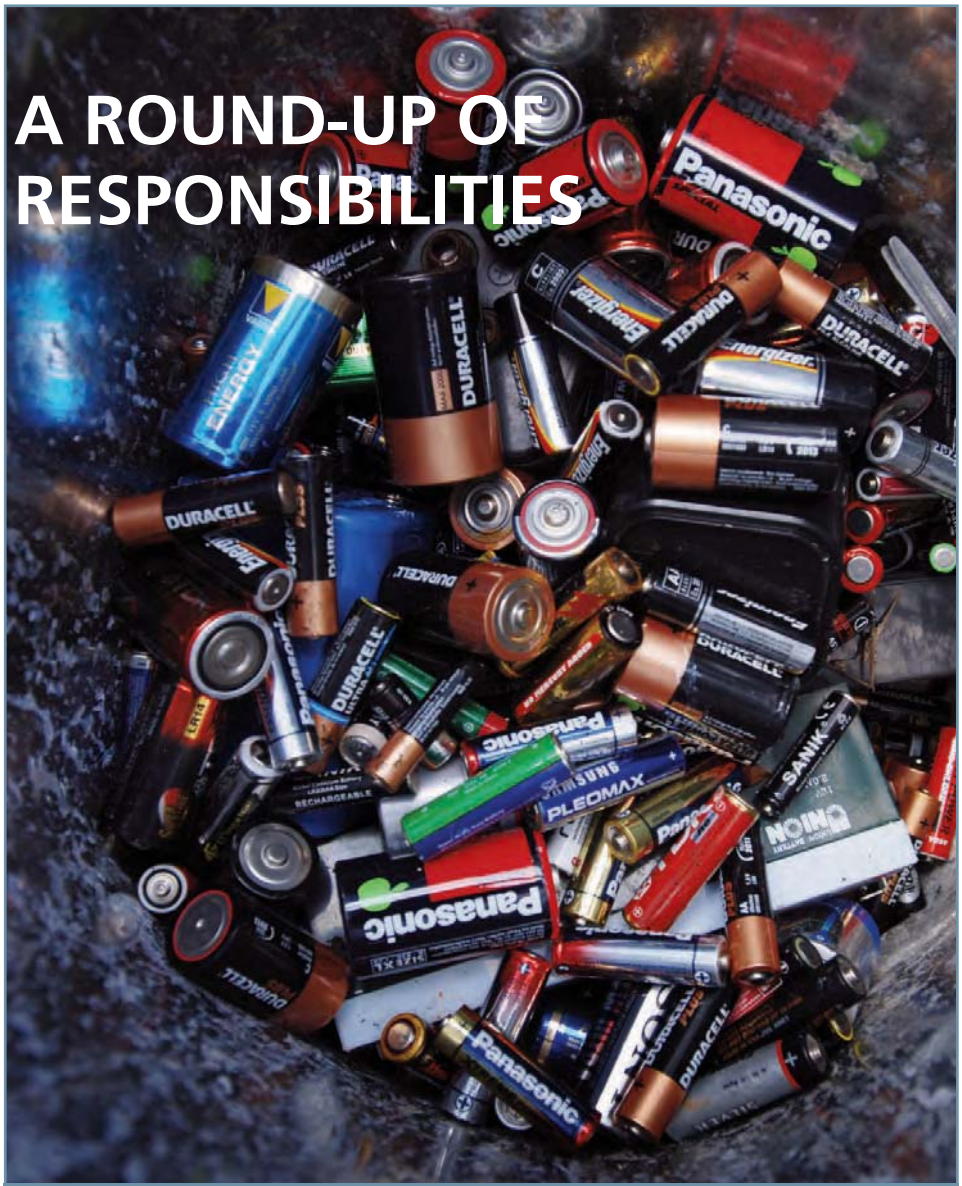

New regulation applies to batteries

Bob Mead, Batteries Implementation Project Manager for the Environment Agency, gives an update on the Batteries Directive and how to ensure stay within the law

THE 'NEW' Batteries Regulations came into force on 5 May 2009 and every producer and distributor must now comply.

It's extremely important to understand the terminology. Vitally, a producer is not necessarily the battery manufacturer. It is whoever places them on the UK market for the first time.

The definition of producer also applies when batteries are included within other products. Many products, from computers to animated toys, contain them. Some are less obvious than others; for example many DVD players and all desktop PCs contain batteries, as does any product with a clock or memory that works when it's unplugged from an electricity supply.

If you incorporate batteries in any product, you need to understand whether they have already been placed on the market. If you purchase them from a wholesaler they probably have, but if you buy them direct from the manufacturer or importer, they may not. This is a particular issue for battery pack manufacturers. If individual battery cells which combine to make up your pack had not already been placed on the market, then you will place them on when you sell the packs, and you are therefore the producer.

The directive seeks to reduce the environmental impact of batteries and accumulators. The UK has targets to increase from recycling 3% of portable batteries (2007 figures) to 25% (7,500 tonnes) by 2012 and to at least 45% by 2016.

This must all be in line with existing controls, including Health and Safety, Carriage of Dangerous Goods, Environmental Permitting and Hazardous Waste.

All batteries placed on the UK market must also be compliant with the requirements of 2008's Batteries and Accumulators (Placing on the Market) Regulations. These set limits for levels of cadmium and mercury in new batteries, and impose marking requirements. Visit **www.berr.gov.uk** for more information.

Every producer must now record the weights and types of batteries they place on the UK market. In January 2010 figures for 5 May to 31 December 2009 must be reported, broken down by weight of lead-acid, nickel cadmium and 'other'.

If you are responsible for automotive or industrial batteries the rules are slightly different. Again, see **www.berr.co.uk** for information.

The Government distinguishes between 'large' and 'small' producers. 'Large' producers place more than one tonne of portable batteries onto the UK market annually. They must pay for the collection, treatment, recycling and disposal of waste batteries in proportion to their market share. If this applies to you and if you have not already done so, you need to join a Battery Compliance Scheme (BCS) as soon as possible, as the deadline for doing so was 15 October. The BSC will also register

you with the appropriate environment agency (England & Wales; Scotland; or Northern Ireland).

'Small' producers place one tonne or less of portable batteries onto the UK market annually. They do not have to pay for the collection and treatment of waste. If this applies to you, you must, however, register with your relevant environment agency within 28 days of the first date on or after 15 October 2009 on which you place any batteries onto the UK market. You can do this quickly and easily online at **www.environment-agency.gov.uk/batteries**.

Battery bank

IN THE CONTEXT OF THE BATTERIES DIRECTIVE, A DISTRIBUTOR IS ANYONE WHO SUPPLIES BATTERIES ON A PROFESSIONAL BASIS TO AN END USER; THIS IS USUALLY A RETAILER BUT IT CAN BE A WHOLESALER FOR ALL OR PART OF THEIR BUSINESS

Following an assessment period we approved seven BCSs to serve the UK market. You can find details on these at **www.environment-agency.gov.uk/batteries**. Producers can join whichever approved scheme best suits their needs, regardless of where in the UK they are based.

In the context of the Batteries Directive a distributor is anyone who supplies batteries on a professional basis to an end user. This is usually a retailer but it can be a wholesaler for all or part of their business.

If this applies to you, and you sell over 32kg of portable batteries annually to end users, you must offer a battery take-back facility from 1 February 2010. Batteries within products do not count when calculating the 32kg. You must accept any portable batteries, not just those you sell, and regardless of any purchase by the depositor. You do not, however, have to pay for their transport and treatment: approved BCSs are obliged to collect them free of charge.

If you sell less than 32kg to end users you do not have to accept waste batteries. BCSs are obliged to accept them from you if you do, but they do not have to collect them. It's advisable to check the situation before deciding to offer take-back facilities. ■

Disposable batteries

FURTHER INFORMATION ON THE BATTERIES DIRECTIVE

You can get further information on portable batteries at **www.environment-agency.gov.uk/batteries** or by phoning our helpline on 08708 506506.

Contact Defra regarding questions about distributor take-back. Email **portable.batteries@defra.gsi.gov.uk**.

For queries on automotive or industrial batteries contact **BIS at batteryconsultation@berr.gsi.gov.uk**

Defra's advice on the regulations are at **www.defra.gov.uk/environment/waste/producer/batteries/documents/0907-advisory-note.pdf**

Professor Alan Burns of the Department of Computer Science at the University of York and **Chris Dale**, events co-ordinator for the Safety-Critical Systems Club, explore the scheduling and timing analysis techniques necessary for safety critical, real-time operating systems

Scheduling and Timing Analysis for SAFETY-CRITICAL REAL-TIME SYSTEMS

THE CAR YOU drive almost certainly depends for its safe operation on concurrent software executing within embedded systems. And as brake-by-wire, steer-by-wire and throttle-by-wire systems become more common-place, your safety will become ever more dependent on the timing properties of these systems.

The diagram in **Figure 1** represents a brake-by-wire system. Its control unit (ECU1) reads the sensors, computes the outputs, writes the outputs to the actuators (e.g. ECU2) and resets the watchdog timer – set to 20ms. This system worked fine, but after a small change in the software the steering mechanism became wobbly. What had happened? The computation task now overran its assumed Worst Case Execution Time (WCET), failing to reset the watchdog timer and resulting in a very frequent full system reset!

This is exactly the kind of problem that scheduling and timing analysis is intended to deal with. And it's easy to see that this problem is a relatively simple one: the car has many more such embedded systems, for engine control, air conditioning, lighting, navigation, etc; fly-by-wire aircraft are even more complex.

Verification of the timing properties of safety-critical real-time systems, and looking at the standard ways of scheduling and analysing the concurrent software that typically executes within embedded systems, are techniques crucial for safety-critical systems, where there is a need for a high level of assurance that timing constraints will be satisfied in all situations, even those that are very rare.

Scheduling and Timing Analysis

So what are scheduling and timing analysis, and how do they differ? You can think of a real-time application as a collection of concurrent tasks that repeat, either periodically or sporadically, and that are either time-triggered or event-triggered. Timing analysis is concerned with the temporal behaviour of individual tasks, while scheduling analysis deals with the temporal behaviour of the whole system.

In scheduling analysis, a scheduling scheme has two features: an algorithm that orders the use of system resources (in particular the CPUs), and a means of predicting the worst-case behaviour of the system when the scheduling algorithm is applied.

Many hard real-time systems are implemented as cyclic executives. This means that the application code is produced as a collection of procedures, each procedure is mapped onto a 'minor cycle', and the minor cycles together make up the complete schedule (or 'major cycle') of the application. The main advantage of cyclic executives is

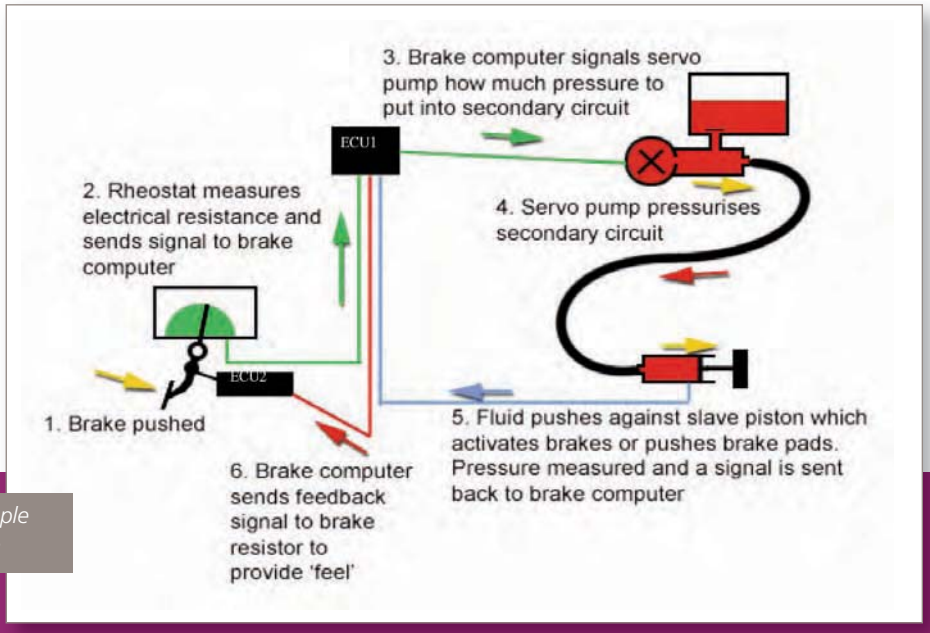

Figure 1: *Brake-by-wire – an example of a safety-critical, real-time system*

Figure 2: *Time-line for a simple real-time system*

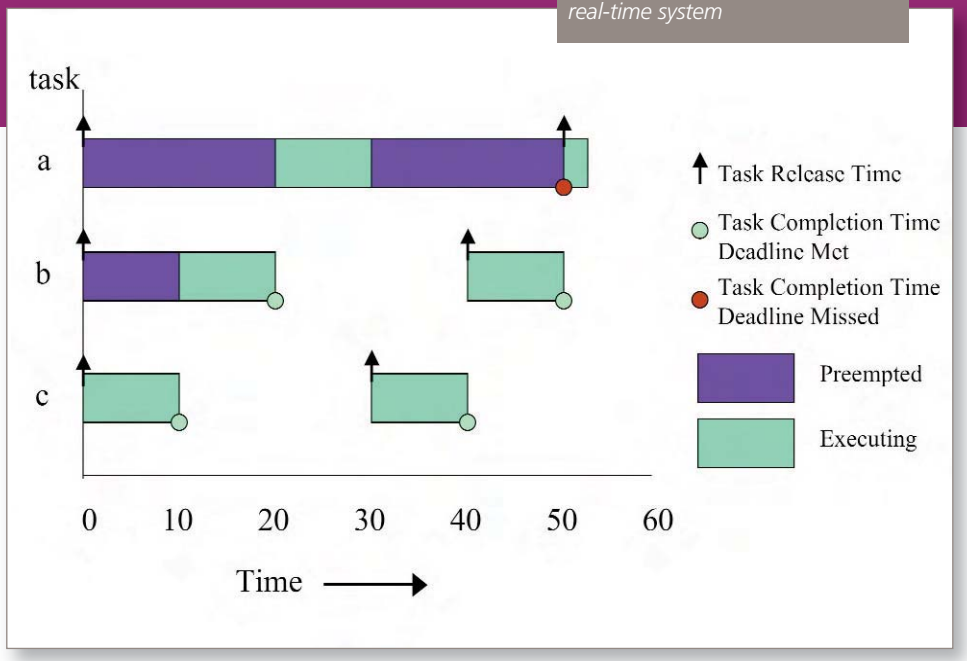

that they are fully deterministic: the cycle times are fixed and individual procedures complete within their respective cycles.

However, this determinism comes at a price: each task period has to be a multiple of the minor cycle time; tasks with long periods are difficult to incorporate; and it is just about impossible to accommodate sporadic activities. This means that cyclic executives are difficult to construct and difficult to maintain. In particular, any task with a sizable computation time must be split into a number of fixed sized procedures; this may cut across the structure of the code from a software engineering perspective and, hence, may be error-prone – not what you want in a safety-critical system. So, although the determinism of cyclic executives may be appealing, predictability of worst-case behaviour is more important, and this can be achieved without having to work within the limitations of cyclic executives.

Non-Deterministic Approach

The most widely used non-deterministic approach is known as Fixed-Priority Scheduling (FPS). Under this scheme, each task has a fixed, static, priority which is computed before run-time, and tasks are executed in the order determined by their respective priorities. In real-time systems, the priority of a task is derived purely from its timing requirements, not from its importance to the correct functioning of the system or its integrity.

But what if a high-priority task is released while a lower-priority one is already executing? There are two possibilities: in a 'pre-emptive' scheme, there will be an immediate switch to the higher-priority task; with non-pre-emption, the lower-priority task will be allowed to complete before the other can start. Because pre-emptive schemes enable higher-priority tasks to be more reactive, they are preferred.

So how does this work in practice? Suppose, for the moment, that an application is made up of a fixed number of tasks, each with its own known period, each independent of all the others, each with a deadline equal to its period and each with a fixed worst-case execution time. We're also going to ignore all the system overheads. Task priorities are assigned according to Rate Monotonic Priority Assignment – i.e. the shorter the task's period, the higher its priority; this is the optimal way of assigning priorities in an FPS scheme with task deadlines equal to their periods. The utilisation of each task is defined as its computation time divided by its period and is the proportion of the time that the task would be executing, if left uninterrupted by others.

Table 1 lists the three tasks that make up a simple real-time application. Task 'c' has the shortest period, so it has the highest priority, and thus executes first – as shown in **Figure 2**; tasks 'a' and 'b' are also ready to execute at time zero, but they are pre-empted by the higher-priority task. Once task 'c' has completed, task 'b' executes, and only then can task 'a' begin. However, before task 'a' can complete its execution, it is pre-empted again, first by task 'c' and then by task 'b', and thus fails to complete before it is due to recommence at time 50. So, in this example, it is impossible to schedule the tasks within the given constraints.

In general, it is not necessary to analyse tasks to the level of detail given in this example, which would be quite onerous even for a moderately complicated application. Instead, tests based on the utilisation levels of tasks can be applied to show whether a given application can be scheduled: provided the total utilisation of all the tasks is less than a threshold value that depends on the number of tasks, they can be scheduled.

In the example above the total utilisation is 0.82, which is greater than the threshold value of 0.78 for three tasks, so this application fails the test. But these tests are sufficient, not necessary – in other words, while passing the test shows that the application can be scheduled, failing the test does not show that it cannot be scheduled. To overcome this problem, a more complicated response-time analysis can be

Task	Period	Computation Time	Priority	Utilisation
a	50	12	1	0.24
b	40	10	2	0.25
c	30	10	3	0.33

Table 1: *The list of three tasks that make up a simple real-time application*

applied in cases of doubt, resulting in a necessary and sufficient test for schedulability.

Classified Tasks

The above description of FPS assumed that the tasks were all periodic, but it can be generalised to include sporadic tasks (those whose arrival times are not regular). However, in many situations the worst-case rates of arrival for sporadic tasks are considerably higher than the averages – e.g. interrupts often arrive in bursts and an abnormal sensor reading may lead to significant additional computation. This means that schedulability calculations made using worst-case figures may lead to very low processor utilizations being observed in the actual running system.

To overcome this, tasks are classified as 'hard' or 'soft': a hard task is one for which execution after its deadline is of no value (e.g. because it leads directly to system failure), and a soft task is one where missing the deadline can be tolerated, albeit with the possibility of decreased service quality. This leads to two rules that can be used to guide schedulability analysis:

1. All tasks should be schedulable using average execution times and average arrival rates.
2. All hard real-time tasks should be schedulable using worst-case execution times and worst-case arrival rates of all tasks.

Rule 1 means that there may be situations in which it is not possible to meet all current deadlines; this condition is known as a transient overload. Rule 2 ensures, crucially, that no hard real-time task will miss its deadline. If it turns out that rule 2 gives rise to unacceptably low utilizations for 'normal execution', then action must be taken to reduce the worst-case execution times (and/or arrival rates).

The above analysis assumed that task deadlines were equal to their periods – i.e. they needed to complete before they were due to restart. In some circumstances, task deadlines may need to be less than their periods, and in these cases a different priority scheme is used: Deadline Monotonic Priority Ordering (DMPO), in which the highest priority task is the one with the shortest deadline.

Unfortunately when deadlines are no longer equal to periods, the utilisation-based method of checking for schedulability is no longer effective. Rather, the response-time of each task must be calculated and then compared with the task's deadline.

In some circumstances, where tasks share resources, it is possible for a task to be suspended waiting for a lower-priority task to complete some required computation. This situation is known as priority inversion and the task waiting on a lower-priority task is said to be blocked. Such a problem nearly led to the loss of the Mars Pathfinder mission: as a shared bus became heavily loaded, critical data was not being transferred; time-out on this data was used as an indication of failure and led to re-boot. In this instance, the problem was resolved through an appropriate patch. In general, response time analysis can be extended to cope with the blocking of resources.

Timing Analysis

In order to be able to apply the scheduling analysis outlined above, we need some knowledge of the timing of individual tasks. But the time a task takes to execute, even on a fixed hardware platform, is not a fixed quantity – it will vary depending on the system state and on the particular data values being processed. Consider the following:

- Best-Case execution time. This is the shortest possible time taken to execute the task. To use this in a scheduling analysis would be grossly optimistic.
- Average execution time. As we saw above, for soft tasks, it can be acceptable to use the average execution time as the basis of scheduling analysis.
- Maximum observed execution time. It is very unlikely that the worst possible (in terms of timing) set of inputs would be encountered during any particular regime of testing. Thus, it would be unwise, and possibly unsafe, to use the maximum observed execution time in a scheduling analysis for hard real-time tasks.
- Worst Case Execution Time (WCET). For any task running on a particular hardware configuration, there is a worst case execution time, but it may not be practically possibly to determine what this is (for reasons such as the complexity of the task). Thus, although this figure would be acceptable and safe to use for scheduling analysis, it may be simply unobtainable.
- WCET Upper Bound. In the absence of perfect knowledge about WCET, some means of obtaining an upper bound on the WCET is needed. This constitutes a pessimistic but safe approach to timing analysis for hard real-time tasks within safety-critical systems.

Detailed guidance about how to obtain WCET Upper Bounds, and further detail on scheduling analysis, can be found in Real-Time Systems and Programming Languages by Alan Burns and Andy Wellings (fourth edition, 2009, Addison Wesley).

Among the tools available for worst-case execution time analysis are those of Rapita Systems Ltd (www.rapitasystems.com), a spin-off from the Real-time Systems Research Group at the University of York. ■

ECO-DESIGN
A CARBON APPROACH WEBINAR

17 February 2010 at 10.00

Register Today!

- A topical and interactive forum for eco design.
- Find out how you could benefit from design in a low carbon economy.
- Live demonstrations will examine low carbon design principles.
- Rob Holdway, our design expert will be on hand to advise on implementing a design strategy.

Register online today

www.envirowise.gov.uk/webinar

envirowise SUSTAINABLE PRACTICES SUSTAINABLE PROFITS

Realistic Implementation of GUIDELINES

Griff Derryberry, Zuken Applications Engineer, in this article looks at each of the roles of the engineer and PCB designer; considers the traditional design process and makes suggestions on how the designer can contribute significantly to improving a design's overall signal integrity whilst saving time

PICK UP ANY PCB or electronics design magazine and probably at least one article will be devoted to some aspect of high-speed digital design; be it routing techniques, signal integrity analysis, electromagnetic compatibility or power integrity. This has now become a mainstream design practice.

Being able to achieve successful designs repeatedly suggests that a honed process is in place – one that captures electrical intent, enables design realization and promotes verification to ensure that the product delivered matches the requirements and specifications. However, having in place such a formal process is often more an ideal than a reality.

Realizing products with significant high-speed digital content requires the engineer to operate in different specialties. Besides performing circuit design, he or she needs to evaluate the design's signal integrity, apply mitigation techniques and all the while keeping in mind the physical limitations of the board. The time an engineer can devote to such activities is usually a fraction of the overall design time, with the majority being spent on researching parts, attending design reviews, testing the design in the lab and participating in numerous meetings. This also usually means that an in-depth, post-layout analysis of the product is compromised.

So what can be done to improve the situation? This article looks at each of the roles of the engineer and PCB designer, considers the traditional design process and makes suggestions on how the designer can contribute significantly to improving a design's overall signal integrity, while simultaneously saving time so that the engineer can focus on more in-depth issues.

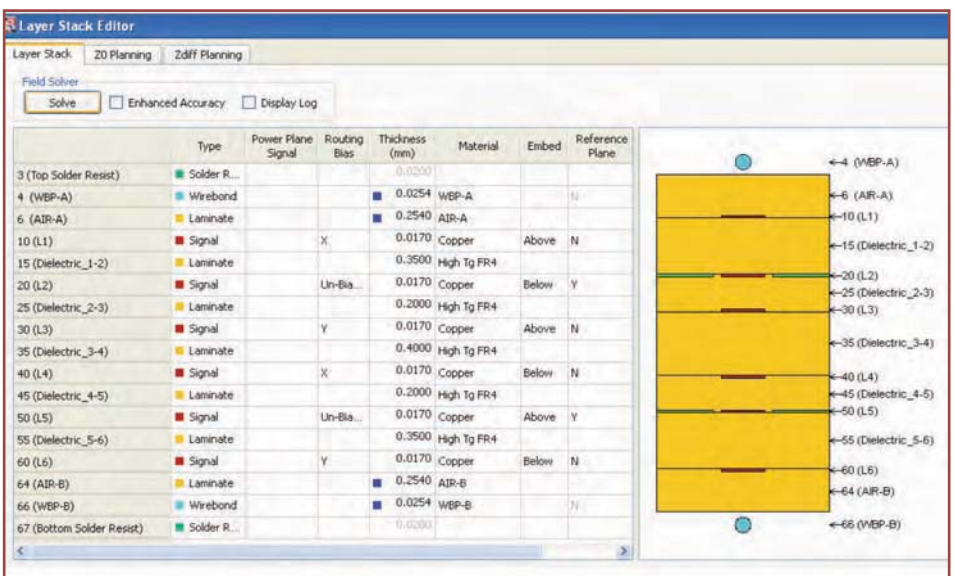

Figure 1: *Screenshot of the stackup editor in Zuken's CR-5000 Lightning solution*

Context

The PCB designer works closely with the engineer who also works under a tight schedule. The designer's role is to physically realize the design as outlined by the engineer vis-à-vis constraints as quickly as possible. Because the engineer often has to manage a constant stream of changes, in turn the PCB designer has to be patient and implement these at whatever point in the design process they are required.

Traditionally, the main urgency to complete the layout is to obtain a prototype so that the engineer can test the product in the lab – getting to this stage is often set as a design process milestone. However, gaining access to probe signals in today's high-speed designs can be difficult, due to component packaging.

One option to help mitigate this issue is to have a signal integrity (SI) analysis tool in the design environment. Yet, having access to an SI analysis tool versus implementing the practice of SI analysis as part of the design process are two entirely different states. Further, performing SI analysis has been perceived historically as a black art, requiring specialist knowledge and skill set. Some companies even have an entire group dedicated to performing this task.

The traditional design schedule does not

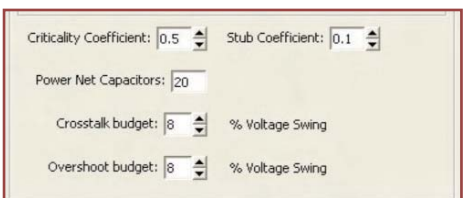

Figure 2: *The designer can constrain the design automatically for crosstalk, overshoot and stub length*

Signal Integrity Screening FOR PCB DESIGNERS

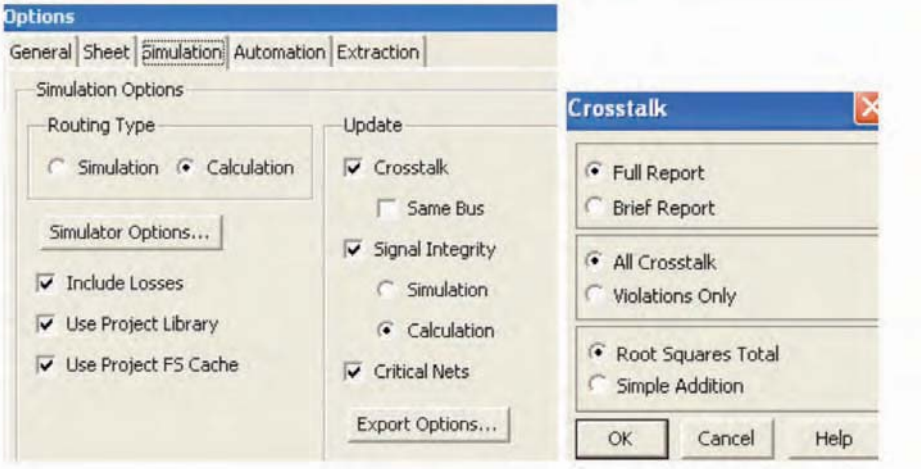

Figure 3: *The designer can configure how the tool will behave when screening is performed*

promote this activity and to incorporate it now into the schedule does not remove the need for examining the product in the lab. Consequently, the perception could be that adding SI analysis to the design process gives the electrical engineer another job to do, in addition to their already long list of actions. But this can be perceived as a somewhat limited view as the benefits of addressing signal integrity up front can be vast, both in time and money savings – enabling the design team to deliver boards that are going to be right-first-time.

PCB Designer Best Positioned to Resolve SI Issues Early

In traditional-style development organizations, the engineer does not look at SI issues until the board is completely placed and routed. Any issues the engineer finds are flagged and sent back to the designer. In order to complete the design, several iterations of this SI review cycle may occur. However, if the design environment includes a PCB design tool that incorporates constraint management with the ability to analyze signal integrity, as in the case of Zuken's CR-5000 Lightning Realize solution, the PCB designer becomes the person best positioned to improve a design from a high-speed design standpoint; as part of the layout process, the designer can perform first-pass signal integrity checking. Through this SI screening, the designer makes the design process more efficient by greatly reducing the number SI review cycles, thereby saving time and money.

The advanced capabilities of today's design tools that are available to the PCB designer mean that in one application, the designer can realize the placement and routing as guided by physical and electrical rules (constraints), and then screen for crosstalk and signal quality without needing a deep understanding of electrical theory. Solutions such as CR-5000 Lightning from Zuken are so sophisticated that the designer can effectively resolve some lesser issues that were once only solvable by the engineer, therefore reducing the number of actions on the engineer's to-do list.

As CR-5000 Lightning is such a powerful environment, the PCB designer is able to perform signal integrity screening as a fluid part of the PCB design process. Issues are presented in a spreadsheet alongside corresponding constraints, and crosstalk hotspots can be highlighted on the PCB design 'canvas'. The designer can then resolve these issues by manual or automated means. In this preliminary look at signal integrity, there is no need to involve the engineer, thereby saving the engineer time that can be used to look at more significant SI issues later in the design cycle.

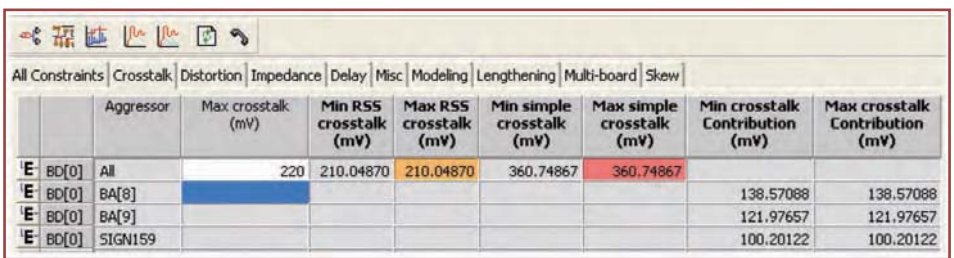

		Aggressor	Max crosstalk (mV)	Min RSS crosstalk (mV)	Max RSS crosstalk (mV)	Min simple crosstalk (mV)	Max simple crosstalk (mV)	Min crosstalk Contribution (mV)	Max crosstalk Contribution (mV)
E	BD[0]	All	220	210.04870	210.04870	360.74867	360.74867		
E	BD[0]	BA[8]						138.57088	138.57088
E	BD[0]	BA[9]						121.97657	121.97657
E	BD[0]	SIGN159						100.20122	100.20122

Figure 4: *Colour-coded results of the screening process are shown in a spreadsheet-based presentation of the Constraint Manager*

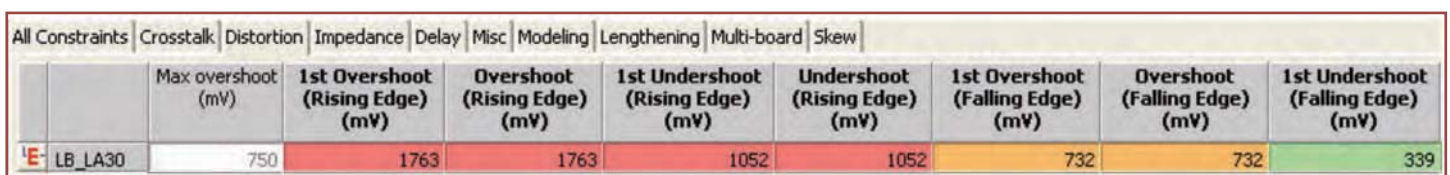

		Max overshoot (mV)	1st Overshoot (Rising Edge) (mV)	Overshoot (Rising Edge) (mV)	1st Undershoot (Rising Edge) (mV)	Undershoot (Rising Edge) (mV)	1st Overshoot (Falling Edge) (mV)	Overshoot (Falling Edge) (mV)	1st Undershoot (Falling Edge) (mV)
E	LB_LA30	750	1763	1763	1052	1052	732	732	339

Figure 5: *The designer can check the impedance of a net based on the layers it traverses*

Why would the PCB designer take on this task of signal integrity screening? Clearly, we are adding to the designer's already compressed schedule. Performing this action is a departure from traditional placement and routing activities. However, when proposing this action to designers, some have actively told us that they want to increase their value in the design process. Others desire a better understanding of the effects of their routing choices, and a few totally embraced the idea on its own merits – the overall time and money savings for the organization.

Again, PCB designers are best positioned to perform the signal integrity screening, as they are the most intimate with the physical realization, knowing which signal names correspond to which traces. They already have the design environment open and, therefore, the screening activity simply becomes an extension of what they are already doing – realizing the PCB.

If design tools are so sophisticated in realizing electrical design intent, why would screening need to take place at all? For crosstalk or multi-drop nets with min/max delays (as opposed to lengths), often it is easier to route nets first without considering constraints and then resolve issues using manual or automated methods.

Resolving min/max delay issues requires multiple passes. For the first pass, a faster calculation mode is used to check delay of each branch. Then at screening time, actual simulation is used to incorporate all of the reflections and logic low and high voltage levels. Screening for overshoot is best performed after placement. At this stage, the design is returned to the engineer with nets flagged that require termination. Therefore, signal integrity screening is really a continual process throughout the board layout phase.

PCB Designer SI Screening Steps

Let us get specific about the PCB designer's actions when performing signal integrity screening using Zuken's CR-5000 Lightning solution as an example.

Confirm Board Stackup: A first step for the designer is to confirm the board stackup, ensuring that the materials, thicknesses and reference planes are correctly established. With this in hand, the designer creates a field solution so that crosstalk and delays can be calculated. In addition, the designer can assess the characteristic impedance for a set of trace widths or conversely, calculate a set of trace widths based upon specified impedance.

A similar process applies to differential signalling, where the variables are trace width, trace-to-trace spacing and differential impedance. **Figure 1** shows a screenshot of the stackup editor.

Set up Default Simulation Models: Ideally, the engineer has assembled SI models for the design that were used when he or she experimented with signalling scenarios prior to schematic capture. In this case, assigning models is straightforward. If this is not the case, the designer can specify driver/receiver models that are appropriate for the technology being used.

Assess Constraints: Next, the designer assesses constraints. Ideally, the engineer has captured the design's electrical intent (mV of crosstalk and ps of delay) directly within CR-5000 Lightning. In contrast, it may be that the engineer has communicated requirements in documentation; it may be incumbent upon the designer to enter them.

In the absence of specified constraints, for the purposes of SI screening the designer can constrain the design automatically for crosstalk, overshoot and stub length based on the rise time of the SI models.

Configure Analysis and Checking: At this point, the designer configures how the tool will behave when screening is performed. The default settings are shown below. However, if the designer also wants to assess overshoot, the button for Simulation under Signal Integrity would need to be selected. In the Physical Editor, the designer specifies how crosstalk will be reported when checking is performed.

Perform Screening and Review the Results in Constraint Manager: After the board has been placed and there is routing to evaluate, the designer starts the screening process. Results are shown in the spreadsheet-based presentation of the Constraint Manager. To investigate a net of interest, the designer uses the "Update Selected" command. If the designer chooses many nets to review, the colour-coded results make it easy to see which nets are in compliance (green), close to violation (orange) or exceeding the constraint (red). In the crosstalk case, shown in **Figure 4**, the contribution of the primary aggressors is included.

If desired, the designer can check distortion, as shown in **Figure 5**, check the impedance of a net based on the layers it traverses (**Figure 6**), check delays (**Figure 7**) and distortion (**Figure 8**).

Perform Checking and Review the Results in the Physical Editor: The designer

All Constraints | Crosstalk | Distortion | Impedance | Delay | Misc | Modeling | Lengthening | Multi-board | Skew

	Layer	Impedance Template	Min Zo (Ohm)	Max Zo (Ohm)	Min Zo (Ohm)	Max Zo (Ohm)	Min track width (mm)	Max track width (mm)
LB_LA30	All		55	65	74	206	0.03000	0.20000
LB_LA30	L1		55	65	91	91	0.20000	0.20000
LB_LA30	L3		55	65	74	74	0.10000	0.10000
LB_LA30	L4		55	65	74	74	0.10000	0.10000
LB_LA30	L6		55	65	91	91	0.20000	0.20000
LB_LA30	WBP-A		55	65	206	206	0.03000	0.03000

Figure 6: *The designer can check delays*

All Constraints | Crosstalk | Distortion | Impedance | Delay | Misc | Modeling | Lengthening | Multi-board | Skew

	Min delay (ps)	Max delay (ps)	Max skew (ps)	Min relative delay (ps)	Max relative delay (ps)	Min Typical Flight Time (Rise) (ps)	Max Typical Flight Time (Rise) (ps)	Typical Flight Time (Rise) (ps)	Min Typical Flight Time (Fall) (ps)	Max Typical Flight Time (Fall) (ps)	Typical Flight Time (Fall) (ps)	Max skew (ps)
MCLK1		800	200			455	722		794	911		267

Figure 7: *The designer can check distortion*

may choose to review the design on the board itself. To do so, the designer checks the nets or the entire board for crosstalk and then locates the crosstalk with an interactive command that steps through each crosstalk victim (see Figure 8).

Resolve Issues and Send Design to the Engineer for Deeper Review: As issues are found, the designer should resolve them, if possible. When the designer saves the design, any colour-coded results will persist in the Constraint Manager window for the engineer to examine more closely.

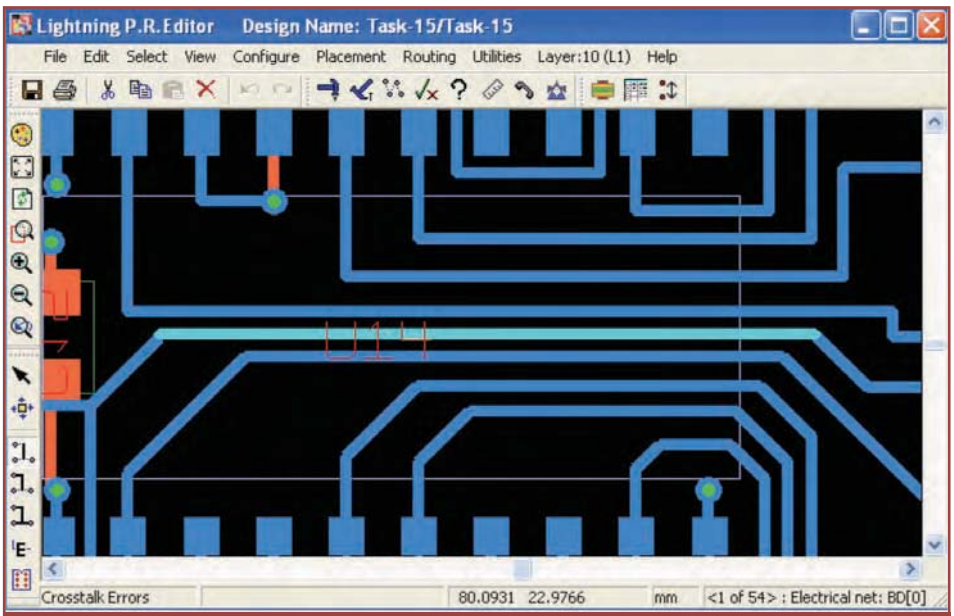

Figure 8: *The designer can check the nets or the entire board for crosstalk and then locate the crosstalk with an interactive command that steps through each crosstalk victim*

Teamwork

The practice of creating a high-speed digital design requires that both agents, the electrical engineer and the PCB designer, work even more closely.

The engineer needs to delve into matters that have typically concerned only the designer, such as stackup, trace widths and spacings, in order to yield the desired characteristic single-ended and differential pair impedances.

The designer needs to understand that routing for delay is different from routing for length and that signals with faster edge rates create more crosstalk than those with slower edge rates. With that in mind and armed with the proper tools, the designer who performs signal integrity screening during layout significantly improves the overall design process.

The designer addresses simple SI issues earlier, freeing up the engineer to focus on tougher problems. This practice can minimize the need for rework and reduce the number of large SI review cycles needed. All of this means that the product can be delivered earlier with higher quality, thereby saving time and money. ■

Improving process development with BETTER DATA

Dr Dirk Ortloff, Co-founder and Chief Technology Officer at Process Relations GmbH, discusses how better data management can benefit engineers in the high-tech process development field

PROCESS DEVELOPMENT

hinges on the effective use of data, whether that is historical or recent experimental data. Historic data, collected over many years and many projects, can provide engineers with a fantastic reservoir of knowledge; proper use of this data can give engineers a good head-start when developing new ideas and new processes.

This is certainly true for the semiconductor industry and for many other high-tech processes, but especially in the MEMS and NEMS fields where a new idea almost guarantees that a new method of manufacture is required. However, the problem is that there tends to be huge quantities of historical data. Moreover, this data has often been poorly managed and is dispersed over many servers and terminals, and filed in a way that can be impenetrable to everyone except the individual engineer who did it in the first place.

Obviously, this situation is a major hindrance to the development of new processes because engineers sometimes only have a vague idea of the results and their contexts. Many projects evaluating promising ideas either don't start properly, or are abandoned early on, because the task of sifting through and evaluating all of the historic data is just too time-consuming. However, many good ideas and exciting new developments can be lost or forgotten about as a result.

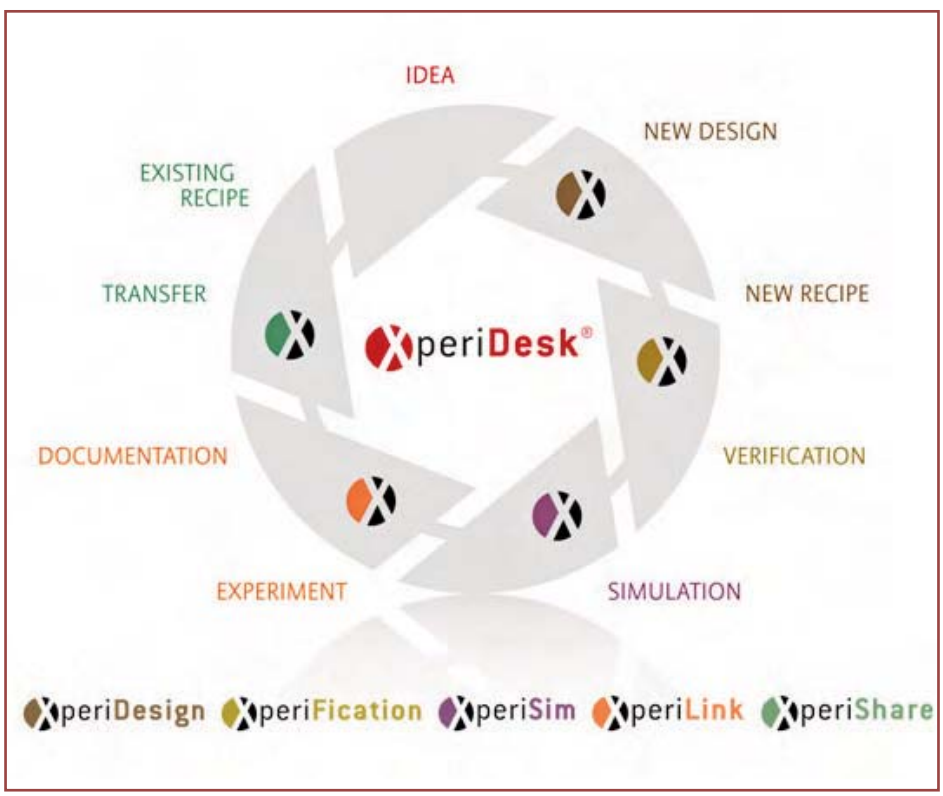

Software for Data Support

Clearly, this situation benefits no one. However, if all of the data and information could be recorded, stored, accessed and analysed easily there would be numerous advantages. These tasks can be easily performed with software. Using software to automate the recording and storage of data delivers the biggest benefit, freeing engineers to focus on the creative side, evaluating and analysing the data to support new ideas.

Indeed, improving data management can effectively result in 'new' knowledge –improved storage and retrieval of data means that it is far easier to analyse and cross reference past recipes and past experiments to see what worked and what didn't, what can be applied to the current project and what could work on other ideas engineers may have. Ultimately, deploying software to remove the burden of tedious data entry work from the engineers allows new ideas, new approaches and new developments to flourish.

To cope with these challenges in fabrication process development new approaches for adequate process design automation are necessary. One solution is a new kind of software – a Process Development Execution System (PDES). This type of software systems covers the complete process development cycle

Figure 1: *Development cycle*

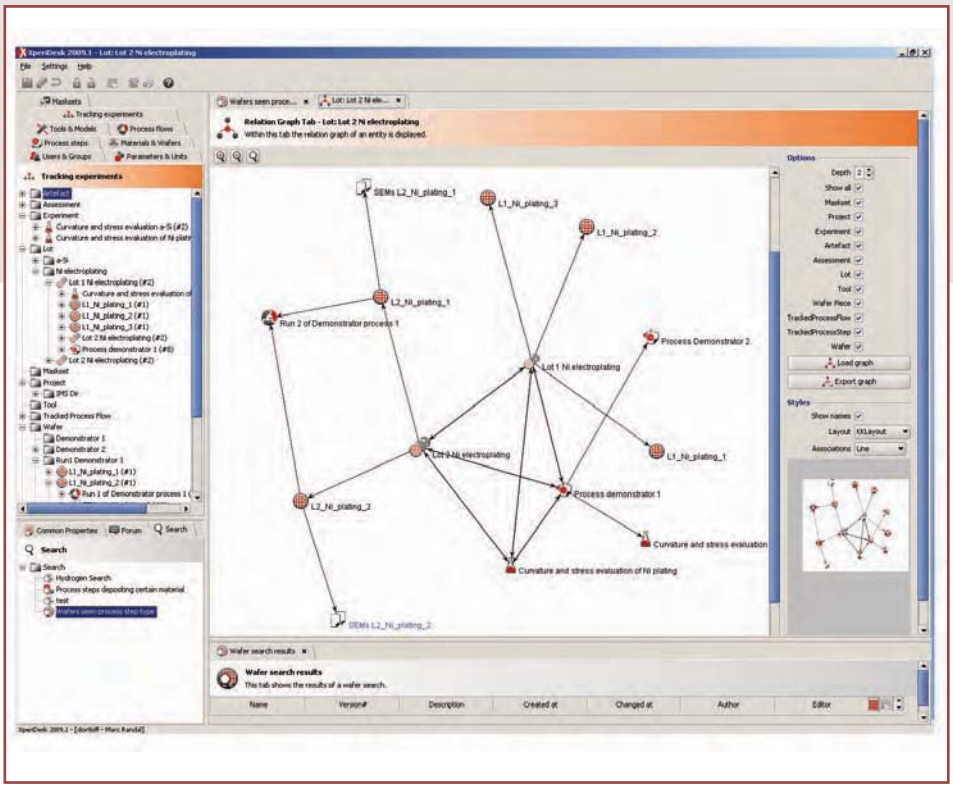

***Figure 2:** Graphical representation of information network*

depicted in **Figure 1** and closes the loop in providing seamless access to historical and current data.

PDES comprises a complete environment for all stages of process development, from initial concept to final experimental success. It provides a framework that can be readily adapted to customer-specific situations and procedures. When starting to develop new processes a PDES can provide a clear overview of available process step options and can automatically detect inconsistencies between process steps used in a process flow. These checks can detect issues like material incompatibilities, temperature budget overshoots, etc. At the simulation stage a PDES can provide a push-button interface to simulation tools and capture the results of simulations automatically.

When simulation results look promising, a PDES again offers a push-button solution to forward the verified process input parameters via a MES-interface to the fabrication facilities. After fabrication, measurement results and other data from process tests can be incorporated into the environment in an automated way. Experimental data can be linked together and relevant data sets are correlated with each other. This provides an excellent overview of the collected and accessible knowledge and information, which now also includes the failed experiments and their results.

Data Analysis Help

Software can also help with the data analysis. By properly recording and documenting experimental data the software tools can help engineers to see all of the comparison, visualization, and cause and effect data. Being able to access all of this information easily means that it is far easier for engineers to see correlations between experiments and between results. In particular, graphical representations of data and information can leverage holistic development data capture.

As an example, **Figure 2** shows how graphical representation allows visual exploration of the dependencies and relations between the different data objects. The data objects are visualized as nodes in an information network. The edges represent the relations between the data objects. Nodes can be selected as new exploration centres for visual navigation through the information network. Additionally, subsets of the information network can be blinded out temporary. Together with the graphical representation of node types this allows the selection of subsets of data for export. Tools such as this, realised in software, can provide invaluable help for engineers in their analysis of the various process steps.

Software can also be used to perform 'quick checks' of ideas via simulations/emulations. Using historic data and recipes, an engineer can quickly create an abstract of a process and run a simulation/emulation to check if an idea is viable. It is then possible to build on this initial abstract with more and more detail, checking each stage with an emulated manufacturing.

These checks mean that engineers have to perform far fewer speculative experiments inside the fabrication facility, saving money and resources, but also ensuring that new ideas are not discarded early on because of the risk of performing real, potentially incomplete experiments. An example of the results of such virtual manufacturing steps can be seen in **Figures 3** and **4**.

Completing The Picture

Collecting the data, managing that data and using the data with software is extremely beneficial, but it is still not the whole picture. The crucial element that needs to underpin these software functions is a powerful search tool that enables easy access to all of this data from diverse perspectives.

As mentioned above, if data cannot be accessed easily, even if it has been recorded and filed correctly, then it is still possible for new ideas to be unnecessarily discarded. One solution is Process Relations's XperiDesk software that has been designed with a powerful search function to ensure that engineers can quickly and easily find the specific data they want. In the phase of assembling process steps into process flows, XperiDesk can help assemble, store and print new process flows. By providing access to previously assembled process flows, a designer can use these as building blocks or modules in the newly developed flow.

The usage of standard building blocks can drastically reduce design time. Indeed, the assembly can be made even easier with

Figure 3: *Results of virtual manufacturing steps*

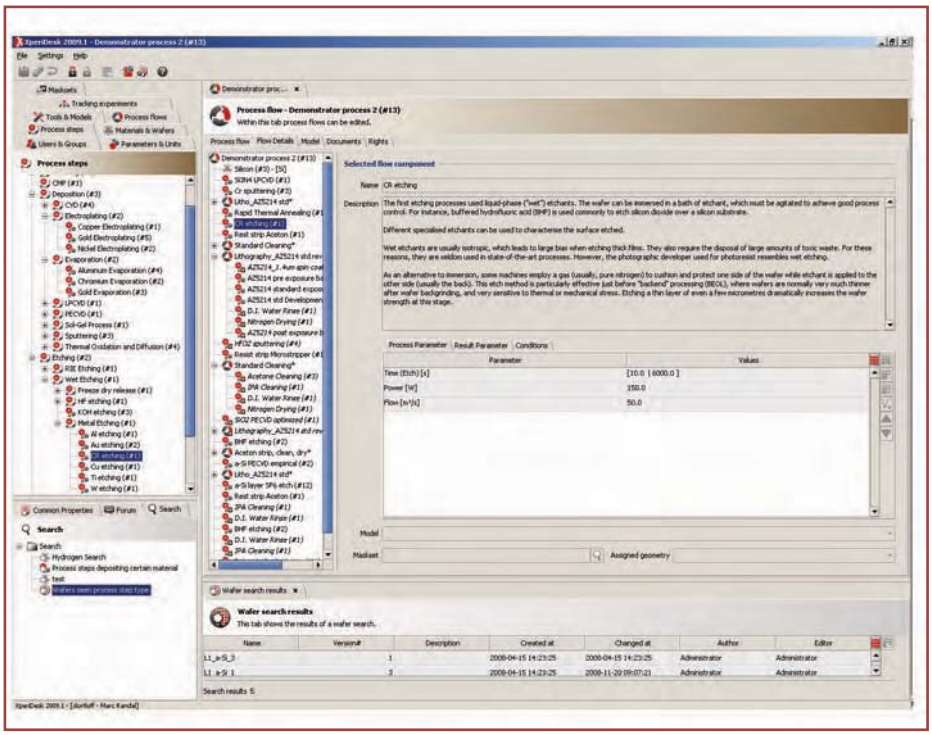

Figure 4: *Graphical representation of the results*

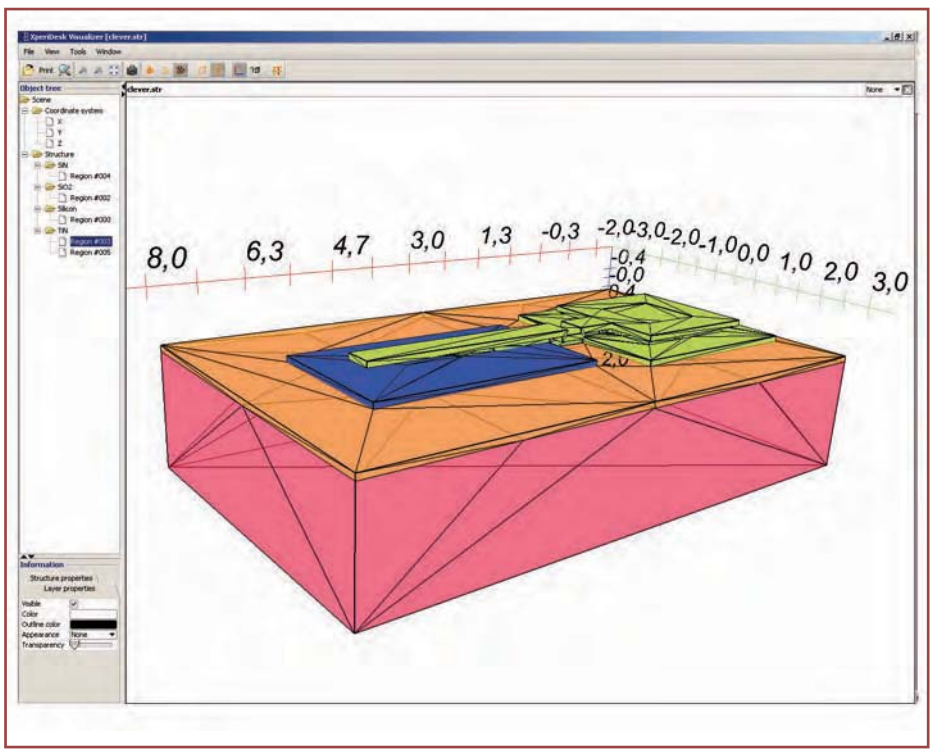

a drag and drop interface. To help with this software needs to be set up to manage data in a far more intelligent way, not just like an ordinary spreadsheet with only the raw numbers. Instead, every data point needs to include the full context in which it was created. This results in a powerful tool, significantly enhancing the search capability and information quality of software. It replaces the lab books and file servers. Information can be retrieved faster and previous results can be found and used more efficiently.

Indeed, a full PDES offers ways to view and search result data (e.g. materials, process steps, machines, experiments, wafers) from different viewpoints, or to categorise data under different aspects. It can also deliver methods to link entities together that belong in a common context and explore the resulting information network.

The limitations of file servers and their one dimensional view is removed such that the system can easily answer questions like "Give me all substrates that have been processed by a process step of a certain category or machine," or "Give me all machines capable of depositing a certain material or category of material." These multi-dimensional search capabilities add significantly to the productiveness of engineers.

Taken together, all of these software tools help to make sure that all of the data created during process development is not wasted and remains accessible. The software tools also ensure that good ideas are given a chance to come to fruition, rather than being discarded. However, the best benefit is that, by using software to make the process more efficient, engineers can concentrate fully on developing new ideas and analysing data, rather than on data collection. ■

Don't just test it... ...Analyse it!

PEAK electronic design ltd

The New Atlas ESR Plus, Model ESR70

This new model of the famous ***Atlas ESR*** offers all the great features of the ESR60 but with extended measurement range and audible alerts.
This is the *Atlas ESR Plus*!

- Capacitance from 1uF to 22,000uF.
- ESR from 0.01 ohms to 40 ohms.
- Great for ESR and low resistance measurements (short tracing etc).
- Automatic controlled discharge function.
- Audible Alerts (for good ESR, poor ESR, open circuit and more).
- Gold plated croc clips.
- User Guide and Battery included.

Atlas SCR - Model SCR100

Connect Triacs or Thyristors any way round.
Auto part identification and pinout display.
Categorises gate from 100uA to 100mA.
Load conditions regulated at 12V, 100mA, even for a dying battery.
Measures gate voltage drop.
Long life alkaline battery supplied.
Supplied with premium probes.

Reduced Price!
was £97
£87.00 inc VAT (£75.65+VAT)

Atlas IT - Model UTP05

Special Offer!

Identify network cabling faults as well as identifying many types of connection configurations automatically.
Great for testing sockets and cables.
Complete with all this:

was £85
£69.00 inc VAT (£60.00+VAT)

Atlas ESR - Model ESR60

Special Offer!

Measure capacitance and ESR!
Capacitance from 1uF to 22,000uF
ESR from 0.01 ohms to 20 ohms
Battery included + Gold plated croc clips

£75.00 inc VAT (£65.21+VAT)

was £87

Atlas DCA - Model DCA55

The famous Peak Atlas, now with fitted premium probes.
Just connect any way round to identify the type of semiconductor, pinout and lots of parameters too.
Complete with battery, user guide and new premium probes.

Darlingtons
MOSFETs
Diodes
Transistors
LEDs
and more...

£53.83 inc VAT (£46.81+VAT)

Atlas LCR - Model LCR40

Passive component analyser.
Automatically identify and measure inductors, capacitors and resistors.
Auto frequency selection.
Removable probes.

1uH - 10H
1pF - 10,000uF
1 Ohms - 2M

£77.31 inc VAT (£67.23+VAT)

Atlas Star Pack (LCR/DCA)

Money Saving Pack: Save £20

Includes the Atlas LCR, Atlas DCA and a premium padded carry case.
Complete with extra spare battery.

£124.99 inc VAT (£108.69+VAT)

Peak Electronic Design Ltd, West Road House, West Road, Buxton, Derbyshire, SK17 6HF.
tel. 01298 70012 www.peakelec.co.uk sales@peakelec.co.uk

Please add £2.00 p&p per order. Prices include UK VAT.
See website for overseas prices.

Special Offer prices for limited period or while stocks last.

LEGAL COMPLIANCE – Software Development to

Mahshad Koohgoli, CEO, **Dr Sorin Cohn-Sfetcu** and **Kamal Hassin** at Protecode discuss what's needed to keep the quality of software intact but also ascertain its legal compliance in today's ever more present open source and large scale outsourcing environments

IN THE AGE of open source and large scale outsourcing, both assuring the quality of software and taking it to market means ascertaining its legal compliance as well. Numerous legal cases in recent years have highlighted the business risks and the enormous costs incurred when this is not done properly.

These costs stem from involvement in judicial procedures, software recalls, fixing legal compliance issues post-release and missed market opportunities caused by delays in the development process. Other consequences include lowered valuations in due diligence processes triggered by customers, potential or existing investors, mergers and acquisitions, and other major transactions.

Software is a pervasive element in most products and processes and, over time, its sources have multiplied. Sources include internal developments, suppliers of sub-systems and chips, outsourced development contractors, open source repositories and the previous work of the developers themselves. Software, unlike hardware, is easily accessed, replicated, copied and re-used.

Open source software has become a significant player in most software development, thanks to the wide availability of source code, its apparent free cost and its high degree of stability and security.

Open source code is generally free on the surface, but it's not without obligations. It comes laden with licensing and copyright conditions that are enforceable by law – sometimes with dire effects for users who are not careful to validate the pedigree of the code in their products; i.e. the origin and any associated obligations of all software components.

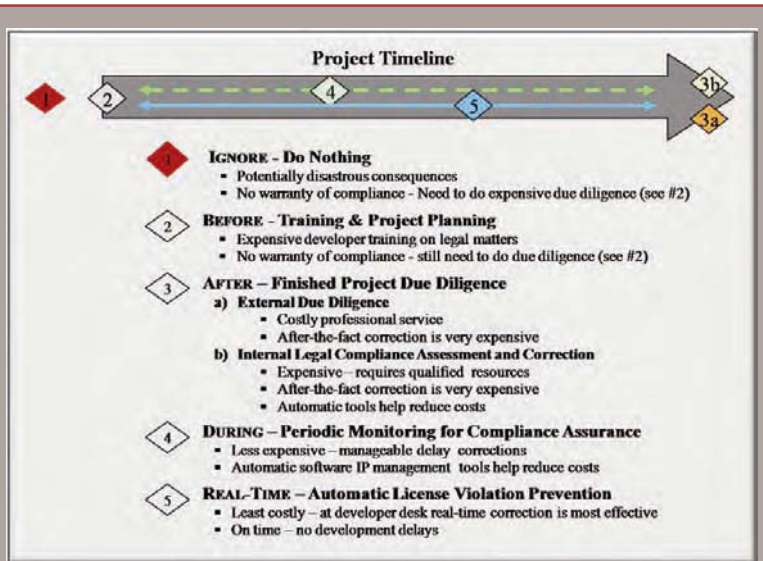

Figure 1: *Project timeline*

This doesn't mean that leveraging outsourcing and/or open source software is to be avoided. The issue is not with the use of open source, but with unmanaged adoption and lack of proper care to the copyright and licensing obligations it entails. It's paramount that industrial managers validate the IP cleanliness of their products and services and ascertain that they meet all legal obligations before they reach the market.

Principal Aspects of Legal Compliance

Assuring compliance to legal obligations implies the following three major aspects:

1. Definition of a corporate (or specific project) intellectual property (IP) policy which must be met by all associated products and services.
2. The auditing of software to determine all implied legal obligations as per associated IP policy.
3. The necessary fixes – legal or development intensive – such that all software components meet said IP policy.

The IP Policy must be defined in accordance with both the business goals of the organization and its engineering processes. Therefore, it requires the involvement of business and engineering managers, as well as the proper legal counsel. The policy must be clear and enforceable. It should be captured for distribution and application within the development and quality assurance departments.

Auditing software for legal compliance is a process that is traditionally only begun just before major commercial or financial events. It's a complex process: preparation, document review, management conferences, designer conferences, analysis, legal consulting and

From Delivery

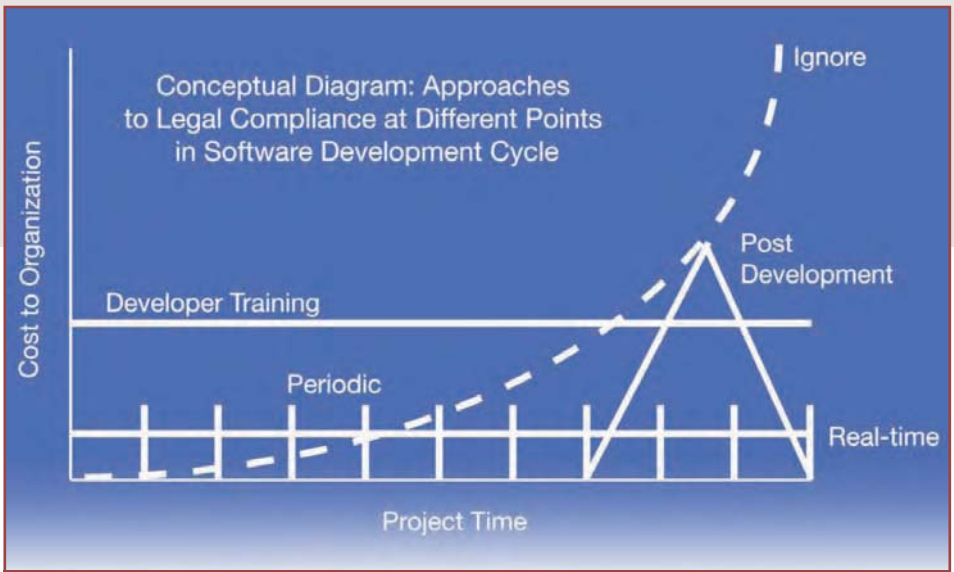

Figure 2: *Conceptual diagram*

reporting. It is time-consuming and expensive as it involves valuable engineering, management and legal resources. Even then, in most cases, the results have been inaccurate as there are usually insufficient records on what is actually in the software. As these problems continue to emerge, automated tools for auditing the software composition and determining legal obligations are becoming a "must-have" component.

The "fixes" necessary to make the software legally-compliant as per IP policy can be complex. Some software components may have to be replaced entirely due to IP infringement. This can be expensive, as new software components have to be found and the overall software needs to be re-tested. In other cases, it may be sufficient to formalize the assumptions of obligations as demanded by license or copyrights.

Bringing Legal Compliance Assurance Into the Development Process

Mitigating business risks associated with software legal compliance is best addressed by building legal considerations into the development process itself. The following options address compliance measures at different points in the development process. Some of the options listed, such as periodic and real-time assessment, can be used in combination for best results.

Do Nothing:

Deciding to ignore the compliance issue carries the lowest up-front cost but bears the highest business risks and largest costs post market introduction.

Developer Training:

Some companies – especially small and mid-size ones – consider that proper training and project planning is sufficient in normal situations, accepting to undertake an audit during imposed due-diligence efforts. Naturally, the more the developers are trained on matters of software legal compliance issues, the more effective the development process can be. This is, however, a rather expensive proposition, given the explosive growth in number of distinct software licenses, the high cost of developer training and the constant churn within the development environment. With this option, compliance depends solely on developers and there is still no assurance of legal compliance before going to market.

Post Development Analysis and Correction:

Taking action later in the project lifecycle can take the form of external or internal auditing and impacts the final stages of testing and the quality assurance process. This option can bear higher costs due to professional services, the cost of any necessary changes to the software after the fact, subsequent re-testing and re-auditing. This option gets results, does not impact development workflow and can be rendered more cost-effective with software tools designed for this purpose. It can, however, prolong the project lifecycle near the end, resulting in delays to the delivery of the final product that are hard to predict.

Periodic Analysis and Correction:

Periodic auditing of software during development involves course corrections along the way if any policy violations are detected. This can be done with automatic tools and is less expensive than waiting until after the development process thanks to the shorter delays in getting the fixes done and re-tested.

Real-Time Preventive:

The most pro-active measure for software compliance assurance is to detect license violations immediately at the developer workstation in real time. The development process is not disturbed and the cost of corrections is minimized as any necessary corrections – which might include justification of selection, code changes or replacement – are done on the spot without involvement of other resources and without need for re-testing. This process can be automated via software tools in ways that are unobtrusive, easy to adopt and, most importantly, do not require developer training in matters of legal compliance. Detecting possible violations in real-time is the most cost-efficient and lowest risk option in the long term.

The later in the software lifecycle such fixes are affected, the more expensive they become. If the legal compliance issues are discovered during the development process, the fixes become less onerous and the business risks are reduced.

Bringing Legal Compliance Into the Software Product Lifecycle

From a business and product management perspective, legal compliance goes beyond the development process and needs to be dealt with at project conception and from a customer standpoint. The critical elements of effective software IP management in an organization are:

- Existence of an IP policy for each project undertaken and a process to disseminate and apply it. Corporate IP policies must be based on the organizations' business goals and they should be clear and enforceable.
- The availability of a central code library, which includes the legacy code in the organization, together with an automated process for ascertaining the pedigree of all components to ensure compliance to all legal obligations.

- Processes and tools for ascertaining the legal obligations and managing the IP of software created and/or acquired in the organization.
- Software Bill of Materials (BoM) that fully records the components in the product, their provenance and the licensing obligations they entail. An adequate software BoM is instrumental in determining the legal compliance of the software.
- Assurance and support for customers concerning the quality and IP cleanliness of software provided.

These elements provide a basis for meeting legal compliance with respect to the lifecycle of the software product from conception to delivery.

Modern software IP management tools allow the implementation of a simple and efficient process for managed open source software adoption which allows developers the freedom of selecting best solutions appropriate with the corporate IP policy. The main stages of such a process are:

1. The definition and tool-capture of an IP policy, together with the actions to be taken in case of violations.
2. Ensuring the legal compliance of legacy and/or acquired code.
3. The real-time detection and fixing of any IP policy violations by new code created or brought in by developers.
4. The automatic verification of any code checked into the organization repository/library.
5. The automatic product build-software analysis and IP policy compliance certification via complete software BoM.

Such a software lifecycle management process ensures automatic compliance with the IP policy without imposing specific pre-approval of open source components. An optional stage dedicated to pre-approval of open source components and the management of a repository of approved open source can be considered as part of Stage 3 above.

With respect to the tools available, modern software IP management applications simplify and enable safe open source adoption, giving developers the freedom to select the best solutions in accordance with the corporate IP policy. For instance, these tools can support pedigree analysis and IP policy violation detection automatically – on demand, on schedule or even in real time within the development process. They can also provide BoM on demand. Taken together, these IP management features deliver higher value and provide customer assurances.

As the critical factors driving the economics of software IP management are the efforts to fix the software IP issues and minimize the associated delays in product introduction to market, everything should be done to ensure its legal compliance throughout its lifecycle for maximized cost efficiencies and minimized risk. As companies continue to leverage third party code, legal compliance issues become increasingly integral to business priorities. Consciously implementing measures for legal compliance in the development process itself, as well as incorporating aspects of effective software IP management into the organization, are now crucial for any entity concerned with software development and delivery. ■

NEW!

PSR30 & PSR60 Series Non-isolated DC-DC Switching Regulators

Features

- Wide 20-53 volt input range
- Power up to 30W (PSR30) & 60W (PSR60)
- Single 5V, 12V, 15V & 24V DC Outputs
- Remote on/off control
- Very high efficiency up 96%
- -25°C to +65°C operating ambient
- 250kHz switching frequency
- Compact SIL package
- Ideal for battery or distributed power applications
- RoHS compliant

POWER SOLVE

www.powersolve.co.uk

Tel: 01635-521858 Email: sales@powersolve.co.uk

VIRTUALIZING Embedded Linux

In this article, **Gernot Heiser**, cofounder and chief technology officer of Open Kernel Labs (OK Labs), explains embedded and mobile virtualization technologies, and explores how they can make Linux-based applications faster, more responsive, secure and IP-friendly.

This is the first article in of a two-part series that continues in the next issue

LINUX IS INCREASINGLY the embedded OS of choice for a wide array of applications, from mobile handsets and network/telecom infrastructure to media-rich consumer electronics devices like portable media players and digital video systems. As many as one third of embedded systems developers already use Linux in current products or are considering doing so in next-generation designs. OEMs find that it is easier to contract developers and to applications for Linux than for proprietary OSes, and that using Linux reduces costs.

Linux Challenges

However, Linux can still present a number of challenges to building and deploying embedded applications. These include the following:

- **Bifurcated Programming:** embedded development task is often divided between Linux-based, high-level programming and coding for low-level RTOS code. For example, in many media-rich consumer applications, Linux hosts high-level application code that is similar – or even identical – to application code running on personal computers. This code is typically developed by application programmers, not experts in programming low-level embedded systems.

  Embedded applications can also require real-time functionality with low and predictable interrupt and pre-emption latencies. In the case of mobile phone terminals, wireless communications subsystems usually impose tight real-time requirements. And while Linux kernel developers have certainly improved Linux responsiveness, small lightweight RTOS kernels still best meet these performance and latency requirements.
- **Security:** in a mobile-phone handset, for example, the communication stack is of critical importance. If subverted, the phone could be turned into a jammer, disabling communications in the whole cell. Similarly, an encryption subsystem needs to be strongly protected from being compromised.

  It's no insignificant challenge to create a secure system comprising a large code base. Today's embedded applications code contain hundreds of thousands, or even millions, of lines of code, with up to tens of thousands of bugs, many of which can compromise system security.

  Increasingly prone to attacks, embedded Linux implementations are large enough to contain thousands of bugs. Because the Linux operating system normally runs in privileged mode, once compromised, all parts of the system will be vulnerable to attack.
- **Hardware-Software Integration:** software may have "nine lives" but hardware designs and sourcing change with increasing frequency. The need to integrate, debug, test and certify new device interfaces and code running on new CPUs injects costs at the OEM and in channel during the life of embedded projects, again with each generation of a product and across product lines.
- **IP and License Separation:** Linux is a frequently deployed high-level operating system. Among its advantages are the royalty-free status, vendor-independence, widespread deployment, and a strong and vibrant developer community. Some OEMs, however, still harbour concerns about code distributed under reciprocal open source licenses (like the GNU General Public License – GPL), with disclosure requirements that accompany deployment. These concerns span hardware and software stacks, from device drivers to CODECs, even to some application code.

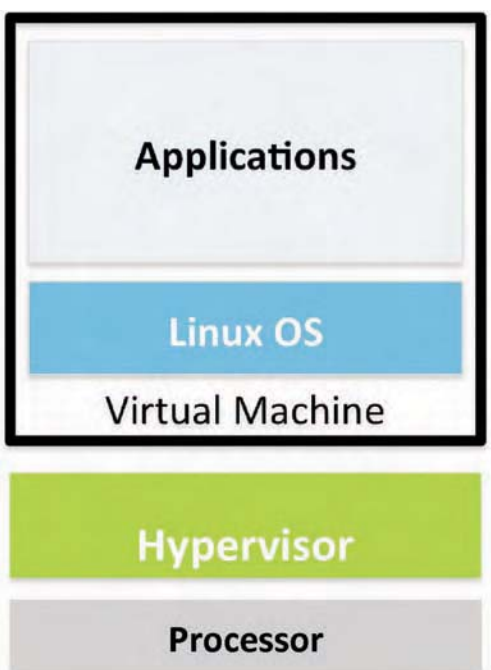

Figure 1: *Bare metal virtualization*

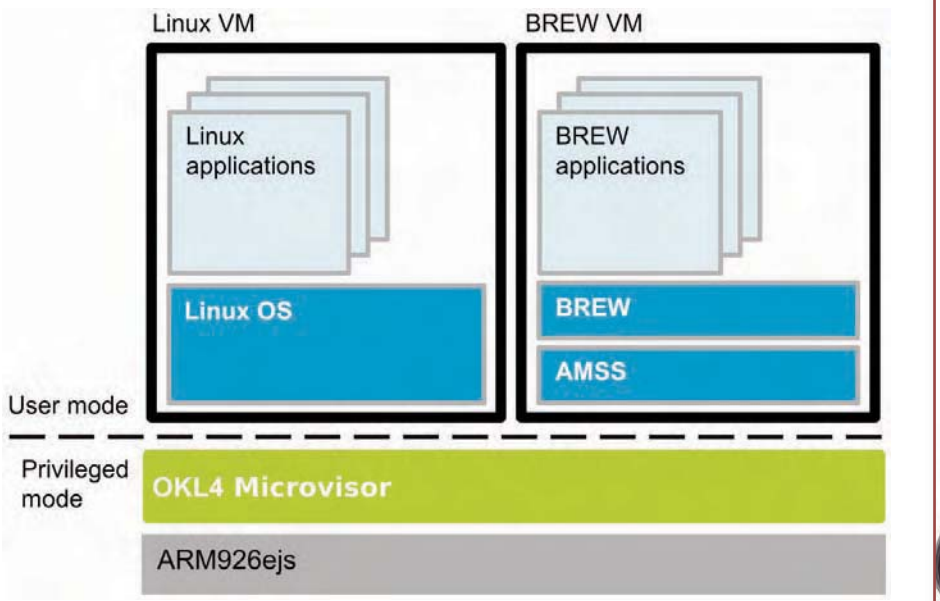

Figure 2: *The Motorola Evoke QA4 virtualization architecture*

Virtualization Basics

Virtualization refers to providing a software environment in which programs (including operating systems) can execute as if on actual "bare" hardware (**Figure 1**).

Such an environment is called a virtual machine. A virtual machine is a container, an efficient isolated simulacrum of actual computer hardware. The software layer that provides the virtual machine environment is called the virtual machine monitor (VMM), or hypervisor, with three essential characteristics:

- It provides an environment for programs that is essentially identical to the original machine.
- Programs that run in this environment exhibit minimal performance degradation.
- The hypervisor is in complete control of system resources.

All three characteristics are important and contribute to making Virtualization viable. The first (similarity) ensures that software that runs on the real machine will run on the virtual machine. The second (efficiency) ensures that Virtualization is practical from the point of view of performance. The third (complete control) ensures that programs cannot break out of their virtual machines.

Efficiency requires that the majority of instructions execute directly on the hosting hardware, avoiding performance-sapping emulation or interpretation that involves multiple instructions replacing single operations of the underlying hardware.

Such near-hardware performance requires that the virtual hardware be almost identical to the physical hardware on which the hypervisor is hosted. Small differences are possible, such as the physical hardware may miss some instructions of the virtual hardware (as long as they aren't heavily used), the memory-management unit may be different, or devices may differ.

An important difference lies in the ability execute privileged instructions. Disciplined resource control requires that instructions impinging upon system resources must access the virtual rather than physical instantiation of those resources. Such privileged instructions must be interpreted by the hypervisor, to avoid violating the definition of an isolated virtual machine.

Specifically, the virtual machine must interpret two classes of instructions:

- Control-sensitive instructions that modify the privileged machine state and, therefore, interfere with the hypervisor's control over resources;
- Behaviour-sensitive instructions that access the privileged machine state.

While these instructions can't change resource allocation, they reveal the state of real resources, specifically when they differ from the virtual resources and therefore break the virtual machine "illusion".

Virtualization Typology

There are two basic ways to ensure that code running in a virtual machine doesn't execute sensitive instructions: "pure" virtualization, which relies on blocking privileged instructions from executing and paravirtualization, where sensitive instructions are removed from software deployed in a virtual machine at compile or build-time.

The classical approach is Pure Virtualization, wherein all sensitive instructions be privileged. These instructions can be executed only when running in a privileged state of the processor (typically called 'privileged mode', 'kernel mode' or 'supervisor mode'). These instructions cause an exception when executed in 'unprivileged mode' (also called 'user mode'). Exceptions enter privileged mode at a specific address (the exception handler), which is owned by the hypervisor. The hypervisor can then interpret (virtualize) the instruction as required to maintain virtual machine state.

Until recently, pure virtualization was impossible on most or all embedded CPU architectures, as they included sensitive instructions executable outside privileged context. For several years now, major enterprise/desktop silicon suppliers, Intel and AMD in particular,

HYPERVISOR TYPOLOGY

System-level virtual machines support sharing of underlying computer hardware resources among different virtual machines, each running its own operating system or other stand-alone software. The software layer supporting this paradigm is called a virtual machine monitor (VMM) or more commonly, a hypervisor.

Hypervisors come in two 'flavours':

Type I Hypervisor: Bare Metal
This is a hypervisor that runs on top of actual machine hardware and mediates execution and resource allocation to all hosted software, in particular to "guest" operating systems and their applications.

Type I Hypervisors have been deployed for decades on mainframe computers and today are the staple of data centre virtualization. Embedded and mobile virtualization also employs Type I Hypervisors to host general purpose OSes like Linux and Android, side by side with RTOSes.

Type II Hypervisor: Application VMM
Hypervisors can also run as environments hosted on operating systems. These application-level virtualization platforms enable workstation and personal computer users to run Linux on Windows PCs or to run Windows over MacOS.

Type II Hypervisors can impose resource and compute requirements on their host systems that make them mostly unsuitable to embedded applications.

have added virtualization extensions to support configuration to force exceptions for sensitive instructions and reduce the cost of virtualization. More recently, licensees of embedded architectures like ARM, MIPS and Power Architecture have begun to field test comparable capabilities.

Despite these advances, pure virtualization may still not suit deployment in embedded applications. One is that advanced hardware features are not available on all processors, especially lower-cost SoCs for high-volume applications. Another is that, computationally, exceptions are expensive. On pipelined processors, exceptions drain execution pipelines, resulting in delays in processing, typically one cycle per pipeline stage. A similar delay occurs when returning to user mode after an exception. Moreover, exceptions (and exception returns) implement branches usually not foreseen by processor branch-prediction, resulting in additional latency.

These effects typically add up to 10 to 20 cycles, more in deeply pipelined high-performance processors. Adding the work required for actual instruction emulation can result in virtualization of a single instruction expending dozens of cycles. Some processors (notably the x86 family) have exception costs that are even higher (hundreds of cycles). These costs create substantial overhead for operating system code, which frequently executes many privileged instructions in a short time.

Paravirtualization

Paravirtualization mandates modification of source code to remove direct access to privileged state information, replacing such accesses by explicit invocations of hypervisor APIs ("hypercalls"). Paravirtualization often involves replacement and consolidation of multiple sensitive instructions by a single hypercall, also reducing the number of (expensive) switches between unprivileged and privileged mode.

When properly implemented, paravirtualization can dramatically reduce virtualization overhead. Variants of paravirtualization have been deployed for years by enterprise platforms from VMWare and Xen. In the last several years, paravirtualization platforms targeting embedded designs have also emerged, such as L4/Wombat from the University of New South Wales, and the commercial system supplied by my company, OKL4 from Open Kernel Labs.

Paravirtualization, however, requires access to and understanding of the source code in question. Open Source OSes, like Linux and Android, are therefore most readily accessible to paravirtualization, but can still demand substantial engineering investment to run as paravirtualized OSes. In fact, enterprise and embedded virtualization suppliers have invested substantial effort in optimizing Linux and other OSes to run well on their platforms (e.g. Microsoft supplying GPL Linux drivers to run on Hyper-V).

To meet both the challenge of access and familiarity for embedded deployment OK Labs delivers ready-to-integrate off-the-shelf versions of Linux, Android, SymbianOS and other guest platforms instead of forcing OEMs, integrators and others to acquire, comprehend and customize those platforms for paravirtualized execution.

Practical Application – A Virtualized Mobile Phone

Last year, handset manufacturer Motorola released the Evoke QA4 messaging phone, the world's first virtualized smartphone. Motorola chose the architecture as depicted in **Figure 2** to support a mix of requirements, including the need to sustain legacy baseband and applications execution on the BREW platform, and desire to present users with a modern user-nterface running on Linux.

In the Evoke architecture, both Linux and BREW, run as de-privileged "guest" OSes, executing in user mode. User-interface (UI) functionality is implemented over the Linux stack, while baseband (AMSS) and multimedia functions like video rendering use real-time software engines running on BREW.

BREW and Linux applications co-exist on the same Linux desktop, integrated into the same UI, and the Linux stack owns the Evoke touch screen and other input sources. When a user requests invocation of a BREW application, Linux communicates with BREW across virtual machines (using secure IPCs) to invoke the application. BREW applications, in return, leverage Linux-based display services for UI presentation.

Virtualization overhead is kept to unnoticeable levels by using a lightweight, high-performance hypervisor, in this case the OKL4 Microvisor. Through intelligent address space management and by thorough-going utilization of the resident ARM9 MMU, the architecture avoids flushing caches and TLBs on context switches, greatly enhancing performance. The boost in context-switch performance in particular benefits the Evoke user interface. The same approach is finding favour with OEMs for rapid bring-up and integration of Android in tandem with legacy baseband and application stacks.

Benefits of Virtualization

Embedded developers can enjoy a range of concrete benefits from adding virtualization to their design and implementation toolbox. These include:

- **Hardware consolidation:** save on Bill-of-Material (BoM) costs by running multiple stacks on a single CPU instead of diverse dedicated silicon.
- **Better security:** confer greater security on embedded designs by cleanly isolating software, especially high-level OSes subject to exploits, and by reducing the Trusted Computing Base (TCB) from millions of lines to just the code within the embedded hypervisor.
- **Ease integration:** add new software to existing designs by encapsulating it within dedicated virtual machines (including legacy OSes) without need to re-architect or port to new platforms.
- **Product line support:** a virtualization platform provides a uniform operating system environment across one or OEM product lines.
- **Future-proofing designs:** facilitate addressing future requirements and design decisions by abstracting hardware dependencies from value-added software, and by enabling preservation of legacy investments in dedicated virtual machines.
- **Low overhead:** a well-crafted hypervisor and paravirtualization induce extremely low overhead. In some cases it can actually improve performance by reducing the need for costly context switches and the overhead associated with them, and through optimized IPCs.
- **IP segregation:** make intelligent choices about how and where to deploy code with highly reciprocal licenses or other restriction, easing concerns from legal departments and observing corporation IP policies.

Meeting Developer Challenges

This article has examined the particulars of embedded virtualization and how it addresses the challenges faced by developers building and deploying Linux-based devices. In the example mobile phone application, as in myriad other scenarios and device types, virtualization can unify bifurcated development, improve security and ease many integration challenges. These and other benefits make embedded virtualization an extremely attractive tool to support both building next-generation intelligent devices, and also for extending the lives of existing ones. By virtualizing embedded Linux, it further broadens the application space for this truly ubiquitous embedded OS. ■

MIND THE GAP – Making the Link from Requirements to

Brian Hooper and **Bill StClair** of LDRA delve into an area that is regularly flagged up by gap analyses of software projects: the troubling field of requirements traceability, a key discipline for safety and security-critical projects, as well as being one of the remaining areas of software development that is yet to be fully automated

GAP ANALYSIS is a technique regularly used in business to ascertain the maturity of working processes and to identify areas for improvement. Companies are forever seeking to evolve and improve whether to reduce overheads or to achieve a demonstrable level of quality such as ISO 9001 or CMM/CMMI. The easiest way to prepare a plan for business evolution is by gap analysis; that is, gather data and then perform analysis on that data to gauge the difference between where the business is currently and where it wants to be.

Gap analysis offers an opportunity to examine operating processes and generated artefacts, typically employing a third party for the assessment, without the pressure and constraints of a formal assessment or certification. The outcome will be notes and findings upon which the company or individual project may act, valuable information that will help improve processes such that the formal assessment or certification, when carried out, is much more likely to be passed at the first attempt.

Companies involved in systems and software development for the aerospace industry, through working to standards such as DO-178B, are familiar with the efforts required to achieve certification for their products from the relevant authority. The need for certification has mandated business evolution such that processes and project plans are documented, requirements are captured, implementation and verification are carried out with respect to the requirements, and all artefacts are fully controlled in a configuration management system.

An ever-increasing reliance on software control has meant that many companies from non-aerospace business sectors, that do not have a traditional requirement for sophisticated software development processes and practices, now find themselves compelled to undertake safety- and security-related assessments to prove the quality of their products. Consider the modern motor car equipped with all manner of safety systems, such as anti-lock braking and traction control, each managed by software running on a networked set of processors; the failure of any of these systems will be a major safety concern and could even lead to recalls and lawsuits. With the need for increased software quality across different industries such as medical, telecommunications and financial, companies are looking outside their own market sector for best practice approaches, techniques or standards. Gap analysis provides an established method by which these companies may isolate the areas in which they need to improve with respect to a standard such as DO-178B; in addition, the results allow the company to correctly and efficiently focus resources in order to achieve that improvement.

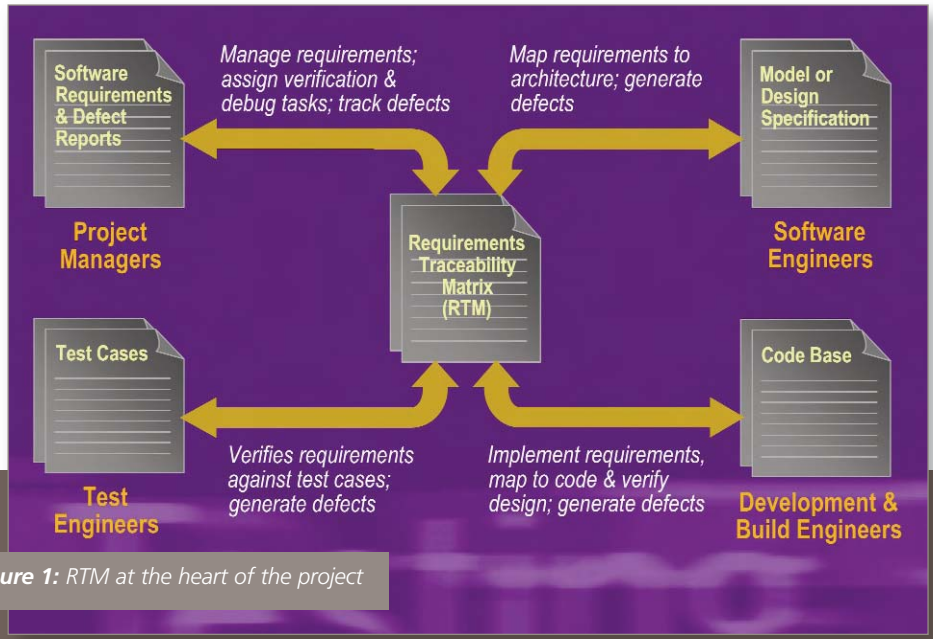

Figure 1: *RTM at the heart of the project*

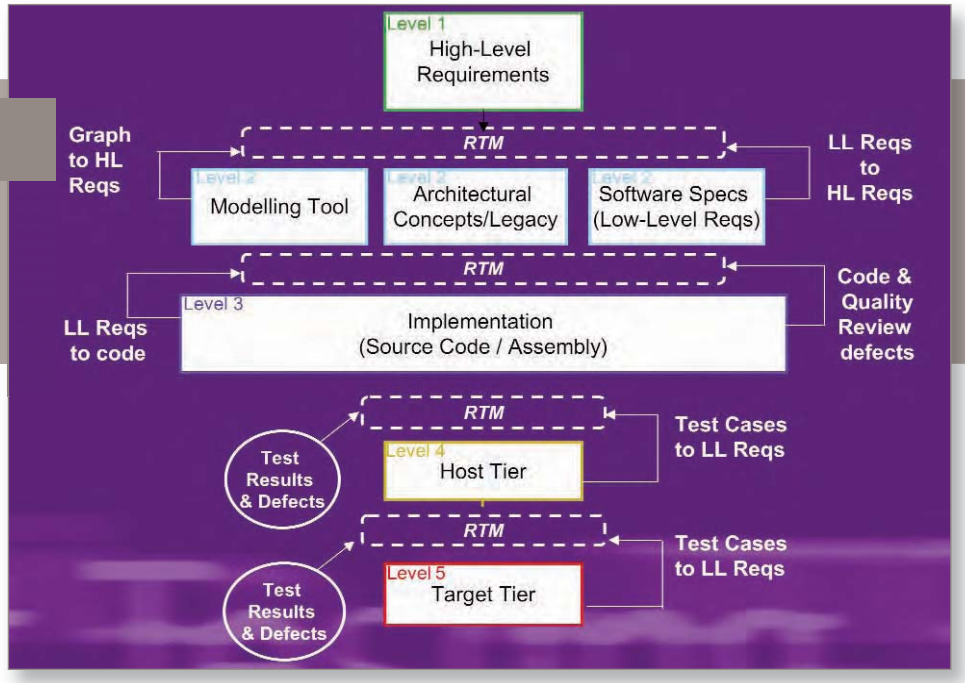

Figure 2: *Development lifecycle model emphasizing the RTM*

Area of Concern

When companies undergo gap analysis, whether for military or commercial projects, the results often show maturity within each individual phase of the software development lifecycle and a high level of investment in tools to efficiently undertake each phase. However, the results regularly reveal one particular area of concern, that of traceability between the software development phases. More often than not, the construction and maintenance of traceability matrixes is performed as a low-priority task and carried out manually, requiring continual human interaction and interpretation regarding what traceability means.

Evidence of traceability is a key deliverable for most project standards, and failings in this area have major repercussions when certification is undertaken. With there being a strong correlation between requirements degradation and software defects, companies are becoming more focused on ways to mitigate this risk through rigorous requirements and traceability management. Meanwhile, as the broader military market is attempting to save costs through the use of commercially-developed components, there is a drive to reconcile the difference between military standards and commercial standards under which products are developed to further trim costs, particularly in the areas of verification and certification which can account for 50% to 70% of the overall development budget.

Requirements Traceability

Requirements traceability is widely accepted as a development best practice to ensure that all requirements are implemented and that all development artefacts can be traced back to one or more requirements. A Requirements Traceability Matrix (RTM) is also a key deliverable within many development standards.

Despite good intentions, many projects fall into a pattern of disjointed software development in which requirements, design, implementation and testing artefacts are produced from stand-alone, silo-like development phases, resulting in tenuous links and references between themselves and the overall RTM.

This pattern is just as evident on projects using state of the art requirements management tools, modelling tools, IDEs and testing tools. Focusing on requirements in particular, the lack of reference to later development phases is attributable to the centralized, database-like architecture and application model of most requirements management tools; there is plenty of functionality to encourage good quality and good management in the requirements domain, yet little to aid the downstream effort where projects are designed, implemented and tested.

In many companies, the requirements manager is a database specialist with responsibility for the RTM on multiple software development projects, usually with limited knowledge of those projects. The RTM is maintained in a lightweight repository with no direct connectivity to the software development artefacts, nor to the status of these artefacts. Management is typically one of manually bringing the RTM updates and the relevant artefacts to the requirements repository, whether on paper or via email or even over the telephone. If any formal structure is in place, it is typically a set of intermediate files managed in tools such as Microsoft Excel, where omissions and mistakes are inevitable.

Maintaining RTM

The tradition view of software development is to show each phase flowing into the next, perhaps with feedback to earlier phases, and a surrounding framework of configuration management and process (e.g., Agile, RUP). Traceability is assumed, via the relationships between phases; however, the mechanism by which trace links will be recorded is seldom stated. The reality is that, while each individual phase may be conducted efficiently thanks to investment in

HOW DO-178B DEFINES COMPLIANCE DETERMINATION

The certification authority determines that the aircraft or engine (including the software aspects of its systems or equipment) complies with the certification basis. For the software, this is accomplished by reviewing the Software Accomplishment Summary and evidence of compliance. The certification authority may review at its discretion the software life cycle processes and their outputs during the software life cycle.

Level 1 High-Level Requirements	A definitive statement of the system to be developed and the functional criteria it must meet. This level may or may not need to be elaborated further.
Level 2 Design	A representation of the design of the system described by Level 1. Above all, this level must establish links or traceability with Level 1 and begin the process of constructing the RTM. Capture low-level requirements, these being requirements specific to the design and implementation domain with no impact on the functional criteria of the system.
Level 3 Implementation	Produce the source/assembly code in accordance with Level 2. Verification activities begin, including code rule checking and quality analysis. Maintenance of the RTM presents many challenges at this level, tracing requirements to source code files may not be specific enough, linking to individual functions may be required.
Level 4 Host-Based Verification	Begin formal verification – the test strategy may be top-down, bottom up or a combination of both. Make use of software stimulation techniques, automated test harnesses and test case generators as necessary. Test cases should be repeatable at Level 5 if required.
Level 5 Target-Based Verification	Specific to embedded software, especially where safety criteria require verification. A further RTM layer, tracing from source code to object code, may be generated where a component has been assessed as safety-critical.

Table 1: *From high-level requirements to target-based verification*

up-to-date tool technology, these tools are likely to contribute to the RTM by accident rather than design. With such a low profile and little support from tooling, it is no surprise that the RTM is poorly maintained over the duration of projects and typically completed as a rush job.

In truth, the RTM sits at the heart of any project (see **Figure 1**). Whether or not the links are physically recorded and managed, they still exist. For example, a developer will create a link simply by reading a design specification and using it to drive the implementation.

This alternative view of the development landscape immediately illustrates how pervasive the RTM is and the importance that should be attached to it. It is, therefore, vital that project managers consider investment in tooling for RTM construction with the same priority and enthusiasm attached to the purchase of requirements management, version control, change management, modelling and testing tools.

Furthermore, the RTM must be represented explicitly in any lifecycle model to emphasize its importance; **Figure 2** shows a way that might be done. With this elevated focus, the RTM will be constructed and maintained efficiently and accurately as an integral part of the development process, thereby avoiding any last-minute panic and associated costs.

For aerospace and military projects developed under the constraints of standards like DO-178B, such as air-traffic control or missile guidance systems, following a model where the RTM has high visibility is commonplace. Gap analyses show that some degree of automation to gather traceability information is often in place, principally across the early lifecycle phases where tool support is strongest. However automatic tracing to implementation and verification artefacts is typically weak, therefore more can be done even by companies who believe their processes and toolsets are generally sufficient.

As for companies in other sectors, the gap between their current working practices and those required to operate under safety-related standards is wide; of course, the upside is that the scope for cost savings is equally huge.

Avoiding Errors

Gap analysis highlights time and again that requirements traceability remains a low-priority, manual task prone to error and omission for many projects, even in the defence industry where there is a long history of rigorous process and standards adherence. Consequently, when a certification case is being assembled, the time and effort required to construct a requirements traceability matrix (RTM) is huge and a serious impact on the costs accrued by a project.

Clearly, the development of a RTM can be easily automated with investment commensurate with that made in other areas of the development lifecycle, such as modelling and version control. The return on investment has been reaped from tools that manage requirements, support design and enable verification; further cost savings and process improvements are also available to project managers in the realm of requirements traceability. ■

PLC with PIC16F648A Microcontroller

Part 16

```
      #include <definitions.inc>    ;basic PLC definitions, macros, etc.
      #include <cntct_mcr_def.inc>  ;Contact & Relay based macros
      #include <dmux_mcr_def.inc>   ;dmux based macros
;
;--------------- user program starts here ------------------------
      Dmux_1_2    I0.0,I0.1,Q0.1,Q0.0          ;rung 1

      ld          I0.2                         ;rung 2
      Dmux_1_2_E  I0.3,I0.4,Q0.7,Q0.6
;--------------- user program ends here --------------------------
```

Figure 1: *The user program UZAM_plc_8i8o_ex26.asm*

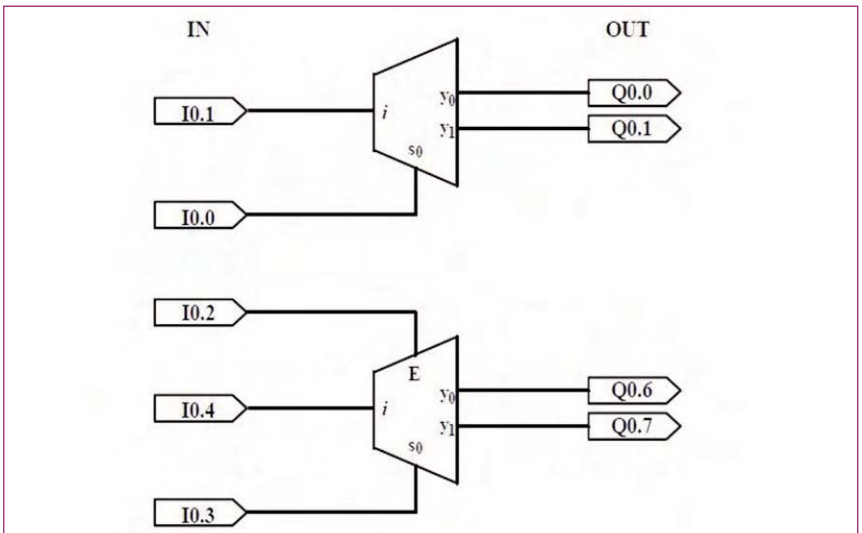

Figure 2: *Schematic diagram for the user program UZAM_plc_8i8o_ex26.asm*

```
      #include <definitions.inc>    ;basic PLC definitions, macros, etc.
      #include <cntct_mcr_def.inc>  ;Contact & Relay based macros
      #include <dmux_mcr_def.inc>   ;dmux based macros
;
;--------------- user program starts here ------------------------
      Dmux_1_4    I0.0,I0.1,I0.2,Q0.3,Q0.2,Q0.1,Q0.0      ;rung 1

      ld          I0.5                                    ;rung 2
      Dmux_1_4_E  I0.6,I0.7,T10,Q0.7,Q0.6,Q0.5,Q0.4
;--------------- user program ends here --------------------------
```

Figure 3: *The user program UZAM_plc_8i8o_ex27.asm*

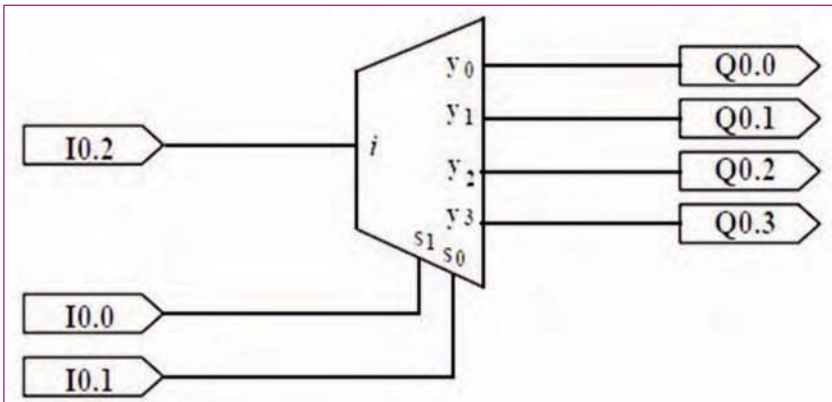

Figure 4: *Schematic diagram for the user program UZAM_plc_8i8o_ex27.asm*

Professor Dr Murat Uzam from Nigde University in Turkey presents a series of articles on a project that focuses on a microcontroller-based PLC. This is the sixteenth article providing four examples to show the use of demultiplexer macros

IN THIS SECTION we will consider four examples, namely UZAM_plc_8i8o_exN.asm, N = 26, 27, 28, 29, to show the usage of demultiplexer macros.

In order to test the respective examples, you can download the files from http://host.nigde.edu.tr/muzam/ and open the UZAM_plc_8i8o_exN.asm, N = 26, 27, 28, 29 program with MPLAB IDE to compile it.

Following that, by using the PIC programmer software, take the compiled file "UZAM_PLC_8i8o_exN.hex" and with a PIC-programmer hardware send it to the program memory of PIC16F648A microcontroller within the UZAM_PLC. After loading the "UZAM_PLC_8i8o_exN.hex", switch the 4PDT in "RUN" and the power switch in the "ON" position. This readies it for testing the example.

To check the accuracy of each program, you are referred to the related information for each demultiplexer macro provided in **Tables 1, 2...6** of the previous article (published in the January issue of *Electronics World* magazine). Note that, in some of these examples, we use the free-running reference timing signal T10 with the "T" timing period 524.288ms.

Examples

The first example program, "UZAM_plc_8i8o_ex26.asm", is shown in **Figure 1**. It shows the usage of two demultiplexer macros: "Dmux_1_2" and "Dmux_1_2_E". The schematic diagram of this user program is shown in **Figure 2**.

In the first rung, the demultiplexer macro "Dmux_1_2" (1×2 demultiplexer) is used. In this demultiplexer, the input signal is "i" = I0.1 and the select input is s_0 = I0.0, while the output lines are y_0 = Q0.0 and y_1 = Q0.1.

In the second rung, the macro "Dmux_1_2_E" (1×2 demultiplexer with active high enable input) is used. In this demultiplexer the input signal is "i" = I0.4 and the select input is s_0 = I0.3, while the output lines are y_0 = Q0.6 and y_1 = Q0.7. In addition, the active high enable input E is defined as E = I0.2.

The second example program, "UZAM_plc_8i8o_ex27.asm" is shown in **Figure 3**. It shows the usage of two demultiplexer macros "Dmux_1_4" and "Dmux_1_4_E". Its schematic diagram is depicted in **Figure 4**.

In the first rung, the demultiplexer macro "Dmux_1_4" (1×4 demultiplexer) is used. In this demultiplexer, the input signal is "i" = I0.2 and the select inputs s_0 and s_1 are I0.1 and I0.0 respectively, while the output lines are y_0 = Q0.0, y_1 = Q0.1, y_2 = Q0.2 and y_3 = Q0.3.

In the second rung, the macro "Dmux_1_4_E" (1×4 demultiplexer with active high enable input) is used. In this demultiplexer the input signal is "i" = T10 and the select inputs s_0 and s_1 are I0.7 and I0.6 respectively, while the output lines are y_0 = Q0.4, y_1 = Q0.5, y_2 = Q0.6 and y_3 = Q0.7. In addition, the active high enable input E is defined to be E = I0.5.

The third example program, "UZAM_plc_8i8o_ex28.asm" is shown in **Figure 5**. It shows the usage of the demultiplexer macro "Dmux_1_8". Its schematic diagram is shown in **Figure 6**. In this example, the demultiplexer macro "Dmux_1_8" (1×8 demultiplexer) is used. In this demultiplexers the input signal is "i" = T10 and the select inputs s_0, s_1 and s_2 are I0.2, I0.1 and I0.0 respectively, while the output lines are y_0 = Q0.0, y_1 = Q0.1, y_2 = Q0.2, y_3 = Q0.3, y_4 = Q0.4, y_5 = Q0.5, y_6 = Q0.6 and y_7 = Q0.7.

The forth and last example program, "UZAM_plc_8i8o_ex29.asm", is shown in **Figure 7**. It shows the usage of the demultiplexer macro "Dmux_1_8_E". The schematic diagram is shown in **Figure 8**.

In this example the demultiplexer macro "Dmux_1_8_E" (1×8 demultiplexer with active high enable input) is used. In this demultiplexer the input signal is "i" = T10 and the select inputs s_0, s_1 and s_2 are I0.2, I0.1 and I0.0 respectively, while the output lines are y_0 = Q0.0, y_1 = Q0.1, y_2 = Q0.2, y_3 = Q0.3, y_4 = Q0.4, y_5 = Q0.5, y_6 = Q0.6 and y_7 = Q0.7. In addition, the active high enable input E is defined to be E = I0.7. ■

```
.
      #include <definitions.inc>    ;basic PLC definitions, macros, etc.
      #include <cntct_mcr_def.inc>  ;Contact & Relay based macros
      #include <dmux_mcr_def.inc>   ;dmux based macros
.
;--------------- user program starts here ------------------------
                                                                ;rung 1
Dmux_1_8    I0.0,I0.1,I0.2,T10,Q0.7,Q0.6,Q0.5,Q0.4,Q0.3,Q0.2,Q0.1,Q0.0

;--------------- user program ends here --------------------------
```

Figure 5: *The user program UZAM_plc_8i8o_ex28.asm*

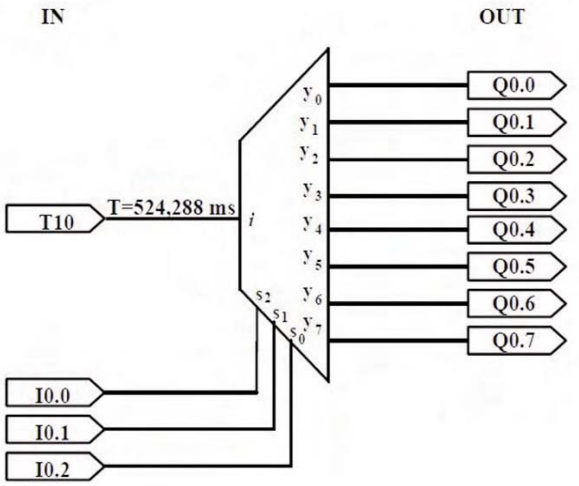

Figure 6: *Schematic diagram for the user program UZAM_plc_8i8o_ex28.asm*

```
.
      #include <definitions.inc>    ;basic PLC definitions, macros, etc.
      #include <cntct_mcr_def.inc>  ;Contact & Relay based macros
      #include <dmux_mcr_def.inc>   ;dmux based macros
.
;--------------- user program starts here ------------------------

ld          I0.7                                                ;rung 1
Dmux_1_8_E  I0.0,I0.1,I0.2,T10,Q0.7,Q0.6,Q0.5,Q0.4,Q0.3,Q0.2,Q0.1,Q0.0

;--------------- user program ends here --------------------------
.
```

Figure 7: *The user program UZAM_plc_8i8o_ex29.asm*

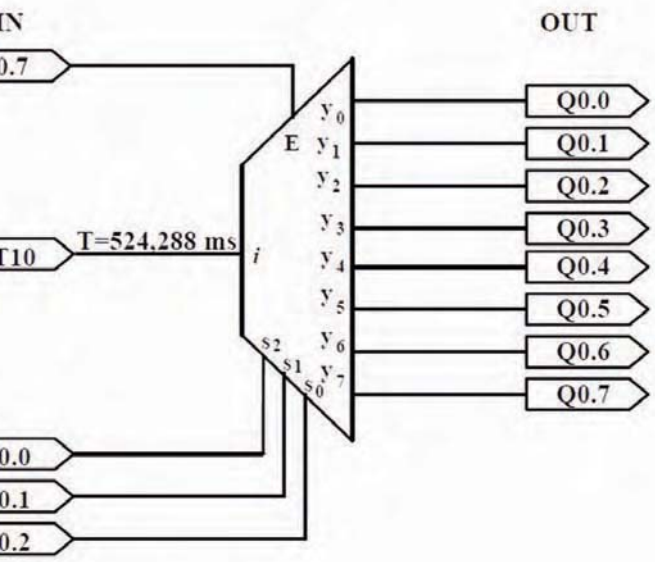

Figure 8: *Schematic diagram for the user program UZAM_plc_8i8o_ex29.asm*

ARE YOU MISSING A BACK COPY OF *ELECTRONICS WORLD* FROM 2006, 2007, 2008?

ELECTRONICS WORLD is pleased to announce its brand new on-line facility enabling you to purchase long out-of-print back issues of the articles in digital format from the years 2006, 2007 and 2008.

Each single copy is priced £4.20; it is fully searchable and is available as a downloadable link sent straight to your inbox with no additional postal charges!

UNIQUE FEATURES INCLUDE:

- **FULLY SEARCHABLE – you can search for particular articles in seconds**
- **DOWNLOADABLE – download specific pages or the whole magazine in PDF format**
- **EASY TO USE – fast page turning and page zoom functionality**
- **IMAGE CROPPING – why not cut out a page and share with friends and colleagues**
- **PAPER-FREE AND ENVIRONMENTALLY FRIENDLY – keep all your back copies anywhere on your PC for quick reference.**

Or why not sign up to receive all the latest issues of **ELECTRONICS WORLD** in digital format via an annual subscription and save ££'s on the printed version?

For 12 issue digital subscription you pay only £36.50!

SIGN UP NOW, VISIT
www.electronicsworld.co.uk

SUBSCRIBE TO *ELECTRONICS WORLD* AND SAVE 10%

.....AND receive a FREE one year subscription to the digital publication – delivered to your inbox (WORTH £36.50)

One year subscription for only **£40** (UK) and **£60** (Overseas)

- **FREE DELIVERY** to your door
- **SAVE 19%** on the Shop Price
- **NEVER MISS** an Issue
- Solutions & Case studies for the professional engineer
- Fresh ideas, tips and tricks
- Analysis of what's coming next in technology
- Features on a wide range of markets from automotive to medical
- Company profiles
- International coverage around the globe

SUBSCRIBE ONLINE:

www.electronicsworld.co.uk

CALL OUR HOTLINE:

UK: 0870 428 1426

International: +44 1635 879361

IR TRANSCEIVER SYSTEM DESIGN WITH ASK MODULATION

IR (INFRA RED) communication is a simple and useful method in many applications such as security control, remote control of electrical and electronics appliances, door and window controls, factory automation systems, smart home design networks and so on.

IR transmitter and receiver circuits use the IR waves in the range of 30-40kHz. There are IR diodes (transmitter and receiver) that work in this band on the market already. In literature, many IR application circuits exist and they can be classified as continuous current IR transmitter and receiver, and pulse or coded-IR transmitter and receiver.

In continuous-current IR transceiver systems, the transmitter IR LED works with a DC current and the communications distance is short; typically under one meter. Coded-pulse IR systems work with square waves in the range of 30-40kHz. These systems send triggering pulses.

The operation of an IR receiver LED is explained as follows: Normally, IR-receiver LED output is at high level (logic 1). When an IR wave is incoming (into the IR receiver LED), the output of that LED changes from logic 1 to logic 0 for a few seconds, and then goes back to logic 1. In other words, the IR receiver LED only detects whether the IR wave is incoming or not. This means that if IR waves continue to be sent, the output of the receiver LED will be at logic 1.

Therefore, the IR receiver LED generates a trigger pulse only in response to an IR wave. In this situation, designers should use a two-button operation (ON/OFF) for the load. One button sends the coded signal for ON, the other button sends the different code signal for OFF.

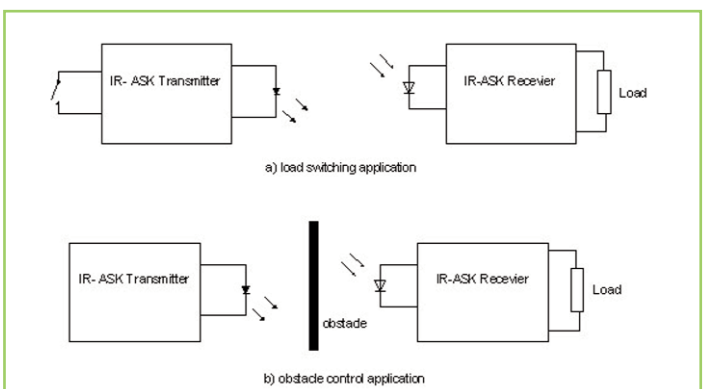

Figure 1: *IR-ASK application systems*

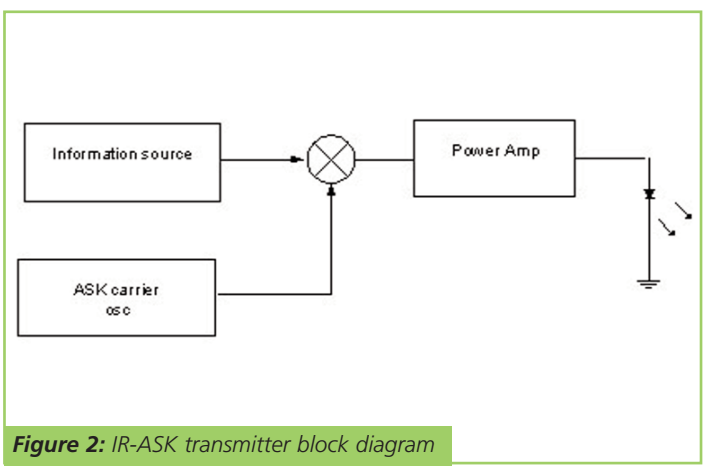

Figure 2: *IR-ASK transmitter block diagram*

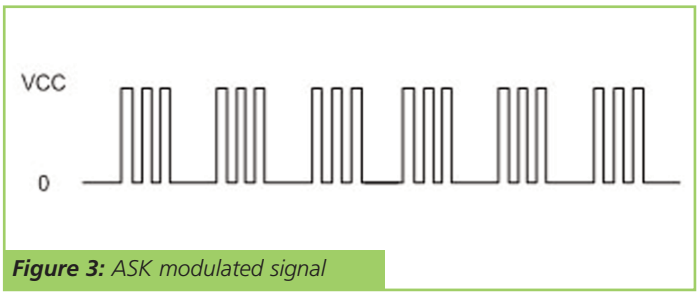

Figure 3: *ASK modulated signal*

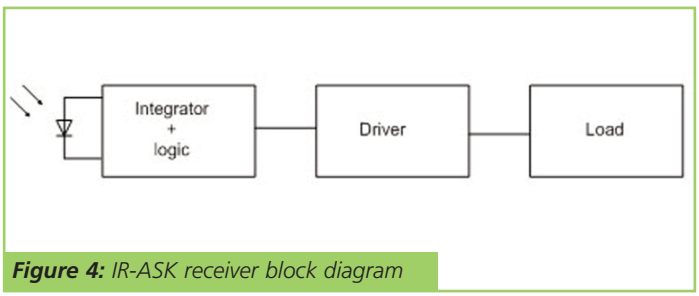

Figure 4: *IR-ASK receiver block diagram*

On the other hand, in the obstacle type of control systems it is important to know if there is an obstacle between the transmitter and receiver. The function of the IR receiver LED described above is not suitable for this.

Figure 1 shows the one-button control and obstacle control mechanism. In this design, ASK (Amplitude Shift Keying) modulation is proposed to increase the communication distance. Communication distance is an important parameter for control systems, such as door control, security, home control applications, remote control and others. The principle of ASK modulation is that the information signal and carrier signal are multiplied.

ASK modulation can be applied to IR communication, as shown in **Figure 2**. The function of the IR receiver LED is as described before, where if a carrier signal is sent in the form of OFF/ON, the receiver LED will generate a continuous logic 1 or 0, according to the IR signal.

In Figure 2, the ASK carrier oscillator works at 38kHz and the information signal is approximately 1Hz. **Figure 3** shows the ASK modulated signal in the time domain. Both sources are square waves. ASK modulator may be realized with an AND gate, and the power amplifier may be used to increase the communication distance.

The block diagram of an IR-ASK receiver is shown in **Figure 4**. IR-ASK modulated waves occur as square waves at the output of the receiver LED. The integrator takes the mean value of this square wave signal. This mean value level is compared to a DC level of the comparator. After the comparator stage, logic 1 or 0 is obtained depending on whether the IR wave has been sent out or not, respectively. These logic signals are sent to the driver circuit.

Transmitter and Receiver Circuits Design

An IR-ASK transmitter circuit is shown in **Figure 5**. A 555 astable multivibrator may be used for both information and carrier sources of

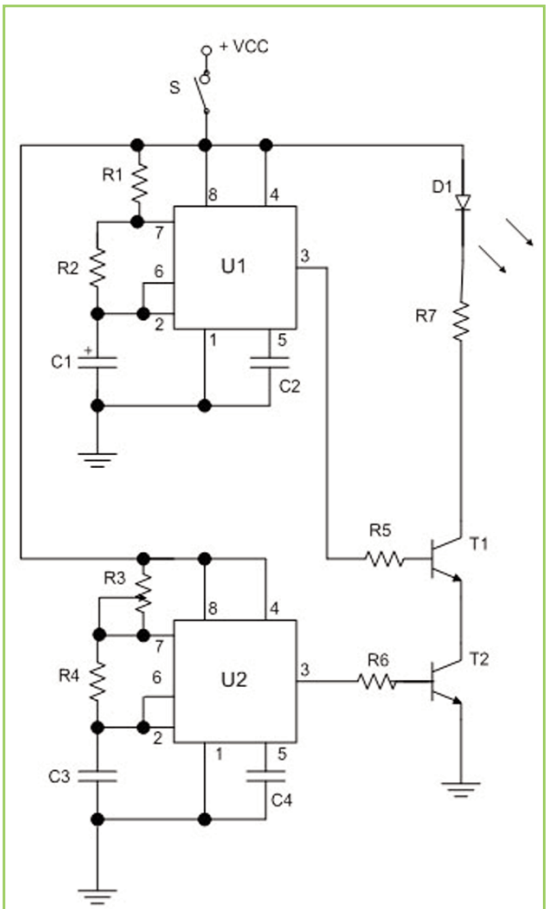

Figure 5: *ASK-IR transmitter*

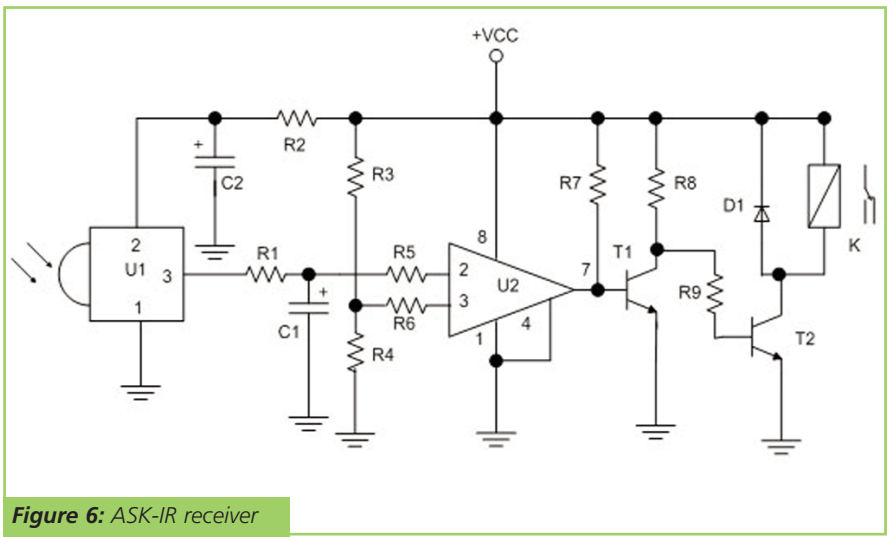

Figure 6: *ASK-IR receiver*

the ASK modulator. U1 is the information source oscillator, at 1Hz, and the U2 is the carrier oscillator, 38kHz. T1 and T2 transistors are used as the ASK modulator. In fact, T1 and T2 work as an AND gate. Because the basic principle of ASK modulation is the switching of the carrier signal, the AND gate is a suitable idea for the ASK modulation. At the same time, T1 and T2 are used as a power amplifier here too.

The part list for the circuit in Figure 5:

U1 = U2 = LM555	T1 = T2 = 2N2222
D1 = IR LED	R1 = 100K
R2 = 180K	R3 = 1K
R4 = 1K5	R5 = 6K8
R6 = 6K8	R7 = 220 Ω
C1 = 4.7 μ F	C2 =10nF
C3 = 10nF	C4 = 10nF
VCC = 5-15V	

The IR-ASK receiver is shown in **Figure 6**. The switch S is used for a one-button control application and VCC is directly applied to the transmitter circuit in the obstacle-type control systems. R1 and C1 are components of the integrator. U2 is the comparator, T1 is the inverter and T2 is the driver transistor. If an ASK modulated-IR signal appears at U1, pin 3 of U1 has a square wave of 1Hz. The integrator R1-C1 takes the integral or mean value of this square wave.

Pin 3 of U2 has a reference voltage, bigger than the mean value of pin 3 but smaller than VCC. Therefore, an IR wave exists at U1, which means the comparator U2's output is logic 0, so the relay K is active and the load is ON. Otherwise, if the IR signal is cut off by the switch S, or due to an obstacle, pin 3 goes to logic 1, and the U2 output goes low, to logic 0.

The part list of the circuit in Figure 6:

U1 = TK19 or SFH506	U2 = LM311
T1 = T2 = 2N2222	D1 = 1N4001
R1 = 1K	R2 = 47 Ω
R3 = 10K	R4 = 10K
R5 = 4K7	R6 = 4K7
R7 = 6K8	R8 = 3K3
R9 = 3K3	C1 = 10 μ F
C2 = 47 μ F	VCC = 5-15V

VCC may be selected between 5-15V in Figures 5 and 6. The values of R5, R6 and R7 in the transmitter circuit, and R7, R8 and R9 in the receiver circuit depend on the range of the VCC value. Therefore, these components should be calculated according to the selected VCC value.

Also, the relay coil voltage and the VCC voltage should be the same. In this design, the relay is preferred as a switching driver component, which can be changed with a suitable triac. ■

I. Hakki Cavdar
Karadeniz Technical University
Electrical and Electronics Engineering
Turkey

READER OFFER WINNER!

The winner from the Microchip PIC32 USB Starter Board competition in April was Martyn Cowell of Made Electronics in Darwen, Lancashire.

ILLUMINATED ROTARY ENCODERS OFFER NEW OPTIONS FOR PRO-AUDIO AND INDUSTRIAL CONTROL

Foremost Electronics announces the availability of the MERP12 range of rotary encoders. This new series of encoders has been designed to replace traditional rotary potentiometers in a wide range of applications and have the added benefit of being illuminated. The new MERP12 range will interest designers currently working on broadcast and professional audio equipment, in addition to next generation building automation systems and test and measurement instruments. Rotary encoders offer significant operational benefits when controlling "soft" menu functions.

MERP 12 encoders feature an LED illuminated actuator which can be single or multi-coloured and can be specified with or without a pushbutton action. MERP12 encoder switch contacts are rated at 5VDC @ 10mA, have a typical lifetime of 30,000 operations, an operating temperature range of -40 to +85C and may be single hole or PCB mounted.

www.4most.co.uk

LCEDI CONNECTORS FOR THE PC MARKET FROM TYCO

Tyco Electronics offers a next generation LCD Coaxial Embedded Display Interface (LCEDI) family of connectors designed to provide exceptional electrical performance in both low-voltage differential signalling (LVDS) and embedded DisplayPort (eDP) applications. This family of connectors is licensed by I-PEX CO Ltd and is fully compatible and intermateable with I-PEX CABLINE-VS connector series, recently selected by VESA (Video Electronics Standard Association) as the global standards connector for LED backlight wide (16x9) panel interface. Its ultra-low profile mating configuration (1.1mm height) makes it ideal for the slim LED backlight LCD panel of advanced notebook personal computers.

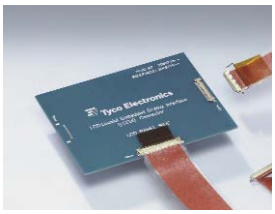

Tyco Electronics's LCEDI connector family accommodates consistent digital data transmission through one, two or four DisplayPort standard lanes at a reduced bit rate of 1.62Gbps or a high bit rate of 2.7Gbps through each lane, and even faster data rates over different wiring schemes.

www.tycoelectronics.com

NEW NEMESIS SC CONNECTORS FROM ITT INTERCONNECT SOLUTIONS CAN BE SEALED AND CLEANED

ITT Interconnect Solutions's new Nemesis SC connectors are sealed but also cleanable, enabling users such as soldiers and medical staff to operate in all conditions. The new, miniature, lightweight connectors incorporate a patented pull back barrel system which facilitates cleaning and allows the user to return to full operational functionality in seconds.

Nemesis SC MIL-DTL-38999-style connectors are blind-mateable and feature ITT Cannon's innovative Pogo Pin/Pad and Breakaway technology, which increases the number of mating cycles to over 10,000. Devices can be cleaned more than 2,500 times. Contacts are rated at 3A.

The first products are based on the key elements of the MIL-DTL-38999 specification; but, a full range of circular and rectangular styles are also planned. Currently available devices can have 7, 14 or 19 pins, feature five locking positions and are colour-coded for simple identification.

www.ittcannon.com

TESTING TIMES FOR POWER CONSUMPTION

Hioki's 'Power HiTESTER' 3332, 3333 and 3334 are ready to meet the new 'Energy' test requirement, measuring electrical power consumption under normal running conditions and in standby mode.

With a greater emphasis on energy conservations electrical equipment manufacturers not only have to monitor the power consumed by their products while in operation but also when in standby mode.

Manufacturers must now comply with new mandatory standards and will also benefit from the growing demand for energy saving products.

The EuP Directive 2005/32/EC covers the eco-design requirements of 'Energy-using Products' (EuP). Regulations controlling the eco-design of some domestic electrical products are already in place, additional product categories will follow, therefore manufacturers need to test and then maybe redesign their products in order to meet these legal requirements.

www.gmciuk.com

KONTRON INTRODUCES ITS FIRST INTEL CORE I7 PROCESSOR-BASED VPX BLADE

Kontron has introduced the Kontron VX6060, an innovative VPX computing blade for parallel data and signal processing applications.

With two independently implemented Intel Core i7 processing nodes linked to a powerful Ethernet and PCIe infrastructure, the Kontron VPX blade VX6060 is the ideal building block for intensive parallel computing workloads where a cluster of Kontron VX6060s can be used in full mesh VPX or switched OpenVPX environments.

Each processing node implements Intel's next generation high performance embedded processor with integrated memory controller and Intel HD graphics – the Intel Core i7 processor – coupled with the highly integrated Intel Platform Controller Hub (PCH) QM57 with numerous Gigabit Ethernet, SATA, USB 2.0 and PCIe channels.

Target applications include radar, sonar, imaging systems, airborne fighters and UAV radar which will use clusters of the Kontron VPX blade VX6060.

www.kontron.com

AVX EXPANDS LED LIGHTING CONNECTOR FAMILY

AVX Corporation has expanded its wire-to-board connector product offering to include board-to-board connectors specifically designed for the LED lighting industry. Designated the 9159 Series, the cost-effective connectors provide design flexibility by offering both card edge (one-piece) and plug-and-socket (two-piece) interconnect options. The connectors feature a small footprint, making them ideal for applications where multiple printed circuit boards need to be plugged together, such as LED lighting strips.

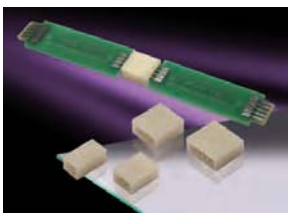

The double-ended card edge version provides a simple and direct connection to both ends of a standard PCB with tin plated pads in 2, 3, 4 and 5 positions. These 2.0mm pitch connectors support 3A current and 250V ratings. The connectors are 5.0mm high and come in both black and white options.

The two-piece connector version is surface mounted on one side of the PCB, which allows the LEDs to be placed on the other side to maintain consistent spacing.

www.avx.com

KONTRON EMBEDDED BOARDS WITH DISPLAYPORT

The new Kontron CP308 3U CompactPCI multicore board with the CP308-MEDIA extension card is one of the first embedded products to feature the new HD digital display interface standard DisplayPort. With S/P-DIF-Out audio and the stereo audio ports for Line In, Line Out and Microphone, the processor board with the innovative Kontron CP308-MEDIA adds extensive multimedia capabilities to embedded computing.

The Kontron CP308 features the high performance 45nm Intel Core 2 Duo processor up to 2.26GHz, a powerful embedded Intel GS45 Graphics and Memory Controller Hub, up to 8GByte of energy-efficient DDR3 RAM and the Intel I/O Controller Hub ICH9M. It is the perfect solution for applications such as digital signage, passenger information and entertainment, process and quality control, surveillance and security in the transportation, aerospace, military and industrial automation markets.

www.kontron.com

TYCO LAUNCHES HYBRID CONNECTORS FOR REAL TIME OPERATION

Tyco Electronics's new real-time hybrid connectors with 8+4 power and signal contacts have been designed for machine automation applications that require a higher performance of up to 10A and a separate power supply. The Power4Net connectors provide perfect connection technology and meet the demanding durability and quality requirements of real-time Ethernet applications in industrial automation.

Both the 8 power contacts and 4 signal contacts of the fast Ethernet interface (and shielding) are robustly built and both the PCB mount and the crimp versions feature multiple contact points. Optimized streamlining of the contacts ensures permanent shock and swing-proof interconnection, providing the foundation for real-time capabilities. An isolation crimp ensures optimal seating and additional strain relief for the power contacts.

Hybrid cabling for the Power4Net connector components is readily available and there are also versions for use with cable carriers.

www.tycoelectronics.com

LOW COST INDUSTRIAL ETHERNET STARTER KITS FROM HARTING

Harting has introduced low-cost, entry-level, industrial Ethernet switch starter kits and patch cables. The Basic and Basic+ Ethernet kits and the new Cat 6 patch cables will be available at substantial discounts until 26th March 2010.

The Harting Basic Ethernet switch kit contains everything needed for a simple industrial Ethernet implementation, including a five-port unmanaged switch, ten "tool-less assembly" RJ45 connectors and 20 meters of Cat 5 Type B Ethernet cable. In addition, the Basic+ kit contains a 24VDC/1.4A power supply and a Harting cable stripping tool.

The RJ45 connectors and the new range of Cat 6 patch cables are part of Harting's Automation IT generic cabling system. The RJ45 has a unique two-part boot design that offers robust handling and bending protection.

They are designed to be robust enough for industrial environments and are heat and flame retardant.

www.Harting.com

CHERRY SENSORS AND SUBMINIATURE SWITCHES FROM ZF ELECTRONICS IMPROVE WATER SOFTENER CONTROLS

Cherry subminiature switches and sensors from ZF Electronics are used to improve salt usage and regeneration in water softeners more effectively than competitive products due to two factors: the switches can be supplied in a space-saving module package and the sensors offer greater dimensional consistency due to a special low-pressure mould process used in their manufacture.

The composition of water varies from area to area, according to the local geological rock type. Water softeners are one of several ion exchange devices used to improve water quality by removing particles and mineral deposits present naturally in many water sources. Their most common use is the removal of calcium and magnesium ions present in "hard" water which cause scale build-up in plumbing, staining of fixtures such as baths and taps, and over the longer term, damage to appliances such as dishwashers and washing machines.

www.cherry.co.uk

OPTICAL SPECTRUM ANALYSER FOR MEASUREMENTS ON LEDS AND LASER LIGHT SOURCES

The new Yokogawa AQ6373 is an optical spectrum analyser designed to carry out measurements over the wavelength range from 350 to 1200nm, including the visible light spectrum from 380 to 780nm.

In addition to its world-class optical performance in areas such as resolution, accuracy, sensitivity, measurement speed and dynamic range, the AQ6373 features a colour analysis function that makes it ideally suited to measurements on LEDs and laser light sources.

Key performance parameters include a wavelength accuracy of ±0.05nm, a wavelength resolution down to 0.01nm and a sensitivity of -80dBm, switchable to a choice of high dynamic-range modes.

High-speed measurements result from a standard sweep time down to one second (0.5s in automatic mode), while the free-space optical input makes the instrument applicable to use with single-mode, multimode and large-diameter core fibres. A built-in light source is also provided to aid optical alignment.

www.yokogawa.com

LECROY EXPANDS WAVEACE OSCILLOSCOPES RANGE

LeCroy has expanded its low-end WaveAce Oscilloscope Series to include 4 channel and 40MHz models for simple and efficient debug, in the price range from £600/EUR695.

The 4-channel models provide 10kpts/ch memory and up to 2GS/s sample rate; the 40MHz model provides 4kpts/ch and a sample rate of up to 500MS/s. All models offer long memory, colour displays, extensive measurement capabilities and advanced triggering to improve troubleshooting and shorten debug time.

With USB host and device ports, plus a LAN connection, the WaveAce oscilloscopes easily connect to a memory stick, PC or printer for saving data or remote control. Combined with the streamlined, time-saving user interface, these features make the WaveAce oscilloscopes the ideal tools for affordable design, debug and troubleshooting from 40MHz to 300MHz.

The high performance and large feature set of the WaveAce is controlled by an intuitive user interface with 11 different languages and streamlined front panel.

www.lecroy.com

NEW PIC32MX4 MULTIMEDIA BOARD FROM MIKROELEKTRONIKA

mikroElektronika has introduced a new PIC32MX4 MultiMedia Board as an addition to its PIC32 development tool product line.

The PIC32MX4 MultiMedia Board is a compact development system for advanced user interface applications development and test. The tool provides a complete, high-quality, multimedia development platform supporting the PIC32MX4 devices.

This board is also compatible with the latest PIC32MX5/6/7 series. With the built-in electronic circuits including TFT colour display 320x240 with touch screen, stereo codec, accelerometer, SerialFlash, SerialROM and many others, the PIC32MX4 MultiMedia Board has many multimedia features for users to develop advanced graphical, audio, memory and storage applications using the most popular components in consumer electronic products.

For more information on the PIC32MX4 MultiMedia Board, visit the mikroElektronika website at www.mikroe.com

mikroElektronika
SOFTWARE AND HARDWARE SOLUTIONS FOR EMBEDDED WORLD

www.mikroe.com

TRANSFORMER MANUFACTURE

WEBSITE DIRECTORY

SCHNEIDER ELECTRIC USES BIOMETRIC TECHNOLOGY FOR ACCESS CONTROL

Schneider Electric has launched a complete standalone biometric switch dedicated to access control for industrial process or building applications, reducing the risks and costs linked with operational errors and increasing levels of security within a machine, process and working areas.

The Harmony Biometric switch has been designed to prevent unauthorised persons from accessing certain areas or controls that may cause them or the running process a danger. A unique pattern is embedded in the device so all control operations can be managed through its intuitive touchscreen menu, and the two levels of rights for user and administrator brings easiness during add/delete operations. The switch can be placed in a standard 22mm cut-out.

So now, instead of a traditional key or passwords which are often forgotten, the switch can be operated by a user with their thumb or fingerprint. Up to 200 unique prints can be stored within the compact unit.

Our panel of commentators says the following on this development:

HAFIDH MECHERGUI, ASSOCIATE PROFESSOR IN THE FIELD OF ELECTRICAL ENGINEERING AND INSTRUMENTATION AT THE UNIVERSITY OF TUNIS, TUNISIA:

Biometrics, known and exploited since antiquity, consist of extracting or calculating the physical or behavioral parameters, specific to each individual, with an aim of being able to identify a person in a reliable way. This technique was monopolized especially by the police force. They used it by the intermediate fingerprints with an aim of identifying an individual and this since 1908, following work of Berlillon, who in 1890 developed the first methods of structured analyses.

Nowadays the spectacular and considerable evolution of technological sciences, especially in the field of electronics and data processing, facilitate the opportunity to develop biometrics. The need to identify people becomes increasingly important, taking into account the terrorist threats, but also to carry out various current operations such as access control or safe monetary payments. In schools it can be applied for automated attendance and registration, school library automation and data protection.

Biometrics is expected to also play a key role in personal authentication in large scale enterprise network environments and digital rights management, health care applications, encryption keys digital signature and other.

The technological innovations make it possible to exploit different physical characteristics more and more in order to identify a person. The different types of biometrics are fingerprint, face recognition, speaker recognition, iris recognition, hand and finger geometry, signature verification and even the proteinic sequences carried by DNA. All these biometric types are automated methods of identifying a person, or verifying the identity of a person based on a physiological or behavioural characteristic.

Indeed, the program launched by Schneider Electric is an exciting challenge but we need to achieve the utilization of strong personal authentication procedures too.

MAURIZIO DI PAOLO EMILIO, TELECOMMUNICATIONS ENGINEER, INFN – LABORATORI NAZIONALI DEL GRAN SASSO, ITALY:

The continuous development in the biometric technology speeds up its popularity in various uses. No doubt soon it will be found in many new applications too.

In the industrial field, where high-level safety is required, by using biometric authentication it makes it safe and practical.

Any biometric access control system consists of biometric access control reader or scanner. But, somehow, this kind of system is not perfect. The fingerprint of a person could change over time with age or through damage, and the scanner could easily become "polluted", which would tamper with the process. The next step would be to overcome these obstacles.

IVOR CATT, ELECTRONICS ENGINEER, UK:

This is a key technology advance into the future, but perhaps at present it has a prohibitive cost. At some time in the future we can expect such a key technology to be in general use.

There was a prior hi-tech facility which became practical ten or twenty years ago and remained unexploited, except for use on my own house. That was the use of a single push-button where the user would tap in his private pattern of numbers in the manner of Morse Code. The single push-button could have been concealed under the front door carpet – or as the doorbell. It would be battery operated.

A decade ago its retail cost would have come down to a little more than fifty pounds, obviating the expensive use of keys or access cards. Time is running out for someone to "think of it" and make a million before fingerprint recognition becomes economic.

BARRY MCKEOWN, RF AND MICROWAVE ENGINEER IN THE DEFENCE INDUSTRY, AND DIRECTOR OF DATOD LTD, UK:

Whereas there is a plethora of fingerprint readers and software on the market for cars, PCs and smartphones, this appears to be the first with both integrated embedded control and image processing data bank. Provided the operating environment is clean and benign, then this biometric switch should prove superior to other means of positive identification for access control, although a difficulty shall arise in process control applications where operator cuts and abrasions are common or dirty, and especially oily hands require cleansing before operating the switch in an emergency, especially for switching items off.

PROFESSOR DR DOGAN IBRAHIM FROM THE NEAR EAST UNIVERSITY IN NICOSIA, CYPRUS:

Security and especially secure and reliable access control to industrial processes and working environments are very important in preventing unauthorized access. Scheneider Harmonic Biometric Switch is based on the well-known and proven fingerprint technology. The user does not have to remember passwords or carry keys and can have access by simple touch to the switch. The unique advantage of the device is that it can store up to 200 unique fingerprints, thus providing access control in large industrial plants and sensitive working areas.

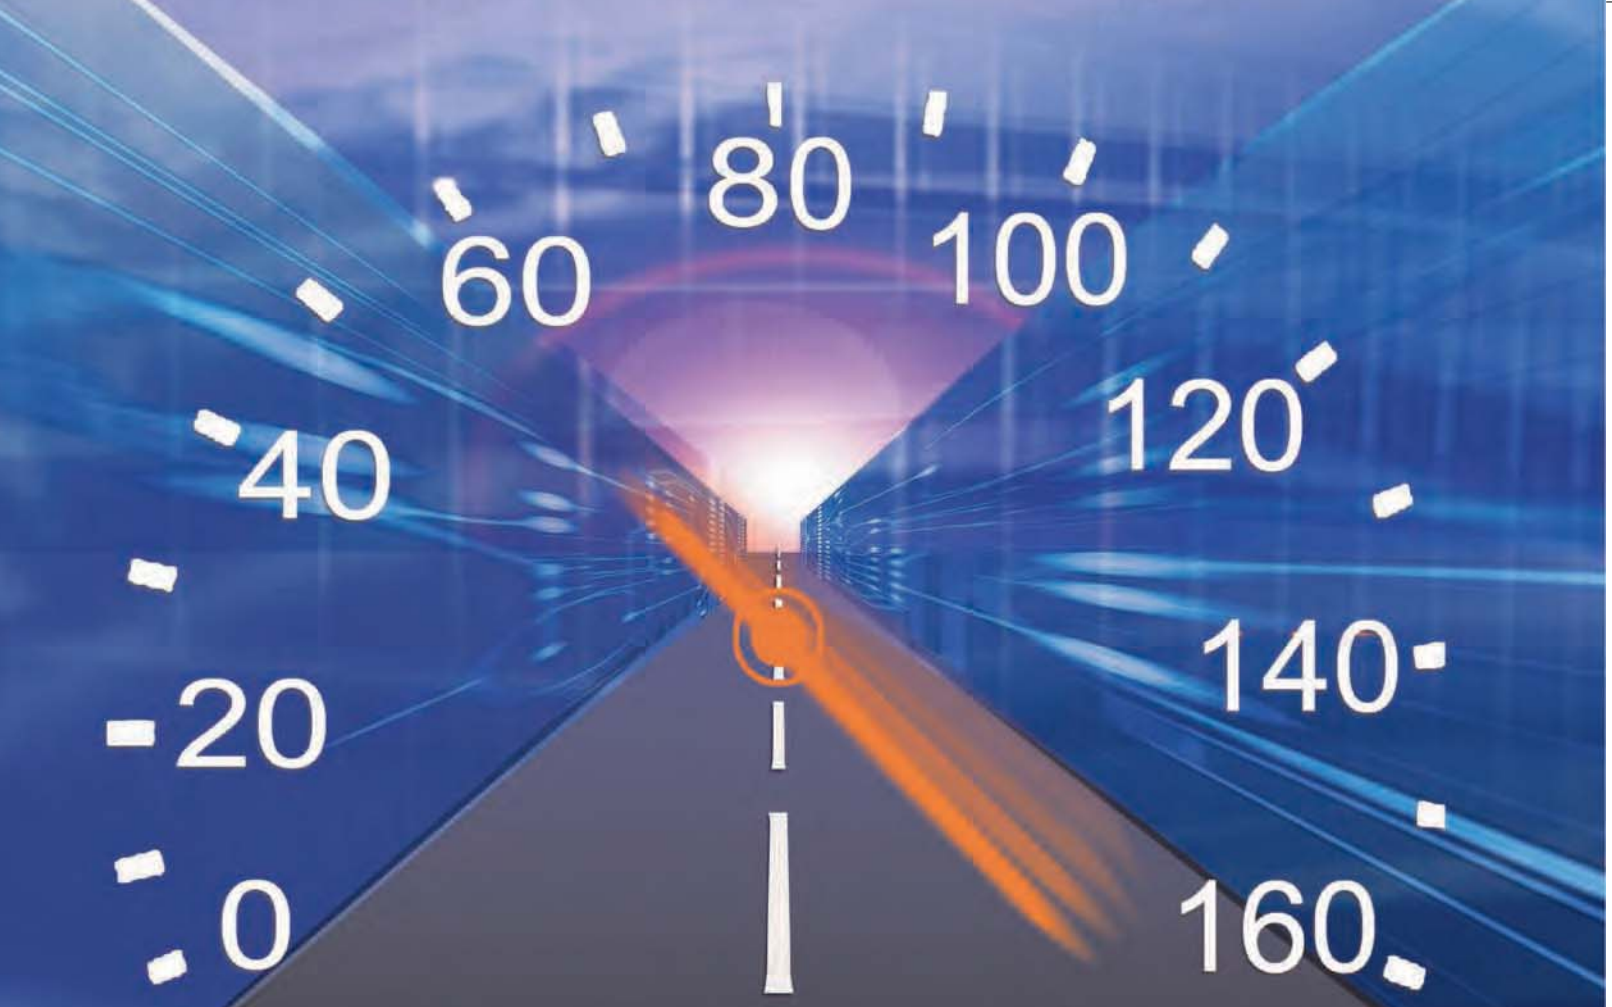

GO FASTER WITH PROTEUS PCB DESIGN

The latest version of the Proteus Design Suite harnesses the power of your computer's graphics card to provide lightning fast performance. Together with unique transparency options it's now easier than ever to navigate and understand large, multi-layer boards.

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++ & 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.

All levels of the Proteus Design Suite include a world class, fully integrated shape-based autorouter 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

ENABLING PERFORMANCE BASED OPERATIONS

avionics europe
Incorporating Defence Electronics

Conference and Exhibition
24th–25th March 2010
Passenger Terminal (City Centre)
Amsterdam, The Netherlands
www.avionics-event.com

PLEASE USE THIS PROMOTIONAL CODE WHEN REGISTERING: EW-0210

www.avionics-event.com

Learn and discover more at the leading conference and exhibition for global avionics and defence electronics

International competition has intensified during a challenging economic atmosphere as companies and governments look to keep pace with next-generation avionics, air traffic management capabilities and defence electronics developments.

Avionics conference track includes:

- Avionics With ATM
- Performance Based Operations
- Regulatory Frameworks for Avionics
- Next Generation Military Avionics

Defence Electronics conference track includes:

- COTS Integration Challenge
- Thermal Management
- Trusted COTS
- Software Defined Radio In Defence Applications

Exhibition Times

Wednesday 24th March: 10:00am – 6:00pm
Thursday 25th March: 10:00am – 4:30pm

Opening Keynote

Open and FREE of charge to all
Wednesday 24th March: 9:00am – 10:00am

Networking Reception

Wednesday 24th March: 5:30pm – 7:30pm

Exhibitor Presentations

Open and FREE of charge to all throughout the exhibition opening hours

For full conference programme or to register online visit www.avionics-event.com

Co-located with: Near Space Security
Flagship Media Sponsors: avionics Intelligence, Military & Aerospace
Owned & Produced by: PennWell
Sponsors: WIND RIVER, LYNUXWORKS
Supporting Organisation: n-C, ADS, NIDV Defensie & Veiligheid, AIRBUS, IET